

A constraint-based approach to phonology and morphology

Kevin Russell

This is a somewhat revised version of my dissertation, defended in April 1993, and presented to the Graduate School of the University of Southern California in June 1993. Chapter 7 is new, as are a few paragraphs throughout the rest of the dissertation, but otherwise changes from the version on deposit with the USC library are largely limited to correction of typographical errors and repagination. In an ideal world, all future references to page numbers for this dissertation would follow the pagination of this version.

Committee: Jean-Roger Vergnaud (advisor), Bernard Comrie, Alicja Gorecka, Mark Seidenberg, Bernard Tranel.

Acknowledgements

I love acknowledgement sections. Whenever I get a new book or thesis, they're the first thing I turn to. It might be that they're the only indication that the author I'm about to read is really a human being after all, or it might be the thought that if even such a famous person needed all this help too, there may be hope for me yet. Since this is my first chance to write a real acknowledgements section, I hope the reader will indulge me.

I can't imagine having written this dissertation anywhere other than USC, with any committee other than this one, and with any advisor other than Jean-Roger Vergnaud. From my first day in his 531a class, he somehow turned phonology from the boring grunt work you had to go through before you got to interesting stuff like semantics and into something I would seriously consider spending the rest of my life doing. Ever since then in my work with him, there has been no idea too sacrosanct to be questioned and no idea too heretical to be considered seriously. Without his constant support, enthusiasm, and bottomless cups of cappuccino, this dissertation would not have been possible.

For their help (and gentle insistence) in bringing stratospheric ideas to their potential (and down onto paper), I am indebted to Alicja Gorecka, Bernard Tranel and Bernard Comrie. With unfailing patience and good humor, they have pointed out counterexamples to my sweeping generalizations, brought to light implications I would never have been aware of, suggested directions for further study, and come up with just the right reference at just the right moment. They have also been instrumental in transforming successive drafts and fragments from their original state into a reasonable facsimile of English. I must thank my external member, Mark Seidenberg, for never letting me forget the wider question of language research and for letting me be a hanger-on in the most stimulating lab environment I know of.

It was well into my third year before I finally and firmly settled on phonology as a specialization. My excellent teachers at USC deserve credit for making it a difficult choice: Elaine Anderson, Joseph Aoun, Doug Biber, Jim Gee, Osvaldo Jaeggli, Elinor Ochs, Maria Polinsky, Barry Schein, and (pulling for the phonology side) Debbie Schlindwein. In other departments, Barry Glassner, Peggy Kamuf, Marsha Kinder, and Dallas Willard kept me realizing that there was more to be learned than feature hierarchies.

For contributions to my growth as a linguist, and for keeping me at exactly the right level of sanity, I must thank my colleagues in the Linguistics Department—Nancy Antrim, Alfredo Arnaiz, Dwight Atkinson, Elabbas Benmamoun, Jose Camacho, Nigel Duffield, Gorka Elordieta, Connie Gergen, Elena Herburger, Kaoru Horie, Matt Hunt, Sue Kalt, Ke Zou, Stephen Matthews, Carla Ponti, Vai Ramanathan-Abbott, Charlotte Reinholtz, Suchitra Sadanandan, Liliana Sanchez, Patricia Schneider-Zioga, and Linda Taylor—and in other departments—David Corina (honorary grad student), Kim Daugherty, Joe Devlin, Dana Murphy, Alan Petersen, and Martine van der Vlugt. Es-

pecially deserving of mention are Heather Goad (for support when being a phonologist at USC was a lonely business and long conversations that kept me from losing my accent) and Robin Belvin (Nisgha syntactician, surfer, pool player, and auto mechanic extraordinaire). The department staff, Laura Reiter, Vivian Smith, Kathy Stubaus, and Linda Williams-Culver, cheerfully suffered through even my most outrageous requests for help and almost made it seem as if the USC bureaucracy didn't exist.

There were many who made southern California, as well as Southern California, such an exciting place to be: at UCLA, Donca Steriade, Pam Munro, and the whole American Indian Linguistics Seminar bunch; Carol Genetti at UC Santa Barbara; at UC Irvine, Terri Griffith, Jim Huang, and many others (special thanks to Bernard Tranel for braving the 405 to come to my defence(s)).

I would never have met any of these people if it weren't for my undergraduate teachers at the University of Manitoba: Dick Carter, John Haiman, Lorna Macdonald, David Pentland, and especially Chris Wolfart, whose efforts to turn me into a linguist have been superhuman. For five years, the U of M has been my home away from home, with stimulating conversations, technical support, and the occasional cheque. Their direct contributions to this work should not be overlooked: their offer of a job in March allowed me to complete and defend this dissertation in a state as near to serenity as it's possible for a grad student to be.

From Steven Bird, Jim Scobbie, and Shelly Waksler, support and advice in a barely explored area of phonology was only an e-mail message away. Steven read and commented helpfully on earlier incarnations and drafts and of this dissertation. For putting up with strange questions about Moroccan Arabic at strange times, I am indebted to Abdesalam Elomari, Laila Lalami, and especially Elabbas Benmamoun. For their hospitality in New Aiyansh, B.C., and their near futile attempts to get me to pronounce Nisgha correctly, I thank Bertha Azak and Sam and Sarah Haizimsque, and for their permission I am grateful to the Nisgha Tribal Council.

I am grateful to the LSA for a fellowship to the 1987 summer institute at Stanford. A doctoral fellowship from the Graduate School at USC and a doctoral fellowship from the Social Sciences and Humanities Research Council of Canada provided me with the time to write this dissertation and have allowed me to finish graduate school with a nearly positive net worth.

Thanks also to Donna Cunningham, Carlyle Hoffman, the Minielys, Arden Ogg, Danielle West, and everyone at United University Church.

This dissertation is dedicated to my parents, without whose constant love, support, and encouragement I would not now be writing these woefully inadequate words.

*June 1993
Venice, California*

Table of contents

Chapter 1. Introduction	1
1.1 Constraints	4
1.1.1 Declarativity	6
1.1.2 Monostratality	7
1.1.3 Monotonicity	8
1.2 Government in phonology	11
1.3 Empty positions	14
1.3.1 Pronounced empty positions	14
1.3.2 Unpronounced empty positions	15
1.4 Morphemes as constraints	16
1.5 Formalism	19
1.5.1 Avoiding ambiguity	20
1.5.2 The representation/description distinction	27
1.6 About this dissertation	28
 Chapter 2. Representing constriction gestures	 33
2.1 Gorecka's constriction model	34
2.2 Dominance models of secondary articulation	40
2.3 Degrees of constriction	43
2.3.1 Steriade's closure and release features	45
2.4 A synthesis	46
2.5 Consonant representations	49
2.6 Vowel representations	51
2.6.1 Low vowels	53
2.6.2 Light diphthongs, contours, and releases	56
2.7 Unresolved questions	57
2.7.1 Laryngeal features	57
2.7.2 Nasals and laterals	58
2.8 An example from Japanese	85

Chapter 3. Outline of the formal system	61
3.1 Phonological structures	63
3.1.1 Some properties of PSs	65
Onset and nuclear licences	70
Coda licensing	74
Composed government relations	76
Sign-level structure	77
3.2 Descriptions of PSs	79
3.2.1 Lexical constraints and the status of the lexicon	83
3.2.2 Autosegments	89
3.3 Sorts	95
3.3.1 Nullness and empty nuclei	101
3.3.2 Underspecifying government arcs	111
3.3.3 Metrical structure	117
Recent approaches to prosodic representation	117
A possible representation of prosody	122
Metrical sorts and full values	125
Characterizing prosodic feet	127
3.4 Local domains and spreading	133
3.4.1 Dissimilation	135
3.5 Phonetic interpretation	136
3.5.1 Phonetic targets	137
3.5.2 Phonetic Event Structures	139
3.5.3 Phonetic Event Structures and phonetic events	142
3.5.4 The mapping principles	143
3.5.5 Phonology vs. phonetics: a cautionary tale for the border patrol	144
3.5.6 Defaults	150

Chapter 4. Locality: harmonies and reduplication	155
4.1 Recursive locality	156
4.2 Symmetric vowel harmony: Hungarian	161
4.3 Asymmetric vowel harmony: Kalenjin	167
4.4 Pasiego	175
4.4.1 Data	175
4.4.2 Feature-changing and feature-adding analyses	178
4.4.3 Analysis	181
4.4.4 Summary: Harmonic neutrality and transparency	189
4.5 Reduplication	190
4.5.1 The basic mechanism	192
4.5.2 “String” reduplication	195
4.5.3 Bases and templates	196
4.5.4 “Copy-back” reduplication	198
4.5.5 Why does reduplication care about prosody?	201
Chapter 5. Constraint-based morphology	205
5.1 Preliminaries	205
5.2 General properties of concatenation	214
5.2.1 Sister alignment	218
5.2.2 The Mother’s Border Principle	222
5.3 Special cases of “concatenation”	225
5.3.1 Suppletion and portmanteau morphs	225
5.3.2 Zero morphemes	226
5.3.3 Infixes	228
5.3.4 Circumfixes	230
5.3.5 Ablaut and allomorph selection	234
5.4 Prosody and morphology	235
5.4.1 The prosodic hierarchy	235
5.4.2 Representing the prosodic hierarchy	236
5.4.3 Prosodic edge requirements	241
A foretaste of Nisgha	246
5.5 Morphological overdetermination	250
5.6 Case study: Nisgha	253
5.7 Summary and implications	361

Chapter 6. Templatic morphology: Moroccan Arabic	273
6.1 Moroccan Arabic	273
6.1.1 Segments	273
6.1.2 Clusters, epenthesis, and syncope	275
6.1.3 Templatic morphology	277
6.2 Government Phonology analyses of MA	280
6.3 The cold nucleus system	288
6.4 The prosodic system	295
6.4.1 Moraic trochees	295
6.4.2 Iambs	297
6.4.3 Syllabic trochees	298
6.5 Roots	299
6.6 Moroccan Arabic “templates”	304
6.6.1 Base forms	304
6.6.2 Active participles	305
6.6.3 Reciprocals	306
6.6.4 Causatives	306
6.6.5 Passive participles	308
Chapter 7. Comparison with other frameworks	313
7.1 Government Phonology	314
7.2 Harmonic Phonology	321
7.3 The Theory of Constraints and Repair Strategies	324
7.4 Autolexical Syntax	328
7.5 Optimality Theory	331
7.6 “Declarative” phonology	338
7.6.1 Attribute-value structures and unification	339
7.6.2 Bird (1990)	349
7.6.3 Scobbie (1991)	355
Conclusion	363
References	368

Chapter 1

Introduction

Most of generative phonology has tended to see the phonological component of grammar as a subroutine in a computer program. The subroutine takes a character string as an input. It executes instructions one at a time in a predetermined order, performing various operations on the string: deleting a character, inserting one, changing one into another, switching the places of two, and so on. When the last of the instructions has been completed, what is left of the character string is the output of the subroutine.

In more recent generative phonology, instead of a single character string, there are several parallel strings hooked up to each other in various interesting ways. There are proposals to limit the set of instructions that can be used in the program or the order that they can be used in. There are even well-formedness conditions that intermediate strings are checked for and, if they fail, will either trigger clean-up instructions or abort the program. But, whatever elaborations have been added, the fundamental concepts and architecture of generative phonology remain tied to the dominant metaphor of phonology as a serial computer program.

The situation is quite different in recent approaches to syntax. Early transformational grammar had the same sort of structure as current phonology: an input, “S”, was transformed into a grammatical sentence by executing instructions, step by language-particular step. But more recent work in the Government Binding or Principles-and-Parameters framework (e.g., Chomsky 1981, 1986) has started to think of a grammatical sentence as any structured object that meets all the conditions that are imposed on sentences by Universal Grammar and by the particular language. What is relevant is not the history of the processes by which a sentence came to be, but simply whether or not it obeys all the conditions. Other syntactic frameworks, usually lumped together under the rubric “unification based approaches” (see, e.g., Shieber 1986), have made constraint-satisfaction the central idea of their theories.¹

¹These approaches include Functional Unification Grammar (e.g., Kay 1982), Lexical Func-

These ideas have been slow in making their way into phonology. Though individual phonologists may not like all the implications of the procedural computation model, it has become so ingrained in the way phonology has been done that it is difficult to see any alternative. Gradually, phonologists have been developing the start of an alternative in their theories of phonological representations, but these ideas have not been pushed as far as they could be. Given a choice between an analysis that crucially uses a difference in representation and an analysis that crucially relies on extrinsically ordering two rules, most phonologists would prefer the representational solution over the procedural. At the same time, phonological accounts have been relying more and more on general constraints on phonological representations, such as the No Crossing Constraint (e.g., Goldsmith 1976, Schein and Steriade 1986; see also Sagey 1988a, Bird and Klein 1990, for attempts to derive the constraint formally), the Obligatory Contour Principle (e.g., Leben 1973, McCarthy 1986, 1988, Mester 1986), or Prosodic Licensing (Itô 1986, McCarthy and Prince 1986, Goldsmith 1990).

There is a definite distaste for explanations that rely on proceduralist assumptions, and an ever-dwindling set of phonological phenomena seem to require them, but very few have taken up the challenge of trying to do away with them altogether. It is the purpose of this dissertation to attempt just that. I take a number of ideas (described below) that have been persuasively argued for individually in the phonological literature, and show how together they allow one to build a formal theory of constraints on **phonological structures (PSs)**. These constraints alone are enough to define the legal words of a language, without any need to specify how or in which order or by which procedure the constraints are applied.

I now briefly describe each of the ideas that I will be using in the course of this paper. Chapter 2 is devoted to the gesturally-based model of segmental content. The rest of the ideas will be discussed more fully in the following sections.

Phonology is

constraint-based: The content of the phonological component of the grammar is a set of constraints on what constitutes a legal phonological repre-

tional Grammar (e.g., Bresnan, ed., 1982), Generalized Phrase Structure Grammar (e.g., Gazdar, Klein, Pullum, and Sag 1985), and Head-Driven Phrase Structure Grammar (e.g., Pollard and Sag 1987). Besides a commitment to the centrality of constraints, these frameworks also share the practice of using feature structures or **attribute-value structures** as representations and unification as the sole, or at least main, operation of the grammar. The notation system of each framework has been similar enough that Shieber (1984) was able to develop the formal language D-PATR as a sort of *lingua franca* for unification-based theories. The work of Johnson (1988, 1991) has made explicit the formal underpinnings of these approaches and showing how their notations can be made equivalent to constraints written in a first-order language whose universe of discourse is linguistic objects. In this dissertation, I shall rely heavily on the ideas of Johnson. One of the results is that much of the framework I propose could be rewritten in a notation similar to that of unification-based syntax, or to closely related phonological frameworks (e.g., Scobbie 1991), if one really preferred to use matrices instead of more familiar phonological diagrams.

sentation. (There are no rewrite rules.) In the literature, being constraint-based usually involves three related properties: declarativity, monotonicity, and monostratality.

declarative: Phonology specifies *only* what counts as a valid phonological structure. It is neutral as to exactly what algorithm one chooses to use in building those structures or in checking their validity.

monostratal: There is only one level of phonological representation. There is no underlying phonological representation that is distinct from the surface representation and needs to be transformed into the surface representation through rewrite rules. Put another way, each part of any phonological representation should have a direct effect on either the content or timing of phonetic events.

monotonic: Information cannot be destroyed. Or, an imposed constraint can't be ignored because some other constraint is "stronger". We shall see that this property cannot hold of a system with default rules, but it will be argued that default rules belong to the system of phonetic interpretation and not to phonology proper.

government-based: PSs are built out of asymmetric relations between atoms or smaller PSs. This follows and extends recent work in Government Phonology (Charette 1988, Kaye, Lowenstamm, and Vergnaud 1990), where there is, for example, an asymmetric relation between a syllabic nucleus and the preceding onset, where the nucleus is said to *govern* the onset.

gesturally based: The representations of sub-segmental structure represent fairly directly the articulatory constriction gestures that are their phonetic interpretations. Gestures are phonologically specified for the primary articulator, the site of the constriction, and the degree of constriction.

partial in principled ways: It is possible for a legal PS to have no segmental content specified for certain positions. Other frameworks have proposed various principles for what should be done with these empty positions and how they should be interpreted. For example, Underspecification Theory (e.g., Archangeli and Pulleyblank 1989) will apply redundancy rules to empty positions to fill them with the features that they will be phonetically interpreted with. Government Phonology proposes that many empty positions can, under the right circumstances, remain phonetically uninterpreted, while in other circumstances they receive default interpretations such as velarity.

All of the above ideas will be combined and framed in a formal language, based on first-order logic, whose job will be to describe what count as legal and illegal phonological representations.

It is worth emphasizing the independence of each of the above ideas. It would be conceivable to have a declarative phonology that uses SPE feature matrices, though it may be extremely clumsy and miss several important generalizations. It is just as possible to have gesturally-based segmental representations as part of a theory that assumes the power of unrestricted rewrite rules. Any novelty in this dissertation comes not from arguments for any one of the ideas, but from suggestions that all of them taken together can make up for the weaknesses each of them has when taken alone.

It is also worth emphasizing that I am not making claims that procedural phonology is empirically inadequate. The procedural devices autosegmental phonology has at its disposal are enough to characterize any fragment of a human language, and a good deal more besides, so it is not possible to disprove procedural autosegmental phonology by pointing to some piece of data that it is not able to handle. Arguments come instead from considerations of theoretical economy: the toolkit of phonology can be stripped down to some fraction of its present size and still do all the work it needs to do. We shall also see throughout this dissertation some examples of phenomena where a constraint-based analysis genuinely appears to be more appropriate. These are not cases that procedural phonology is inherently incapable of dealing with, but they *are* cases where procedural accounts can be convoluted and seem to miss the right generalizations—they are not insurmountable problems, but they are still problems. On the other hand, a constraint-based account can deal with these cases easily and naturally—they are the kinds of phenomena that a constraint-based account would lead one to expect should exist.

1.1 Constraints

Phonology has for quite some time been moving away from the view that the best or the only way to express generalizations concerning phonological representations is by means of transformational rewrite rules that can in principle be extrinsically ordered. Little by little, much of the work formerly done by rules has been assigned to enriched representations or to (preferably universal) conditions on the well-formedness of representations. Chomsky and Halle (1968, henceforth “SPE”) had already discussed Morpheme Structure Rules, though these were still treated as a kind of rewrite rule. Others extended the idea of Morpheme Structure Conditions as constraints on the phonotactics of the underlying form of morphemes. Kisseberth’s (1970) discussion of rule “conspiracies” in Yawelmani made it clear that some similar phonotactic constraints also applied to surface forms.

In the early 1970s, many phonologists embarked on a research programme that aimed to capture generalizations without extrinsic ordering (e.g., Koutsoudas, Sanders, and Noll 1974). They still used transformational rewrite rules, and the rules still applied at different stages in a derivation, but the order of rule application was supposed to be predictable from universal principles. Many of these ideas were taken up in the framework of Natural Generative Phonology (e.g., Hooper 1976).

The no-ordering research programme got lost in the shuffle as phonologists discovered

the delights of autosegmental phonology. But the initial work of Williams (1976), Leben (1973), and Goldsmith (1976) was prompted by essentially the same concerns: using richer representations to avoid the need for accounts using powerful transformational devices (and lexical or segmental diacritic features regulating how those devices applied). McCarthy (1979) placed much of the burden of Semitic root-and-pattern morphology on enriched representations (such as templates) rather than transformational rules (e.g., gemination rules, vowel substitution rules). Marantz (1982) extended McCarthy's ideas to handle reduplication, arguing forcefully against an approach based on transformational rules. The autosegmental phonology tradition of seeking representation-based explanations to replace rule-based ones continues in today's work on feature hierarchies.

A few more thoroughgoing attempts to forego the procedural apparatus have appeared in the literature. Hudson (1980) is an early attempt to do much of morphophonology without derivations. The Categorical Phonology of Wheeler (1981) shares many of the goals of recent constraint-based phonology. While most of Halle and Vergnaud (1987) talks about metrical structure in terms of procedural algorithms ("first build line 1 constituents, then..."), the section that lays out their theory formally uses a mostly declarative characterization of legal metrical structures that is deliberately modelled on Prolog programs. By far the most sustained attempt to rework phonology in a constraint-based framework has been recent work by often labelled "Declarative Phonology". Starting with researchers such as Bird (1990), Waksler (1990), and Scobbie (1991), who borrow much the formal apparatus of unification-based approaches to syntax, a large body of work has grown that tries to capture phonological regularities in a non-procedural way. (The papers in the volume edited by Bird (1992) are representative of this trend.)

The basic idea of any constraint-based approach is simply there are conditions on PSs that come from a variety of sources, and any PS (or perhaps the smallest PS) that simultaneously satisfies all of them can instantiate a legal word of a language.²

There are universal constraints on what can count as a licit PS. The most obvious kind simply specifies what sorts of entities it is that phonology has truck with. A PS can be built out of syllables or skeletal slots or morae or features or whatever one's favourite set of primitive entity-types is, but Wh-traces, lambda operators, and telephone numbers have no place in a phonological structure. Of course, universal grammar does set some more stringent conditions than these on the set of valid PSs. The exact membership of the set of universal constraints is largely an empirical question, but some promising candidates would seem to include constraints such as "A syllable has at most one onset" or "Feet are maximally binary."

²This is not a necessary feature of theories of phonology that have constraints as central components. The Theory of Constraints and Repair Strategies of Paradis, for example, remains strongly proceduralist. The Optimality Theory of McCarthy, Prince, and Smolensky also retains many aspects of proceduralism forced by their continuing acceptance of the hypothesis that morphemes are made up of pieces of representation that need to be actively joined together. For comparisons between the framework proposed in this dissertation and these and other constraint-centred theories, see chapter 7.

Individual grammars also place constraints on what PSs can occur in their languages. Some frequently encountered language-particular constraints are: “Any coda consonant must have the same place of articulation as the following onset consonant”, “Any word final obstruent is voiceless”, “Syllable rhymes are maximally binary”.

Besides the purely phonological constraints on PSs that we have been discussing so far, there are also constraints on possible sound-meaning pairs. Morphemes are essentially licences that allow certain phonological structures to be associated with certain syntactic and semantic structures. Thus, while *bnik* is ruled out as a word of English by the constraints on possible onsets, *blik* meets every purely phonological constraint on PSs perfectly well—its downfall is simply that there happens to be no syntactic or semantic object that the English lexicon licenses it to be paired with.

Obviously we need to be more exact about what these constraints are. It’s all well to say “Any PS that is to count as an instance of the past tense morpheme must look like this...”, as long as we have an explicit and adequate way of spelling out what the “like this...” is. A fuller outline of my proposal for accomplishing this is given in chapter 3.

1.1.1 Declarativity

The concept of declarativity comes from computer science, where declarative strategies are opposed to procedural (see, e.g., Kowalski 1979). **Procedural** problem-solving spells out exactly the steps and operations that are needed to construct a solution to the problem. There might be several different procedures that can be used to arrive at the same solution. Someone writing a procedural program must choose just one of these as the “correct” one, usually based on considerations such as simplicity and efficiency. This is the situation of most generative phonology. There are usually several different derivations that can produce a given surface form from a given underlying form, and phonologists are forced to choose between them, using much the same criteria as the computer programmer uses.

Declarative problem-solving, on the other hand, does not care in the slightest what procedures are used in order to construct a solution; it simply specifies what a correct solution would have to look like, what properties it would have. This description of the correct solution is independent of how we might decide to go about finding it. Again, there are many possible procedures that can check a candidate solution to see if it fits the description of the correct solution, and there are many possible procedures to generate candidate solutions. But the declarative specification of the correct solution is neutral with respect to these different procedures.

It is not clear exactly what claims generative phonology makes concerning its procedural underpinnings. As in any field of the generative enterprise, phonologists would undoubtedly subscribe to the idea that their task is to characterize “what a speaker knows when she knows a language.” But it is less than obvious what “characterize” should involve. Is this slogan to mean that a speaker “knows” rule ordering conventions and so forth in any psychologically real way? If this is the case, generative phonologists

have been avoiding their responsibility to demonstrate that speakers actually choose among the myriad of procedural models in the same way that phonologists have chosen among them, that speakers actually care about factors like simplicity, efficiency, and formal elegance. Or, perhaps, is the claim of characterizing speakers' knowledge the weaker claim that the phonologist's procedural model discriminates the same class of linguistic objects as does the speaker's competence? If the task is only to single out the right linguistic objects, there is a heavy burden of proof on generative phonology to show that a model that both describes these objects *and* spells out a procedure to construct them is more adequate for this task than one that only describes them. Or, put another way: given that constraints will be needed in any case, it must be shown that a model that uses only constraints is clearly inferior to one that needs constraints and transformational rules besides (for the purpose of discriminating the set of legal forms).

In order to accomplish this task of describing valid phonological linguistic objects (what I shall call **phonological structures**, or **PSs**), we need a theory of what PSs are and a formal method for describing them.

1.1.2 Monostratality

Because most linguistic frameworks have declared themselves fairly clearly and explicitly concerning the number of linguistic levels that they hold to exist, the property of monostratality has seemed to be much more straightforward than it in fact is. Much of the problem centres around the unclearness of the word "level". Arc Pair Grammar (Johnson and Postal 1980), for example, denies that it has any need for a derivation with several different levels of representation and claims only one level of representation, the R-graph. But at the same time, it explicitly acknowledges the existence of "strata" within the R-graph, only one of which has any direct bearing on the phonological form of the sentence.

Frameworks such as Arc Pair Grammar show that there is no necessary relationship between monostratality and other theoretical choices such as declarativity. Under reasonable interpretations of the terms, Arc Pair Grammar can be said to be declarative but multistratal. It holds linguistic representations to be sequences of strata, but characterizes these valid sequences declaratively using a language based on first-order predicate calculus. The work of Stabler (1990, 1992, 1993) in the formalization of Government-Binding (and its use in parsing) is another excellent example of a rigorously logical and declarative approach to a multistratal grammatical model. Similarly, there is nothing preventing a theory from characterizing monostratal representations procedurally. Examples of this are harder to come by among hard-core linguistic theories, but this seems to be a fair way to classify the Marcus parser (Marcus 1980) and Augmented Transition Networks (Woods 1970).

So, though the property of monostratality is not as watertight as one might have thought it to be, it still is useful as a rough-and-ready tool for distinguishing various approaches to declarative phonology. The Cognitive Phonology model of Lakoff (1988),

for example, or its connectionist implementation by Touretzky and Wheeler (1989), fairly clearly involves two strata. The model tries to relate declaratively an underlying representation with a surface representation, albeit directly and without the numerous intermediate stages of procedural models, but there are two strata nonetheless, only one of which is relevant to the phonetic interpretation of the form. Similar remarks could be made about the two-level morphological models of Koskeniemi and Kartunnen (Koskeniemi 1983).

Unlike Cognitive Phonology or two-level morphology, I do not propose two levels of phonological representation; rather, constraints are all brought to bear on a single phonological structure. In this sense, my proposal can be said to be monostratal, like those of Bird and Scobbie. Using a stronger definition of monostratality, whether only a subset of the formal structure is relevant to phonetic interpretation, the question becomes more complex. In an obvious sense my proposal, recognizing like other “sign-based models” that there are syntactic and semantic structures parallel to the phonological structure which have no direct effect on phonetic interpretation, would no longer qualify as monostratal. Even within the phonological structures, I admit that there are designated “null” positions (for example, those corresponding to the empty nuclei of Government Phonology), which by definition receive no *overt* phonetic interpretation, though they continue to contribute to the temporal ordering of the other positions with which they stand in government relations.

1.1.3 Monotonicity

Monotonicity is a property of some constraint systems that involves how persistent the constraints are. Using the terminology of **hard** and **soft constraints**, all constraints in a monotonic system must be hard. No hard constraint can be ignored because some other constraint is stronger. New information does not cause any revision in existing assumptions; as new morphemes are added to a word, the amount of knowledge about its phonology can only grow, never shrink.

Logical systems can be seen as systems that can infer or prove theorems on the basis of a set of axioms. Looked at this way, if a monotonic system can infer a theorem from the set of axioms {A,B,C}, it can also infer the theorem from the set {A,B,C,D}. No theorem of a monotonic system will be made false by the addition of a new axiom. It may not be immediately apparent how this relates to phonology, but the logical description language I propose can be seen as treating constraints as axioms and the properties of phonological structures as theorems that can be derived from those axioms. Applying a new constraint to a form, i.e., adding to the set of axioms, cannot reverse the effect of constraints that have already applied, i.e., cannot falsify an existing theorem.

This has consequences for what lexical constraints must be like, if lexical constraints are hard constraints. A lexical constraint may require that a certain consonant position in all forms of a certain verb must be filled by a velar stop. If one form has a voiced velar stop while another has a voiceless, “velar stop” is in fact the most that the lexical

constraint can say for the verb as a whole. It cannot require, say, that an “underlying form” have a voiced velar stop, expecting some other constraint to cancel or overrule this requirement in the second verb form. The most a phonological constraint can specify about a set of words is their “greatest common denominator”, what properties all the words have in common.³ Any property not shared by all the words, even a property that is unpredictable from the common properties of the set, can be specified only for those words that it actually belongs to.

There is one area of phonology where the desirable property of monotonicity breaks down. That area is default specification. A default specification is an inference rule that lets one draw a conclusion in the absence of evidence to the contrary. For example, in the absence of any information about what vowel should fill a certain slot, a default rule can allow one to conclude that it is an *i*. It is the required absence of contrary information that makes default rules non-monotonic. It is possible that the constraints on one form of a verb may make no claim on the identity of a certain vowel, so the default rule will fill something in. But in another form of the same verb, that particular vowel slot may be spoken for, say by an affix, so that the default rule can no longer apply. In effect, the addition of a new axiom (the constraint from the suffix) has falsified what used to be a theorem (that a vowel position is filled with the default vowel). Bird (1990) and Scobbie (1991) have also remarked on the non-monotonic properties of systems with default rules and the implications for phonology.

It appears likely, however, that all the effects of this non-monotonicity can be localized in the phonology-phonetics interface. That is, the default rules have no effect on what does and does not count as a legal phonological structure. Their only role is in phonetically interpreting legal PSs. It is unsurprising that the principles relating PSs and physical phonetic events are non-monotonic, that most of the constraints of the interface are—to some degree or other—soft. How well a phonetic event satisfies a phonological structure is a gradient measure, not absolute. A PS will be well-satisfied by a phonetic event where the sub-events corresponding to unstressed vowels are clearly articulated close to those vowels’ prototypical values. But the PS might be satisfied almost as well by a phonetic event where the parts corresponding to unstressed vowels are reduced and schwa-like. The constraints governing the interpretation of unstressed positions are relatively weak, and violations of them are not particularly costly. The addition of another type of non-monotonic principle (default rules) to this already non-monotonic set of phonology-phonetics mapping principles is not a serious defect, especially if it will allow the set of *phonological* constraints, those that define legal PSs, to remain monotonic. Section 3.5 discusses some of the properties of the phonology-phonetics interface.

A second area in which phonology falls short of a perfectly monotonic ideal involves selecting between rival candidate PSs that all meet the constraints imposed on them. For example, all the constraints that are brought to bear on a certain word might be equally well satisfied if one of the vowels were either short or long. Most languages in

³The effect of using only their common properties to constrain a set of words is often similar to using archiphonemes to represent the underlying form of the set.

most situations will allow only the short vowel alternative as the legitimate form of the word. Generally, languages do not take kindly to superfluous phonological structure, i.e., structure that is not required by any of the constraints that apply to a word. In addition to constraint satisfaction, we will need some mechanism to choose among the alternatives that satisfy the constraints, perhaps as part of the phonology-phonetics interface. Although much more study needs to be done in the area, it seems probably that this selection can be done by referring simply to the “size” of the competitors, as defined by some straightforward metric, without the need to refer (somehow) to the “intent” of the constraints.

A more serious challenge to the property of non-monotonicity comes from cases where the most elegant analysis of a phenomenon would appear to involve a deletion rule. The question is more one of aesthetic inconvenience than empirical inadequacy, since monotonic frameworks are perfectly able to handle apparent deletion phenomena. Bird (1990), for example, relies on Hudson’s (1980) arguments for allophonic “alternation with zero”. Hudson uses disjunction as the logical connective involved in these cases, though conditionals would also work. The lexical entry for a morpheme could specify unconditionally everything that all its forms have in common, and then specify conditionally any other properties that only some of its forms have. Such a constraint might look like the phonological equivalent of donkey-anaphora: “If there is another consonant after this, it is a *k*.”

One possible way of looking at this is that the undeleted material in an alternation is not just inert phonological content that happened to have survived unscathed through a dangerous derivation; rather, it actively marks, however irregularly, the morphological content of the form it appears in. It is not a question of the monotonic analysis missing a generalization that the non-monotonic one captures. Rather, it a question of choosing which generalization to express. A deletion analysis essentially says “X occurs in all forms of this word, except in those it doesn’t (which have Y in common).” A monotonic analysis says “All forms of this word have this in common. If additionally X occurs in environment \bar{Y} , it means...”⁴ There is some evidence that speakers can and do choose the analysis that monotonic systems are forced to use.

Hale (1973) discusses an alternation in Maori between active and passive verb forms (data from Bynon 1977):

⁴It should be noted that the additional complexity in spelling out the environment \bar{Y} as opposed to Y is not as great as is usually assumed.

(1.1)	<i>Verb</i>	<i>Passive</i>	
	awhi	awhitia	‘embrace’
	hopu	hopukia	‘catch’
	aru	arumia	‘follow’
	tohu	tohuŋia	‘point out’
	mau	mauria	‘carry’
	wero	werohia	‘stab’
	patu	patua	‘strike, kill’
	kite	kitea	‘see, find’

There are (at least) two possible analyses of this data, an elegant one and an ugly one. The elegant one uses deletion. There are underlying representations that may end in a consonant *awhit*, *hopuk*, *arum*, *patu*, a rule deleting word-final consonants, and a passive suffix that alternates between *-ia* after a consonant-final and *-a* after a vowel-final stem.⁵ The ugly analysis assumes vowel-final stems (*awhi*, *hopu*, *aru*, *patu*), each of which idiosyncratically selects an allomorph of the passive suffix: *-tia*, *-kia*, *-mia*, *-a*, etc. The second, ugly solution (or one along the same lines of allomorphy) is the one that a monotonic framework is more or less forced to adopt. This might seem like evidence that the monotonic framework is inferior to one that allows deletion rules, but Hale argues that native speakers also choose the second, ugly analysis. Some of the evidence that there is suffix allomorphy, with the unmarked allomorph being *-tia*, includes the fact that the *-tia* group is attracting stems from the more marked suffixes and that foreign loan words ending in a vowel, for which speakers can have no evidence of an underlying consonant, invariably take the *-tia* form of the passive.

The Maori example should provide a note of caution for proponents of transformational rules. Even when it seems that a deletion analysis is more elegant than a monotonic alternative, it is by no means a foregone conclusion that it is more appropriate. But many deletion rules proposed in the literature do not even have the advantage of being more elegant. Rather, they are used simply because they are already conveniently lying in the phonologist’s toolbox when there is a need to undo the damage done by some (not entirely plausible) assumptions about the nature of phonology or morphology. In our discussion of morphological overdetermination in chapter 5, we shall see that a large class of apparent deletion rules are in fact handled far more elegantly by a monotonic constraint-based approach.


1.2 Government in phonology

Recent approaches to phonology have been developing the idea that the “structure” of phonological structures comes from a certain kind of asymmetric relationship between the

⁵This could be further simplified by assuming an underlying representation *-ia* for the suffix and another deletion rule removing *i* between vowels.

building blocks, a relationship where one of the things joined can be seen as dependent on the other. Autosegmental phonology has been moving in this direction, replacing the unstructured feature bundles of SPE with feature geometries organized using the asymmetric relation of dominance between nodes (and the asymmetric relation of linear order along tiers). Opposed to the relatively unstructured representations of Particle Phonology (Schane 1984), there are theories of segmental content whose central idea is that the segmental primitives enter into asymmetric relations, e.g., **dependency** relations in Dependency Phonology (Anderson and Ewen 1988)⁶, **government** relations in Government Phonology (Kaye, Lowenstamm, and Vergnaud 1985).

The major difference between autosegmental phonology and approaches such as Dependency and Government Phonology lies in the way asymmetric relations are exploited. In the latter approaches, the primitives of segmental content enter directly into asymmetric relations with each other. In autosegmental phonology, the “primitives” (binary features) seldom enter into dependencies with each other,⁷ rather they bear their asymmetric relations to higher-level nodes that are supposed to represent some common aspect of the primitives that can “dock” onto them, such as the region of the vocal tract they all occur in. For example, in Sagey’s (1986) feature hierarchy, [+high] would not enter into a relation with [-low], as in (1.2a), a structure central to frameworks like Dependency and Government Phonology; rather, the two would be sisters, each standing as a dependent in its own dominance relation with the higher Dorsal node, as in (1.2b).




This use of asymmetric relations requires a commitment to some degree of ontological reality for the “higher” nodes like Dorsal.

The difference between the two approaches becomes clearer when they deal with structure above the segmental level, such as syllable structure. A not-untypical autosegmental treatment, translating Fudge’s (1969) proposal more or less directly (cf. Steriade 1982), would have a suprasegmental structure for the English word *drank* along the lines of:


⁶Dependency Phonology also allows two primitives to stand in a symmetrical relation, or at least to mutually depend on each other with two asymmetric relations. See den Dikken and van der Hulst (1988) for a discussion of some of the problems with this kind of structure.

⁷though they do so in a few proposals, such as Mester (1986), Goad (1991).



Accepting a structure like (1.3) commits the autosegmental phonologist to accepting the nodes involved as autonomous entities which should behave as such, and to adopting one of a limited number of possibilities regarding which autosegmental tier each of the nodes lives on. Much of the recent autosegmental research in syllable structure has been prompted by the failure of various nodes in (1.3) to behave the same way subsegmental nodes in similar dominance relations do.

Other approaches to syllable structure see the relations between parts of a syllable as relations that hold directly between the segments involved, without the mediation of higher-level nodes of uncertain status. A Dependency Phonology representation of the syllable structure of *drank* might look like (1.4), where a line indicates a dependency relation in which the segment lower on the page is the dependent.



Government Phonology's views of syllable structure (e.g., Kaye, Lowenstamm, and Vergnaud 1990, Charette 1988) are a compromise between the two above extremes, handling some relations as direct dependency between segments as in (1.4), some as sisterhood relations as in (1.3), and some redundantly as both.

The framework I propose here more closely resembles Dependency Phonology in assuming all syllabic structure to be the result of government relations holding directly between the root nodes of the segments involved. As in Dependency Phonology, and as for some types of government in Government Phonology (e.g., coda licensing, onset licensing), a government relation between two positions does not require or create an independent object at some higher level of the representation. In this sense, government relations are primitive and not structurally defined.

There is another sense in which governments are not primitive. The basic government relations can of course be listed; I shall propose, among others, asymmetric relations holding between:

- a nucleus and its onset
- an onset and the preceding coda
- a nucleus and the next nucleus
- a segment and the specification of its active articulator
- a segment and the specification of the site of its articulation
- a segment and its stricture specification
- the stop portion and the fricative portion of an affricate

But there is more to be said about government relations than a simple list like this. Specifically the different types of government relation form natural classes, e.g., they can pattern together for the purposes of some constraint. The possibility of there being natural classes of government types opens up the possibility of underspecifying them. The lexical entry for a morpheme might require that a relationship between two points simply be one of the members of a certain natural class, and let the choice of which member be determined by its interaction with other constraints. (For example, see the analysis of Rotuman metathesis in section 3.3.2.) The possibility of natural classes also suggests that government relations themselves may have a compositional rather than a primitive structure. Just as natural classes of phonemes have led phonologists to posit just a few primitive segmental features that combine into intricate structures, it is possible that the whole bestiary of government relations too can one day be shown to be combinations of a small handful of primitives.

1.3 Empty positions

1.3.1 Pronounced empty positions

The idea that there is some level of phonological representation where not every feature has to be specified has been around for a long time. One possible, though not entirely accurate, interpretation of the structuralists' archiphoneme is as a device for making underlying representations only as specific as they had to be. Chomsky (1965:87) proposed that the lexicon held only those

- aspects of phonetic structure that are not predictable by general rule (for example, in the case of *bee*, the phonological matrix of the lexical entry will specify that the first segment is a voiced labial stop and the second an acute vowel, but it will not specify the degree of aspiration of the stop or the fact that the vowel is voiced, tense, and unrounded).

Chomsky and Halle (1968:163–171) fleshed out this suggestion with lexical entries that could contain underspecified features like [0tense] and Lexical Redundancy Rules to fill these zeroes in with the appropriate value, [+tense] or [-tense], to result in the underlying representation. Because of criticisms by Lightner (1963) and Stanley (1967) against the power of a system that allowed three values for a feature (+, -, and 0), it was generally assumed that the possibility of 0-features was not available in the phonology any time at or later than the underlying representation.

This assumption was challenged by new theories of underspecification (e.g., Archangeli 1984, Pulleyblank 1983) that allowed all occurrences of a particular polarity of a feature to be absent in underlying representation and to become filled in during the course of a derivation. Supporters of underspecification theories have shown that many interesting results can be obtained by assuming that some specification does not yet exist at the time some phonological rule applies.

The possibility that certain positions in a PS can be empty, but still receive specifications by default, is important. But there is a small problem raised by traditional underspecification theories for a monostratal framework. In a theory that assumes derivations through time, it is possible for a position that is [α F] on the surface to be underlyingly empty and have the [α F] filled in by a default rule. Somewhere in the middle of the derivation, a rule can apply that we would normally expect to affect the [α F] segment, but since the default rule has not yet applied, the segment is not subject to the rule. Later, after the default specification has been filled in, another rule (or perhaps the same one) might apply that *does* affect the [α F] segment as expected. The [α F] specification is thus allowed to behave both as if it were there and as if it were not. In a framework with no derivations and only one level of phonological representation, a specification can only be or not be in the environment for a phenomenon, not both. If a position in a PS is empty, the default interpretation it receives in the phonetics should have no effect whatsoever in the phonology.

Because of this, it is useful to distinguish between **default rules** and **redundancy rules**. Default rules are the phonetic principles that give an interpretation to an empty position in a PS. They may refer to properties of the environment of the PS in deciding what the default interpretation may be, but crucially they cannot affect what is and is not a legal PS of the language. Redundancy rules, on the other hand, are constraints on PSs and affect the legality of PSs like any other constraint. They require certain properties to co-occur in PSs, i.e., any PS that has a configuration A (say a round vowel) must also have configuration B (backness for the vowel).

1.3.2 Unpronounced empty positions

Not all empty positions need to receive an overt phonetic interpretation. One of the features of the framework developed here is that PSs can have empty positions that are designated as **null** and will have no phonetic realization. Default rules will not apply to positions that are designated null. These null positions can be seen as a kind of

“placeholder” in PSs.

This idea comes more or less directly from Government Phonology and its arguments for the existence of empty nuclei. GP’s treatment of empty nuclei and the way it is adopted here will be treated in more depth in chapter 3, but I would like to give some idea now of how allowing syllable nuclei with no phonetic content can greatly simplify the expression of constraints.

Consider the alternation in Moroccan Arabic between *ktəb* ‘he wrote’ and *ktbu* ‘they wrote’. If we assume more or less traditional syllabifications of these forms, [k \emptyset] [tə] [b] and [kə] [t \emptyset] [bu], it would be difficult to express what the two verb forms have in common in terms of which syllabic positions their consonants occur in. Both *t* and *b* move back and forth between onset and coda in the two forms. The problem is especially pressing in a monostratal system that has no place for resyllabification processes. Such a system is forced to express the syllabification constraints on the *b* in terms general enough that they are satisfied by *b* being in either a coda or an onset.

But empty nuclei suggest another possible syllabification for the forms: [k \emptyset] [tə] [b \emptyset] and [kə] [t \emptyset] [bu]. It is now easy to see what the two words have in common. The consonants do not jump between coda and onset—they are always in onsets. The only property that changes between the two forms is which nuclei are designated as null and which contain vowels. It will in fact turn out that all of the nuclei, except for the third person plural suffix *-u*, are empty, and what changes between the words is which of the empty nuclei get default interpretation as schwa and which do not. The principles constraining the legal distributions of null positions throughout a word will be discussed in 3.3.1.

1.4 Morphemes as constraints

One of the dominant tensions in modern phonology and morphology is the proper division of labour between representations and rules. Should phonological generalizations be captured by assuming a rich theory of representations which can be complicated and potentially quite abstract together with a very restricted inventory of rules with limited power (a bias towards representations), or should they be captured by phonological rules (and rule inventories) of great power and complexity, possibly with very simple representations (a bias towards rules)? In one form or another, this tension has been present throughout the twentieth century, as can be seen in Anderson’s (1985) survey, admittedly biased towards the rule end of the spectrum.

For the most part, current autosegmental phonology is biased toward representational accounts. Intricate structures of feature hierarchies and prosody have greatly limited the amount of transformatory work that needs to be assigned to rules, and have made it conceivable to have a very limited inventory of possible rules (e.g., Clements 1989, Archangeli and Pulleyblank 1992).

The same tension between reliance on rules and reliance on representations has characterized most of the history of morphology as well. But there is a significant difference.

In phonology, there is general agreement that representations of some kind and rules of some kind are both necessary parts of the theory; any disagreement involves the emphasis that should be given to one or the other. In morphology, on the other hand, what is at stake is the ontology of the subject matter itself. The question is aptly summed up by a section heading in Spencer (1991): “Morphemes: things or rules?”

On the one hand, there are many who view morphemes (or at least the phonological content of morphemes) as pieces of phonological representation. Morphology is responsible for sticking together these pieces of representation, preferably in a manner that is as much like simple agglutinative concatenation as possible. Ideally, the principles that produce the surface representations from these smaller pieces would have the status of universal conventions. In reality, there is usually a need for a battery of clean-up rules to come in and fix up the representations in ways that universal conventions cannot. Many of these clean-up rules are triggered by explicitly morphological properties. In many presentations, it is not clear what the status of these morphologically conditioned rules is.

On the other hand, there are some researchers who believe that morphemes are best seen as *processes* or rules that operate on base representations to produce new representations. A classic example is the English past tense morpheme. For most verbs, the content of this morpheme is the operation of adding *-ed*; for the verb *take*, it is the operation of replacing the /ei/ with an /U/. Anderson (1992), though he disavows the term “morpheme” for his “word formation rules”, is one of the most complete expositions of this point of view.

One of the weaknesses of the rule-based approach is that it has little to say on the issue of the restrictiveness of the rule inventory, one of the primary concerns of the representation-based approach. Anderson (1992: 172), for example, proposes the following unabashedly transformational word formation rule for Potawatomi that exchanges en masse the features of two layers of a morphosyntactic structure (the outer layer representing the subject, the inner the object):

$$(1.5) \quad \left[\begin{array}{l} +\text{Verb} \\ \left\{ \begin{array}{l} +\text{Obv} \\ -\text{me} \\ -\text{you} \end{array} \right. \left\{ \begin{array}{l} [+ \text{Obv} \\ + \text{Anim}] \\ \left[\left\{ \begin{array}{l} +\text{me} \\ +\text{you} \end{array} \right\} \right] \end{array} \right. \right\} \end{array} \right] \\ \left[+\text{Verb} 1 [2] \right] \rightarrow \left[+\text{Noun} 2 [1] \right] \\ /X/ \rightarrow /Xuko/$$

If word formation rules have the power to perform operations of this complexity, it is difficult to imagine what they could *not* do. Specifically, a representationally biased morphologist would wonder what would prevent a grotesque, but formally comparable, rule like:

$$(1.6) \quad \left[\begin{array}{l} +\text{Verb} \\ +\text{Past} \end{array} \left\{ \begin{array}{l} +\text{Obv} \quad \left[\begin{array}{l} -\text{me} \\ +\text{you} \end{array} \right] \\ +\text{me} \\ -\text{Anim} \quad \left[\left\{ \begin{array}{l} -\text{you} \\ +\text{Obv} \end{array} \right\} \right] \end{array} \right\} \right] \\ \left[+\text{Verb} 1 [2] \right] \rightarrow \left[+\text{Future} 1 [2 [+\text{me} [2 [1]]]] \right] \\ /X/ \rightarrow /XuXko/$$

In Anderson's framework, the word formation rules that are morphemes (besides being this complicated) are also extrinsically ordered into highly intricate patterns. Given a framework that not only allows, but *requires*, this level of expressive power, it is difficult to see how any of the insights of the representation-based approaches on the restrictiveness of rules could be integrated.

The central issues of this debate were presaged in Hockett (1954), who characterized and compared two models of morphology, which he called **Item-and-Arrangement** (or **IA**, the representation-based model) and **Item-and-Process** (or **IP**, the rule-based model). Hockett offered formalized definitions of both models (IP had never been explicitly formalized before). Pointing to problems caused for Item-and-Arrangement by examples like the English past tense *took*, Hockett decided in favour of the Item-and-Process model.

In fact, the version of representation-based morphology usually assumed nowadays is more extreme even than the Item-and-Arrangement model discussed by Hockett. While Hockett's characterization of IA countenanced such indiscrete behaviour on the part of morphs as blurring together at their edges (portmanteaux morphs being an extreme example), most modern researchers adopting an IA stance tacitly make much more stringent demands on the representations that form their morphemes. Put bluntly, morphemes act like bricks: they can be placed end to end, but under no circumstances can two bricks occupy exactly the same place at the same time. I shall refer to this assumption as the "Physical Integrity of Morphemes" hypothesis.

A large part of the work in autosegmental phonology can best be seen as an attempt to maintain the Physical Integrity of Morphemes hypothesis in the face of blatantly uncooperative data from languages. Cases that seemed to involve two morphemes living in the same string of segments (for example, a tense morpheme that involves a distinct sequence of tones superimposed on the vowels of a verb root) were instead analyzed by segregating the information belonging to the two morphemes onto separate tiers. Using the brick metaphor, the bricks may be stacked on top of each other, but it is still impossible for two of them to occupy exactly the same space. Even in a system like Semitic, where morphemes seem to be inextricably fused, analyses like that of McCarthy (1979) managed to separate them onto independent tiers with some degree of success. The concept of morphemic planes allowed autosegmental phonology to handle those cases where even segregation onto phonologically motivated tiers was not sufficient to maintain the Physical Integrity of Morphemes hypothesis, for example, if two morphemes both

consisted of features of exactly the same type attached to the same or to interleaved skeletal slots.⁸

This dissertation will argue for an approach to morphology that is different from both the rule-based and representation-based ones. The central idea of this approach is that the phonological content of a morpheme is a constraint, that is, a description of a representation that can be framed in the same description language used for other phonological constraints. Morphemes as a whole are licences on possible associations of phonological, syntactic, and semantic structures. The basic passive constraint-satisfaction mechanism remains the same. Morphemes can be seen as similar to cooccurrence constraints in phonology: an English phonological constraint might say that if a representation has [+round] it must also have a [+back] on the same segment; the morphemic version would say that if a ⟨phonology, syntax, semantics⟩ 3-tuple has a semantics that looks like *bake'* it must also have a phonology that looks like /beek/, or else the 3-tuple is not a legal linguistic object of English. The formal description language will allow us to spell out exactly what looking like /beek/ involves.

When we acknowledge the representation/description distinction, and accept that morphemes are descriptions rather than representations, the problems concerning two morphs living in the same place dissolve. In a framework where morphemes are pieces of phonological representation that obey the Physical Integrity of Morphemes assumption, it is hard to imagine how a single node can “belong” to two different morphemes. In a constraint-based morphology, there is no problem. The same node can easily satisfy two different morphemic constraints simultaneously, just as it presumably satisfies segmental content and syllabification constraints simultaneously.

Chapter 5 will discuss the general ways in which more than one morphemic constraint apply to a representation, resulting in effects like concatenation, infixation, and templatic morphology.

1.5 Formalism

A major psychological disadvantage awaiting those who approach much of the research published within constraint-based (or unification-based) approaches to phonology is that it is bristling with formalism, much of it unfamiliar to the large majority of phonologists. The papers seem dense, difficult to read, and the central ideas can get lost in the forest of strange symbols. In some sense, this is perhaps a cosmetic problem with the field. It might be possible to lessen the forbidding demeanour and make phonologists' initial exposure to the central ideas easier. Unfortunately, this cosmetic problem is not the only difficulty in reading work in constraint-based phonology, and even if it were removed, a larger problem would remain: the ways of thinking required by the central ideas of

⁸A fair amount of effort has gone into trying to demonstrate that these suspiciously convenient planes are needed by phonology even in cases that do not involve separate morphemes, and thus are independently motivated mechanisms of phonological theory (e.g., McCarthy 1986, 1989a).

the research programme are often incommensurable with the ways of thinking that are usually at work in mainstream generative phonology.

This second, non-cosmetic source of difficulty is largely unavoidable. It is not easy to stop thinking of phonological representations as things that are actively built by rules in the course of a derivation that takes place in (perhaps abstract) time, and to start thinking of representations as things whose provenance is unknown (and irrelevant), that are born fully-grown and are subjected to legality-checking constraints whose only power is to accept the entire representation or throw the entire representation away. Difficult as it may be, it is absolutely necessary to change to this way of thinking (at least temporarily) in order to understand at all what constraint-based phonology is about. I do not know of any rhetorical techniques that can make this transition easier.⁹

I shall try as much as possible to avoid the first, cosmetic source of difficulty—a dense, unreadable text resulting from the constant use of formal symbols of questionable utility. A glance through the rest of this dissertation will show that I have not been entirely successful. The worst of the formalism has been exiled to an appendix, which the reader may skip without missing the main ideas. I have tried to keep down the level of unfamiliarity by not inventing a new language to formalize the constraints, instead taking over wholesale first-order logic, with which most linguists have at least a passing acquaintance. I hope the abbreviations I have chosen will make the formulae easier to read, at least once the reader is familiar with them—though $v_1 \xrightarrow{n} v_2$ may at first appear more daunting than $arc(v_1, n, v_2)$, I believe it expresses more iconically the proposition that a representation has a line labelled “n” joining v_1 and v_2 .

Even these hopefully straightforward formulae will be avoided as often as possible. Given this policy, some justification is in order for the large number of such formulae I have chosen to keep. Researchers in constraint-based frameworks are often accused of using “formalism for the sake of formalism”. There are, however, two excellent reasons for the liberal use of a formal apparatus in spelling out exactly what the claims of one’s theorizing are. (Admittedly, we have often allowed these excellent reasons for using formalism to outweigh the aforementioned good reasons not to.) I take them up in turn.

1.5.1 Avoiding ambiguity

First, formalism used judiciously helps linguists avoid ambiguity when stating the content of their theories. Linguists are in the business of making claims that certain facts about


⁹Indeed, the reader may notice that I have not entirely made the transition myself. I often speak of constraints applying in quasi-temporal terms (“Constraint A forces X to be true of a representation, but this then triggers constraint B which forces Y to be true, which then...”). This is not because I am trying to sneak derivations in through the back door to handle something the purely constraint-based model is unable to account for. The discussion could always be reframed in terminology that is strictly constraint-based, though perhaps much more tedious to read. These lapses occur simply because I still often feel that I have a better intuitive grasp of what is going on if I can explain it to myself step by step, as if each step occurred after the other.

languages follow as consequences from rules or stipulations or, even better, from the very architecture of their theories. But exactly what these rules, stipulations, or architectures are often not spelt out clearly. Other linguists trying to evaluate the claims can only rely on intuitions about whether it is plausible that some conclusion could follow from a rather vague set of premises. This situation is not a problem—until different linguists start having different intuitions.

Many syntacticians formalize their accounts enough that questions about consequences can be answered more or less objectively by generally agreed-upon rules of inference. Faced with disbelievers who cross their arms and stubbornly insist “A should be able to govern B in that tree,” a syntactician can often answer, “No it can’t. That follows from my definition of government on page 63.” If the syntactician has done her job correctly, it should be undebatable whether or not the definition on page 63 actually applies. Even if the situation is so complex that intuitions begin to break down, an answer can be arrived at by a more or less mechanical application of inference rules that even the stubborn disbelievers should agree with.


This ideal is seldom the case in phonology. Claims are usually evaluated almost entirely by intuition. As already said, this is not a problem unless different linguists have different intuitions. Phonological intuitions are largely based on spatial relations and geometry, and, since phonologists are born into the same almost-Euclidean universe endowed with comparable capacities for visual perception, our intuitions about simple cases seldom disagree.

For example, a main workhorse of autosegmental phonology, the ban against crossing association lines, is seldom given formally (but see Sagey 1988a, Bird and Klein 1990). Rather it is presented as a diagram of the unwanted situation, with a star in front of it:



Whether or not a particular phonological representation violates this constraint is a judgment phonologists make by drawing on their topological intuitions about spatial relations. But there is no way of ruling out cases that violate the unspoken assumptions that the intuitions are based on. For example, it is assumed without discussion that the association lines in (1.7) are perfectly straight. But this is not uniformly assumed elsewhere in phonological theory. Hayes (1989), for instance, proposes that the following may be a possible representation of a syllable like *yi* that does not violate the Obligatory Contour Principle:


(1.8)



Regardless of whether or not one believes this is the correct structure for *yi*, most phonologists have not rejected it out of hand on the grounds of some supposed geometric incoherence.

Now, if association lines do not have to be completely straight, Bird (1990) asks, only partly tongue-in-cheek, then exactly how crooked can they be? Can we avoid violations of the No Crossing Constraint by having representations like (1.9)?

(1.9)



Obviously, Bird does not think that (1.9) is a possible phonological structure that plays any role in natural language. Neither do I. Our intuitions in this matter agree with the rest of the phonological community. But they are only intuitions. There is no principled way of excluding (1.9). If we were faced with disbelievers who crossed their arms and stubbornly insisted that the structure in (1.9) was possible, we would have no definition on page 63 that we could use to change their minds.

Concerns about formalism might seem misplaced when we look only at trivial examples like this. Intuition is clearly enough. But very little in modern phonology is this trivial. Representations are much larger and more complex than (1.9), usually built on several different planes by intricate derivations through time that rely on the subtle application of several constraints and conditions. In the face of the vaster degree of complexity that phonologists usually have to deal with, topological intuition begins to break down.

I shall look at one case where reliance on an intuitive notion of crossing association lines has led a phonologist to make claims about a theory that may or may not be incorrect, but are at the very least questionable and impossible to evaluate. I hope that by choosing as scapegoat Bruce Hayes, a researcher whose work is exemplary, it will be clear that the following discussion is not about the failings of any individual phonologist (with the implication that the problem could be avoided by anyone who was simply more careful) or any individual theory. Rather, it is about the dangers of a way of theorizing that relies too heavily on intuition.

The phenomenon Hayes (1989) addresses is a problematic asymmetry: deleting a vowel can cause compensatory lengthening of the preceding vowel, but apparently can never cause compensatory lengthening of the following vowel. Hayes wants to explain this asymmetry as a logical consequence of the representations used in moraic theory and the No Crossing Constraint. The possible failure of Hayes' explanation was pointed out by Jean-Roger Vergnaud, and has been discussed in more detail in Schindwein (1989).


There are several cases of a deleted vowel causing compensatory lengthening in the preceding syllable, e.g., the Middle English historical change in (1.10). But there seem to be no cases of the mirror image process in (1.11).

(1.10) talə → ta:l

(1.11) əla → la:


Hayes (1989) derives the forms by the following (historical?) changes. *tal* begins with the expected moraic representation:

(1.12)




The word final schwa is deleted:

(1.13)




This causes parasitic delinking (and deletion) of the syllable node, though the mora node remains on the moraic tier:

(1.14)



The free mora now reassociates to the *a* of the first syllable.


(1.15)



Finally, the former onset *l* docks as a coda consonant onto the second mora.


For *la*, on the other hand, the initial representation is:

(1.16)




Schwa deletion and parasitic delinking will result in:

(1.17)



Unlike *tal*, however, the newly free mora here cannot simply dock onto the vowel of the neighbouring syllable because, so Hayes claims, this would result in crossed association lines:


(1.18)




This is exactly the point where problems are caused by having to rely on intuitions about multi-dimensional geometry. Because the σ -to-*l* association line and the attempted μ_i -to-*a* association line both seem to occur on the same plane in the diagram, that is, on the surface of the piece of paper, Hayes assumes without comment that they must

likewise occur in the same plane in the abstract phonological representation. But nothing in his formalism requires this, indeed much in his formalism militates against it.

Let us make the situation clearer by diagramming the structures from a different perspective and introducing an angle at the moraic tier. An onset-less bimoraic syllable would look like:




Hayes argues that the root node of the consonant attaches directly to the syllable node. Taking him literally at his word, we make the association line between σ and the onset consonant go *directly* between them, without stopping off for a rest at the moraic tier:¹⁰




Given this structure, we can see that the σ -to- l association line and the μ_i -to- a association could not possibly cross in the derivation of la from la , so that the derivation could not possibly be ruled out by the No Crossing Constraint.

¹⁰It might be objected that the association line in this diagram is not on a plane. There is no reason to suppose this. Without the extra stipulation that planes are perfectly flat and the strange supposition that the human linguistic capacity includes the whole apparatus of Euclidean geometry, there is nothing to prevent two different planes being bounded by the same two lines, either in Hayes (1989) or as far as I am aware in any other work in autosegmental geometry. Extending Morris Halle's analogy, phonological representations would not resemble a spiral notebook so much as a spiral notebook with some of the pages stuck together along the outside edges.



Hayes' explanation of the vowel-loss asymmetry works only if the association lines joining syllables and root nodes are crucially forced to intersect with the moraic tier. Hayes needs a structure like the following, where I have marked the intersection of the association line and the moraic tier with an ω :



Even this is not enough to rule out the $la \rightarrow la$ derivation. We also need to ban any metathesis of μ and ω , strongly suggesting that ω has a certain degree of independent reality. This in turn raises serious questions: what is the nature of this ω ? Why do μ s contribute to syllable weight while ω s do not? Why can we not have reduplicative templates like XXX, where X ranges over ω and μ ? It is not hard to see that this interpretation of moraic theory is no moraic theory at all. All the desirable properties that were supposed to follow for free from segregating onset slots and morae no longer follow for free once we allow onset surrogates (ω s) back onto the same tier as morae.

So Hayes' partly-formalized version of moraic theory is ambiguous. It leaves open two interpretations, one with two planes, one with a single plane plus ω s. The latter has essentially no explanatory superiority over theories with a single skeletal tier. The former keeps some of the desirable properties, but is unable to derive the claims Hayes has made concerning compensatory lengthening.

The main point of this discussion has not been the inadequacy of moraic theory. Rather, it has been the inadequacy of a definition of moraic theory, or indeed any theory, that relies too heavily on informal intuitions. Hayes' article does not give us the unambiguous statement of a framework that would be needed in order to evaluate it objectively. If Vergnaud, Schindwein, and I were to cross our arms and stubbornly insist that we were right, Hayes would be hard pressed to convince us otherwise.

Formalism is not a panacea. The problem with Hayes' account is not that it is clearly wrong, but that it is impossible to tell whether it is wrong. If Hayes had offered a rigorous definition of association line, of tier, of what it means for lines to cross, of what it means for an onset to be connected directly to σ , of what exactly the No Crossing Constraint prohibits, then it is not a foregone conclusion that he would have avoided the alleged problem discussed above, nor is it a foregone conclusion that anyone would have noticed it. But we would be assured that if anyone *did* claim to notice a problem, there would be a clear and objective way to decide who was right. Bruce Hayes would be able to say, "But that follows from my definition of association line on page 63, and here's how..." and Vergnaud, Schindwein, and I, being rational people, would be forced to agree.

Without doubt, this dissertation also fails to meet the standards implied above. It is certain that at least some of the analyses of individual languages will contain inconsistencies, as might the set of constraints I propose as universals. It is also possible that sometimes it will not be decidable whether one of the informally sketched analyses could be made to work. These are unarguably failings. Until they are remedied, I have proposed no more than a research programme.

1.5.2 The representation/description distinction

Mark Johnson (1988, 1991) has argued for a firm distinction between linguistic objects (or linguistic representations) and descriptions of those linguistic objects.¹¹ In the phonological domain, the linguistic objects are the familiar phonological representations, the things with nodes, features, tiers, and association lines. Constraints are descriptions that any phonological structure must conform to in order to be well-formed. Johnson also introduces two quite different ways of expressing objects and constraints. Objects can continue to be represented by drawing them on the page as we are used to doing. Constraints, though, are framed in a formal description language based on first-order logic.

Adopting this distinction, and the accompanying notational differences, would have a very visible effect on papers written about phonology. We have heretofore made no notational distinction between representations and constraints, between object and description, but have used exactly the same kinds of drawings for both. Opening a phonology article to a randomly chosen diagram, one must usually search through the surrounding text to determine if the diagram is supposed to stand for a representation (underlying, surface, or somewhere in between), a subset of a representation calling attention to only those details that are of interest, a constraint or a condition, or perhaps neither constraint nor representation but just a type of situation that may arise in the course of a derivation.

¹¹In the context of attribute-value structures and unification-based approaches to grammar, the distinction was first made by Kaplan and Bresnan (1982).

The advantage of using a (perhaps unfamiliar) formalism to express constraints lies in constantly reminding the reader of some of the fundamentally different claims being made by constraint-based approaches to phonology. There would likely be little confusion if we restricted our attention to the kinds of constraints already widely used in phonology. There is little danger that a phonologist would mistake a picture of the No Crossing Constraint for a piece of representation that is hauled out of storage and inserted into the representation under construction at the appropriate point in time. But constraints play a much wider role in a constraint-based framework. Specifically, the lexicon is also built out of constraints. Phonologists will often assume that a lexical entry contains, among other things, a piece of representation that can be hauled out of storage at the appropriate time and inserted into the larger representation under construction. In contrast, it is argued here that lexical entries consist entirely of constraints on the possible pairings of phonological and semantic structures (or triples of phonological, syntactic, and semantic structures). A phonological structure (in conjunction with syntactic and semantic structures) can *instantiate* or *satisfy* a morpheme, but it cannot *be* the morpheme. This is a different conception of morpheme from the usual ones. Though theoretical discussions of morphology will often make just this distinction, it is usually not reflected in the day-to-day practice of phonology. The distinction can be easily forgotten if the constraints that made up the morpheme were expressed using the same sorts of diagrams used to express representations. The use of description language formalism to express lexical constraints may result in less clarity about what the constraints do, but should result in more clarity about what they actually are.

As an aid to visualizing the effects of the constraints, lexical entries discussed in this dissertation will as often as possible be accompanied by a representation-like diagram, to which no theoretical status should be attributed. The actual claims made about the nature of the lexical entries are to be found in the constraints themselves, which will be framed in the formal description language introduced in chapter 3. I hope in this way to have the best of both worlds. The diagrams should allow an easier understanding of what the constraints do, an understanding that might not have been as possible if only pure formalism had been used. But the presence of the formalism as well should serve as a constant reminder of the different way of thinking about phonology and morphology that is demanded by a constraint-based framework.

1.6 About this dissertation

Almost all of the ideas discussed in this dissertation have already been argued for by other researchers. The pieces may never have been put together into precisely this constellation before, but the pieces for the most part already exist. The largest part of this dissertation will be devoted to showing that these pieces taken together result in a coherent and powerful approach for explaining the sound structure of language.

Along the way there will be a handful of novel proposals. Some of these are: vowel height and consonant frication form a natural class, and are represented by the same

feature, which shows the expected properties of a feature, e.g., it can spread. Reduplication and autosegmental spreading (as found in vowel harmonies, for example) are not separate grammatical mechanisms, but manifestations of the same underlying principle. The government relations between nodes can be underspecified, with effects that in a more standard framework would seem to involve a piece of “underlying representation” alternating between an association line and a tier adjacency line.

How this dissertation differs from previous work in constraint-based phonology is in trying to deal with as wide a range of phenomena as possible within a unified model. Though I may not offer a convincing account of any of them, I hope I have at least suggested some of the ways that are open to a constraint-based framework for dealing with many of the topics that have been of the most interest to phonologists.

I also hope that this dissertation can begin to fill what I feel are some real needs in the theory of syntax. Other researchers have made great strides toward a theory of the phonological component that would be compatible with syntactic theories like Head-driven Phrase Structure Grammar (e.g., Bird and Ellison 1992). But I am aware of nobody who has put forward a rigorous constraint-based morphology compatible with many of the assumptions of GB. With Chomsky’s recent proposal of a “minimalist” programme (Chomsky 1992), it seems an opportune time. The assumed architecture of the syntactic component is more monostratal than ever. But at the same time, proceduralism is regaining force, with increasing reliance on constraints on derivations and a theory of morphological “spell out” that accepts too uncritically the conclusions of Bromberger and Halle (1989). There is a need for a coherent, non-derivational picture of the relationship between syntax and phonology, before a large number of researchers commit themselves almost irrevocably to the assumption that syntax and morphology can only be done with constraints on derivations rather than constraints on representations. I hope the present work can be useful in suggesting some steps toward such an alternative.

Outline of the dissertation

Chapter 2 outlines the model that segmental structure of will be adopted in this dissertation. The model is somewhat different from the usual hierarchical organization of binary features of autosegmental phonology. It is rather more similar to proposals of Gorecka (1989) and Browman and Goldstein (1989, 1990) where segmental structure represents more directly some of the properties of articulatory constriction gestures.

Chapter 3 sketches the basis of the rest of the formal system. As chapter 2 discussed the model of sub-segmental structure, section 3.1 will discuss the model of phonological structure above the level of the segment, including proposals for syllabic structure that incorporate many of the ideas of Government Phonology. Section 3.2 takes up Johnson’s distinction between representations and descriptions, and outlines a formal language for describing representations in which phonological constraints may be written. I explore some of the implications of the representation/description distinction in accounting for

autosegmental phenomena (e.g., the behaviour of tones) without the need for multiple tiers. While acknowledging that many languages require information about tones and segments to be segregated, I argue that this segregation is accomplished in the descriptions and not in the representations—two different symbols of the description language can refer to the same node in a representation. Chapter 3 also contains the beginnings of a proposal for prosodic and metrical structure and an analysis of Rotuman metathesis that does not rely on separate planes for consonants and vowels. In section 3.5 I sketch briefly some properties of the phonetic component and the phonology-phonetics interface. I argue that some apparent counterexamples to a monotonic framework are not in fact phonological processes at all, but phonetic effects. Using the ideas of section 3.5, I discuss some properties that a phonetic effect should have and show that Chumash sibilant harmony (often argued to be evidence for feature-changing non-local spreading) displays all of them.

Chapter 4 is devoted to exploring the implications and applications of the theory of phonological locality that was introduced in chapter 3. It is proposed that two nodes can be required to share their features only if they stand in the right type of government relation (or if, recursively, their respective parents or grandparents or greatⁿ-grandparents stand in the right type of government relation). These principles defining locality are illustrated with the vowel harmony systems of Hungarian, Kalenjin, and Pasiego Spanish. The final section shows how many types of reduplication can be handled using exactly the same principles that govern vowel harmony.

Chapter 5 deals with morphologically complex linguistic structures, specifically with the question of what happens when morphemes combine. In other areas of constraint-based grammar, there is no need to be concerned with constructing legal structures, only with judging the legality of candidate structures. The same situation holds in morphology. There is no need for a module of grammar to actively combine the phonological contents of two morphemes. All that is needed is the ability to judge a candidate structure to see if both morphemes are satisfied and if the boundaries are aligned in accordance with specific and general constraints. I propose a set of simple alignment constraints. Each possible choice from within the limited degrees of freedom allowed by these alignment constraints will result in a different type of morpheme combination: prototypical concatenation, infixation, templatic morphology, etc. I shall also deal with the interaction between morphology and prosody. An extended examination of determiner clitics in Nisgha will illustrate both the interaction between morphology and prosody and the analysis of morphological overdetermination.


In chapter 6, an extended analysis of the templatic morphology of Moroccan Arabic will further illustrate the morphological principles introduced in chapter 5 and the theory of empty nuclei borrowed largely from Government Phonology. The complex consonant clusters and apparent resyllabifications of Moroccan Arabic are handled without the need for such procedural apparatus as epenthesis or syncope rules.

Chapter 7 compares the framework presented in this dissertation with other approaches to phonology (or grammar more generally) that are also concerned crucially

with constraints. These include Government Phonology, Optimality Theory, Harmonic Phonology, the Theory of Constraints and Repair Strategies, Autolexical Syntax, and Declarative Phonology as exemplified by Bird (1990) and Scobbie (1991).

I conclude with by considering briefly some of the implications of the proposals, such as the fate of underlying representations and how a constraint-based grammar can itself be constrained.

(2.2)



In the matrix versions of the diagrams, this will be represented with an embedded matrix.

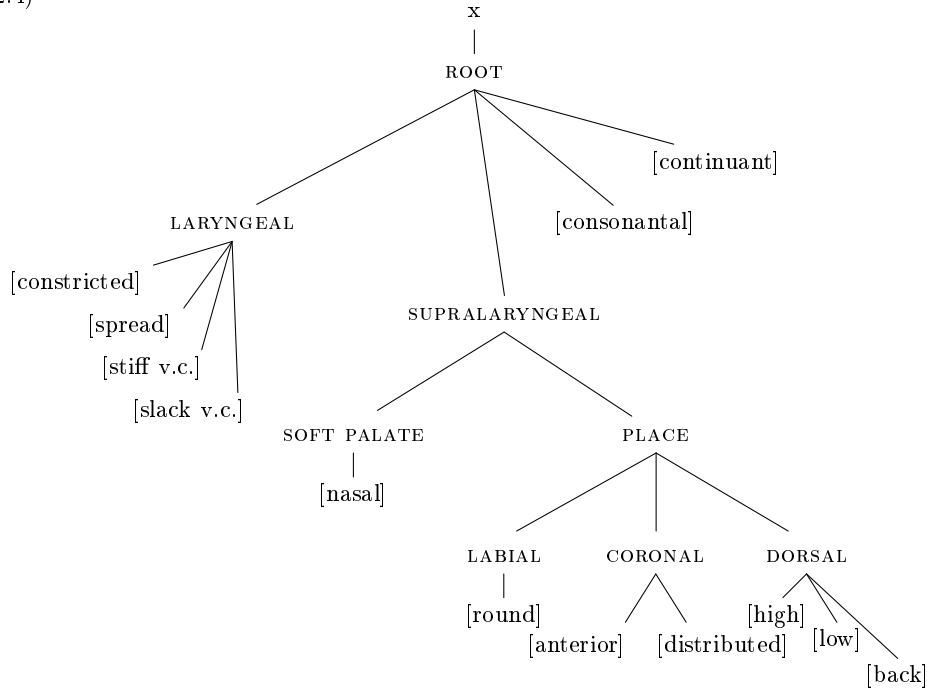
$$(2.3) \quad \left[\begin{array}{l} \text{a: Tongue Body} \\ \text{s: Velar} \\ \text{d: Closure} \\ \text{2: } \left[\begin{array}{l} \text{a: Lower Lip} \\ \text{s: Upper Lip} \\ \text{d: Approximant} \end{array} \right] \end{array} \right]$$

In section 2.1, I outline the major points of Gorecka's (1989) model of constrictions. In the next section I discuss some of the recent models that have relied on dominance relations rather than sisterhood relations to represent aspects of segmental structure such as secondary articulations. In section 2.3, I discuss aspects of the degree of closure of constrictions, especially those models such as Browman and Goldstein's that represent it as an inherent part of each constriction. I also sketch the use Steriade (1993a–d) has made of degree of closure specifications in an account of complex segments. Section 2.4 offers a synthesis of these various ideas. Proposals for the representation of consonants and vowels are the subject of the following two sections. Perhaps one of the most novel proposals of this chapter is that fricatives and high vowels from a natural class and are represented by the same degree of articulation interpreted differently for consonants and vowels. The final section offers evidence from Japanese affrication for this proposal.

2.1 Gorecka's constriction model

Most researchers in the American autosegmental tradition have assumed a version of the feature hierarchy that resembles more or less closely that of Sagey (1986):

(2.4)



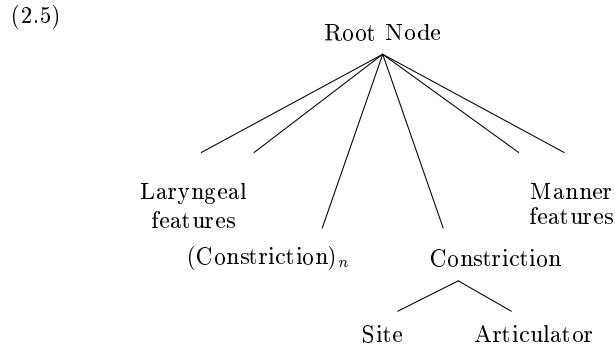
Phonologists differ concerning which of the non-terminal nodes they believe exist, and in the position of some of the features, especially the manner features. But one of the constants throughout all the proposals is the existence of a place node that dominates, at least, nodes for each major active articulator of the mouth: Labial (the lips), Coronal (the tongue tip and blade), and Dorsal (the tongue body).¹ For many languages, this simple partition into three major class nodes is enough to handle all the phonologically relevant differences in place of articulation in the consonant inventory. Where the simple division has not been sufficient, the finer distinctions have been handled by adding further dependent features (usually binary) underneath these major nodes. For example, if a language contrasts dental and palatal consonants, this will be represented by the feature $[\pm\text{anterior}]$ underneath the Coronal node. Vowels, which following Sagey are usually assumed to be Dorsal,² similarly have their differences encoded by the dependent features $[\pm\text{high}]$ and $[\pm\text{back}]$.

¹A common addition to this set is Radical (the tongue root). See, e.g., McCarthy 1989b.

²The major exception to this are front vowels, which many hold to be Coronal, e.g., Mester and Itô(1989), Pulleyblank (1989), Clements (1990), Lahiri and Evers (1991).

What each of these treatments seems to be missing is any principled treatment of the *site* of a constriction. Sageyan models are designed to specify easily which active articulator is involved in a constriction, but treat exactly *where* the active articulator forms the constriction as a secondary property, leaving no choice but to represent sites with seemingly ad hoc dependent features. This is in marked contrast to the taxonomy that underlies the IPA consonant chart, where the site of a constriction is one of the most important organizing principles.

The fundamental insight of Gorecka (1989) was that both the articulator and the site of a constriction need to be specified, and that many analyses become much clearer once this is done. Gorecka's (1989: 112)³ proposal for the hierarchical organization of features is:



The Articulator and Site nodes in (2.5) can have the following values:

(2.6)

Possible Articulators:	Possible Sites:
Lower Lip	Upper Lip
Tongue Blade	Anterior
Tongue Body	Palatal
Tongue Root	Velar
	Pharyngeal

(The Tongue Blade articulator includes the tongue tip.) Much of Gorecka (1989) is devoted to showing that each physically possible combination of Articulator and Site is used phonologically in some language. As well, each Articulator and Site specification shows behaviours expected of any autosegment, such as assimilatory spreading and dissimilation under pressure from the Obligatory Contour Principle.

³Future references to Gorecka (1989) will be of the form “PA 112”.

Before continuing with general discussion of the model, I shall mention, following Gorecka's discussion, some of the segments that can be found inhabiting each of the combinations of Articulator and Site.

Lower Lip — Upper Lip

bilabial sounds: *p, b, m*,
and glides: *u, o, ü, w*.

The site of the labial constriction can be significant. Before the passive morpheme /-wa/, bilabials (Upper Lip site) but not labiodentals (Anterior site) are velarized in Venda and palatalized in Zulu (PA 41–52).

Lower Lip — Anterior

labiodental sounds: *f, v*.

Anterior Site segments can form a natural class. In Standard Thai labiodental *f* patterns with anterior coronals rather than with bilabials in undergoing velarization before *i* (PA 57–9).

Tongue Blade — Upper Lip

linguolabial sounds, found in some Austronesian languages (Maddieson 1988).

Tongue Blade — Anterior

standard anterior coronals: *t, d, s, z, n*, etc. Gorecka assumes that the feature [±distributed] depends on the Tongue Blade articulator, so this is also the constriction specification for laminal sounds such as .

Tongue Blade — Palatal

palatal coronals: *c, č, š, ž, ç*, etc.

Note that there is no more need for the SPE feature [±anterior]. It has been replaced by the distinction between two sites: Anterior and Palatal.

Tongue Body — Palatal

non-low front vowels: *i, e, ü, ö*.

In this way, Gorecka captures both the Dorsal properties of front vowels, which come from the Tongue Body articulator they share with other dorsal segments, and their “Coronal” properties, such as triggering palatalization, which come from the Palatal site, which they share with palatal consonants. Feature hierarchies without Site specifications that wanted to capture this relation between front vowels and palatal consonants (e.g., Lahiri and Evers 1991) could only do so by arguing somewhat counter-intuitively that they shared a Coronal major articulator.

Tongue Body — Velar

velar consonants: *k, g, x*,

Tongue Body — Anterior

Gorecka argues on the basis of both acoustic and phonological evidence that some languages (especially those with more than two high vowels) can have front vowels with

constrictions in the dental-alveolar region and which should be represented as having Anterior site. E.g., Standard Thai, perhaps Russian and Polish. (PA 138–145)

Tongue Body — Pharyngeal

æ, α


Gorecka argues on phonetic and phonological evidence for two types of “a”, one formed in the upper pharynx with the tongue body (the dorsum) and subject to tongue body harmonies, often transcribed /α/, and one formed in the lower pharynx with the tongue root and neutral to such processes, the more central /a/. (PA 173–6). Front /æ/ involves the same tongue body–pharyngeal constriction as /α/.

Tongue Root — Pharyngeal

/a/, that is, the /a/ referred to in the last paragraph. This combination is also the secondary constriction in pharyngealized segments. As well, it plays a role in uvulars, which Gorecka analyzes as having two constrictions—one tongue body–velar and one tongue root–pharyngeal—which merge to give a uvular constriction.

Gorecka’s constriction representations for vowels bear closer examination, since they are less immediately obvious and more controversial than the consonantal representations.

/i/ is represented with a constriction articulated by the Tongue Body at the Palatal site.⁴ It also has the dependent feature [+high] attached to the Tongue Body node:




/u/ usually has a velar constriction.⁵ There is also a secondary labial constriction.

⁴except for those languages in which /i/ is an anterior vowel as mentioned above.


⁵though in some languages, such as German, Gorecka argues for a [+high] Tongue Body constriction with an (upper) Pharyngeal site.

(2.8) /u/




Mid vowels are usually like their corresponding high vowels, but with the addition of a Tongue Root constriction at the Pharyngeal site and [-high] rather than [+high] as the dependent feature. The difference between /e/ and /ɛ/ or /o/ and /ɔ/ lies in the feature [±tense] or perhaps [±ATR], whose exact identity or position in the hierarchy Gorecka (1989) does not address.

(2.9) /e/



The simple low vowel /a/ is formed by a Tongue Root–Pharyngeal constriction. But the other /a/-like vowels, /æ/ and /ɑ/, though they still have Pharyngeal site, have the Tongue Body as the articulator rather than the Tongue Root.

(2.10)



There are a couple of criticisms that can be made of Gorecka's model. One of them is the continuing need for so many binary features dependent on Articulator specifications. We saw that the work of the SPE feature [±anterior] could be done by Site features in a sufficiently worked out constriction system. Many of the remaining

binary features in Gorecka's system seem to be the result of not taking the constriction model as far as it could be taken. For example, I shall argue later that the need for the dependent feature [\pm distributed] is caused by not having an adequately fine-grained set of Articulator specifications.



Another weakness is the model's inexplicitness concerning its relation to some fundamental principles of autosegmental phonology. While there is some excellent motivation for the segment-internal hierarchy shown in (2.5), there is little discussion of what this organization means for relations *between* segments. Does each value of an Articulator or Site specification define its own tier? What tier does the constriction node live on? More pressingly, what tiers do the optional extra constriction nodes in a segment live on?⁶ These kinds of questions are especially crucial in a theory that wants to account for dissimilation effects as the result of the Obligatory Contour Principle operating on adjacent items on a tier. The inexplicitness in defining the relations between constriction nodes and between a constriction node and the segment's stricture features (represented under the root node) also makes it difficult to determine what claims the theory makes about the differences in representation between, for example, a labiovelar stop, a labialized velar stop, a velarized labial stop, and a click.

2.2 Dominance models of secondary articulation


Not all researchers have assumed that the relevant features for vowels are related to each other as sisters (or some degree of cousin) under a common major articulator node. Another line of thought, starting with Archangeli's (1984, 1985) analysis of Yawelmani vowel harmony and continuing with Mester (1986), has explored the possibility that features can depend on other features.

Ngbaka has a requirement that all vowels in the same morpheme that agree in height must also agree in backness and ATR (Churma 1984). Mester (1986) analyzes this by having [back] and [ATR] specifications dependent on the specifications for [high] and [low]. If there are two different heights, as in *pEpu* in (2.11), each may have a different [back] specification. If two vowels are linked to the same height feature, as in *liki* in (2.12), they are also necessarily linked indirectly to the same backness feature.

⁶Gorecka (1989:120) remarks that "an instruction to spread the Constriction Node will result in the spreading of all qualifying nodes", but there is no discussion of what kind of rule it is that can spread several pieces of a tree at the same time, how this rule knows what qualifies and what does not, and where exactly the qualifying nodes spread *to*.



The only way there could be two backness feature for a single height would be to have the structure in (2.13), which Mester rules out as a violation of the Obligatory Contour Principle:




Mester argues that the dependency relations of tiers need not be universal, but may be parametrized. He offers an analysis of Ainu dissimilation in which [high] is dependent on [back], the opposite of Ngbaka. Others have not accepted this aspect of tier-dependency models, arguing instead for some particular universal geometry of dependency relations (e.g., Goad 1991, 1993).

Not all frameworks that rely on dominance use SPE-style binary features, nor do they all use their hierarchical organizations solely for the purpose of suggesting what is and is not a natural phonological process. Dependency Phonology, for example, has long used different dependency relations between primitive elements to represent phonetic differences between segments (Anderson and Ewen 1987). It would be plausible that in a framework like Mester’s, a [-high] dependent on a [-back] would have the same interpretation as a [-back] dependent on a [-high], though we would never find both in the same language. In Dependency Phonology, the difference in dominance between A→I and I→A can be used to distinguish /ε/ and /e/ within the same language. The problem with such systems is that as the number of primitive elements

grows, and the number of dependency contexts in which they can occur grows, the compositionality of the interpretation principles breaks down.⁷

An interesting hybrid of autosegmental and privative dependency frameworks is offered by Selkirk (1991). She proposes that the major class nodes of Sagey (Labial, Coronal, and Dorsal) with the addition of Pharyngeal (or Radical) should be treated much the same way as privative primitive elements in frameworks like Dependency Phonology, Government Phonology, or Particle Phonology (Schane 1984), and should be allowed to enter into dependency relations with each other. A discussion of Selkirk's proposal can show the problems that are inherent in all similar privative dependency models and autosegmental models. Her proposed process for the coalescence of *a* and *i* into *e* is a good place to start the discussion:



The Pharyngeal node in (2.14) delinks from its root node in step (i) and reassociates to the Coronal node of the segment to its right, in the process changing its status from head to dependent. The Pharyngeal constriction at first glance appears to have jumped tiers, moving one storey downstairs. The first problem is the nature of the tiers that allow Pharyngeal to spread in this way, especially given that a segment may contain more than one Pharyngeal node. Selkirk devotes a fair bit of attention to the question of tiers, especially in Selkirk (1988), but the complexity of the required mechanics seems out of proportion to the complexity of the phenomena they are supposed to handle. It is likely that any hierarchy that allows or requires the presence of more

⁷The phonetic interpretation of structures in Dependency Phonology cannot be fully predicted from the values of the elements and general principles. Government Phonology offers a limited compositional semantics for its vowel representations (Kaye, Lowenstamm, and Vergnaud 1985), though there is no simple way it can be extended to consonants. Van der Hulst (1990) perhaps has the most consistent interpretation principles. (Strictly speaking, van der Hulst's dependency relations are not tier-dominance relations like Mester's.) Even in these last two cases, though, the interpretation principles only provide a translation from the formal structures of the new frameworks into SPE-style feature matrices. This may be a convenient paedagogical tool to help phonologists grasp the intent of the formalism, but the SPE-style "interpretations" themselves will eventually have to receive interpretations in terms of the properties of articulatory and acoustic events, which, judging from the volume of literature in phonetics, is no trivial task.

than one node of a given sort within a segment will suffer from this problem, whether the duplicates are sisters like Gorecka's constriction nodes or potentially mother and daughter in dominance models like Selkirk's. Given the other problems that the notion of tier creates, such as the massive redundancy and potential inconsistency of temporal ordering information (cf. Bird and Klein 1990, Scobbie 1991), it seems clear that a new way is required of representing the ways in which two specifications can be "local" to each other. This question will be taken up in section 3.4 and in chapter 4.

The second problem suggested by (2.14) is the different interpretations the Pharyngeal node receives as head and dependent.⁸ Pharyngeal has no single phonetic correlate, in some positions being interpreted as a vowel height feature, in others as some sort of pharyngeal region activity: tongue root advancement, pharyngeal constriction tensing, or centralization. The fact that some of the interpretations of Pharyngeal are contradictory is a symptom of the fact that no node and no relation between nodes has a consistent phonetic correlate. Not only is the interpretation non-compositional, the same structure in different languages can have radically different phonetic interpretations. It is conceivable that a Dorsal dependent on a Coronal could be what distinguishes a velarized coronal in one language, a retroflex in another, a high front glide in a third, and perhaps a click in a fourth.

Dominance-based models have given us the idea that secondary articulations are hierarchically dependent on their primary articulations, and have often given strong evidence from phonological processes for this structural relationship. Where most of them fail is their inability to define exactly what it is that the representation of a secondary articulation contributes to the interpretation of the segment.

2.3 Degrees of constriction

Much of the weakness in Selkirk's proposal is due to the fact that she specifies only the active articulator and has no principled way of representing the site of a constriction gesture. But part of the problem, which is shared by Gorecka's work, is that there is no principled way of relating the part of the representation that encodes the place of the constriction gesture to the stricture features under the root node that control many other crucial properties of the gesture. The information that a segment is a stop can be quite far away from the information on where the stop is to be formed.

The inability to capture this relation has been a failing of feature hierarchies since Clements (1985) and Sagey's (1986) work on complex segments. Like Clements, Sagey segregated the features related to the degree of aperture of a constriction from those related to the place of the constriction. This has no major consequences, as long as there is only one constriction per segment and the relation between stricture and place is unambiguous. But once we try to represent complex segments using feature

⁸In this particular case the problem could be avoided by also interpreting the Phary of /a/ as simple lowness, but this way out is not always available.

hierarchies, there is an indeterminacy in how the stricture features, usually represented only once under the root node, are to relate to the place specifications, of which there can be several in a segment. Sagey's solution, drawing an arrow from the root node to the place specification for the primary articulation, is an unsatisfactory ad hoc device; at best, the arrow notation says nothing about how the non-primary places are to be interpreted in terms of stricture features.⁹

A different approach is suggested by the work of Browman and Goldstein (1989, 1990), who propose a framework based much more closely on phonetic reality and give little, if any, weight to matters of purely phonological concern. One of the levels of representation that Browman and Goldstein argue for is the **gestural score**, a more developed version of the idea discussed in Goldsmith (1976).¹⁰

The gestural score represents a group of temporally coordinated constriction gestures. Each gesture is specified not only for the articulator and the site (or "location" in Browman and Goldstein's terminology), but also the degree of the constriction, the stiffness of the articulator, and (for the tongue tip and body) the shape of the articulator. Possible **articulators** for Browman and Goldstein are: lips, tongue tip, tongue body, tongue root, velum,¹¹ and glottis. Possible values for the **location** are: protruded, labial, dental, alveolar, postalveolar, palatal, velar, uvular, and pharyngeal. The possible constriction **degrees** are: closed, critical (the degree of aperture for a fricative), narrow (the degree of aperture for approximants and high vowels), mid (high-mid vowels) and wide (low and low-mid vowels). Browman and Goldstein do not discuss **shape** settings in as much depth, but shapes are involved in contrasts between apicals and laminals and in the representation of laterals. **Stiffness** is a parameter that is more closely tied to Browman and Goldstein's modelling of physical articulatory gestures than with any phonological property. They do however mention that it may be phonologically relevant in some cases. They suggest that the difference between /i/ and /j/ and between /u/ and /w/ might be only in the stiffness parameter. The value might also be relevant in controlling such things as taps, trills, and pitch. I shall not deal with the possibility that stiffness should be represented in phonology.

Browman and Goldstein intend their gestural scores to be one level closer to physical reality than pure phonological representations, but they do not discuss exactly what they believe phonological representations are and how they differ from gestural scores. It is clear that gestural scores themselves are inadequate as phonological representations.¹² They contain much information about temporal coordination that is irrelevant to phonology, and they say nothing about differences in status that affect

⁹Sagey (1988b) argues that, at least within any given language, stricture features for secondary articulations do not phonologically contrast.

¹⁰See Bird (1990), chapter 3, for a constraint-based model of phonology that incorporates Browman and Goldstein's ideas fairly directly.

¹¹There is no location specifiable for the velum, simply two closure values—closed and wide.

¹²See Scobbie (1991:11–16) for arguments that phonology is not just "tidied-up phonetics".

which gestures are and are not accessible to phonological operations such as spreading. The geometry I shall propose abstracts away from irrelevant details of temporal ordering and represents differential accessibility to phonological operations using a gesture's hierarchical position in a system of dominance relations like those in section 2.2.

Many of the relevant ideas have also appeared in autosegmental literature. Clements (1989), for example, recognizes that vowel height often functions to define major classes analogous to obstruents and approximants. Clements represents these vowel height differences using a separate aperture node in his hierarchy, though he continues in the tradition of keeping these features that control the degree of aperture segregated from the features that control where in the vocal tract the degree of aperture features will be realized. Padgett (1991) proposes a model closer to the present one, arguing that each place specification carries its own stricture features.

2.3.1 Steriade's closure and release features

One of the more interesting uses of constriction degree features in phonology is Steriade's account of consonants, especially complex onsets or complex segments like affricates. Steriade (1993a–d) proposes that there are three phonologically relevant degrees of aperture for consonants, which she terms A_0 (complete closure), A_f (critical closure, the degree of aperture for fricatives), and A_{max} (the degree of aperture for approximants).


The power of Steriade's proposal comes from the hypothesis that each consonant involving a complete interruption of oral airflow can correspond to two aperture specifications, one A_0 for the **closure** of the segment, and another for the **release**. Three classes of segments are distinguished by the choice of whether there is a release position and which aperture feature occurs there:

$$(2.15) \quad \begin{array}{lll} (\mathbf{A}_0 \ \mathbf{A}_{max}) & (\mathbf{A}_0 \ \mathbf{A}_f) & (\mathbf{A}_0) \\ \text{plain released} & \text{affricate or} & \text{unreleased} \\ \text{stop} & \text{aspirated stop} & \text{stop} \end{array}$$

Segmental content may also be associated to the release position. The difference between an aspirated $/t^h/$ and an affricated $/t^s/$ is whether or not the release position shares the coronality specification of the closure:


$$(2.16) \quad \begin{array}{ccc} t^h & \text{COR} & t^s & \text{COR} \\ & | & & \wedge \\ & (\mathbf{A}_0 \ \mathbf{A}) & & (\mathbf{A}_0 \ \mathbf{A}) \end{array}$$

The two positions also allow us to represent four types of nasality/orality that a segment may have:



Steriade argues that pre- and post-nasalized fricatives do not exist because fricatives project only one, rather than two, aperture positions. The same sort of representations for pre- and post-glottalized segments could be obtained by replacing the nasal specification, N, in (2.17) with a glottal specification. (A doubly connected glottal closure would be a reasonable representation for an ejective. Similarly, a plausible representation for a click would be a doubly connected velar closure.)

The segmental material associated with the release position does not necessarily have to be shared with the closure position in all languages. If a release position can have segmental content of its own, the result is a complex onset. For example, an English complex onset *pr* might have a representation along the following lines:¹³



2.4 A synthesis

I propose that each phonologically relevant constriction gesture is represented by a node dominating three specifications: the articulator, the site, and the degree of aperture. In diagrams, the lines joining the gesture node and these specifications will be labelled **a**, **s**, and **d**.

There is one main difference between the sets of articulators I propose and those in the works reviewed so far. Analogous to Gorecka's division of "coronal"-related sites into Anterior and Palatal, I divide the coronal articulator itself into the Tongue Tip, involved in apical sounds, and the non-tip part of the Tongue Blade, involved in laminal sounds. If it turns out that there is a need to define coronal as a natural class, then the formalism of the "sort" introduced in section 3.3 provides a way to do


¹³The claim that fricatives have only one aperture position is obviously incompatible with analyzing *fr* similarly. I would rather abandon the claim concerning fricatives than the analysis of complex onsets.

that. But I know of no argument in the literature that addresses itself specifically to showing that apical and laminal articulators ever pattern together for the purposes of some phonological phenomenon.¹⁴

Table 2.1 lists the sets of articulators, sites, and degrees that I will be using throughout this dissertation.

I adopt Steriade's three aperture degrees, A_0 , A_f , and A_{max} , though I will use Banner Inouye's (1989) notation: **0**, **1**, and **2**. The more neutral numeral symbols are more appropriate, since I argue that they are phonetically interpreted as closed, critical, and approximant only when they are specifying consonant gestures. When they are specifying vowel gestures with a Dorsal articulator, [d:1] is the aperture of a high vowel (not a fricative) and [d:2] that of a non-high vowel (not an approximant). Whether the segment will be treated as a consonant or vowel will depend on its position in the syllabic structure.¹⁵ An example of how consonant and vowel degrees of constriction can be used is given in section 2.8's analysis of Japanese palatalization. I also propose two special degrees of closure that apply only to the Radical articulator: **A** for Advanced Tongue Root, and **R** for Retracted Tongue Root. These will be discussed more fully in section 2.6.

As in most of the dominance-based models discussed in section 2.2, secondary articulations will be represented by a node that is dominated by the node that represents the primary articulation. A simple /k/, for example, would be represented with a tongue body articulator (Dor), velar site (Vel), and a degree of complete closure (0), as in (2.19).



¹⁴There are of course many arguments that operate on the assumption that it is sufficient to show that Anterior and Palatal *sites* pattern together. The evidence that has the best chance so far of being converted into an argument for the grouping of Apical and Laminal *articulators* is probably “coronal” assimilation in Arabic. In the dialect I am most familiar with, Moroccan Arabic, there is no good reason why this phenomenon should not refer simply to the Laminal articulator (rather than the union of Laminal and Apical), but it is possible that some more eastern dialects contrast the two articulations, yet treat them identically for the purposes of coronal assimilation. (Alleged morpheme structure constraints on Semitic roots are also often assumed to be evidence for a unified coronal articulator. For arguments against the existence of these morpheme structure constraints, see Pierrehumbert (1993) and Paradis and Prunet (to appear).)

¹⁵More precisely, it will depend on a primitive property of the root node that usually (and in many languages always) correlates with syllabic position.

Table 2.1: **Gesture specifications****Sites:**

Lip	Lip	upper lip
Ant	Anterior	dental-alveolar region
Pal	Palatal	palato-alveolar and palatal regions
Vel	Velar	
Pha	Pharyngeal	

Articulators:


Lab	Labial	lower lip
Apc	Apical	the tongue tip, apical and sub-laminal regions
Lam	Laminal	the non-tip part of the tongue blade
Dor	Dorsal	tongue body
Rad	Radical	tongue root

Degrees:

	in consonants	in vowels
0	complete closure (A_0)	—
1	critical closure (A_f)	high
2	approximant (A_{max})	non-high

A	Advanced Tongue Root (ATR)
R	Retracted Tongue Root (RTR)

Here, all the specifications for the primary velar articulation are immediately dependent on the root node, which also serves as the “skeletal slot” of the segment. There is, of course, no temporal ordering relations among the different specifications of a constriction. A labialized /k^w/ would have a second constriction node dependent on the root constriction node, as in (2.20). Just as the lines between the root node and its articulator, site, and degree are labelled a, s, and d, the line between the root node and its secondary articulation is labelled 2:¹⁶



The only temporal ordering relation implied by a secondary articulation is some degree of overlap between the primary and secondary constrictions. The exact coordination of the two constrictions is a matter left to phonetic interpretation, which itself often allows considerable freedom. If a segment contains more than two constriction gestures, the secondary articulation nodes involved may be sisters. Thus, there may be more than one secondary articulation node linked to the same root node by a “2” line.

The relation that a specification bears to its constriction node (i.e., articulator, site, degree, secondary) is indicated by the label on the association line connecting them. It is *not* a function of what tier the specification is on. In this framework, there are no tiers, except that of the root nodes. Simply because two site specifications belong to neighbouring segments, this does not automatically mean that they are local to each other. Whether specifications of neighbouring segments are local to each other and subject to spreading is determined by the type of the government relation that joins the two segments. This will be discussed in more depth in section 3.4 and chapter 4.

2.5 Consonant representations

Most of Gorecka’s combinations of articulators and sites, as sketched in section 2.1, can be easily translated into the present framework with the addition of a degree specification. The velar constriction gesture for /k/ is simply [a:Dor, s:Vel, d:0], that is, a gesture with a Dorsal articulation, a Velar site, and a complete degree of closure (zero

¹⁶The [2:] that labels secondary articulations should not be confused with the [d:2] that represents an approximant or mid-vowel degree of aperture.

aperture). /f/, which Gorecka would represent as having the Lower Lip as articulator and Anterior as site, becomes [a:Lab, s:Ant, d:1], where [d:1] represents the fricative degree of aperture.

The situation is more complex for coronals. Now, as well as having two sites, [s:Ant] and [s:Pal], corresponding roughly to SPE [+anterior] and [-anterior], there are two articulators, [a:Apc] and [a:Lam], corresponding roughly to SPE [-distributed] and [+distributed]. The four possibilities given by combining articulator and site are shown in the following table.

(2.21)


	[s:Ant]	[s:Pal]
[a:Apc]	apico-alveolars true apico-dentals	retroflexes some languages' palatal fricatives probably English <i>r</i>
[a:Lam]	laminal dentals inter-dentals	palatals laminal post-alveolars

Some languages, e.g., Australian languages, use all four possibilities distinctively. The four combinations should be enough to capture *all* the phonological distinctions made among coronals in any given language. [s:Ant], for example, is indeterminate about the exact location of the constriction, as long as it is made somewhere between the bottom of the upper teeth and the alveolar ridge. No language is known to contrast place within this region without also showing a contrast in apicality vs. laminality. Similarly, [a:Apc, s:Pal] covers both those retroflexes made with the apex of the tongue and those made with the sublaminar surface. Again, no language has been shown to contrast the two types of retroflexes, though Ladefoged and Maddieson (1986) speculate that, depending on how one interprets Emeneau's (1984) description, Toda may be such a language.

Adopting Steriade's idea of closure and release positions allows us to represent more complex segments. The closure features of a segment are to be represented on the root node. The release features are specified on a second node associated to the root by a line labelled **release** or **r**. One of the few substantive differences between the release relation and the secondary or 2 relation is that the former is phonetically interpreted as involving a temporal precedence relation between the root and dependent constrictions, rather than an overlap.

The affricate t^s can be represented by a closure position specified [d:0], i.e., complete blockage of the airflow, and a release position specified [d:1], interpreted in consonants as the degree of aperture for frication. The two positions share their articulator and site specifications:


(2.22)



Recall that there is no sense in which *a* and *s* specifications are ordered on tiers or with respect to each other, so (2.22) does not involve any crossing of association lines.

Since “branching” or complex onsets are simply the result of specifying the release position with different segmental content from the root position, the English onset *fr* would be represented as:

(2.23)



2.6 Vowel representations


The tongue root articulator is different from those of the oral cavity, since it has only two phonologically contrastive degrees of constriction, though the exact narrowness of each phonological degree may vary slightly from language to language. While the oral articulators are capable of performing complete closure, or [d:0], constrictions, the same possibility does not seem to be available to the tongue root. I also argue that the tongue root has no neutral constriction-less position. Thus, the tongue root must always be at one of [d:A] or [d:R]. A language can, of course, make one of these a default. The narrower of the two, [a:Rad, s:Pha, d:R], corresponds to the pharyngeal constriction active in the vowel /a/, and is also the constriction that distinguishes Retracted Tongue Root, or RTR, segments. The wider constriction, [a:Rad, s:Pha, d:A], distinguishes Advanced Tongue Root, or ATR segments.¹⁷

¹⁷The phonetic interpretation of these two constrictions will typically be enhanced, to various degrees in various languages, by other articulatory gestures such as raising of the larynx, raising the body of the tongue, or independently narrowing the pharynx (cf. Lindau 1979). These other gestures have to my knowledge never been shown to operate independently of “pure” ATR/RTR gestures. (See Trigo (1991) for an argument for the phonetic independence of laryngeal height and an admission that the evidence for phonological independence is ten-

The role that Gorecka (1989) gives to pharyngeal constrictions in the representations of mid vowels is problematic. For /e/ and /o/ in languages that use [ATR] contrastively, it is difficult to conceive how the tongue root could simultaneously be advanced and execute a pharyngeal constriction.¹⁸ As Selkirk (1991) more or less admits when distinguishing the dominant and dependent statuses of PHAR, vowel height and pharyngeal constrictions are essentially independent properties. This idea is also implicit in Gorecka's use of the almost redundant binary feature [\pm high], whose plus value is required to represent a natural class.

In a framework that represents the degree of every constriction, there is a natural way to express vowel height distinctions. As mentioned above, degree specifications are interpreted somewhat differently for vowel gestures articulated with the tongue body than they are for consonants. [d:1] represents a high vowel, [d:2] a mid. The resulting representations for *i*, *e*, and

(2.24) *i*




uous.) I shall assume that these gestures are never more than phonetic enhancements of the phonologically relevant ATR/RTR distinction.


Because of the rather drastic differences between the closure possibilities of tongue root gestures and those made with the rest of the tongue, I have thought it best to use completely different symbols for tongue-root degrees. It would of course be possible to force the tongue-root degrees into the more general model: the narrower constriction [d:R] could be represented as [d:1] and the wider [d:A] as [d:2]. But the phonological evidence for this identification is slim. Indeed, what phonological evidence there is suggests that if any identification is to be made at all, it should be the other way round: [d:A]=[d:1] and [d:R]=[d:2]. High vowels ([d:1]) tend to correlate with ATR, mid vowels ([d:2]) with RTR. A different example of such evidence from Moroccan Arabic will be discussed in a footnote to section 2.8.

¹⁸Gorecka no longer believes that /e/ involves a pharyngeal gesture.

(2.25) e



(2.26) ε



2.6.1 Low vowels

I adopt Gorecka's proposal for the difference between the vowels /a/ and /α/. The first, /a/, is a pure tongue root-pharyngeal constriction: [a:Rad, s:Pha, d:R]. The pharyngeal constriction of the second, /α/ is made with the tongue body rather than the tongue root: [a:Dor, s:Pha]. The effects of these representations will become clearer.

In considering the proper representation of /æ/, we come across two problems that solve each other. The first is how to represent the “frontness” of the /æ/ segment without the need for the binary feature [±front] that Gorecka used. If we keep with Gorecka's claim that /æ/ and /α/ have their pharyngeal constrictions made with the tongue body ([a:Dor, s:Pha]), then we cannot simply add a secondary palatal articulation of [a:Dor, s:Pal], since this would require the tongue body to move in two opposite directions at the same time.


Another conceivable solution for representing frontness by palatality would be to say that the articulator for the palatal site in /æ/ is not the tongue body [a:Dor], but the tongue blade [a:Lam]. This is not a very satisfactory solution, as it would make /æ/ the only vowel whose [s:Pal] gesture is performed by [a:Lam], and would lead us to expect the same possibility to be open for the other front vowels. But even if the [a:Lam, s:Pal] solution were formally desirable and the [a:Dor, s:Pal] solution formally possible, there is some doubt that they would be empirically desirable. One of Gorecka's

strongest arguments for the feature [-back] is that, while /æ/ often patterns with the clearly palatal vowels in vowel harmonies, it never triggers palatalization.

The other problem raised by the claim that /æ/'s pharyngeal constriction is [a:Dor, s:Pha] is that this leaves the tongue root unaccounted for. If the tongue body is making a pharyngeal constriction, there is nothing immediately obvious stopping us from expecting that the tongue root could be performing an ATR gesture independently. But this runs counter to the well-known claim that “a”-vowels cannot be [+ATR]. Kaye, Lowenstamm, and Vergnaud (1985) have even made this into one of the central facts that is supposed to fall out automatically from the very organization of the formal system. In direct contradiction to Gorecka, these researchers argue that what gives /æ/ its frontness is exactly the same element that gives other front vowels their palatality.

The answer is: both are right.

Gorecka is right that a true low /æ/ is not palatal. Rather, /æ/ and /α/ are ATR and RTR counterparts of each other:



Kaye, Lowenstamm, and Vergnaud are right that there is a low vowel that cannot have an ATR counterpart by its very definition, namely /a/. A segment cannot have two [a:Rad, s:Pha] constrictions, one of them [d:A] and one of them [d:R]. The structure of /a/ is simply a bare RTR specification:

$$(2.28) \quad /a/$$

o	a: Rad
Rad Pha R	s: Pha
	d: R

The prediction that the present model makes is that there can be no vowel system that contrasts / ϵ / and / æ /, but in which / æ / functions as a palatal, [s:Pal]. If a segment that appears to be / æ / behaves as a front vowel, then the system can have only three distinctive vowel “heights”. In this case, the / æ / will have the representation that / ϵ / has in other languages, a representation whose phonetic interpretation is ambiguous between a low-mid and a low vowel height:

$$(2.29) \quad /æ/\sim/\epsilon/$$

o	2	
a		
s	/	
d	\	
Dor	Pal	2
		o
		a
		s
		d
		Rad
		Pha
		R

a: Dor			
s: Pal			
d: 2			
2 :			
<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">a: Rad</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">s: Pha</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">d: R</td> </tr> </table>	a: Rad	s: Pha	d: R
a: Rad			
s: Pha			
d: R			


Hungarian is a good example. The vowel whose orthographic symbol is short *e* behaves as the front counterpart of orthographic *a*, a rounded low to low-mid vowel. In an essentially three-height system, Hungarian *e* is pronounced as either [ϵ] or [æ], depending on the dialect. The feature [-back] is not needed to describe these systems, Palatal site is sufficient.

In a language such as English, with four vowel “heights”, / æ / can only be represented as the ATR counterpart of / α /, Tellingly, many dialects of English have no genuine / a /, suggesting that for these dialects *all* vowels, even the low ones, are constrained to have an [a:Dor] specification. Moroccan Arabic is a language where the / æ /~/ α / alternation can be seen more clearly to be the result of the tongue root constriction varying between [d:A] and [d:R]. Moroccan Arabic consonants may be either pharyngealized—[2:[a:Rad, s:Pha, d:R]]—or not—[2:[a:Rad, s:Pha, d:A]], which is presumably the default. These specifications are shared by all segments in a given domain. In a pharyngealized domain, the low vowel is realized as / α /, otherwise as / æ /.

The last category of languages consists of those that have their low vowel's pharyngeal constriction made by the tongue root instead of the tongue body, i.e., the / a / is a genuine [a]. This is the category of language (e.g., Kpokolo) that prompted the claim by Kaye, Lowenstamm, and Vergnaud that / a / cannot have an ATR counterpart. The / a / in these languages behaves consistently as RTR, since it in fact consists of nothing but a bare RTR specification.

2.6.2 Light diphthongs, contours, and releases


Heavy diphthongs are simply sequences of two root nodes, and will be dealt with using the concepts outlined in chapter 3. Light diphthongs pose more of a problem. A common treatment of light diphthongs in the autosegmental literature is to have two root nodes linked to a single timing slot on the skeletal tier or moraic tier, as in (2.30):



The suprasegmental framework outlined in the next chapter, however, does not allow any “timing slot” that is not also a root node. (Put slightly differently, a mora and the root node of its vowel are the same object looked at from different points of view.) Furthermore, there are no tiers independent of the timing or moraic tier along which dependents can be ordered. The usual treatment of light diphthongs cannot work in such a framework.

Instead of having the two members of a light diphthong being sisters under the same timing slot and ordered with respect to each other along the root node tier, I propose that they stand in a hierarchical relationship. Specifically, I propose that the relation between the two members of a light diphthong is exactly the same as that between a consonant and its release. Although it is not entirely accurate to call this a “release” (or “r”) relation, I shall continue to do so, in order to emphasize their unity (and because all the other good letters have already been taken). So the light diphthong $\hat{i}a$ would be represented:¹⁹

¹⁹It is debateable whether the i should dominate the a rather than vice versa. On the one hand, this organization preserved the temporal ordering interpretation of releases as used in consonants. On the other hand, there may turn out to be some processes that unambiguously treat the most sonorous member of a light diphthong as its head, that is, would treat $\hat{i}a$ as if it were a . If this is in fact the case, it might be better to propose a kind of anti-release, an *arrest* or an *attack* node hierarchically dependent upon its root node but temporally preceding it. These attacks or arrest might also be relevant for consonant representations. For example, Azra (1992) argues within Government Phonology that while there are good reasons in languages like Italian for treating sC sequences as coda-onset clusters, in French there is evidence that the s should be treated as the dependent in a branching onset. In the present framework this would have to be represented using such arrests or attacks. I shall occasionally remark in passing some of the other phenomena that might be more perspicuously treated with such arrests or attacks, but I shall not develop a proposal for incorporating them into the present



In section 3.2.2, I shall discuss the possibility that contour tones occurring on short vowels should also be represented using releases.

2.7 Unresolved questions

2.7.1 Laryngeal features

Lombardi (1991) argues that all the relevant laryngeal contrasts for consonants can be expressed in a system using three privative laryngeal specifications: one for glottal closure, one for voicing, one for aspiration. (There may be no laryngeal specification at all, as in an unvoiced, unaspirated, unglottalized stop.)

There would seem to be an easy and tempting correlation between Lombardi's three privative features and the three possible degree specification. We might say that the glottis is an articulator (or site) like any other, and set up the following assignments.

- (2.32) [a:Glo, d:0] glottal closure
 [a:Glo, d:1] voicing (glottal aperture of critical degree)
 [a:Glo, d:2] aspiration (approximant aperture)

Unfortunately for the easy and tempting correlation, aspiration often seems to have more in common phonologically with oral fricatives than with oral approximants (though not always). The analogy between the critical degree of closure needed to produce the aerodynamic properties of a fricative and that needed to produce voicing seems strained. Furthermore, the hypothesis leaves no obvious way to handle other laryngeal phenomena, such as tones or perhaps Piggott's (1988, 1989) proposal on spontaneous voicing. It may be that different interpretations of [d:1] and [d:1] for consonants

framework or for explaining their relationship to releases.

Despite the formal possibility of "anti-releases", there are often good reasons for using releases instead, even in seemingly counter-intuitive cases that involve an on-glide on the root node governing a more sonorous vowel as a dependent release. The analysis of Rotuman metathesis in chapter 3 crucially uses releases for light diphthongs, rather than arrests or attacks. There is also some evidence that in the Pasiego dialect of Spanish, it is the glide portion of a light diphthong that is on the root node and available for spreading.

and vowels, analogous to their different interpretations in oral gestures (e.g., [d:1] as frication for consonants and height for vowels), will be able to handle these different laryngeal phenomena (as well as the fricative/approximant problem just mentioned). Another promising approach might be to incorporate Browman and Goldstein's idea of articulator stiffness as a phonologically relevant property.

Obviously, laryngeal specification is an area that needs more work. In the meantime, for the purposes of this dissertation, I shall assume that these classes of laryngeal gestures *are* results of their degree values, but that, as with the tongue root degrees, A and R, these values need not be the same as other oral degree values. For example, voicing might be represented as [a:Glo, d:Voi] — or in a more neutral shorthand as [Voi]. This treatment of voicing, etc., as primitive should be regarded more as a practical expedient than as a theoretical claim.

2.7.2 Nasals and laterals

Again, for expediency, I shall assume nasality is controlled by the Nasal articulator (Nas), i.e., the velum, which can have two degrees of aperture: closed or open. For want of evidence identifying these with [d:0] and [d:2], I shall use the labels [d:O] and [d:N].

Laterals pose something of a challenge for this sort of framework. One possible representation for an English light /l/ would be to have the tongue tip making a complete closure [a:Apc, s:Ant, d:0], while the tongue blade is only at an approximant degree of closure, [a:Lam, s:Pal, d:2]. In order to distinguish this representation from that of, say, /tʰ/, we would have to assume that the laryngeal specification could have as one of its values something along the lines of spontaneous voicing, as proposed by Piggott (1988, 1989), that is, voicing that is indefinitely sustainable. The incompatibility between a spontaneous voicing specification and a stop or [d:0] specification could conceivably result, by some sort of interpretive convention, in the kind of partial occlusion characteristic of laterals. This type of representation also stands a good chance of being extendible to more exotic sorts of laterals, e.g., the velar lateral described by Capell (1969) for some languages of Papua New Guinea, by using different sites and articulators for the stop and approximant components. This is, so far as I can tell, a reasonably compositional representation for voiced laterals that will certainly fit into the present framework without the need for further modification. It remains to be seen, however, if it is the representation that most accurately reflects laterals' behaviour in phonological processes, or whether we will have to add another type of feature to our gestures, such as the shape parameter proposed by Browman and Goldstein. One type of segment the simple proposal probably will not work for is voiceless lateral fricatives in languages like Nisgha. In the discussion of Nisgha in chapter 5, I shall use the abbreviation [Lat] to stand for whatever the representation of laterality will turn out to be.

2.8 An example from Japanese

Perhaps the least convincing part of the representations I have been proposing is the treatment of degree features, especially the claim that the same feature value is interpreted differently in vowels and consonants. Ideally, there should be evidence that degrees act like other specifications, e.g., that they spread. As well, there should be evidence that high vowels and fricatives can form a natural class. The Japanese phenomenon of affrication offers evidence on both counts.

Japanese has a process where *t* and *d* are affricated to *ts* and *dz* before the high vowels *i* and *u* (actually unrounded [u]). The onset is further palatalized to *ç* and *ǰ* before palatal *i*.

According to the phonetic descriptions in Vance (1987), Japanese *t* and *d* are laminal dentals, meaning their articulator specification is [a:Lam]. Normally, their site specification is anterior [s:Ant], but before *i* they share the vowel's site, resulting in palatals, [a:Lam, s:Pal]. The affrication component of the phenomenon is more difficult to account for. It is not immediately obvious what high vowels and affricates have in common that should allow them to interact in this way. But, as suggested above, high vowels and fricatives are both characterized by the intermediate degree specification [d:1].

The simplest account is that, in this environment, the release of the onset and the following vowel share their degree specifications. If the following vowel is high, that is [d:1], the onset will be [d:0, r:[d:1]], or in Steriade's notation ($A_0 A_f$), that is, an affricate. Otherwise the onset will be [d:0, r:[d:2]], or ($A_0 A_{max}$), that is, a plain released stop. The constraint that the vowel and the release must share features in this way can be expressed using the formalism developed in section 3.4 and chapter 4.

A standard account, on the other hand, has no way of explaining exactly what it is that is going on here. A simple rule like (2.33) may be descriptively adequate.

$$(2.33) \quad \begin{array}{c} [-\text{cont}] \\ | \\ \text{COR} \end{array} \rightarrow \begin{array}{c} [-\text{cont}] \quad [+ \text{cont}] \\ \backslash \quad / \\ \text{COR} \end{array} / \text{ ______ } [+ \text{high}]$$

But (2.33) has nothing to say about why the transformation should occur in the environment of a [+high] vowel rather than the highly unnatural, but equally easy to express, environment of [-high].

The affrication phenomenon of Japanese thus provides evidence both that the classes characterized by [d:1] and [d:2] form natural classes and that degree specifications can spread independently of other constriction specifications. In section 3.3.4, we shall see more examples of degree spreading and of the same degree specification, [d:1], being realized differently depending on whether it is interpreted as a consonant or a vowel. Further evidence on the differential interpretation of degree features can be found in the discussion of the sorts *consonantal* and *vocalic* in section 3.3.

Chapter 3

Outline of the formal system

Phonological structures (PSs) are formal objects. As with any formal objects used as a scientific model, we are posed with two tasks: specifying what might be called the “syntax” and the “semantics” of PSs.

We must specify the “syntax” of PSs: that is, characterize what are possible PSs and what are not. We do this by means of constraints on the properties of mathematical objects known as **graphs**.

We must also provide a “semantics”. Since our aim in creating formal objects in the first place is to model some aspect of reality (part of the linguistic behaviour of human beings), simply marking out a certain class of formal objects is not enough. We must spell out how these relate to the phenomena we are trying to model. The class of graphs we isolate could conceivably be interpreted in various ways: as the family trees of Egyptian dynasties, as possible outcomes of a game, as maps of the London underground system. The only interpretation we are interested in is the one where PSs correspond to types of events in the real world, types of events that involve certain configurations through time of a human vocal tract or certain patterns of sound waves.

Phonology can be seen as the theory of the possible formal objects. Phonetics can be seen as the theory of articulatory or acoustic event-types. The phonology-phonetics interface is thus the set of principles that determines how the formal objects of phonology are to be interpreted as the event-types of phonetics. The mapping between phonological structures and phonetic event types can also be done in a constraint-based way, though it will contain elements that are non-monotonic. Though a full discussion of the phonology-phonetics mapping is beyond the scope of this dissertation, in section 3.5 I shall try to give a general idea of what I believe the principles to be.

The rest of this chapter will have an alternating structure. In each section, I shall first present an aspect of the formalisms that have been used in many constraint-based approaches to language—a piece of technology, so to speak. In the second part of each section, I shall apply the formalism to one or more aspects or problems of phonological

theory, that is, I shall give examples of how that piece of technology can be used to deal with the sorts of questions phonologists are interested in. It should be emphasized that my specific phonological proposals do not follow from the formalisms as logical necessities. When one of my proposals turns out to be empirically inadequate, this will not invalidate constraint-based approaches to phonology in general, any more than the empirical inadequacy of some proposed transformational rewrite rule could invalidate the entire research programme of generative phonology.

Section 1 introduces the mathematical concept of the graph—a collection of nodes and lines joining them—as an appropriate formalism for the diagrams used in both generative and constraint-based phonology. My particular proposals for phonological structure (which owe a great deal to Government Phonology and Dependency Phonology) will be framed in the terminology of graphs. Nodes represent articulatory constriction gestures. Lines between nodes represent the government or dependency relations that the gestures stand in.

Section 3.2 introduces Johnson's (1988, 1990) fundamental distinction between linguistic objects and *descriptions* of linguistic objects. The PSs introduced in section 1 are linguistic objects; constraints are descriptions that any valid PS must conform to. Following Johnson, a formal description language, based on first-order predicate logic, is proposed for describing (and hence constraining) linguistic objects. I then discuss some of the consequences of taking lexical entries, as well as more traditional constraints, to be descriptions of linguistic objects rather than themselves linguistic objects. I also show how many of the effects that Autosegmental Phonology explained by segregating information on different tiers can be obtained in a “no tier” system—using tone as an example, I argue that “association lines” are the effect that arises when two different symbols in a description language statement refer to the same node of a PS.

Section 3.3 discusses types of PS objects or *sorts*. A node can belong to various sorts, such as the onset sort, the nucleus sort, or the sort of positions that are heads of iambs. Lines or arcs representing dependency relations also belong to sorts, which determine exactly what type of dependency the relation is, e.g., the relation between a nucleus and its onset, the relation between the head of a trochee and its dependent. Section 3.3.1 discusses one of the most interesting sorts of nodes in PSs: null positions. Drawing heavily on work in Government Phonology, I propose a typology of the various sorts of empty positions in PSs and discuss some of the constraints on where they may occur. Section 3.3.2 uses the technology of sorts to accomplish the “underspecification” of government relations, and shows how this is useful for explaining a consonant-vowel metathesis phenomenon in Rotuman without recourse to a metathesis rule or segregating consonants and vowels onto different planes of the representation. Section 3.3.3 shows how the framework developed so far can be applied to metrical structure. I propose a model based on the metrical theory of Halle and Vergnaud (1987), where the asymmetrical relations between elements within a constituent are represented by government arcs and headship is represented by sorts.

Section 3.4 introduces no new “technology”, but shows how the system developed so far can deal with spreading processes and the notion of locality in the absence of tiers. I argue for the strong hypothesis that locality is a primitive property of (some) government relations. Whether two PS positions are local to each other for the purposes of spreading is a property of the arc or line that joins them. It follows that *only* nodes that are in a dependency relation with each other can form a local domain.


The final section sketches a model of phonetic structure and of the phonology-phonetics interface.

3.1 Phonological structures

One of the main arguments of this dissertation is that PSs form a definable sub-class of the class of labelled graphs.¹


Some simple graphs are illustrated below:

(3.1)



As can be seen, graphs are essentially collections of points and lines between the points. The points are called **vertices** or **nodes**. The lines are called **edges** or **arcs**. Both nodes and arcs can have **labels**. Arcs may also have a direction—they may originate from one node and terminate at the other:


(3.2)



The trees of generative syntax are graphs with two additional constraints: they are acyclic (no path of arcs leading away from a node ever ends up back at that node), and each node has only one incoming arc—in syntactic terms, only one mother:

¹See, for example, Behzad et al. (1979) or Berge (1983) for an introduction to the basic concepts of mathematical graph theory.

(3.3)




Like generative syntacticians, I shall usually suppress the arrows on the arcs and adopt the convention that arcs originate from the node that is vertically higher in the diagram. Unlike syntactic trees, however, there is evidence that many nodes in phonological graphs have more than one parent (for example, a geminate has both of its syllabic positions for parents).²

When evaluating the representations in this dissertation and the machinery for constraining them, it is useful to keep in mind that any formal version of autosegmental phonology would need to do more or less the same thing. Autosegmental diagrams are also graphs consisting of nodes (with labels such as [-high], σ , and COR) and arcs. The most obvious kind of arc is the association line, but arcs would also be the easiest way of explicitly representing the temporal or tier orderings that are almost always represented implicitly by the left-to-right distribution of symbols on the page. A typical if very underspecified autosegmental representation, using a to label association lines and t to label tier or temporal orderings, might look like:

²Despite no longer being trees, PSs are probably still acyclic, though Johnson (1988) shows that imposing an acyclicity *requirement* results in no advantages of mathematical simplicity.

(3.4)



Readers are encouraged to refer back to this diagram whenever they feel a PS that I propose is entirely too complex to be plausible.

In chapter 2, I used two different notations for representing articulatory gestures as if they were interchangeable. They are. I shall continue to use abbreviated matrices (e.g., $[\text{a:Dor, s:Pal, d:1}]$) in the text, though diagrams will use the graphic notation only from now on.³

3.1.1 Some properties of PSs

Before moving on to outline how properties of PSs are formally stated using constraints, I shall discuss informally what I believe some of the central properties are. I assume that the appropriate representation of segmental content is as discussed in chapter 2, couched in the appropriate graph representations. I devote this section to a sketch of


³This equivalence between matrix and graph representations has deeper consequences. It will probably turn out that all the graphs that are needed to express phonological phenomena are in the class of **directed acyclic graphs**. This is exactly the type of formal object that unification-based approaches to syntax use for the formal underpinnings of their attribute-value matrices. Johnson (1988, 1991) has investigated in more depth this relation between matrix representations and graphs, and the logical foundations of both.

some of the aspects of intersegmental structure, that is, the relations between segmental positions.

Most of the ideas concerning intersegmental structure are borrowed from work in Government Phonology (GP). The basic girders of a PS are governments, or asymmetrical relations that hold between positions. Such relations are represented by an arc or edge from the governor to the governee. To motivate my choice of which ideas to take from GP, which to leave behind, and which to extend further than they were originally intended, a brief detour through X-bar theory will be helpful.


Several attempts have been made to import X-bar schemata from syntax into syllabic theory (e.g., Levin 1986, Charette 1988, Bures 1989). What most of them have in common is taking a coda consonant to be the complement of the nuclear head. Levin argues further that onsets occur in specifier position:

(3.5)




Government Phonology, on the other hand, does not have onsets within the maximal projection of V at all, and argues that it is coda consonants that are in specifier position (the complement position being reserved for the second skeletal slot of a long vowel or diphthong). Paraphrasing their notation somewhat:

(3.6)



GP also argues for a relationship between a coda consonant and the following onset: the onset governs (or *licenses*) a preceding coda:

(3.7)



This licensing relationship can only exist under certain conditions. The exact nature of these conditions need not concern us here (see, e.g., Kaye, Lowenstamm, and Vergnaud 1990)—they essentially boil down to a requirement that the governor not be more “sonorous” than the governee, and that the governor be more complex than the governee (in terms of the segmental representations that Kaye, Lowenstamm, and Vergnaud propose). For example, *rt* is a legal coda–onset sequence, since *r* is more sonorous than *t* and its KLV representation is simpler. *tr* is *not* a legal coda–onset sequence, and cannot have the structure in (3.7), since it violates both these requirements. If a language has sequences such as *tr*, there can only be two possible representations: it may be a branching onset, or if the language does not allow branching onsets, it may be two onsets separated by an empty nucleus. That is, it may have the following structure:

$$(3.8) \quad \begin{array}{cccc} \text{C} & \text{V} & \text{C} & \text{V} \\ | & & | & | \\ \text{t} & & \text{r} & \dots \end{array}$$


GP has strong constraints on the distribution of these empty nuclei. These constraints are discussed throughout the GP literature, but perhaps most completely and accessibly in Charette (1988, 1989) and in the contributions to the 1990 theme volume of the journal *Phonology*.

I push this idea of empty nuclei one step further and argue that *all* apparent coda positions are in fact onsets followed by an empty nucleus, not only those that violate GP’s requirements. While admitting there can be a special relationship between two consonant positions (which we can for now call “coda licensing”), there is no reason to assume that these positions are not also “separated” by an empty nucleus. We need not assume that the addition of an extra government relation between the two consonants results in a radically different structure along the lines of (3.7).⁴

⁴Actually, GP has excellent reasons for denying this possibility, since the two consonants really would be separated given two additional assumptions, neither of which I accept: a) that empty nuclei are not really empty, but rather contain the lax high back unrounded vowel [ɪ] which remains unpronounced under the appropriate circumstances, and b) that there is a tier of phonological structure on which the consonant positions and the “empty” position’s vowel are adjacent. The phonetic realization of empty nuclei is handled somewhat differently in the present framework, so that these problems will not arise if we do away with (3.7) in favour of a strict CVCV structure where some positions may be empty.

Even within the GP tradition, there have been moves in this direction. Guerssel and Lowenstamm (in preparation) propose a strict CV syllabic structure for Semitic. In order to avoid the problems caused by Finnish long vowel–geminate sequences for GP’s rhyme binarity theorem, Gibb (1992) analyzes the first element of a Finnish geminate as standing in onset position and introduces an onset-to-onset government relation. Given that such an onset-to-onset relation is necessary anyway within GP, it would seem plausible that all coda licensing relations are onset-to-onset relations (as argued here), and very strong arguments would be needed to show that a second type of relation, as in (3.7) is also needed.


GP puts the onset outside the maximal projection of the nucleus. Levin puts the onset inside. But despite their disagreements over what exactly is the specifier and the complement of a nucleus, both agree that the relation between one nucleus and the next is of an entirely different stripe and outside the scope of the phonological X-bar schema. This claim loses some of its attractiveness when we recall that neighbouring nuclei often show behaviour (e.g., vowel harmony) that suggests they are at least as intimately related to each other as the positions that are supposedly within the maximal projection of the nucleus. If we are treating apparent coda consonants along the lines suggested in the last paragraph, that is, as onsets followed by empty nuclei, this frees up one of the two major X-bar relations, suggesting a revision of Levin's diagram:




The only possible complement of the nucleus is the maximal projection of another nucleus.

But reflecting honestly on (3.9), there are no compelling reasons for making the onset the specifier and the next syllable the complement rather than the other way around. Furthermore, nucleus-to-onset and nucleus-to-nucleus are far from the only two relations that a nuclear position participates in. There are also the segmental specifications (articulator, site, degree, etc.) and a small number of others that will be introduced later. All of these considerations cast doubt on the strict applicability of the X-bar schema to phonology.⁵


X-bar syntax tries to characterize structural asymmetries in terms of constituency: both head and dependent are members of a separate higher-level constituent which has all the category features of the head and a higher bar level, as in (3.10). As discussed in section 1.3, another alternative is to characterize the asymmetry as a hierarchical relation that holds directly between the head and the dependent, as in (3.11):




⁵Its relevance to generative syntax is no longer unanimously assumed either, cf. Speas (1990).



The difference between the two approaches is illustrated by the ways Autosegmental Phonology and Dependency Phonology tend to deal with syllabic structure, as we saw in the following two sample representations discussed in chapter 1. Autosegmental Phonology relies on the constituency approach and can only represent an asymmetric relation between two segments using the mediation of a higher level node such as Onset or Rhyme (or μ or σ), which one of the segments is taken to be the head of:



Dependency Phonology, using the dependency approach, represents the asymmetry between two segments directly with a hierarchical relation:




The dependency approach seems more appropriate to the properties of phonology. As needed, a node can participate in several different kinds of hierarchical relationships without our having to choose one or two of them as the most basic, the ones upon which X-bar trees should be constructed. The resulting representations will also contain fewer nodes and those nodes that remain can be demonstrably justified—they will behave as nodes ought to behave, unlike some of the more questionable artefacts of the constituency approach.

Onset and nuclear licences


Instead of trying to force all relations into the mould of the X-bar schema, I shall treat PSs as if the relations between a nucleus and its onset and between a nucleus and the next nucleus down are as primitive and direct as the relationship between a segmental node and its articulator specification. To handle the suprasegmental skeleton of PSs, we will need two different kinds of labelled arcs representing the two fundamental kinds of relations that we tried in the last subsection to make into the complement and specifier of an X-bar schema. The nucleus-to-nucleus government that I proposed held between a nucleus and the maximal projection of the next nucleus down will be called the **nuclear licence**, and represented in PS diagrams by the label **n** on its arc. The nucleus-to-onset government will be called the **onset licence**, and represented by government and **o** and represented in diagrams by the label **o**.

(3.14)



As one device to improve the readability of diagrams, I shall label nucleus nodes by **v** and onset nodes by **c**. Subscripts will be used to distinguish between nodes of the same type. Subscripts can also be used to distinguish between different arcs of the same type, since they are formally distinct and nameable objects, though the need to do so will not arise as often. (3.14) will now look like:


(3.15)




I shall call the line of connected nuclear skeletal positions the **nuclear spine** of the PS.

If each of these nodes is specified for segmental content, the above structure results in a perfect CVCV alternation. Deviations from such a strict alternation come about as a result of nodes that have no segmental content.⁶ These empty positions are represented in (3.16–3.17) with a superscript null symbol. An apparent coda consonant is the result of its governing nuclear position being unpronounced (3.16). A long vowel or a heavy diphthong is the result of an onset node not being specified (3.17).

(3.16) PS skeleton of a “CVC syllable”




(3.17) PS skeleton of a “CVV syllable”



A branching onset does not arise from these sorts of principles. Rather, as argued in chapter 2, it is the result of giving articulators and sites to the release position of an onset:

(3.18) English onset *pr*



For languages where coda consonants result in heavy syllables, the spine of nuclear positions in (3.15) corresponds quite closely to the moraic tier in those moraic theories


⁶To be more precise, the relevant property is not the absence of segmental content, which because of default interpretation principles need not result in an unpronounced node. Rather, as we shall see, there is a positive requirement that a particular node have no content and receive no phonetic interpretation.

where onsets attach to the mora (e.g., Hyman 1985) rather than to the syllable (e.g., McCarthy and Prince 1986, Hayes 1989). Other than this, there are no easy equivalences to be drawn between the suprasegmental structure in (3.14) and those of other frameworks of phonology. Specifically, there is no skeletal tier mediating linear order, or tier of root nodes performing more or less the same function. Writers in Government Phonology (e.g., Charette 1988, Kaye 1990a) have noted that, given a representation of linear order—the skeleton—and a rigid structural definition of government, the government relations in a PS are predictable. Turning this observation on its head, it is equally valid that, given a set of primitive government relations and strict principles for phonetically interpreting them, temporal order is predictable. It is my position that there is no need for an explicit representation of temporal ordering in PSs, no need for a skeletal tier, and certainly no need for an independent statement of ordering on every autosegmental tier of the PS. Rather, only government relations are relevant for phonology, and temporal order is the province of phonetics and the principles of the phonology-phonetics mapping.⁷

I believe the proposal that onsets hierarchically depend on their nuclei captures the good points of both the arguments against and the arguments in favour of Onset as a phonological entity (that is, as a structural position that can be uniquely referred to, can be required to be present, can be required not to branch).

On the one hand, researchers in moraic theory such as McCarthy and Prince (1986) and Hayes (1989) offer good arguments against earlier conceptions of the skeleton, what might be called the “one segment = one timing slot” theories, such as the CV skeleta of Clements and Keyser (1983) or the X skeleta of Kaye and Lowenstamm (1984) and Levin (1985). First, onsets never count for the purposes of syllable weight, a result that needs some extra mechanisms in a skeletal framework that treats onset segments and moraic segments as equal members of the skeletal tier, but which seems to follow naturally from a model that posits morae as the only syllabic structure standing between σ nodes and root nodes, as in (3.19).

(3.19)




⁷Scobbie (1991) has argued convincingly that there are no separate tiers for each feature and node, rather the relative linear order of any two nodes of the same class is determined entirely by the relative order of their root nodes. I remove even the necessity for imposing a linear order on the set of root nodes by treating linear order as a phonetic effect of independently needed government relations.

Furthermore, phonology does not seem to deal with arbitrary sequences of timing units (C/Vs or Xs). As McCarthy and Prince (1986) point out, no language has a reduplicative prefix template of XXX that will simply copy the first three segments of the word regardless of their prosodic structure, with results as below:

$$(3.20) \quad \begin{array}{lcl} \text{badupi} & \rightarrow & \text{bad-badupi} \\ \text{bladupi} & \rightarrow & \text{bla-bladupi} \\ \text{adupi} & \rightarrow & \text{adu-adupi} \end{array}$$

Again, the kind of representation in (3.19) accounts for the non-occurrence of such prosodically arbitrary XXX templates essentially by banishing onset segments from the timing tier, the tier of things that can be counted. The price for such a pure moraic theory is the inability to refer to onsets as characteristic positions in the syllable. It can refer to individual onset consonants indirectly (as root nodes that are directly dominated by the syllable node), but not to the entire set of onset consonants. There is no principled means available to the theory for determining whether or not a given syllable even has an onset.

This price is too high for a number of researchers who point to cases where phonology needs the ability to refer explicitly to onsets as onsets. The growing number of arguments to this effect include Schindwein (1989), San Duanmu (1990), and especially Piggott (1991). Prosodic Morphology itself often sneaks references to onsets in through the back door, for example by using the prosodic constituent of the “core syllable”, σ_c , consisting of exactly one C and one V, though without any indication of how we are to check whether one and only one C is in fact present. Interestingly, under one of the interpretations discussed in chapter 1, even the purest of pure moraic theories, Hayes (1989), requires and offers the ability to refer to onsets, in the form of the ω of the diagram reproduced below:




But allowing ω s onto the same tier as μ s like this vitiates most of the benefits that were supposed to follow from postulating μ s in the first place.⁸

⁸As mentioned in chapter 1, another version of moraic theory is still tenable, though a weaker one that is unable to live up to some of the claims made on its behalf.

So we have two requirements for any syllable theory: it should be able to refer to onsets easily and explicitly as syllabic constituents, but it should keep them out of areas where they could interfere with syllable weight or become part of prosodically arbitrary templates. Making onsets hierarchically dependent on their nuclei accomplishes both these tasks. First, onsets are still discrete and identifiable parts of a PS. We can determine whether one is present or absent, we can require one to be present, we can impose constraints on its content, and so on. But at the same time, the onset is kept segregated from the main line of the PS—the spine of nuclear nodes, the moraic tier—which determines syllable weight. Because onsets and nuclei do not occur on the same level, they cannot be generalized over with impunity; there is no mechanism by which a template could say, “Give me three segments, I don’t care if they’re consonants or vowels.” The three XXX instantiations of (3.20), which would have the structures below, have no common property by which they could be subsumed into a single template.


(3.22)



Coda licensing

Government Phonology allows (indeed, requires) a special type of government relation to hold between the onset of a syllable and a coda consonant in a preceding syllable. This was illustrated in (3.7), repeated here:


(3.23)



Several properties are usually proposed for this government type. Like any government type in GP, the head in a coda licensing relation must be more “charmed” than the

dependent, or else it must be more complex.⁹

In GP, coda licensing does much of the work of onset rules or onset principles in other theories.




Most other frameworks explain the illegality of this structure either by appealing to the relative strength of an onset principle or to the relative extrinsic ordering of an onset rule. In GP, on the other hand, a coda licence is literally a licence. The very existence of a coda consonant is made dependent upon the existence of a following onset.

It will be useful to have an analogue to this kind of licensing relation in the present framework. I propose that a consonant onset-licensed by a null nucleus may (and in some languages must) also be the dependent of a coda-licence arc, abbreviated *cl* in PS diagrams. This coda-licence arc has all the properties of other arcs. For example, “spreading” may occur along it, that is, the two nodes it joins may be required to share certain features, resulting in geminates or homorganic nasal-stop sequences.

To illustrate, the sequence *andi*, in a language where head and dependent of a coda-licence arc are required to share place of articulation (i.e., share their root node gesture features such as [s:Ant]), might have the structure:

⁹“Charm” is a primitive of the segmental theory of GP developed in Kaye, Lowenstamm, and Vergnaud (1985, 1989), Harris (1990), etc. Segments can have three charm values—positive, negative, or neutral—predictable from the charm values of the elements the segment is built from. In the context of coda licensing, the requirement that the head be “more” charmed than the dependent usually boils down to a requirement that the onset be negatively charmed (prototypically a segment such as a released stop) while the coda is neutral (usually unreleased stops or fricatives) or perhaps positive (a nasal). The alternate requirement, that the onset be more complex than the coda, is measured in terms of the number of elements it takes to construct the segment in KLV’s segmental theory. Since it is stipulated that elements shared with the onset do not count towards the complexity of the coda, this often has the effect of encouraging shared structures such as geminates or homorganic nasal-stop clusters.


(3.25)



Composed government relations

In addition to the primitive government relations discussed so far (nuclear-licence, onset-licence, coda-licence, release, secondary, and the gestural relations, articulator, site, and degree), there is another class of government relations: **composed** government relations,¹⁰ or the relations that hold between a nucleus and a position two government arcs away. For example, *o·r* (which we might call the release-licence) is the relation that holds between a nucleus and the release position of its onset, that is, the position that can be reached by first following the *o* arc, then the *r* arc, as illustrated in (3.26a). This composed government is what makes possible the Japanese affrication discussed in section 2.4.4. Similarly, *n·o* is the relation that holds between a nucleus and its following onset, i.e., the position reached by following first the *n* arc then the *o* arc, as in (3.26b). Assimilatory effects are often found in this environment as well. *n·n* can hold between a nucleus and the nucleus two steps away, as in (3.26c). This allows a relation to hold directly between two full nuclei, while in effect “skipping over” an intervening null nucleus, a property often exploited in vowel harmony and metrical systems.

(3.26)




¹⁰The name comes from one of the ways of looking at these relations: as the result of functional composition of two primitive functions. For example, the function $M(x)$, or “mother of x ” can be composed with $F(x)$, “father of x ”, to give the composed function $M \cdot F(x)$, “paternal grandmother of x ”.

Composed governments should be distinguished from **paths**, an abbreviatory device to be introduced later. Using a path like $n|n|o|r|2|s$ will let you, if you would ever want to, refer to the site specification of the secondary articulation of the release of the onset that is two nuclei down from the present nucleus, without having to refer explicitly to each of the intervening nodes along the path. But this remains just an abbreviation for the convenience of the linguist; the formal mechanisms underlying the full path are still as long and complicated as if you had spelt them out in full.

Composed government relations, on the other hand, seem to have an independent status and behave as single unitary arcs just like other arcs. Feature spreading can occur between nodes that are joined by a government arc, and nowhere else. Like other arcs, but unlike arbitrary paths, feature spreading can occur between two nodes joined by a composed government relation. Furthermore, as will be seen, government relations belong to natural classes, and it is possible to underspecify an arc using these natural classes. Like other arcs, but unlike arbitrary paths, composed government relations can belong to these natural classes.

I shall assume that at least the composed government relations $o\cdot r$ and $n\cdot o$ are always present where possible, though they may have no independently observable effects on a PS. A stretch of PS would thus look more like this:

(3.27)



Though I assume their presence, I shall usually not include composed government relations in a PS diagram unless they are crucial for the discussion at hand (e.g., the analysis of reduplication in chapter 4).

Sign-level structure

While the foregoing relations are all useful in handling purely phonological phenomena, PSs do not exist in a vacuum. The phonological structure of a word or sentence is just one part of a larger linguistic structure that also contains information about syntax and semantics. It has been something of a common practice in generative grammar to

represent the lexical entry of, say, a verb using a matrix like:


$$(3.28) \quad \left[\begin{array}{l} \text{PHONOLOGY: } /gIv/ \\ \text{SYNTAX: } [+V,-N] \\ \quad +[\underline{\quad} \text{ NP NP}] \\ \quad < \text{Ag Go Th } > \\ \text{SEMANTICS: } \text{give}' \end{array} \right]$$

This practice is made explicit in some unification-based approaches. Pollard and Sag (1987), followed by Bird (1990), propose that a sign (a lexical item or a phrase) is an attribute-value structure where the main attribute-value pairs are specifications of the relevant phonological, syntactic, and semantic information.

$$(3.29) \quad \left[\begin{array}{l} \text{PHON} \quad [\] \\ \text{SYN} \quad [\] \\ \text{SEM} \quad [\] \end{array} \right]$$

Keeping in mind the fact that attribute-value matrices and directed acyclic graphs are notational variants, we can take the graph counterpart of (3.29) literally as the top level of a linguistic structure:¹¹

(3.30)



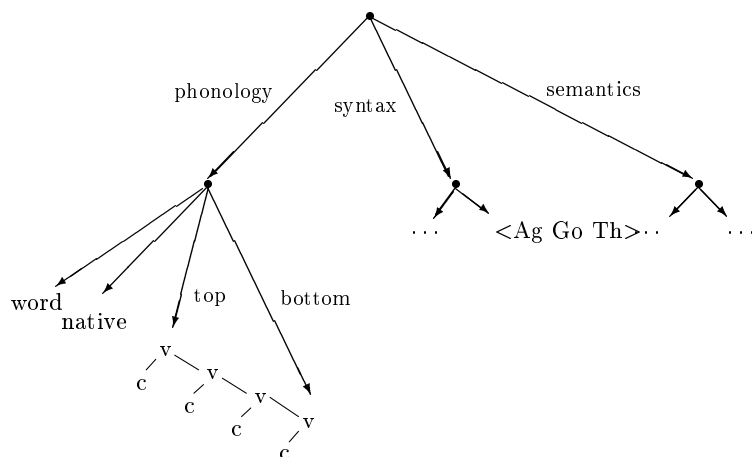
The **phonology** branch of the tree, as well as containing the PSs we have been discussing so far, is the most reasonable place to represent some other properties. One such property would be the prosodic “level” of the PS: whether it represents phonological word, a clitic group, an intonational phrase, or so on. This might also be the most natural place to put such “morphological diacritics” as declension class. Another property represented at this point in the linguistic tree would be whether or not a word is part of a significant sub-class of the vocabulary that behaves coherently with regard

¹¹The questions of whether syntax and semantics are distinct branches (cf. Pollard and Sag, in press) and whether other branches such as pragmatics or morphology could exist (cf. Sadock 1991) are far beyond the scope of this dissertation. Since little in this dissertation hinges on the exact nature of representations outside of phonology, I shall continue to talk as if the tripartite division were a given. I will, however, explore in chapters 5 and 6 the possibility that morphology can be dealt with by directly relating syntax and phonology without the need for an independent level of morphology (cf. Lieber 1992).

to some phonological constraints or processes. A distinction between native vocabulary and (possibly several) loan-word subclasses is made over and over again across languages, with significant effects on the properties of their PSSs, e.g., [\pm Latinate] in English, [\pm Sanskrit] in Malayalam (Sadanandan 1990), [\pm Sino-Japanese] in Japanese (McCawley 1968). Also originating from the **phonology** value are arcs that serve as pointers to the **top** and **bottom** of the PS, that is, the first and last nuclear positions. (It might also be the case that many languages utilize a pointer to the nucleus that bears primary stress.)

The kinds of phonological structures we have been discussing, embedded in a higher sign-level structure, look something like:

(3.31)



3.2 Descriptions of PSSs

The graph representation for PSSs is all well and good, but in and of itself it does not give us much help in the way of expressing constraints on possible PSSs. Many languages, for example, have a constraint that if a consonant is in coda position (in present terms, if it is governed by a null nucleus), it must be identical to the following onset, that is, it must be the first member of a geminate. With graphic representations of PSSs, we could represent this in a somewhat clumsy fashion as:

(3.32)

If (3.32) were to be interpreted as a constraint, it would share the weaknesses of many of the diagrammatically represented constraints or conditions in autosegmental phonology (cf. Bird 1991). While being intuitively appealing, it is not obvious how the representation in (3.32) could be brought to bear on a candidate PS. This inexplicitness in how exactly the diagram is to be interpreted becomes more and more of a problem as the constraints to be imposed on PSs become more complex and more numerous.

In order to characterize the class of valid PSs, we need a formal way of describing them. Following Johnson (1990), I propose a special description language based on first-order predicate calculus. The description language will need both a syntax and a semantics. The syntax can be borrowed wholesale from the usual syntax of first-order languages. The semantics provides a set of principles for relating statements in the description language to properties of the world—the world for the description language being the world of PSs.

The basic predicate of the description language is *arc*, a ternary relation.¹² $arc(a, g, b)$ is true of a PS if there is an arc named by g going from the node named by a to the node named by b .

(3.33)

The usual notation in semantics for the interpretation of a language symbol uses double brackets.

(3.34) $\llbracket a \rrbracket = a'$

means that the description language symbol a refers to the PS node, a' . A more phonologically oriented example is:

(3.35)

¹²Some other important predicates will be introduced in later sections.

Table 3.1: Logical expressions of the description language

$a \wedge b$	a and b. (A PS satisfies the complex description $a \wedge b$ if it satisfies both description a and description b .)
$a \vee b$	a or b. (A PS satisfies $a \vee b$ if it satisfies a or b or both.)
$a \vee_{ex} b$	a or b (exclusive or). Exactly one of a and b is true: a is true, or b is true, but not both.
$a \rightarrow b$	if a then b. (Whenever a is true, b is also true. If a is false, $a \rightarrow b$ is vacuously true. Any PS can satisfy $a \rightarrow b$, except one that satisfies a but fails to satisfy b .)
$a \leftrightarrow b$	b if and only if a. (A PS must satisfy either both a and b , or neither. One cannot be true while the other is false.)
$\sim a$	not a. (a is not true.)
$\forall x(F)$	for all x, F. (When F is a description, either simple or complex, with a variable symbol x in it, any choice of a value for x will make F true.)
$\exists x(F)$	for some x, F. (There is at least one choice of a value for x that will make the description F true.)

That is, for any v_1 and v_2 that are adjacent on the nuclear spine where v_1 is null, any onset v_1 has, v_2 must also have, and vice versa.

The structure consisting of two successive nuclei and their onsets often occurs in diagrams and their descriptions, e.g., (3.16), (3.17), the PS left of the arrow in (3.32). I shall abbreviate the description of series of skeletal slots as in:

(3.39) Structure	Abbreviation	Full description
	$c_1 \cdot v_1 \cdot c_2 \cdot v_2$	$v_1 \xrightarrow{o} c_1 \wedge v_1 \xrightarrow{n} v_2 \wedge v_2 \xrightarrow{o} c_2$
	$c_1 \cdot v_1 \cdot c_2 \cdot v_2 \cdot c_3 \cdot v_3$	$v_1 \xrightarrow{o} c_1 \wedge v_1 \xrightarrow{n} v_2 \wedge v_2 \xrightarrow{o} c_2$ $\wedge v_2 \xrightarrow{n} v_3 \wedge v_3 \xrightarrow{o} c_3$

A final abbreviatory convention will be useful in associating segmental content with a node:

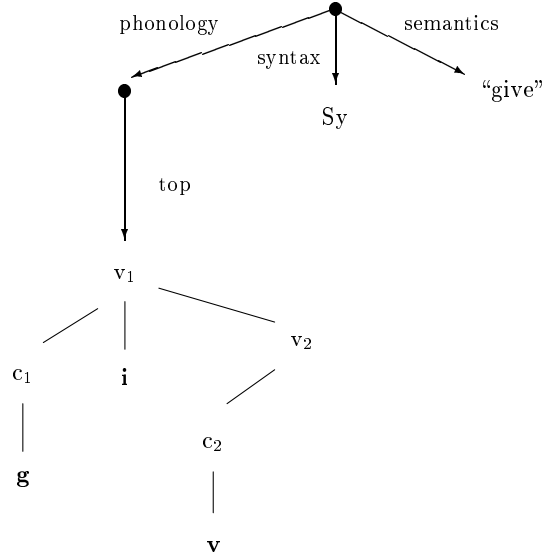
$$(3.40) \quad c_3 \approx \mathbf{t}$$

This will mean that the node $[c_3]$ stands in as many gestural relations (articulator, site, degree, secondary) as are necessary to minimally specify the phoneme /t/ in the language in question, e.g., $c_3 \xrightarrow{s} Ant \wedge c_3 \xrightarrow{d} 0$.

3.2.1 Lexical constraints and the status of the lexicon

Morphemes, the licences for sound-meaning pairings, are simply conditions on fuller linguistic structures that have PSs as one part along with representations for syntactic and semantic properties. The content of a morpheme, say the verb *give*, might be represented graphically as:

(3.41)



The exact nature of the syntactic and semantic substructures is a question beyond the scope of this dissertation, although some tentative assumptions about syntactic structure will be outlined in chapter 5. It is important to keep in mind that the morpheme itself is *not* a graph structure as in (3.41), but the *description* of that graph structure. It licenses a graph containing the relations diagrammed in (3.41) as a valid linguistic structure. The morpheme give is more accurately the following description:¹⁴

$$\begin{aligned}
 (3.42) \quad \forall w \ w \xrightarrow{\text{semantics}} \text{"give"} &\rightarrow \exists Sy, c_1, v_1, c_2, v_2 \\
 &w \xrightarrow{\text{syntax}} Sy \wedge \\
 &w \xrightarrow{\text{phonology|top}} v_1 \wedge \\
 &c_1 \cdot v_1 \cdot c_2 \cdot v_2 \wedge \\
 &c_1 \approx \mathbf{g} \wedge v_1 \approx \mathbf{i} \wedge c_2 \approx \mathbf{v}
 \end{aligned}$$

where the syntax specification Sy would be further specified with such information as grammatical category and argument structure. The possibility of homonymy means

¹⁴The vertical bar abbreviates **paths**, or sequences of arcs. $x \xrightarrow{\text{phonology|top}} y$ can be expanded to $\exists z \ x \xrightarrow{\text{phonology}} z \wedge z \xrightarrow{\text{top}} y$.

that we cannot in general make the conditional in (3.42) a biconditional.

Just as there may be parts of a PS that are not directly traceable to the requirements of any particular morpheme (e.g., epenthetic segments), it is quite possible for there to be syntactic or semantic content that has no effect on the related PS. These are often called zero-morphemes and have been the cause of much consternation among linguists. The following quotation from Hoeksema (1985:18) is a typical example of morphological angst:

For example, the postulation of zero morphemes makes necessary certain arbitrary decisions about e.g. their position in the word: are they prefixes, or suffixes, or perhaps even infixes? Such questions are impossible to answer.

Indeed, in a constraint-based framework, such questions are impossible even to ask. A “zero morpheme” is simply a piece of syntactic or semantic structure that happens not to impose any additional constraints on the form of the associated phonological structure.

A word is in order on how to deal with some irregularly inflected words. Since we have a requirement of monotonicity, that is, a constraint cannot be ignored once it is imposed, the description of a verb with partially suppletive allomorphs can only be a general specification of what all the forms have in common. Put roughly, the description of the verb stem for ‘take’ cannot be /teik/, since there would then be no way that /tUk/ could count as a form of the verb. The best the description can do would be /t_k/, the maximal generalization across all the forms. The other properties of each form would be filled in by lexical descriptions of the full words. The lexicon would contain descriptions for each of the relevant forms: a morphemic generalization for the lexeme TAKE as in (3.43); a description of the basic present tense form *take* (3.44); and a description of the past tense form *took* (3.45). Under each morphemic description, I have included a schematic diagram of the kind of PS that might satisfy that description (of course, the diagrams are so incomplete that they would fail to satisfy many other well-formedness constraints of English).¹⁵

¹⁵The schematic diagrams should not be mistaken for underlying representations. They are not strictly speaking PSs at all. It might be helpful to think of a schematic diagram as being drawn on an overhead projector transparency. It can be laid over a real PS diagram drawn on paper in order to check if the PS corresponds to the schema. The PS may have more structure than the schema, but the schema cannot have any pieces that do not correspond to the PS.


$$(3.43) \quad \forall w \ w \xrightarrow{\text{semantics}} \text{"take"} \rightarrow \exists Sy, c_1, v_1, c_2, v_2$$

$$w \xrightarrow{\text{syntax}} Sy \wedge$$

$$w \xrightarrow{\text{phonology|top}} v_1 \wedge$$

$$c_1 \cdot v_1 \cdot c_2 \cdot v_2 \wedge$$

$$c_1 \approx \mathbf{t} \wedge c_2 \approx \mathbf{k}$$

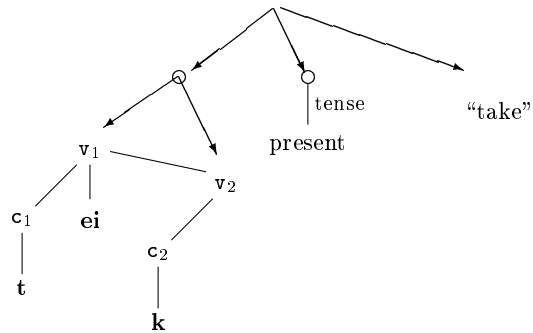


$$(3.44) \quad \forall w \ w \xrightarrow{\text{semantics}} \text{"take"} \wedge w \xrightarrow{\text{syntax|tense}} \text{present} \rightarrow \exists c_1, v_1, c_2, v_2$$

$$w \xrightarrow{\text{phonology|top}} v_1 \wedge$$

$$c_1 \cdot v_1 \cdot c_2 \cdot v_2 \wedge$$

$$c_1 \approx \mathbf{t} \wedge v_1 \approx \mathbf{ei} \wedge c_2 \approx \mathbf{k}$$

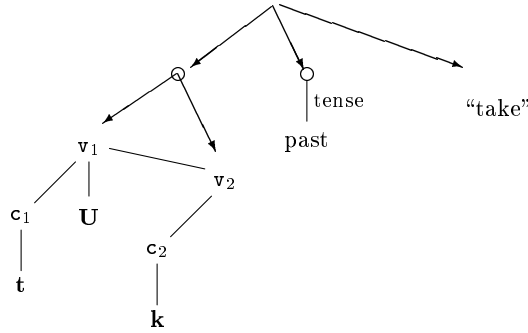


$$(3.45) \quad \forall w \ w \xrightarrow{\text{semantics}} \text{“take”} \wedge w \xrightarrow{\text{syntax|tense}} \text{past} \rightarrow \exists c_1, v_1, c_2, v_2$$

$$w \xrightarrow{\text{phonology|top}} v_1 \wedge$$

$$c_1 \cdot v_1 \cdot c_2 \cdot v_2 \wedge$$

$$c_1 \approx \mathbf{t} \wedge v_1 \approx \mathbf{U} \wedge c_2 \approx \mathbf{k}$$



One question that might be asked is whether the most appropriate choice for representing these lexical items is by having three separate entries or constraints, one for the lexeme in general and two for the fully inflected words. Perhaps all the information can be combined into the entry for the lexeme itself, making a description like (3.46).

$$(3.46) \quad \forall w \ w \xrightarrow{\text{semantics}} \text{“take”} \rightarrow \exists Sy, c_1, v_1, c_2, v_2$$

$$w \xrightarrow{\text{syntax}} Sy \wedge$$

$$w \xrightarrow{\text{phonology|top}} v_1 \wedge$$

$$c_1 \cdot v_1 \cdot c_2 \cdot v_2 \wedge$$

$$c_1 \approx \mathbf{t} \wedge c_2 \approx \mathbf{k}$$

$$(Sy \xrightarrow{\text{tense}} \text{present} \rightarrow v_1 \approx \mathbf{ei}) \wedge$$

$$(Sy \xrightarrow{\text{tense}} \text{past} \rightarrow v_1 \approx \mathbf{U})$$

This question is something of a red herring. Since the ultimate fate of the descriptions in (3.43–3.45) is to be conjoined together, along with all the other lexical entries, into one large all-embracing constraint, they are the same as the more compact description in (3.46): not only would grammars that use either version mark out the same class of PSs, the two versions are logically equivalent and can be derived from each other using the logical counterparts of the distributivity laws in arithmetic that tell us that $5 \times (3 + 2) = (5 \times 3) + (5 \times 2)$.

If the above picture of lexical entries is correct, it raises some questions about the status of the lexicon. There are two dominant metaphors of the lexicon that have permeated research in generative linguistics: the lexicon as a place, and the lexicon as a time. The lexicon is often conceptualized as a place in the head, a warehouse where underlying representations, argument structures, selectional restrictions, and so forth, are all stored until they are needed by lexical insertion into a syntactic tree. Often this warehouse comes equipped with a small factory, where words can be brought out of storage and operated on to create new words, which are put back into storage. The other dominant metaphor sees the lexicon as a time, a certain stage during the derivation of a surface form, e.g., what happens before the postcyclic rules of phrasal phonology. Under either interpretation, it is assumed that one can draw a clear box around the lexicon in one's diagram of the system of grammar, with arrows between it and the other boxes it interacts with. There then ensue fierce arguments over whether certain unprototypical phenomena, such as clitics or phrasal idioms, should go inside the box or outside.

As in the case of zero-morphemes, a constraint-based framework makes such questions hard to ask, let alone answer. Lexical entries cannot be seen as existing at either a certain time of derivation or a certain place. A lexical entry is, simply, a type of constraint, and made of the same formal stuff as any other constraint, including purely phonological constraints like coda gemination, purely syntactic constraints like binding principles, or purely semantic constraints. The most one could say about the lexicon is that it is that set of constraints on linguistic structure that all share a certain form.

But even this raises problems, since there appears to be no single form of constraint that would define all and only those that researchers have wanted to admit to membership in the lexicon. One of the most characteristic features of lexical constraints is that they relate two or more levels of linguistic structure. But this is not an exclusive property of lexical constraints. If temporal order is represented only in phonology, we could also imagine a head-complement or head-adjunct ordering principle (along the lines of, e.g., Venneman 1972, Lehmann 1973, Gazdar et al. 1985) that related two levels of syntax and phonology without being a lexical constraint. The principles that relate the phonological property of contrastive stress to the syntactic or semantic representation of sentence focus also cross levels without in any clear sense being lexical. Researchers in prosodic phonology (e.g., Nespor and Vogel 1986, Selkirk 1986) propose principles that define prosodic constituents with reference to syntactic constituents, though we would not want to call these lexical constraints. Another possible criterion might be based on the fact that a large number of lexical entries tend to deal with a PS that represents a phonological word. But if this were taken as the criterion for membership in the lexicon, it would controversially include many cliticized words and controversially exclude compounds and semantically non-compositional phrases like *kick the bucket*, *put up with*, and *and so on*, and so on. More seriously, it would also exclude any form that contains a phonological word boundary, such as any English word with the suffix *-ness* (cf. SPE, p.85).

In short, there are clear examples of lexical constraints with the following prototypical properties: they relate two or more levels of linguistic structure, they deal with phonological words, the associations between levels tend to be arbitrary in the sense of de Saussure, their semantics will often be non-compositional. There are also clear examples of non-lexical constraints that hold across large portions of the language: purely syntactic constraints such as binding principles, purely phonological constraints such as syllable structure. But there is a long continuum of constraints in between, whose population is likely to exceed that of the endpoints. A constraint-based framework is capable of expressing all the constraints along this continuum and treating them as full and natural parts of the grammar. There is no need to make arbitrary decisions about which side of a theoretical line a particular constraint belongs on.

3.2.2 Autosegments

This section is devoted to illustrating how the separation between linguistic representations and descriptions of representations can be used to achieve many of the effects that non-linear phonology has taken as evidence for segregating information within representations onto different autosegmental tiers.


Contour tones were one of the most pressing problems that led to the development of autosegmental phonology. SPE-style phonology had no acceptable way of representing changes in pitch on a single segment. Using special features like [+falling] missed the generalization that falling tones often behave as high tones with respect to their left environment and low tones with respect to their right environment. Representing sequences of level tones within a single segmental matrix resulted in incoherent formalisms like:

$$(3.47) \quad \begin{bmatrix} +\text{back} \\ +\text{low} \\ -\text{high} \\ [+H][-H] \end{bmatrix}$$

Nevertheless, the sequence approach seemed to be the most promising. Tones often show restrictions simply on the sequence of their level pitches, regardless of how these level pitches are distributed among syllables. Leben (1973) points out that *níkílí*, *nyáhá*, *nyáá*, and *mbā* are all illegal Mende lexical entries, a fact which can be stated more clearly by referring only to tone levels: *any* sequence of HLH is illegal, regardless of how many syllables it is spread over. The problem with expressing this lies in the mismatch between the sequence of tones and the sequence of segments.

Other types of phenomena, such as “tonal stability”, also suggested mismatches between the two sequences. It was argued that languages can delete tones from the tonal sequence but leave the segmental sequence untouched, and delete segments from the segmental sequence but leave the tonal sequence untouched. In the latter case, the tones would appear to surface on a different vowel.

The increasingly obvious solution, synthesized in Goldsmith (1976), was that information about segments and information about tones should be segregated into two different levels of representation, called “tiers”, with temporal coordination between elements on different tiers marked by association lines:



Rules and constraints could refer to one level without necessarily referring to the other.

One consequence is the ability to deal with “floating tones”, a phenomenon noted by many earlier researchers but which could not be appropriately formalized in an SPE framework. Many morphological oppositions in tone languages are marked by tone alone in a way that makes it seem the morpheme in question consisted of nothing but a tone specification. Instead of needing morphologically conditioned rewrite rules, these oppositions could now be represented by tonal morphemes, prefixed or suffixed to the tonal tier of the base and hooked up to the appropriate segment by general principles.¹⁶


In the rest of this section, I show how the same kinds of behaviour can be accounted for in a model that does not keep tones and segments segregated on different tiers.

Ideally, it would be good to have a clear idea of the internal structure of tone specifications. Unfortunately, I have nothing intelligent to say about the representational difference between high and low tones, let alone the various levels of mid. As with many features in chapter 2, I shall simply treat tonal specifications as unanalyzed wholes that are dominated by a nucleus' root node via a *tone* arc (which is probably a sub-species of the secondary articulation relation).




The question arises as to how many tones a nuclear position can bear. The usual representation in early autosegmental phonology for a contour tone was a single vowel on the segmental tier linked to two or more tones on the tone tier.


¹⁶In fact, the rules needed to make sure the docking happened correctly were usually as complicated as the rules that would have been needed to transform the base directly into the new form. The analysis with floating tones was thus more the result of a commitment to an Item-and-Arrangement model of morphology than something forced by the nature of the data.



A simple translation of this into the present framework would be:



(3.51) takes advantage of the general possibility that a language may allow more than one secondary articulation arc for a segment. But if the internal structure of tone is such that “H” and “L” are contradictory, (3.51) would be illegal. There is also the possibility that the tonal specifications I have been abbreviating “H” and “L” are not incompatible—they might, for example, together represent some form of mid tone. But the point remains that (3.51) would be an inappropriate structure for a *contour* tone. The illusion of ordering in (3.51) is purely the result of representing the PS with a two-dimensional diagram. It is formally identical to:




Since there is no independent ordered tier for tones, the only way phonology can talk about a difference in ordering is through hierarchical position. But the structural position of the two tones in (3.51) is identical. Rejecting (3.51) would imply that each position along the nuclear spine of a PS (roughly, each mora) could have at most one tone. This is the position of Woo (1969). It makes the prediction that all vowels with contour tones must be long, a prediction which is unfortunately contradicted by several languages (cf. Leben 1973).¹⁷

To represent monomoraic contour tones, we need a representation that is like (3.51) but avoids its problems. Fortunately, we already have a type of arc that is similar to, but not quite identical to, the secondary articulation arc: release. Though releases were originally intended to handle consonantal phenomena such as affrication, glottalization, and branching onsets, I have already mentioned the possibility that they

¹⁷San Duanmu (1990), though, reasserts the argument that each mora can bear only one tone, and discusses many of the empirical counterexamples.


are the appropriate representation for light diphthongs. If we specify tones on vocalic releases, we could get a PS for a falling tone like the following:

(3.53)



The two tone specifications in (3.53) are phonologically distinguishable by the type of government relation they bear to the root node, and phonetic interpretation will be able to assign them the correct temporal ordering. The segmental content of a monophthong with a contour tone would be shared by the root node and the release node. A short *i* with a falling tone would be represented as:


(3.54)



This is similar to the structure of an affricate (as proposed in chapter 2), where all gestural features except degree are shared between the root node and the release.

There is another benefit of this kind of representation. Autosegmental phonology has no general way of preventing associations like:

(3.55)




It is mysterious in frameworks that allow multiple linking of tones to positions why such linkings are generally limited to a maximum of two.¹⁸ If (3.53) is the correct

¹⁸Bao (1990:61–2) discusses some counter-examples to this generalization, but all such convex and concave contours on short vowels result from tone sandhi. Bao concludes that all contours are underlyingly two-level (though he extends this conclusion to long vowels as well as short). It is likely that the present framework will also be able to maintain that only two

PS for a monomoraic falling tone, the answer is clear: there are only two possible distinguishable relations that could hold between a tone and the root node.

Furthermore, contours on long vowels are limited to three levels (cf. Yip 1980, Bao 1990):

(3.56)



It would seem that the first vowel here should also be able to have a release with a tone, resulting in a four-level contour. The failure of this to occur is parallel to the behaviour of releases in consonants: in most languages, the first member of a geminate consonant also has no release. It may be possible to subsume both limitations under the same constraint.

The most salient difference from autosegmental phonology in the framework I propose is that there is no tier specifically for tone specifications. Tones cannot be ordered along such a tier independently of the rest of the PS. What is phonetically interpreted as temporal ordering between two tones is mediated by exactly the same hierarchical structure that mediates the phonetic ordering of any two segmental specifications. The one “tier” that temporal order is defined on is the nuclear spine, or the “moraic tier”. (Cf. Archangeli and Pulleyblank 1992.)

This seems to run counter to the original arguments of Williams, Leben, and Goldsmith that in many African languages, phenomena such as tonal stability and floating tones require that information about tones be segregated from information about segments. It is the central hypothesis of autosegmental phonology that this segregation happens in the phonological representation itself, that tonal and segmental features live (physically, so to speak) in different places. But there is more than one way to segregate information.

I propose that the segregation of information lies not in the PSs, but in the *descriptions* of PSs. Tones and segments live in the same place, they just give their addresses differently.

This is possible because of the semantics of the description language. It is possible for two different description language symbols to refer to the same object in the universe

tones can ever associate to one position in a PS, and explain apparent counter-examples in tone sandhi environments as similar to vowel coarticulation, the result of the indeterminacy in the exact timing of extended phonetic events implied by the theory of the phonology-phonetics interface outlined in section 3.5.

of discourse. Lexical constraints use existentially quantified variables to refer to skeletal positions; it is possible to use two different variables, one for tone, one for segmental material, that end up denoting the same node in the PS.

Many tone languages will not use this possibility. For example, the relevant part of the lexical constraint for *mā* in Mandarin Chinese probably looks like:

$$(3.57) \quad \exists v_{17} \quad v_{17} \approx \mathbf{a} \wedge v_{17} \xrightarrow{t} \text{“H”}$$

where the description requires that the high tone be on the same nucleus as the vowel *a*. In many African languages, however, the variables denoting the skeletal slots of segmental content are less tightly tethered to those denoting the skeletal slots of tones. (3.58) might be part of the lexical constraint for the hypothetical stem *tómù* in a hypothetical African-like language.

$$(3.58) \quad \dots \exists v_{21}, v_{22}, v_{51}, v_{52} \dots \\ v_{21} \approx \mathbf{o} \wedge v_{21} \xrightarrow{n} v_{22} \wedge v_{22} \approx \mathbf{u} \wedge \\ v_{51} \approx \text{“H”} \wedge v_{51} \xrightarrow{n} v_{52} \wedge v_{52} \approx \text{“L”}$$

In many forms of the verb, it may turn out that v_{21} and v_{51} will denote exactly the same object in the PS, as will v_{22} and v_{52} :

$$(3.59) \quad \begin{array}{c} \begin{array}{c} v_1 \\ / \quad | \quad \backslash \\ c_1 \quad \mathbf{o} \text{ “H”} \quad v_2 \\ | \quad \quad | \quad / \quad \backslash \\ \mathbf{t} \quad \quad c_2 \quad \mathbf{u} \text{ “L”} \\ \quad \quad | \\ \quad \quad \mathbf{m} \end{array} \end{array} \quad \begin{array}{l} \llbracket v_{21} \rrbracket = v_1 \\ \llbracket v_{51} \rrbracket = v_1 \\ \llbracket v_{22} \rrbracket = v_2 \\ \llbracket v_{52} \rrbracket = v_2 \end{array}$$

But in other forms, the variables may denote different positions. Say, the tonal variable v_{51} is identified with the vowel of a prefix, while v_{52} coincides with v_{21} :

$$(3.60) \quad \begin{array}{c} \begin{array}{c} v_1 \\ / \quad | \quad \backslash \\ c_1 \quad \mathbf{i} \text{ “H”} \quad v_2 \\ | \quad \quad | \quad / \quad \backslash \\ \mathbf{c} \quad \quad c_2 \quad \mathbf{o} \text{ “L”} \quad v_3 \\ \quad \quad | \quad \quad | \quad \quad | \\ \quad \quad \mathbf{t} \quad \quad c_3 \quad \mathbf{u} \\ \quad \quad \quad \quad | \\ \quad \quad \quad \quad \mathbf{m} \end{array} \end{array} \quad \begin{array}{l} \llbracket v_{21} \rrbracket = v_2 \\ \llbracket v_{51} \rrbracket = v_1 \\ \llbracket v_{22} \rrbracket = v_3 \\ \llbracket v_{52} \rrbracket = v_2 \end{array}$$

It remains to be seen if, as a general rule, all the possible co-references of variables will fall out naturally from the interaction of lexical constraints and all the other independently needed constraints of a language, or if some additional (possibly universal) constraints analogous to autosegmental phonology's Association Principles will be needed.

One relevant possibility for the mismatch of tonal and segmental information arises if some of the information is conditional. For example, the final *u* of v_{22} in our hypothetical *tomu* might be a conditional segment, that is, it occurs only in some forms of the verb and thus morphologically marks those forms (like the final consonants that morphologically mark feminine gender in French, or the idiosyncratic epenthetic-like consonants that help mark passive voice in Maori). In forms where the morphological conditions are not met and the *u* does not occur (i.e., there is no requirement for v_{22} to exist), the lexical constraint's clause for v_{52} might still require that a low tone appear in the verb. This is what underlies the effect of tonal stability under "deletion" of segments.

Giving an analysis of even a small piece of a real tone language would involve having a more realistic theory of the internal structure of tone specifications, an analysis of underspecification in that language (both segmental and tonal, cf. Pulleyblank 1984), and a clear idea of the language's morphology. This is all beyond the scope of this dissertation and makes it impractical to use a real tone language to illustrate the points I have been discussing. But I hope that I have succeeded in showing some of the resources and strategies that would be available to the present framework in dealing with autosegmental phenomena and that I have laid some of the groundwork for a full-fledged analysis of tone.

3.3 Sorts

Any fully explicit theory of autosegmental phonology would want to make sure that the right kinds of phonological entities inhabit the right niches in phonological structures. Nobody would want to admit the following structure:

(3.61)	COR	—	mora tier
	μ	—	root node tier
	σ	—	supralaryngeal tier
	[-high]	—	place node tier
	[+cons -cont]	—	[atr] tier

Less explicit versions of autosegmental theories usually give a few cursory remarks about the distribution of some of the entities, and trust the common sense of phonologists not to try anything outlandish.

The resources of our description language allow us to formulate explicit constraints on what kinds of entities can occur in what kinds of positions and what kinds of entities can enter into various government relations. We can specify what kind of entity a particular PS object is, or its **sort**, using predicates in the description language. $onset(x_1)$ would be true if x_1 is of the sort **onset**, the sort of object that is appropriate for the skeletal slot of an onset.¹⁹

Constraints can be written using these sort predicates. For example, it is reasonable to suppose that the primitive elements—Lip, Ant, Pal, Vel, Pha, Lab, Apc, Lam, Dor, Rad, 0, 1, 2, A and R—belong to the same sort, which we can call **atom**. $atom(Vel)$ would be true, but $atom(v_4)$ would be false if v_4 is a nuclear skeletal slot. We can also say that nodes that represent articulatory constriction gestures are of the sort **constriction**. We can now express the constraint that the objects that enter into an articulator relation, i.e., the head and tail of an **a** arc, must be of the appropriate sorts, as:

$$(3.62) \quad \forall x, y \ x \xrightarrow{a} y \rightarrow constriction(x) \wedge atom(y)$$

though we would want to be more specific than just this—not allowing $x \xrightarrow{a} Vel$, for example. (3.62) would be one of the set of appropriateness conditions that ensure that **constriction** nodes, and only they, can have specifications for segmental content (e.g., articulator, site, degree, secondary, release). Further descriptions can enforce other conditions we would like to impose on PSs, such as the requirement that a secondary articulation of a constriction node should be another constriction node and not, say, a phonological atom or the syntax's representation of the word's argument structure,

$$(3.63) \quad \forall x, y \ x \xrightarrow{2} y \rightarrow constriction(x) \wedge constriction(y)$$

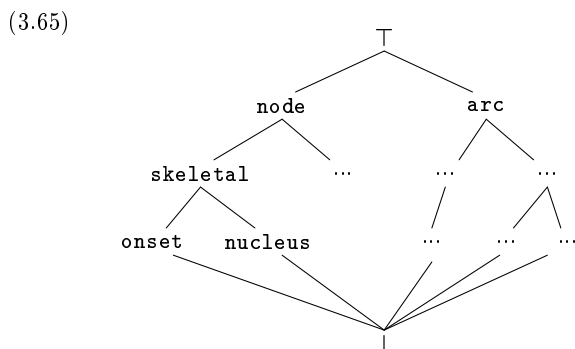
or the condition that atoms like Pha and Apc cannot be the origin of any arc (ruling out, for example, $Apc \xrightarrow{s} Pha$):

¹⁹ An alternative to making each sort a description-language predicate would be to make them constants and express the sort membership of a node by the predicate $sort(x_1, onset)$. I shall continue to use the more concise notation, though I take this to be an abbreviation for the more long-winded version using constants and the predicate $sort$, which will be used in the appendix detailing the formal system. The distinction becomes crucial in section 3.4, when we will want to be able to predicate things of sorts (like the fact that two positions share all arcs of a given sort). If sorts were really themselves predicates, this would require a second-order language. Only the constant-and-*sort* version can accomplish this within a first-order language.

$$(3.64) \quad \sim \exists x, g, y \text{ arc}(x, g, y) \wedge \text{atom}(x)$$

A more complete list of the (universal) appropriateness conditions that will be needed can be found in the formal appendix.

It is clear that sorts are not exclusive of each other. It is possible, indeed the usual case, for a position of the sort `onset` also to be of the sort `constriction`. While `onset` and `constriction` are two independent sorts that happen to overlap in many cases, there is another possibility that is very interesting: one sort may be a subsort of another. For example, we might want to refer to nodes that are either onset positions or nucleus positions, as opposed to secondary articulations or releases, using a supersort called `skeletal`. In this case, `onset` and `nucleus` would be subsorts of `skeletal` and any node that was of either sort would automatically also be of sort `skeletal`. It is common practice to diagram these relations using a sort lattice:²⁰



In these diagrams, \top (or “top”) is the symbol for the universal sort, the sort that all objects in the PS belong to. \perp (or “bottom”) is the sort that no object can belong to. Sorts are diagrammed between these two extremes in terms of how encompassing they are. Sort membership is enforced by constraints like:

$$(3.66) \quad \forall x \text{ skeletal}(x) \leftrightarrow \text{onset}(x) \vee_{ex} \text{ nucleus}(x)$$

As diagram (3.65) suggests, arcs as well as nodes can have sorts. One natural application of this would be to express government types by means of sorts, such as `onset-licence`, `articulator`, etc.. Now we can see how to spell out the iconic arc abbreviation in our description language. For example:

²⁰These can be seen quite literally as instances of the algebraic structure “lattice”, though I shall not dwell on the point. See, e.g., Bird (1990) or Scobbie (1990) for a fuller exposition of sort lattices.

$$(3.67) \quad v \xrightarrow{o} c \equiv \exists g \text{ arc}(v, g, c) \wedge \text{onset-licence}(g)$$

The fact that arcs can belong to sorts opens up the possibility that they too can enter into subsort-supersort relations. This allows us to represent natural classes of government types. For example, we might want to group the government types related to segmental content (**articulator**, **site**, **degree**, **secondary**) into a supersort, say, **content**. Now several of the constraints similar to (3.62) can be expressed more succinctly:

$$(3.68) \quad \forall x, g, y \text{ arc}(x, g, y) \wedge \text{content}(g) \rightarrow \text{constriction}(x)$$

The ability to form natural classes of government types also opens up the possibility of “underspecifying” government arcs. If **secondary** and **release** form a natural class, it would be possible for a lexical entry to specify only the supersort $\{2, r\}$ for a particular arc, leaving the choice of which specific sort of arc appears in the PS to be determined by the interaction of other constraints. This could result in the arc being **secondary** in one inflected form of the word and **release** in another.²¹ Section 3.3.2 will present an example from Rotuman of this kind of underspecification behaviour on the part of government relations.

Consonantal and vocalic

We have proposed the sorts **onset** and **nucleus** to distinguish between the appropriate positions of the skeleton. There remains the problem from chapter 2 of how the phonetics is supposed to know if a node is supposed to be interpreted as a vowel or as a consonant. The choices give different effects for degree features. [d:1] is interpreted as a critical degree of closure (resulting in frication) for consonants but as a high vowel for vowels. [d:2] is interpreted as an approximant for consonants and as a mid vowel for vowels.²² It is quite possible that it is the sorts **onset** and **nucleus** that determine which interpretation is given.

I tend to believe, however, that this difference in interpretation depends on two other sorts, which can be called **consonantal** and **vocalic**. It is true that in most languages and most environments, **consonantal** will coincide with **onset** and **vocalic** with **nucleus**. But there are differences in **degree** interpretation for which the simple onset/nucleus distinction is insufficient. The best example is the interpretation of release positions, which are neither **onset** nor **nucleus**. Releases in fact show a great deal of variation across languages, and sometimes within languages, as to whether their gestures are to be interpreted as consonants or vowels. Lapsing momentarily into a more procedural way of talking: if an *i* [s:Pal,d:1] spreads onto its onset’s release

²¹This might be the appropriate analysis for an alternation between monophthong and light diphthongs, for example, in Spanish verbs.

²²To be more specific, for the dorsal gesture of vowels.

via the *o-r* arc (the mechanism for doing this is dealt with in section 3.4), the result will depend on whether it lands there as **consonantal**, giving the affricate [t^s], or as **vocalic**, giving [t^v].

Moreover, there are cases where onsets and nuclei themselves can take the more marked value of **consonantal** or **vocalic**. Onsets can be **vocalic**, particularly for glides. If a vowel spreads via the *n-o* arc to a following unspecified onset, in many languages the resulting onset will be **vocalic**, for example, the [d:1] palatal gesture of *i* in *i+a* is interpreted as a vocalic [d:1], giving [iya]. But this is by no means universal, as there are also cases where the gesture spread from the preceding nucleus can be interpreted as a consonant.

Hua (Papuan, New Guinea) is a one such language, where the “glides” that fill otherwise empty onsets are fricatives (Haiman 1980:42–43). The empty onset between an *i* and a following vowel can optionally be filled with the glide *z*, phonetically ranging from [z^v] to [δ^v]:

- (3.69) *hi + e* → *hize* ‘he did it’
 io + roga → *izoroga* ‘down in the grasslands’

The fricative *v*, phonetically [β], can occur between a *u* and a following vowel:²³

- (3.70) *hua* ~ *huva* ‘Hua’
 ua’ ~ *uva’* ‘overripe’
 zatura ~ *zatura* ‘house site’

As in other cases of inter-vocalic glides, in Hua the *n-o* government between a high nucleus and the following unspoken-for onset may be a local-domain creator for the gestural features of the nucleus. This includes the high-vowel specification for the onset, [d:1]. But the following onset is **consonantal** rather than **vocalic**, so this [d:1] is not interpreted vocally as height (which would result in a glide, *y* or *w*, but consonantally as a critical degree of closure, resulting in fricatives.

Moreover, whether glide-like onsets are **consonantal** or **vocalic** may vary even within a language. Nisgha (Tshimshianic, British Columbia) shows a clear alternation between glides and fricatives, specifically between the palato-velar fricative /x/ and the glide /y/ and between the labialized dorso-velar /x^w/ and the glide /w/ (Tarpent 1987:82–91). The fricatives and glides are to some extent in complementary distribution, the fricatives occurring word-finally after stress and before consonants and the glides initially before a vowel. There are morphophonemic alternations between the

²³Because of the existence of some *u+a* sequences that cannot be broken up by a *v* (e.g., *rua* ‘competition, quarrel’, *zua* ‘dish’), Haiman concludes that there is no rule of *v* glide insertion, but that all *vs* are present underlyingly and may undergo an optional rule of *v*-deletion. In the present framework, the relevant *n-o* arcs in the non-alternating words may be lexically marked as not creating a local domain, and the glide “insertion” account can be generalized.

3.3.1 Nullness and empty nuclei

Dell (1973) presents the following data illustrating schwa deletion in French. Each of the boldfaced *es* in the first line may be either unpronounced or pronounced as a schwa:

(3.74)	<i>envie</i>	de	te	le	<i>demande</i>	‘desire to ask it to you’
	a.	də	tə	lə	dəmande	
	b.	də	t∅	lə	d∅mande	
	c.	d∅	tə	l∅	dəmande	
	d.	də	tə	l∅	dəmande	
	e.	d∅	tə	lə	d∅mande	
	f.	d∅	tə	lə	dəmande	
	g.	də	t∅	lə	dəmande	
	h.	də	tə	l∅	dəmande	
	i.	də	tə	lə	d∅mande	

The generalization here is that any combination of pronounced and deleted schwas is possible, except one with two deleted schwas in a row.


Dell, and other generative phonologists since him, accounted for the data by postulating underlying schwas and an optional transformational rule deleting them. The requirement that there not be two deletions in a row is encoded either in the structural description of the deletion rule itself or in some extra condition. Most of these analyses also require some further resyllabification rules to apply to clean up the ill-formed mess left by the schwa deletion rule.

Using the concepts of Government Phonology, Charette (1988) takes a different tack in analyzing the pattern in (3.74). Since GP allows for empty nuclei that can be either pronounced or unpronounced, there is no need for any resyllabification. In the *te* in (3.74b), for instance, the *t* whose schwa has been “deleted” does not need to become the coda of the preceding *de*, it can remain an onset whose following nucleus happens to be empty. The only aspect of the GP representation that changes in (3.74) is the distribution of a particular type of government relation known as **proper government**. In GP, there is government from right to left between nuclei.²⁵ Some of these nucleus-to-nucleus government relations may (in French, optionally) be *proper*. A nucleus that is a proper governor must be phonetically realized, the proper governee is unrealized. All of the possible schwa positions in (3.74) are in fact empty nuclei—those that are the governee in a proper government relation are unpronounced (or


²⁵More recent versions admit the possibility that the direction of government between nuclei is a parameter that can vary between languages. Since there is no immediately obvious increase in empirical adequacy in parametrizing this property, I prefer to continue with the assumption that nuclear government is uniformly left-to-right, i.e., the phonetic interpretation of the head is temporally ordered before the phonetic interpretation of the dependent.

“deleted”), the rest are pronounced as schwas. For example, (3.74b) and (3.74c) have the following structures, where the arrows represent proper government.

(3.75) de t̸ le d̸mander



(3.76) d̸ te l̸ demander



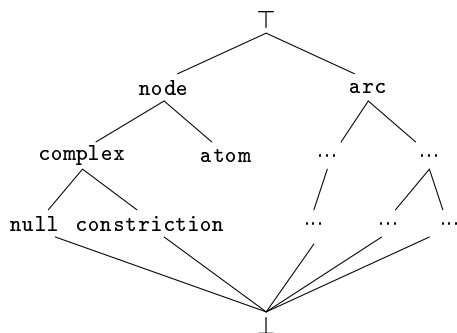
All the possible patterns can be obtained by freely distributing proper government relations, subject to the condition that the governor must be pronounced. This condition accounts for the absence of two deleted schwas in a row: if one nucleus is unpronounced, it cannot possibly properly govern the nucleus to its left, which therefore must be pronounced.

Although I am uncomfortable with this appeal to the eventual semantic (phonetic) interpretation in defining syntactic (phonological) well-formedness, I find Charette’s analysis more convincing than procedural deletion analyses. It is easily translated into the present framework. I argue that those nuclear positions whose phonetic fate is to be unpronounced (in GP terms, licensed by proper government to be unrealized) belong to the sort *null*. This sort has obvious effects on the interpretation of nuclear positions, but it is not defined in terms of those effects. *null* is a phonological sort like any other phonological sort and enters into subsort relations with them of the kind diagrammed in (3.65).

Specifically, among nodes that are not segmental atoms, *null* is the complement of *constriction*. Put explicitly, every node (as opposed to arc) is either an *atom* (Apc, Vel, 1, etc.) or *complex*—this can be thought of roughly as a division between terminal and non-terminal nodes—and every *complex* node is either a *constriction* or *null*.

$$(3.77) \quad \begin{aligned} \forall x \text{ node}(x) &\leftrightarrow \text{atom}(x) \vee_{ex} \text{complex}(x) \\ \forall x \text{ complex}(x) &\leftrightarrow \text{constriction}(x) \vee_{ex} \text{null}(x) \end{aligned}$$

(3.78)



It follows from *null* and *constriction* being mutually exclusive and from appropriateness conditions like (3.62) that a null position can never have any segmental content specified. It would be incoherent to say that v_4 is specified for the vowel *o* but just happens to be null so that the *o* cannot be pronounced.²⁶ A node that is of sort *null* is a quite different animal from a node that is of sort *constriction* but happens not to have any segmental content specifications. For example, the latter, but not the former, will receive default specifications in phonetic interpretation.

The condition that there cannot be two null nuclei in a row can be expressed rather straightforwardly as:

(3.79) Adjacent nulls constraint (preliminary version):
 $\sim \exists x, y \ x \xrightarrow{n} y \wedge \text{null}(x) \wedge \text{null}(y)$


Another version, taking advantage of the complementarity of sorts, would be a positive condition:

(3.80) $\forall x, y \ x \xrightarrow{n} y \rightarrow \text{constriction}(x) \vee \text{constriction}(y)$

It is reasonable to extend this constraint to types of government arcs other than nuclearlicences. For example, we do not find null nuclei licensing null onsets. In general, adjacent nulls are the type of thing we should like to avoid, otherwise we could end up with PS where vast stretches were null, as in:

²⁶This is by no means a necessary consequence of the framework. Letting a node be both *constriction* and *null* could be a not implausible analysis of apparent deletion phenomena. But this would be a much more radical step towards multistratality and I feel the disadvantages of its potential unconstrainedness outweigh any advantages. I shall not pursue the possibility further.

(3.81)



The strongest hypothesis would be that a null cannot govern another null through *any* type of government arc. It remains to be seen if this strongest version can be maintained. The brute force ban of the Adjacent Nulls Constraint may seem less explanatory than the proper government concept of Government Phonology, but I believe its attractiveness increases when the prohibition against adjacent null *nuclei* is seen as a special case of a more general constraint against overly sparse PSs.

Many other phenomena refer to whether or not the neighbourhood nucleus is null. Dialects of French, for example, differ in whether or not they allow a word-final null nucleus to licence release specifications on its onset. Standard Parisian French allows full onset releases in this environment, resulting in word-final clusters like the *tr* of *quatre* or the *bl* of *table*. Many dialects of Quebec French, on the other hand, do not allow such releases to have specifications, resulting in forms like [kat] and [tab] (cf. Walker 1984:108–111). My dialect of English goes one step further and bans not only segmental content on these releases, but also most release nodes themselves, with the result that word-final stops are unreleased.

There are many languages which make no use of the null type and many others for which constriction and null as I have described them so far are adequate. But many languages need a more intricate system of null and constriction sorts. In the following subsections, I shall present a somewhat richer typology of node sorts. Nulls will be sub-divided into **simple nulls** and what, for want of a better term, we can call **extra-nulls**. Constriction nodes will be divided into **specified** and **unspecified**. These subsorts enter into other supersort relations with each other that cross-cut the division between **null** and **constriction**.

Simple nulls

All the nulls we have seen so far have been simple nulls. This is the type of null that respects the constraint in (3.79) against adjacent nulls, and which we should now reformulate:


$$(3.82) \quad \text{Adjacent nulls constraint (revised version):} \\ \sim \exists x, y, g \quad x \xrightarrow{g} y \wedge \text{simple-null}(x) \wedge \text{simple-null}(y)$$

This is not yet the final version of the Adjacent Nulls Constraint. An extra condition will have to be added that *x* and *y* are within the same phonological word.


This can be done using the concepts concerning the prosodic hierarchy that will be introduced in chapter 5.

Extra-nulls


Not all empty nuclei seem to respect the Adjacent Nulls Constraint. In more traditional phonological terminology, these are the cases of consonants that do not respect the usual syllabification constraints of the language. A notorious example is the English word *sixths*, which would have a PS in which no less than three empty nuclei violate the Adjacent Null Constraint:





Not all examples of this kind need to be morphologically complex, as the word *child* shows:



Again, while Moroccan Arabic generally shows a perfect alternation of pronounced and unpronounced empty nuclei, many nouns do not follow that pattern. Beside well-behaved nouns like *nmer* 'tiger' and *qfel* 'lock', there are nouns that phonetically end in two consonants (so phonologically end in two empty nuclei): for example, *kelb* 'dog', *wezn* 'weighing', *ražl* 'man'.²⁷ These would have the structures:



²⁷The claim that the penultimate consonants in these forms are actually onsets followed by empty nuclei is supported by the fact that in their broken plural forms they appear as onsets followed by full nuclei: *kluba* 'dogs', *ržal* 'men'.



The last two nuclei in these forms do not respect the Adjacent Nulls Constraint.

Furthermore, there are other dialects of Moroccan Arabic where not even the verbs discussed in the last subsection show the pronounced/unpronounced alternation of empty nuclei. Instead, all empty nuclei remain unpronounced. Instead of the contrast between *kteb* and *ketbu*, these dialects have simply [ktb] and [ktbu].²⁸ These forms would have the same structure as *kteb* and *ketbu*, except that all empty nuclei everywhere remain null, without regard for the Adjacent Nulls Constraint.

These data can be dealt with if we assume that the sort **null** has two mutually exclusive subsorts, **simple-null** and **extra-null**. As reformulated in (3.82), the Adjacent Nulls Constraint applies only to **simple-nulls**. Extra-nulls are free to occur next to simple-nulls, and even next to other extra-nulls.²⁹

While the distribution of extra-nulls is free with respect to each other and to simple nulls, languages that allow them generally impose other kinds of restrictions on them. They may be restricted to the edges (the top or the bottom nucleus) of a PS. There may be constraints on the types of onsets they can license. There may be restrictions on the kinds of morphological environments they can occur in. English extra-nulls, for example, can only license coronal, or rather anterior [s:Ant], onsets, e.g., the *d* of *child* or the final *s*, *,* and *s* of *sixths*. English also generally restricts extra-nulls to occur as the bottom nucleus of a morpheme. The first dialects of Moroccan Arabic also restrict extra-nulls to the bottom nucleus, but in addition allow them to occur only in words that are nouns. (An additional possibility for extra-nulls, dominating the first member of a geminate consonant, will be discussed shortly.)



Of course, some languages seem to impose no language-particular constraints at all on the distribution of extra-nulls. The dialects of Moroccan Arabic that allow

²⁸Kaye (1990a) mentions these dialects, though he has since denied their existence. They are, however, independently reported by Heath (1987).

²⁹In order not to sneak sparse structures like (3.81) in through the back door, some kind of cooccurrence constraints will have to be introduced for extra-nulls, or else they will have to be restricted to certain sorts of nodes, say, only nuclei.

consonant sequences like [ktb] would appear to be among these. Other likely candidates for this class of language include Berber (Dell and Elmedlaoui 1985) and Bella Coola (Bagemihl 1991).

Government Phonology models using proper government instead of the Adjacent Null Constraint run into the same difficulty with cases like *child*. But beyond the problems shared with Government Phonology, the present framework faces additional difficulties because of its hypothesis that apparent coda consonants are also onsets followed by empty nuclei. While GP could use (3.88) for the structure of the Moroccan Arabic word *xedma*, the present framework must use (3.89).



Again, there is an apparent violation of the Adjacent Null Constraint.


For these cases, it would be possible to invent yet another sort, say **coda-null**. But the evidence that the nuclei dominating coda consonants act qualitatively differently from other types of nulls is thin. In Moroccan Arabic, there is nothing to be gained by taking these nuclei to be anything other than **extra-nulls**. This involves adding a phonological condition to the list of Moroccan Arabic environments permitting extra-nulls, though we should note that one of the distinctive tendencies of extra-nulls still holds: they impose restrictions on the type of onsets they dominate—in this case, the onset must be the first member of a geminate.

The status of “coda” nuclei as **extra-nulls** is by no means universal. In French, for example, the nulls dominating coda consonants behave just like **simple-nulls**. Consider the contrast in the schwa that must be realized between *fortement* with a coda-onset cluster and the possible deletion in *sagement* with no such cluster:

- (3.90) *forte* *fortement* ‘strong/strongly’
sage *sagøment* ‘sensible/sensibly’


Extra-nulls are licensed word-finally, so the adjectival form *forte* can end in what phonetically appears as a two-consonant cluster [fort]. The “coda” consonant *r* is dominated by a **simple-null** nucleus, *t* by the word-final **extra-null**. The effects of the Adjacent Nulls Constraint are therefore escaped:

- (3.91) *forte* [fort]



But when the adverbial marker *-ment* is suffixed, *t*'s nucleus can no longer be extra-null, but it cannot be simple-null either without clashing with the *r*'s simple-null and violating the Adjacent Nulls Constraint, as in:


- (3.92) *fortement* as *[fortmã]



The only option is for it not to be null at all, but realized as a schwa. Charette (1988) needs two completely different theoretical devices to explain the coda case here and the *envie de te le demander* case discussed earlier. The present framework can handle both cases with the Adjacent Nulls Constraint, given only the assumption that French coda consonants are dominated by **simple-null** nuclei. In other words, the presence or absence of a coda-licence arc is irrelevant to the distribution of null sorts in French.³⁰

³⁰For the most part, it also seems sufficient for English coda consonants to be dominated by **simple-nulls**, but a handful of stubborn but by no means unnatural counterexamples suggest that a perfectly general grammar would need to use **extra-nulls**, e.g., *extra*, less obviously nativized Latinate forms like *abstract*, forms historically but not synchronically derivable from other words containing schwas, such as *curtsy* and the [kamftɹb!] variant of *comfortable*.

Although there is no need for a separate sort such as **coda-null**, it will be convenient to have a simple way to refer to the second nucleus of a CVV or CVC structure. I propose that a nucleus and its onset will both have the sort **rhyme-dependent** whenever either one of them is **null**. I choose the name because most, though not all, of the cases for which this sort will be useful correspond to what phonological frameworks have traditionally used the dependent position of a rhyme to represent. In PS diagrams, **rhyme-dependent** nodes will be represented using a superscript **rd**.




Unspecified

Unspecified nodes share with null nodes the property that they cannot have gestural features as dependents. An unspecified node could not be specified [s:Pal], for example. But they differ from null nodes in that their ultimate destiny is to be filled in with default values by the phonology-phonetics interface and pronounced as an articulatory constriction gesture. What default values will be used for **unspecified** nodes depends on the default rules of the particular language.

It should be emphasized that the term **unspecified** refers only to the lack of gesture features on the node. It does not mean that the lexicon (or rather lexical constraints) says nothing about the node. It is entirely possible for a lexical constraint to *require* a node to be **unspecified**, in the sense of not having gesture features yet not being null; the effect of this for the vowel is like having the vowel lexically specified as a schwa (or whatever the default for the language is). As will be illustrated later, having gesture features and being the object of interest of a lexical constraint are orthogonal properties. **Unspecified**, as a node sort, should be understood only to mean the absence of the first property (gestural features), and not the absence of the second (interest by the lexicon).³¹

This is the sort of node that alternates with null in the French and Moroccan Arabic examples. On the basis of their commonality of behaviour in this and several other cases, we can group **null** and **unspecified** together into a supersort, called **cold** after the “cold vowel” v^0 of Government Phonology, whose distribution in that theory is roughly the same as that of the **cold** supersort of the present framework. For example, the nucleus of the French words *de*, *te*, *le*, etc., is probably lexically specified as **cold**.

³¹ “Underspecified” may perhaps have been a better choice of term for this property, though not without problems itself.

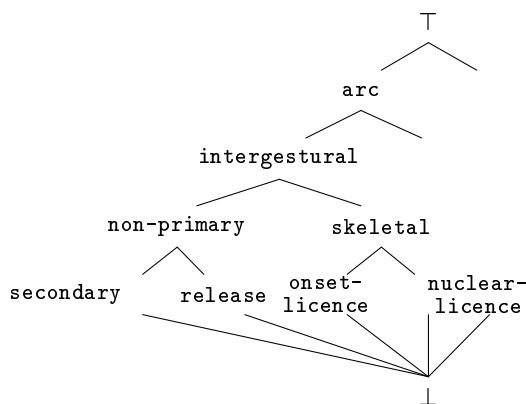
(3.97) *secours* [skur] or [səkur]

The difference lies in what the morphemes for the two words require of the PSs. For *score*, the morpheme demands that the initial nucleus be null, i.e., the morpheme contains the description *null*(v_1). *Secours* on the other hand makes no such demand on the first nucleus, which can alternate between *null* and *unspecified* according to the non-lexical principles of French.

3.3.2 Underspecifying government arcs

I have claimed that government relations as well as nodes can enter into subsort-supersort relations of the kind that have been diagrammed using sort lattices. For example, we might expect to find the *secondary* and *release* sorts of government arcs grouped together into a supersort called, say, *non-primary*. A partial sort lattice representing this situation might look like:

(3.98)



In this section, I offer evidence that such a situation actually exists. Indeed, there will also be evidence for an even higher supersort than *non-primary*, one including *secondary*, *release*, and *nuclear-licence*, or $\{2, r, n\}$. The alternation between two forms of Rotuman words can best be accounted for by proposing that each lexical

item specifies a particular government arc only as belonging to this supersort, leaving the decision of which subsort to the interaction of other constraints. A full analysis of Rotuman will need to use the concepts of locality and spreading introduced in section 3.4. In this section I shall only briefly describe some of the salient data and sketch how underspecification of government arcs is applicable to the solution.

Rotuman (Churchward 1940) is an Oceanic language closely related to Tongan, Maori, Samoan, and Fijian. Rotuman's main claim to fame among phonologists is a morphological process that seems to involve consonant-vowel metathesis. Attempts to find a representational solution that does not involve the massive transformational power necessary to carry out metathesis have been less than fully successful. McCarthy (1986), for example, needs to propose that Rotuman consonants and vowels are segregated onto different planes, despite the fact that this segregation has absolutely no morphological basis in Rotuman, unlike Semitic. The lack of a satisfactory autosegmental account has led researchers such as Hoeksema and Janda (1988) and Anderson (1992) to present Rotuman as another piece of clinching evidence that morphology needs the power to carry out transformational processes.

Almost every Rotuman word in a major lexical category has two forms, or in Churchward's terminology, "phases": a **complete** phase and an **incomplete** phase.³³ The incomplete phase of a word is generally predictable from its complete phase, though not vice versa. Three seemingly different processes are used to derive incompletes from completes:

- a) deletion of the final vowel
- b) deletion of the final vowel with umlaut on the preceding vowel
- c) metathesis of the final vowel and the preceding consonant

These three processes are illustrated in the following examples (the transcription has been adapted from Churchward's):

(3.99)	Deletion		
	Complete	Incomplete	
	<i>haga</i>	<i>hag</i>	'to feed'
	<i>tokiri</i>	<i>tokir</i>	'to roll'
	<i>hoto</i>	<i>hot</i>	'to jump'
	<i>hele'u</i>	<i>hele'</i>	'to arrive'

³³The phases have several uses in Rotuman and are often subject to semantically arbitrary syntactic rules and restrictions. Among perhaps the most prototypical uses of the complete phase is to mark definiteness in nouns and perfectiveness in verbs. The incomplete phase can be used for indefiniteness and imperfectivity, though it seems to be the semantically unmarked member of the pair and is used more frequently than the complete phase.

(3.100) Deletion and umlaut

Complete	Incomplete	
<i>mose</i>	<i>mös</i>	'to sleep'
<i>futi</i>	<i>füt</i>	'to pull'
<i>t</i>	<i>t</i>	'to sweep'

(3.101) Metathesis

Complete	Incomplete	
<i>seseva</i>	<i>sesēav</i>	'erroneous'
<i>hosa</i>	<i>hōas</i>	'flower'
<i>pure</i>	<i>pūr</i>	'to rule, decide'
<i>tiko</i>	<i>tīok</i>	'flesh'


From Churchward's phonetic description, it is fairly clear that the metathesized forms contain a light diphthong.


Informally, we can state the generalization deriving the incomplete phase from the complete phase as follows. If the final vowel is lower than the second last vowel, metathesis will take place. If it is a back vowel of the same or higher height than the second last vowel, it will be deleted. If it is a front vowel of the same or higher height, it will trigger umlaut and be deleted.

I shall assume that Velar site and [d:1] (i.e., highness) are default values for Rotuman vowels.

I shall not deal with the deletion class of (3.99). The generalization we can make about this class is that the final vowel, when present, is identical to the second-last vowel in both site and degree, or else has default values for site and degree. I propose that in this class the complete phase's final vowel has no independent lexical specification, but receives its values either by default rules or by vowel harmony. While this class is interesting for these reasons, it is irrelevant for the present purpose of illustrating the underspecification of government arcs.

Let us consider the umlaut class in (3.100). The complete phase form *futi* would have the PS in (3.102), while the incomplete phase *füt* would have the PS (3.103).





In *futi* in (3.102), v_1 is specified for its roundness. The velarity and height of u will be filled in by default rules. v_2 is specified for Palatal site. v_1 and v_2 are joined by an **n** arc.



In *füt* in (3.103), v_2 is now **null**. Nullness of the final nucleus is one of the chief characteristics of the incomplete phase. v_1 is still specified for roundness, but since there is now a secondary articulation of Palatal site, it will no longer receive velarity by default. It will still receive height ([d:1]) by default. It seems as if the node carrying [s:Pal] had jumped from the now-null v_2 to become a secondary articulation of v_1 .

We must ask ourselves what (3.102) and (3.103) have in common that would allow us to formulate a general lexical constraint for the verb-stem morpheme ‘pull’. Clearly the consonants remain constant, as does the root node of v_1 carrying labiality. The two phases also have in common a node bearing [s:Pal], though the exact position of this node is different. In (3.102), it bears a **nuclear-licence** relation to the root node of v_1 . In (3.103), it bears a **secondary** relation to the root node of v_1 . Positing a supersort encompassing the **nuclear-licence** and **secondary** sorts (let us for convenience call this supersort $\{2, n\}$) will allow us to state the lexical constraint of the verb root simply and directly. We call the moveable node x and require it to be joined to v_1 by an arc which belongs to the supersort $\{2, n\}$.

$$\begin{aligned}
 (3.104) \quad \forall w \quad w \xrightarrow{\text{semantics}} \text{“pull”} &\rightarrow \exists v_1, v_2, c_1, c_2, x \\
 w \xrightarrow{\text{phonology|top}} v_1 &\wedge c_1 \cdot v_1 \cdot c_2 \cdot v_2 \wedge \\
 c_1 \approx \mathbf{f} \wedge c_2 \approx \mathbf{t} &\wedge \\
 v_1 \xrightarrow{\mathbf{a}} \text{Lab} &\wedge \\
 x \xrightarrow{\mathbf{s}} \text{Pal} &\wedge \\
 v_1 \xrightarrow{\{2, n\}} x &
 \end{aligned}$$

In the complete phase, v_1 will have an **n** arc to x , so x will necessarily be identified with v_2 . In the incomplete phase, where v_2 must be **null**, v_1 and x can only be joined by a **2** arc.

The situation with the metathesis class is similar. The PS for the complete phase *tiko* and the incomplete phase *tïök* of ‘flesh’ are:³⁴



Again, everything is the same between the two PSs except for the position of a single node, which we can call x , and its dependent, a secondary ATR articulation. In the complete phase x and v_1 are joined by an n arc. In the incomplete phase, they are joined by an r or **release** arc, resulting in a light diphthong. Using a similar strategy, assuming a supersort containing both **nuclear-licence** and **release**, we can formulate the general lexical constraint for the morpheme ‘flesh’:

$$(3.107) \quad \forall w \quad w \xrightarrow{\text{semantics}} \text{“flesh”} \rightarrow \exists v_1, v_2, c_1, c_2, x, y \\
w \xrightarrow{\text{phonology|top}} v_1 \wedge c_1 \cdot v_1 \cdot c_2 \cdot v_2 \wedge \\
c_1 \approx \mathbf{t} \wedge c_2 \approx \mathbf{k} \wedge \\
v_1 \xrightarrow{s} Pal \wedge \\
x \xrightarrow{d} 2 \wedge y \xrightarrow{a} Rad \wedge y \xrightarrow{d} A \wedge x \xrightarrow{2} y \wedge \\
v_1 \xrightarrow{\{n,r\}} x$$

³⁴Recall that light diphthongs, like complex onsets, are represented using a release node dominated by the root node via an r or **release** arc.

Having seen the two supersorts, $\{2, n\}$ and $\{n, r\}$, we might ask if we are not dealing with a single supersort $\{n, 2, r\}$. In fact, it is completely predictable which stems will use 2 arcs in the incomplete phase and which will use r . Better yet, the prediction is an automatic consequence of the segmental model outlined in chapter 2. The segmental model disallows recursive secondary articulations. That is, the following structure is illegal:


$$(3.108) \quad * \quad \begin{array}{c} x \\ |2 \\ o \\ |2 \\ o \end{array}$$

The generalization covering the distinction between the umlaut and metathesis classes is this: If the x connected by the underspecified arc is simplex, that is, if it is just one node and does not dominate any secondary articulations, the arc will be specified as 2 in the incomplete phase, resulting in an umlaut class word. If, on the other hand, the x is complex and has its own secondary articulation, specifying the underspecified arc as 2 in the incomplete phase would result in an illegal structure with recursive 2s, so the only choice is to specify the arc as r , resulting in a metathesis class word.

To sum up this brief analysis of Rotuman, we have seen that the lexical constraint of a Rotuman word may (but need not) require a dependent joined to the second-last nucleus via an underspecified arc of the supersort $\{n, 2, r\}$. In the complete phase of the word, this clause of the lexical constraint will be satisfied by an n arc. In the incomplete phase, it will be satisfied by either a 2 arc or a r arc, the choice being predictable from independent principles of universal grammar.

This analysis shows an advantage of the present framework over more standard frameworks. In a framework where relations between nodes are uniformly represented by government arcs and “morphemes” are constraints, it is an easy matter for a morphemic constraint to generalize over the different types of government relations necessary. In a more standard framework, where morphemes are actual chunks of phonological representation, there is no principled way for the representation to be underspecified in such a way that a particular piece of it can be sometimes a tier adjacency relation and sometimes an association line. If one relies on theories of representations and of morphemes that are unable to make the needed underspecifications, one has little choice but to use a transformational rule of metathesis, as for example Hoeksema and Janda (1988) specifically argue for Rotuman. This in turn raises serious questions about the constrainedness of the inventory of possible rules.

Another pleasant effect of the Rotuman analysis is the support it gives for a necessary consequence of the proposals on tone in section 3.2.2. We need exactly this kind of underspecification to account for alternations between two level tones spread out over two vowels and a contour tone compressed onto one vowel:



In languages with this kind of behaviour, the lexical constraints for the relevant morphemes will need to underspecify the relation between v_1 and v_2/x as belonging to a supersort that contains (at least) **nuclear-licence** and **release**. Rotuman demonstrates that this is not a peculiarity of tone. The same phenomenon can affect segmental melodies as well.

Representing natural classes

The mechanism of subsort-supersort subsumption seems somewhat too powerful to be the best solution to the underspecification behaviour of government arcs. It would be formally possible to take a random set of government sorts and group them into a supersort. I do not believe this type of power is generally possible for languages. Instead, it seems likely that the inventory of natural classes of government types is universally limited.

Ultimately, an ideal theory would be able to account for the natural class behaviour of government relations in exactly the same way it deals with natural classes of phonemes, i.e., with features. Instead of being a primitive, unanalyzable entity, a government relation would have an internal structure made up of these features. A language that wished to refer to the supersort $\{\mathbf{n}, \mathbf{2}, \mathbf{r}\}$ could simply refer to the feature that **nuclear-licence**, **secondary**, and **release** had in common. Of course, before we could devise a feature system like this, we would need a better idea of what natural classes of government relations actually play a role in the grammars of languages. In the meantime, I shall use the mechanism of sorts as a temporary expedient to describe these natural classes.

3.3.3 Metrical structure

Recent approaches to prosodic representation

One of the central questions in recent metrical theory has been the appropriate type of representation for stress patterns. Two of the most examined alternatives have been the metrical **grid** (e.g., Liberman 1975, Dell 1984, Prince 1983, Selkirk 1984) and the metrical **tree** (e.g., Liberman and Prince 1977, Hayes 1980, 1991, Hammond 1984).

In a metrical grid, each potentially stress-bearing position in a string can have a number of projections on higher levels, the number of projections determining the relative degree of stress. For example, the English place name *Apalachicola* might have the following grid representation:


(3.110)

					*
*	*	*			
*	*	*	*	*	*
Apalachicola					

An asterisk on a level marks a stress-bearing position that has been projected to that level. The *o* that bears main stress, being projected to three levels, is more prominent than the initial *a*, which is only projected to two.

Many stress systems have phenomena, such as stress clash avoidance, that seem best handled by rules that simply move asterisks around in the metrical grid.

The representation of *Apalachicola* in the tree-based system of Hayes (1980) would be:



The tree is made up of binary constituents, which can be recursively embedded. Each branch of a constituent is labelled as to whether it is strong or weak; each constituent must have one of each.³⁵ The relative prominence of a syllable is determined by an algorithm that uses the dominating node labels.

More recent work in tree-based metrical theory has developed a typology of the possible types of metrical constituents or **feet**. Hayes (1985), followed by McCarthy and Prince (1986), proposes that universally there are only three possible types of foot: the **iamb**, consisting of a light syllable followed by a heavy syllable, the **syllabic trochee**, consisting of two syllable where the second is not heavier than the first, and

³⁵There is also the possibility for degenerate constituents, having only one branch which is by convention considered strong.

the **moraic trochee**, consisting of a heavy syllable or two light syllables, i.e., two morae, regardless of the number of syllables they are in.

(3.112)	iamb	$[[\sigma\mu][\sigma\mu\mu]]$	right-headed
	syllabic trochee	$[\sigma\sigma]$	left-headed
	moraic trochee	$[\mu\mu]$	left-headed, $[\sigma\mu\mu]$ or $[\sigma\mu][\sigma\mu]$

McCarthy and Prince (1986, 1990a, 1990b) discuss the role these three foot types play in morphology. Hayes (1991) and Kager (1992) have argued for collapsing the two types of trochees into a single foot type, the generalized trochee, the difference between them being determined by whether the generalized trochee is built on the mora level or the syllable level.

A tree formalism allows a clear expression of rules that clearly depend on metrical constituents as constituents (e.g., a vowel harmony rule that operates only within a foot), but it makes it much more complicated to deal with things like stress clash avoidance that the grid formalism handles easily.

A sort of middle ground is taken by Halle and Vergnaud (1987), who recognize the importance of both metrical constituency (like tree formalisms) and a direct representation of headship (like grid formalisms). Their representation of *Apalachicola* is:

(3.113) * .
	(* . * . *).
	(* *) (* *) (* *)
	Apa lachi cola

On any line of the grid, constituency is represented by parentheses. The head of each constituent is projected to the next level of the grid, and marked there by an asterisc. A central claim of Halle and Vergnaud is that constituency and headship are separate, though mutually constraining, phenomena. It is worth discussing here some of those aspects of the Halle-Vergnaud system that will play a role later in the section.

Line 0 is the foundation of the metrical grid. Every segment of the string that is potentially stress-bearing is projected on this line. Line 0 positions are grouped into constituents and the head is projected to line 1. Line 1 positions are then grouped into constituents, the head is projected to line 2, and so on. Apart from the choice of which positions qualify for line 0 status, Halle and Vergnaud argue that the basics of all stress systems can be derived from a handful of parameters. Their proposed parameters are:

- (3.114) $[\pm\text{HT}]$ head terminal: is the constituent head adjacent to one of the constituent boundaries?
 $[\pm\text{BND}]$ bounded: is the head separated from the constituent boundaries by at most one element?
 $\left\{ \begin{array}{l} \text{left to right} \\ \text{right to left} \end{array} \right\}$ direction of foot construction
 $\left\{ \begin{array}{l} \text{left} \\ \text{right} \end{array} \right\}$ headedness: if $[\text{+HT}]$, are the heads at the right or left boundary?

Each line of the metrical grid can have different parameter settings. For instance, line 0 constituents could have bounded and left-headed constituents constructed right to left, $[\text{+HT}, \text{+BND}, \text{left}, \text{right-to-left}]$, while line 1 has an unbounded right-headed constituent, $[\text{+HT}, \text{-BND}, \text{right}]$. (In this case, assuming there is no extrametricality, the result will be a penultimate stress system.)

The basic pattern of constructing the metrical grid is a three-part rule for each line i :

- (3.115) a. parameter settings for line i are...
 b. construct constituent boundaries on line i
 c. locate the heads of line i constituents on line $i + 1$

Interspersed among these three basic statements, a language may include rules for any special processes that have to be carried out in order to get the right structure: marking something extrametrical, placing higher-line asterisks according to principles that have nothing to do with foot construction (e.g., lexical accents, accents for heavy syllables), or conflation (deleting a lower line, leaving only the next higher line's head marked).

For example, Latin would have the following rule battery:

- (3.116) a. Mark the final syllable extrametrical.
 b. Assign line 1 asterisks to syllables with branching rhymes.
 c. Line 0 parameter settings are $[\text{+HT}, \text{+BND}, \text{left-headed}, \text{right to left}]$.
 d. Construct constituent boundaries on line 0.
 e. Locate the heads of line 0 constituents on line 1.
 f. Line 1 parameter settings are $[\text{+HT}, \text{-BND}, \text{right-headed}]$
 g. Construct constituent boundaries on line 1.
 h. Locate the heads of line 1 constituents on line 2.
 i. Conflate lines 1 and 2.

Agricola 'farmer (nom.)' would have its metrical structure built as follows. First, (a) would mark the final syllable extrametrical. Halle and Vergnaud use angle brackets to represent this in their diagrams:

(3.117) * * * -- line 0
Agrico <1a>

Since *agricola* has no branching rhymes, (b) does not apply. (d) constructs line 0 constituent boundaries using the parameter settings in (c)—maximally binary constituents, built right to left:

(3.118) (*)(* *) -- line 0
Agrico <1a>

(e) projects the heads (the left elements) of line 0's constituents to line 1:

(3.119) * * . -- line 1
(*)(* *) -- line 0
Agrico <1a>

Using the parameters in (f), (g) now constructs an unbounded constituent on line 1:

(3.120) (* *). -- line 1
(*)(* *) -- line 0
Agrico <1a>

And (h) projects the head (the right-hand element) to line 2.

(3.121) . * . -- line 2
(* *). -- line 1
(*)(* *) -- line 0
Agrico <1a>

The conflation instruction in (i) now applies to delete line 1,³⁶ leaving behind only the head-marking on line 2, which marks the main stress of the word:

(3.122) . * . -- line 2
(*)(* *) -- line 0
Agrico <1a>

In contrast, the word *agricolārum* 'farmer (gen.pl.)', stressed on the long \bar{a} , has a branching rhyme that rule (b) marks with a line 1 asterisk before any construction begins on line 0:

³⁶Another way of looking at this would simply be to say that line 1's structure has no phonetic effect.

(3.123) . . . * -- line 1
 * * * * -- line 0
 Agricol a <rum>

Constituent construction in (d) must respect this pre-marked line 1 head:

(3.124) . . . * -- line 1
 (*) (* *) (*) -- line 0
 Agricol a <rum>


After (e)-(h) have applied, *agricolārum* has the structure:

(3.125) . . . * -- line 2
 (* * . *) -- line 1
 (*) (* *) (*) -- line 0
 Agricol a <rum>

Conflation in (i) will delete line 1, leaving behind the line 2 asterisc that marks the word's main stress.

A possible representation of prosody

At first glance it seems it would be very difficult to incorporate any representation of metrical or prosodic information into the kind of phonological structures I have been arguing for. One could of course create new node types—say, foot (F) and superfoot (SF)—and use them to build a metrical tree over the nuclear skeleton, somewhat as in (3.126):




Each metrical node in a PS would have an arc to the **head** and to the **dependent** or **complement** of the constituent it represents.

There are a few unattractive aspects to such a decision. Foremost among these is the fact that the PS nodes, F_1 , F_2 , and SF, are the only ones we have seen so far which cannot be understood as representations of an articulatory gesture. While every other sort of node—nucleus, onset, release, secondary, etc.—can at least potentially be specified for articulator, site, and degree features, these proposed metrical nodes crucially could never be.

At various points, I have touched on how the representation of asymmetric relations can be accomplished in two different ways. For example, the asymmetric relations between parts of a syllable could be represented either through the mediation of superordinate nodes like Onset, Coda, and Rhyme, as in (3.12), or directly, as in (3.13), the choice of Dependency Phonology. We have more or less been using the Dependency Phonology approach for relations below the level of the syllable, but we have also taken some steps towards using it for the supra-syllabic asymmetries found in metrical systems as well. The basic linear relationship between one “syllable” and the next is not mediated by any special σ node on its own tier. Rather, we have been using a **nuclear-licence** arc that directly connects the *heads* of the two syllables, that is, the primary constriction gestures of the nuclei. It is worth exploring whether metrical asymmetries can be handled the same way.

Tentatively using **f** to label the relation between head and dependent of a foot, and **sf** for a superfoot, (3.126) could be redrawn as:


(3.127)



In such diagrams, as with all the others we have been using so far, the “head” position of a constituent stands in for the whole constituent for the purposes of relating it to other constituents.

A diagram like (3.127) is hard to read. It is easier to see what the relations are if we modify it slightly, stretching out into rectangles the points that represent the nuclei:

(3.128)



The similarity between (3.128) and more familiar representations of the metrical grid as in (3.110) should be obvious.

In a PS such as (3.128), it is not necessary to have a nuclear position “project” onto some other level of representation or linked up autosegmentally to an independent entity (e.g., SF) on another tier. The v_1 that foot-dominates v_2 is the same node that superfoot-dominates v_3 . It simply behaves as a different sort of position for the purposes of different sorts of government relations—and the word “sort” here can be used in its technical sense. That is, v_1 , as well as having the sorts **nucleus** and **constriction**, also has the sorts **foot-head** (or **line-1**) and **superfoot-head** (or **line-2**).

This type of approach also captures Halle and Vergnaud’s insight that metrical constituency and headship are separate, though mutually constraining, aspects of phonological structure. The property of the PS that is relevant to answering the question “Do A and B form a constituent?” is whether or not A and B are joined by the appropriate metrical government arc. “Is A a line- i head?”, on the other hand, questions whether or not the primary constriction gesture of A has the appropriate sort. Constituency deals with prosodic arcs, headship with node-sorts.

In what follows, I shall assume that nuclear positions can receive sorts such as **line-0**, **line-1**, **line-2**, and so on. In PS diagrams, I shall designate a node with sort **line-0** with a superscript 10 and a node with sort **line-1** with a superscript 11.

Languages may differ in the sorts of positions they allow to be sorted **line-0**. In all, a **line-0** position must be **nuclear**.³⁷ Some languages may further restrict line-0 to nuclear positions that are not **rhyme-dependent**, as in:


$$(3.129) \quad \begin{array}{ccccccccc} v^{i0} & \text{---} & v^{rd} & \text{---} & v^{i0} & \text{---} & v^{i0} & \text{---} & v^{\emptyset,rd} & \text{---} & v^{i0} \\ & & \swarrow & & | & & | & & \swarrow & & | \\ & & \mathbf{a} & & \mathbf{u} & & \mathbf{e} & & \mathbf{c} & & \mathbf{o} \\ & & & & & & & & & & | \\ & & & & & & & & & & \mathbf{d} \end{array}$$

This results in a quantity-insensitive stress system—loosely speaking, one built strictly on syllables rather than morae.³⁸ It is also possible to allow **line-0** to be **constriction**


³⁷Everett and Everett (1984) have argued that stress assignment in the Amazonian language Pirahã must take segmental properties of the onset into consideration. Though I see no elegant straightforward way of accounting for Everett’s data, it seems unlikely that even Pirahã will require onsets to be sorted **line-0**.

³⁸Another (I believe, better) analysis of quantity-insensitive systems using trochees is to have all nuclei, including **rhyme-dependent** ones, as **line-0** positions, and build syllable-bound moraic trochees (that is, feet whose dependent are constrained to be **rhyme-dependent**) on line 0. This in effect gives each “syllable” a single line 1 position, which can then be used to build the quantity-insensitive trochees on line 1. This analysis is easier to incorporate into a parametric theory of metrical structure of the kind discussed in Russell (1993).


but not null:



This results in a system where CVV syllables are bimoraic, while CVC syllables are monomoraic. Imposing no restrictions on line 0, thus allowing null nuclei to be **line-0** as well, results in a system where coda consonants are also moraic.



It will be convenient to have government arcs that represent adjacency on metrical lines. With these, a line 1 position can immediately dominate the next line 1 position, in effect skipping over any line 0 dependents the two may have. In PS diagrams, I shall label these relations **10g**, **11g**, and so on. With these relations included, (3.128) will look like:



Metrical sorts and full values

It should be pointed out that sorts like **line-1** are first-class citizens of PSs. It is not the case that the only effect of having a metrical sort on a node is its degree of phonetic stress. Metrical sorts can also affect the well-formedness of other parts of the PS. The clearest example of this will be the prosodic morphology phenomena

discussed in chapters 5 and 6. Another common effect of metrical sorts is in controlling the realization of a node's lexically determined full segmental value.

Lexical constraints do not necessarily specify gesture features directly by using the *arc* predicate. Rather, they may specify a gesture indirectly and conditionally, using the predicate *full-value*, a three-place relation taking a **constriction** node, an arc-sort, and a gestural atom. If the constriction node receives the sort **full**, then it will have all the arcs specified for it by its *full-value* predicates; otherwise it will likely end up as **unspecified**. This is done by the constraint:

- (3.133) A node x will have an articulatory gesture arc to some y if and only if y is x 's full-value for that arc type and x is full.
 $\forall x, y \ x \xrightarrow{g} y \wedge \text{gestural}(g) \leftrightarrow \text{full-value}(x, g, y) \wedge \text{full}(x)$

Just as I have used $c \approx \mathbf{k}$ to abbreviate the set of description language *arc* clauses needed to minimally specify c for the phoneme /k/ in some language, I shall use $c \overset{f}{\approx} \mathbf{k}$ to abbreviate the set of *full-value* clauses needed to accomplish the same thing.

In languages where not every node is automatically a full position, it is usually metrical structure that determines which are and which are not. A common restriction is that all and only **line-1** vowels are **full**.³⁹

For example, the lexical constraint of the English lexeme *photograph* would specify its vowels as:

- (3.134) $v_1 \overset{f}{\approx} \mathbf{o} \wedge v_2 \overset{f}{\approx} \alpha \wedge v_3 \overset{f}{\approx} \mathbf{æ}$

In the plain form *photograph*, v_1 and v_3 are heads on line 1 (v_1 being a head on line 2 as well); v_2 is not. So only v_1 and v_3 are **full**. Only they have arcs going to their full values specifications, [o] and [æ], in the resulting PS. v_2 is left unspecified, and is phonetically realized with default schwa.

In the suffixed form *photographer*, however, the metrical structure is different. Now only v_2 is **line-1** and **full**. Only v_2 will have gestural features in the PS. v_1 and v_3 will be unspecified and receive default schwa in phonetic interpretation.⁴⁰

Other languages with stress-related full/non-full alternations include Tonkawa and many dialects of Ojibwe.

³⁹I assume that all onsets and secondary positions are **full**, though some languages may make this conditional on the dominating nuclear position being full as well.

⁴⁰Constraint (3.133) might be weakened to a simple conditional, so that it only requires a full value feature on full positions and does not prohibit full values on non-full positions. Apparently, this is a matter of some variation. In my dialect, at least, schwa realization of non-full positions is obligatory: [fotágræfr] is not a careful pronunciation, it is simply not English.

Characterizing prosodic feet

Just as heads in Halle and Vergnaud's system can be handled by giving a node a sort like **line-1**, constituents could be handled just as straightforwardly. For instance, there might be an arc of sort **line-1-constituent** between two nodes of sort **line-1**, with the dominant node also being of sort **line-2**. This is certainly a possibility, but ideally we should like to have a model that also incorporates the insights of Hayes' limited inventory of feet and can use the same mechanism to handle prosodic conditions of the kind described by McCarthy and Prince.

The easiest foot type to integrate into such a system is the moraic trochee: $[\sigma\mu\mu]$ or $[\sigma\mu][\sigma\mu]$. A moraic trochee is simply a pair of adjacent **line-0** nodes, where the first is the head. The trochee is "formed" by a special sort of government relation between the two elements of the foot. Since there is already a government relation between the two nodes (10g), we can treat this **moraic-trochee** government as a subsort of 10g government. In PS diagrams, I shall label this arc $\mu\mathbf{t}$.

$$(3.135) \quad v^{i0,i1} \xrightarrow{\mu\mathbf{t}} v^{i0}$$

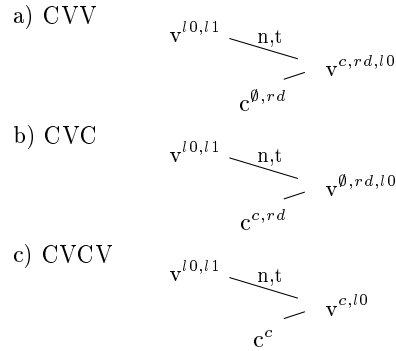
The requirement of headship can be formalized by the following appropriateness condition:

$$(3.136) \quad \text{If two line-0 nodes are joined by a } \mathbf{t} \text{ arc, then the first is of sort } \mathbf{line-1} \\ \text{and the second is not.} \\ \forall v_1, v_2 \quad v_1 \xrightarrow{\mu\mathbf{t}} v_2 \rightarrow \mathbf{line1}(v_1) \wedge \sim \mathbf{line1}(v_2)$$

Note that this constraint also enforces binarity. The only way to have a ternary trochee would be to have a chain of $\mu\mathbf{t}$ arcs. But this would require the middle node, as an arc-head, to be **line-1** while at the same time requiring it, as an arc-tail, not to be.

A number of structures can satisfy the requirements for forming a moraic trochee. A single long vowel will suffice, as in (3.137a)'s CVV syllable. If a language also counts the empty nucleus that dominates a coda consonant as **line-0**, then a CVC syllable can be a trochee too, as in (3.137b). And a sequence of light syllables, CVCV, can also count as a trochee, as in (3.137c), though individual languages may rule out this possibility by requiring the dependent to be **rhyme-dependent**. In the following PS diagrams, a superscript \emptyset designates a null position, a superscript c represents a constriction (i.e., non-null) node, and a superscript rd a **rhyme-dependent** node.

(3.137)



It should be noted that both nuclear positions of a moraic trochee, even a CVC trochee, belong to the sort **line-0**. This is a more liberal definition of line 0 than Halle and Vergnaud's, who would limit membership to potentially stress bearing positions and would therefore rather assign the entire trochee a single line 0 position and specially mark it by rule with a line 1 asterisc. It is true that languages frequently require a **rhyme-dependent** nucleus to be the dependent in a moraic trochee relation as well, though languages that can stress the second mora of a long vowel show that this is not a universal requirement.⁴¹ We do not need a special rule to project a line 1 asterisc over branching rhymes (which formally speaking could just as easily have been a rule to project a line 1 asterisc over non-branching rhymes). The presence of a line 1 asterisc, or the sort **line-1**, is an automatic consequence of constructing the trochee.

With this theoretical equipment, we are now ready to deal with the syllabic trochee. Kager (1992) argues that the difference between the syllabic trochee and the moraic trochee lies in whether the process that constructs the feet scans the moraic tier or the syllable tier. I propose that moraic trochees are constructed on line 0 and syllabic trochees are constructed on line 1.⁴² There are several cases in Halle and Vergnaud (1987) that require a left-headed binary constituent on line 1 built over moraic trochees on line 0. If the language imposes no restrictions on the dependent of a line 0 trochee, a line 1 trochee built over this foundation will simply result in one of these familiar cases, whose phonetic effect is typically an alternating pattern of secondary stress. If, on the other hand, the language requires the dependents of its line 0 trochees to be

⁴¹While CVV is possible, we seem never to find CVC. A null nucleus apparently cannot be a line 1 position under normal circumstances. (Abnormal circumstances may include the edge of a word, where in Moroccan Arabic a null can be a line-1 position, as we shall see in chapter 6.) Ideally, the fact that null and line-1 usually do not cooccur should follow from a theory of the natural classes of node sorts.

⁴²This account presupposes the refinement suggested in a previous footnote, where all nucleus nodes, including **rhyme-dependent** nodes, are allowed to be **line-0**.

rhyme-dependent, then building line 1 left-headed constituents over this will result in syllabic trochees. I shall use the label $\sigma\mathbf{t}$ for any arc forming a left-headed binary constituent on line 1.

Hayes (1991) and Kager (1992) also argue for a mode of foot construction using a “generalized trochee”, where either a moraic trochee or a syllabic trochee is built, depending on the environment. Though I offer no analyses of their examples here, I think it is well-established that $\mu\mathbf{t}$ and $\sigma\mathbf{t}$ often act as members of a super-sort, which we can call simply **trochee** or \mathbf{t} . In fact, if we accept the existence of the sort **trochee**, we need no longer treat $\mu\mathbf{t}$ and $\sigma\mathbf{t}$ as primitives. Rather, $\mu\mathbf{t}$ is simply the intersection of **trochee** and $10\mathbf{g}$, and $\sigma\mathbf{t}$ of **trochee** and $11\mathbf{g}$. For clarity, I shall often continue to use the more specific labels.

Representing the iamb in the present model poses more of a challenge. Unlike the trochee, the iamb cannot simply be treated as two elements joined by an arc. There is crucially a third element present at some level: the element that makes the right-hand syllable of the iamb heavy. The role played by this third element is lost if the iamb is expressed as a relation between one syllable head and another. (This cannot be avoided by taking the iamb to be a relation between two syllables, since in the direct-government framework argued for here, a relation between two syllables *is* a relation between their heads.)

It would seem that we must constrain the second member of an iamb to have a rhymal complement. Yet the only way to keep the iamb binary would be to banish this rhymal complement from line 0. Under this proposal, a stretch of the nuclear spine of a PS from an iambic language would look like:

$$(3.138) \quad \mathbf{v}^{i0} \xrightarrow{10\mathbf{g},i} \mathbf{v}^{i0,i1} \xrightarrow{n} \mathbf{v}^{rd} \xrightarrow{n} \dots$$

While a rhymal complement in an iamb cannot, under this account, be of sort **line-0**, we have seen that a rhymal complement in a trochee should be. (3.138) would be an acceptable solution if all languages using iambs kept rhymal complements off line 0, while all languages using trochees kept them on. There are, however, languages that use both. McCarthy and Prince’s (1990b) account of the broken plural in Classical Arabic, for example, crucially relies on the ability to parse a portion of a string into either an iamb or a trochee. Each CA nucleus must then be **line-0** to be available for parsing as a trochee, but then (3.138) could not possibly be the structure of a CA iamb.

Fortunately, there is a way to require both that the second element of an iamb be heavy and that the position that makes it heavy be of sort **line-0**. If the second member of an iamb is itself a moraic trochee, it must of necessity be heavy. Thus the head of an iamb enters into two prosodic relations: a trochaic relation with the line 0 position to its right and an iambic relation with the line 0 position to its left.⁴³

⁴³It would seem more natural for the head of the iamb also to be the origin of the **i** arc.

$$(3.139) \quad v^{l0} \xrightarrow{i} v^{l0,l1} \xrightarrow{t} v^{l0,rd}$$

There is an immediate problem with this solution: it is possible for a moraic trochee to consist of two light syllables, as in (3.137c), as well as one heavy syllable, as in (3.137a,b). So it should be possible for an iamb to have the unlikely form CVCVCV. It is certainly possible for a language to place extra restrictions on its *t* government, requiring its dependent to be **rhyme-dependent**. But it would be suspicious if every language using iambs happened to do this. Fortunately, there *are* languages that can have feet like CVCVCV. This type of foot—three morae with the central one as head—has been called a **ternary foot** or an **amphibrach**.


Following Levin (1988), Halle and Vergnaud (1987:25–28) offer an analysis of the Bolivian language Cayuvava that uses amphibrachs, the only kind of ternary feet their system of parameters allows. In Cayuvava, stress falls on every third mora counting from the end of the word. This can be accounted for by marking the last mora extrametrical (in the present system: $word \xrightarrow{\text{bottom}} v \leftrightarrow \sim line0(v)$) and constructing amphibrachs from right to left.⁴⁴ The parameters resulting in this state are [-HT] (not head terminal, heads need not be adjacent to a constituent boundary) and [+BND] (bounded, each dependent must be adjacent to the head, or the head must be at most one element away from the boundary). The words *cáadiróbo maráhahaéiki* receive the following structures:

$$(3.140) \quad \begin{array}{ll} * . . * . . * . & \text{-- line 1} \\ (9\ 8)(7\ 6\ 5)(4\ 3\ 2)\langle 1 \rangle & \text{-- line 0} \\ \text{ca a dirobo Bururu ce} & \end{array}$$

$$(3.141) \quad \begin{array}{ll} . * . . * . & \text{-- line 1} \\ (7\ 6\ 5)(4\ 3\ 2)\langle 1 \rangle & \text{-- line 0} \\ \text{maraha ha e i ki} & \end{array}$$

Under the present proposal, each foot would have the general structure:

This would result in a structure like





I have chosen the first method, without great conviction, purely to keep PSs acyclic. A fuller consideration of the implications of the model of reduplication proposed in section 4.5 will likely require the cyclic version of the iamb.

⁴⁴See Halle and Vergnaud (1987:27–28) for arguments in favour of using extrametricality as opposed to constructing dactyls, feet with the form (*..).

$$(3.142) \quad v^{l0} \xrightarrow{i} v^{l0,l1} \xrightarrow{t} v^{l0}$$

which is, in all relevant respects, identical to the schema for the iamb in (3.139). Using a grid-like diagram as in (3.128), the two words would have the following structures:



The amphibrach is simply a special case of the generalized iamb—or perhaps more accurately, the iamb is a special case of the generalized amphibrach. The stress pattern of Cayuvava can thus be accounted for without the need to introduce a new kind of foot, an otherwise unneeded parameter ($[\pm\text{Head Terminal}]$), or a new procedural device.⁴⁵

It should be emphasized that the foregoing is only one possible way that metrical information could be incorporated into a constraint-based approach to phonology. Other alternatives are conceivable. Regardless of the merits or demerits of the particular model I have outlined in this section, there are still excellent reasons for an approach along these general lines. I shall end this section with a brief discussion of some of the general advantages that a constraint-based approach to prosody has over a proceduralist construction approach.

First, there is no need for lexical entries to contain diacritic marks whose purpose is to guide the derivation of irregular or unpredictable items. If a word has an accent, in the sense of Halle and Vergnaud, the lexical description of this fact uses exactly the same vocabulary as any other part of the grammar that deals with stress. If Halle

⁴⁵Hayes' (1991) account of Cayuvava uses moraic trochees, but relies crucially on the procedure that builds them having the ability to skip a mora after every construction, leaving the skipped position without prosodic structure. See Russell (1993) for some arguments against this analysis.

and Vergnaud's device of allowing lexical entries to come with some line 1 asterisks pre-attached seems slightly ad hoc, it is perfectly natural in the constraint-based adaptation: in the model of this section, the description language predicate describing such a situation is *line1*, and this predicate is just as available for use in "lexical" constraints as it is for any other class of constraints in the grammar. The constraint merely checks for the presence of a line 1 position, it does not have to create one.

Secondly, as in other areas of phonology, many stress systems seem to defy analysis in terms of a linear deterministic application of ordered rules. There are cases where even a proceduralist model must construct two different candidate solutions in parallel, then decide between them by means of a condition. For example, Halle and Vergnaud (1987:24) offer an analysis of the stress pattern of Yidin^y that contains the following two ordered rules, from their example (46):

- a. Line 0 parameter settings are [+HT, +BND, left to right] and [right] (that is, right-headed) if the word contains an even-numbered syllable with a long vowel; otherwise, [left] (that is, left-headed).
- b. Construct constituent boundaries on line 0.

But, as they remark:

the setting of the parameter that determines whether constituents are left- or right-headed requires information about the position of long vowels relative to constituent boundaries—in other words, information that is not available until rule (46b) has applied. There is no contradiction here. We shall assume that (46a) is formally implemented by constructing metrical constituents on two planes simultaneously and deleting the inappropriate one by a subsequent rule. As we shall see below, construction of metrical constituents on two planes simultaneously is required in a number of other languages such as Tiberian Hebrew...and Pirahã...

The formal implementation of constructing simultaneous metrical planes is not spelt out, nor is the rule that decides between the competitors and deletes the undesired one. It is not immediately clear how such a process *could* be formalized within the framework of autosegmental phonology. There is no discussion of the constraints on when a derivation can and cannot pursue multiple paths, or on how it is ensured that only one derivation path survives, since the point at which the candidate paths are judged can conceivably be quite distant from the point at which they split.

This is just the kind of situation that constraint-based frameworks excel in. A constraint-based model is concerned only with judging candidate forms. Unlike a proceduralist model, there is no need to worry about *constructing* the candidates as well. Yidin^y, Tiberian Hebrew, and Pirahã do not force our model to add a device that creates unnatural states where a linguistic item hovers, like Schrödinger's cat, between two possibilities, waiting for a constraint to come along and choose between them. There is no need to add such a device because, in a sense, this is already the natural state of *all* items in a constraint-based framework.


3.4 Local domains and spreading

The work of Bird and Klein (1990) and Scobbie (1991) has shown the inadequacy of traditional conceptions of tiers with regard to the information about temporal ordering that they contain. We saw in chapter 2 the inadequacy of traditional conceptions of tiers when it comes to juggling multiple specifications of the same kind attached to the same segment. Yet the notion of tier gives a tempting account of how two specifications interact with each other: they can interact because, at some level of representation, they are adjacent.

If we are to abandon tiers, we will need some other way of expressing this “adjacency” of interacting nodes. Fortunately, one way immediately suggests itself: two positions in a PS can interact with each other if they stand in a government relation. Obviously, not every government arc in every PS induces adjacency effects. By and large, the segmental contents of the nuclei in an English word might as well be in separate universes for all they interact. Yet in other languages, such as those with rich vowel harmony systems, the nuclei show a high degree of interaction, sharing each other’s features.

I argue that these domains of interaction are the result of a primitive property of government arcs. Normally, two positions will have their own separate specifications. But if the government arc linking them has the primitive property of inducing a local domain with respect to some government type, the two positions will share their specifications for that government type. In PS diagrams, the set of government types that an arc creates a local domain for will be indicated in braces after the arc’s label. For example, in (3.145) nuclear arc n_7 creates a local domain for articulator, site, but not for degree.

(3.145)



In the description language, this primitive property of arcs is indicated by the predicate *local-domain-creator*, a two-place relation taking as arguments an arc and a sort of arc:⁴⁶

⁴⁶ It might seem that we are using a predicate (e.g., *site*) as the argument of another predicate (*local-domain-creator*) and that we would need a second-order description language to handle this. But recall that the description language notation *site(x)* is just an abbreviation for *sort(x, site)*. That is, sorts are not really predicates, but constants denoting objects of the universe of discourse. The relation *local-domain-creator(g, site)* holds between two objects— not between an object and a predicate—so the description language remains first-order.

$$(3.146) \quad \textit{local-domain-creator}(g, S)$$

This statement is true if the arc $\llbracket g \rrbracket$ creates a local domain with respect to specifications of sort $\llbracket S \rrbracket$, that is if the two positions joined by $\llbracket g \rrbracket$ share any $\llbracket S \rrbracket$ specifications they have.⁴⁷ For example, the PS in (3.145) satisfies the expression *local-domain-creator*(n_7 , *site*).

If two nodes are joined by a local-domain-creating government arc, we can say that they are local to each other for the appropriate arc sorts. Indeed, we can make this a biconditional—since the main claim of this section is that locality between two nodes can arise *only* when they are joined by a local domain creator:

$$(3.147) \quad \textit{Locality condition (preliminary version)}$$

$$\forall x, y, g, x \xrightarrow{g} y \rightarrow \textit{local-domain-creator}(g, S) \leftrightarrow \textit{local}(x, y, g)$$

(In fact, this requirement will be weakened slightly in the next chapter to allow local domains to be formed between the children of two nodes in another local domain. This is why we shall eventually want to express the locality of two nodes (using the predicate *local*) somewhat independently of their being joined by a local domain creator.)

The substantive effect of being in a local domain is enforced by the following universal constraint:

$$(3.148) \quad \textbf{Spreading Constraint — universal}$$

$$\forall x, y, z, S \textit{ local}(x, y, S) \rightarrow (x \xrightarrow{S} z \leftrightarrow y \xrightarrow{S} z)$$

The effect of this constraint is similar to the result of bidirectional spreading in autosegmental frameworks. If some constraint, lexical or otherwise, requires that x must have an $[S:z]$ specification, y must also have an $[S:z]$ specification, and vice versa. It is clear that the effect of the Spreading Constraint is transitive. If a language requires all nuclear arcs n to be *local-domain-creator*(n , *site*), then every nuclear position in a word will share the same site, even though local domains are actually only created pairwise between adjacent nuclei.⁴⁸ This is the basic mechanism underlying vowel harmony systems, which will be discussed in more detail in the next chapter.

⁴⁷I capitalize the sort variable S in order to avoid any confusion with the abbreviation s for the arc-sort constant **site**.

⁴⁸When the intent is clear from the context, I shall sometimes speak loosely of all nuclei in a word being in a single local domain.

3.4.1 Dissimilation

I shall not have much to say about dissimilation in this dissertation, but this section will point out one of the means available to the framework for handling dissimilation that occurs under some kind of governmental adjacency. This will be able to deal with local dissimilation, such as constraints on the sharing of labiality between onsets and nuclei, but it will not be able to deal with long-distance non-governmental cases, such as Latin liquid dissimilation.

The usual way of dealing with such local dissimilation in autosegmental phonology is by application of the Obligatory Contour Principle (OCP). The OCP bans the occurrence of two identical phonological objects adjacent to each other along a given tier.

(3.149) Obligatory Contour Principle (autosegmental version)

$$* \text{ — } \alpha \text{ — } \alpha \text{ — }$$

Whenever such a structure is found, it is eliminated, either through deleting one of the offending objects or by fusing them. In the fusion option, if two identical objects come into contact with each other (as the result of morpheme concatenation, for example), they are fused to produce a branching structure:

$$(3.150) \quad \begin{array}{ccc} x & x & \rightarrow \\ | & | & \swarrow \searrow \\ \alpha & \alpha & \alpha \end{array}$$

The OCP does not apply everywhere. Some failures are universal (nobody, for example, has proposed that it applies on the moraic tier). Others are language specific. Some frameworks need it to apply only at certain times or at certain stages of the derivation. (For instance, a radical underspecification framework could not have the OCP apply so early that it fuses the underlyingly specified features of possibly non-adjacent segments before the appropriate default rules have applied to the positions between them.)

In any event, the OCP can provide only part of the story for most dissimilation processes. Besides banning sequences of the dissimilating feature (3.149), multiply-attached structures as in (3.150) must also be banned, thus preventing the fusion option and forcing the deletion option. (Cf. Mester's (1986) constraints against many-to-one mappings in languages that show OCP effects.) For example, in discussing Taiwanese labial cooccurrence constraints, Lin (1989) must invoke the OCP on the labial tier, but must also propose the following constraints against multiple association:

$$(3.151) \quad * \left[\begin{array}{c} \\ \backslash / \\ [\text{labial}] \end{array} \right]_{\text{final}} \quad * \left[\begin{array}{c} [\alpha\text{cons}] \quad [\beta\text{cons}] \\ \backslash / \\ [\text{labial}] \end{array} \right]$$

The treatment of local dissimilation in the present framework incorporates the same two steps: a universal ban on sequences of identical specifications that are not shared, and then (language specifically) bans on certain sequences of identical specifications that *are* shared.

I propose to accomplish the first step with the following translation of the OCP into the present framework. It requires that any (governmentally) adjacent nodes that have the same value for some feature must be in a local domain for that feature:


(3.152) **Obligatory Contour Principle** (constraint-based version)

$$\forall x, y, z, g, S \quad x \xrightarrow{g} y \wedge x \xrightarrow{S} z \wedge y \xrightarrow{S} z \rightarrow \\ \text{local-domain-creator}(g, S)$$

It remains to be seen if a constraint of this strength can be made universal.

As an example, if a language represented *u* with the labial gesture (rather than the velar gesture) on the root node, then in any syllable like *pu* the onset-licence would have to be a local domain creator for site and articulator, resulting in a PS like:

(3.153)



This is the first step, banning unshared identical structure. Now in order to accomplish a dissimilation, all the language would have to do is forbid such an onset-licence to be a local domain creator—the second step, banning shared identical structure as well.

(3.154) $\forall g \text{ onset-licence}(g) \rightarrow \sim \text{local-domain-creator}(g, \text{site})$

3.5 Phonetic interpretation

This section discusses the sort of phonetic interpretation component that I assume to be part of the present framework. No discussion of this length could do justice to the subject of phonetics and its relation to phonology. The present section can only hope to sketch the broad outlines of just one conception of phonetics that is compatible with constraint-based phonology.

The first three sections describe a possible kind of phonetic representation (called Phonetic Event Structures or PESs) and the relationship of these representations to

actual phonetic events. The discussion suggests some starting points for an eventual formal theory of the competition between “ease of articulation” and the need to articulate distinctions clearly enough that a listener can reconstruct the intended phonological structure. Section 3.5.4 briefly addresses the principles that map phonological structures (PSs) to phonetic event structures (PESs).

The next section addresses the issue of distinguishing between phonetic effects and phonological effects. Mistaking phonetic effects for phonological processes can greatly complicate theories of phonology. It is argued that Chumash sibilant harmony, often taken to be evidence for feature-changing and non-local spreading, is one such phonetic effect and shows the distinctive characteristics of competition between different phonetic principles.

The role of default rules is discussed in section 3.5.6, two types of which are distinguished. One type of default rule (those that decide what certain articulators should be doing when a PS says nothing about them) can have no effect at all on the legality of a PS. The other type (those that fill in missing or `unspecified` values on a node) might turn out to require that its default value *not* be present in a PS, in order to avoid ternary power. The statement of this type of default rule and its effects on PSs are formalized. The application of default rules in the phonology-phonetics interface to create default values in the phonetic event structure is touched on, but no formal model is developed.


3.5.1 Phonetic targets

Most of the ideas in this section on the nature of phonetic events owe their origin to the work of Keating (1988a,b) on phonetic underspecification.⁴⁹ One of Keating’s main points is that many apparent occurrences of a feature, nasality for example, are not really present phonologically but remain unspecified even at the level of phonetic representation.

For example, English vowels preceding a nasal consonant are accompanied by a certain degree of nasal resonance. But this is not the result of the feature [+nasal] spreading from the consonant to the vowel at any level of phonology. It is not even the result of [+nasal], or its phonetic stand-in, spreading at some level of phonetic spell-out. Rather, the degree of nasalization that occurs on the vowel is the result of the articulator making as smooth a transition as possible from one required state to another. The articulation of an English syllable of the form CVN might be diagrammed as follows, where the vertical axis represents the degree of nasalization:

⁴⁹See also related work by Boyce, Krackow, and Bell-Berti (1991). The reader is warned that this section does not purport to be a faithful summary of Keating’s work, nor of course is Keating responsible for the immoral purposes to which these stolen ideas have been put.

(3.155)




The phonology-phonetics interface does not particularly care what the velum is doing during the vowel. It does care about having a raised velum (and hence blocked nasal airflow) during the first consonant and a lowered velum during the second, nasal consonant. The intermediate and increasing amounts of opening during the vowel are the result of the velum making as smooth and easy a transition as possible from the required raised state to the required lowered state.

We call these points at which a certain articulatory state (e.g., raised velum, lowered velum) is required articulatory **targets**. The smooth transition from one target to another is a process Keating calls **interpolation**. Interpolation usually results in intermediate values of the parameter in question. These intermediate values are simply a byproduct of the way articulators do their job, they are not significant at a phonological or even a phonetic level of representation.


These concepts are useful in distinguishing between phonological assimilation and purely phonetic variation in place of articulation. As an example, Keating (1988b) discusses the behaviour of the Russian velar fricative. Before *i*, /x/ palatalizes to [ç]. Before *a*, it becomes a uvular fricative [χ]. The following diagrams showing the front/back position of the segments are abstracted from second formant of actual spectrograms that Keating provides.

(3.156)



In a context like the above where /x/ is the first segment produced, both assimilations seem to be absolute and equally phonological. But in order to determine whether an assimilation is phonological or phonetic, we need to look at the behaviour of the assimilating segment as a *transition* from a preceding target to a following one. Intervocally, the two assimilations are quite different:

(3.157)



Here, ζ must have the same frontness as i , χ does not have the same frontness as a . The phonology-phonetics interface cares about the frontness of ζ . It receives its own target (represented by a solid circle), which the articulators must respect. On the other hand, the interface does not particularly care about the frontness or backness of χ . It has no target of its own. Whatever backness it has it owes to interpolation, the smooth transition of the articulators from the target for i to the target for a . Palatalization here is a phonological assimilation, and the spreading of the frontness feature from $/i/$ to $/\zeta/$ results in a new phonetic target. Uvularization before a is not a phonological assimilation. There is no spreading, no new target, and, Keating argues, $/\chi/$ remains underspecified all the way through to the final level of motor implementation.

3.5.2 Phonetic event structures


The phonetic representations developed here do not represent phonetic events, but types of phonetic events. Actual phonetic events in the real world may correspond more or less closely to phonetic event types. It is the phonetic event types that are the subject matter of the phonology-phonetics interface.

Phonetic event types are represented by Phonetic Event Structures (PESs). Any characterization of phonetic events will inevitably involve such temporal notions as intervals of time, temporal precedence, and temporal overlap. To accomplish the task, we can draw on the resources of temporal logic. Indeed, we can characterize PESs as structures in an interval-based temporal logic (see, e.g., van Benthem 1983, 1988).

The basic building block of PESs are **intervals**. Two temporal relations can hold between intervals: **precedes**, symbolized in the temporal logic by \prec , and **overlaps**, symbolized by \circ . There is also the **subinterval** or **inclusion** relation, symbolized by \sqsubseteq .⁵⁰ In the diagram below, we can imagine time as the horizontal axis and intervals as regions marked by boxes. A represents the interval diagrammed the solid-lined box, B the interval with the dashed-line box.

⁵⁰Though this relation is in fact definable in terms of overlap. Refer to van Benthem (1983) for details on temporal logic, and in particular interval-based temporal logics.

(3.158)




The main work done by intervals in this framework will be to characterize the temporal properties of certain facts. We can say that some fact is true *in* (or perhaps *of*) some interval. What kinds of facts are we interested in? Following the discussion of Keating's ideas in the last section, I propose that a PES interval characterizes the *target* of an articulatory constriction gesture.

Thus, if the gestural features of a phonological structure are [a:Ap_c, s:Ant, d:0], the phonetic interpretation of this will be the gesture of an anterior coronal (apical) stop. In some sense, we can see this gesture as involving all the motor activities needed to coordinate the tongue tip's ballistic approach towards the dental-alveolar region, its contact, the momentary blockage of airflow, and the tongue's subsequent removal. But not all of this is strictly relevant for the skeleton of the PES. The central object of the PES will be the *target* of the entire gesture, that is, it will be an interval during which it is true that the tongue tip is in complete contact with the anterior region.

It is these target intervals that are ordered relative to other targets. For example, if a nucleus node governs an onset node via an onset-licence arc, then the target interval corresponding to the onset's root node gesture will temporally precede the target interval corresponding to the nucleus' root node gesture. It is important to note that this temporal ordering applies only to the target intervals. It does not imply that every phonetic situation "belonging" to the onset terminates before any phonetic situation "belonging" to the nucleus comes to be. As a concrete example, consider a syllable *nu* where the onset's root node represents the gesture of lowering the velum (nasality) and the nucleus' root node represents the gesture of lip rounding. The precedence requirement applies only to the targets of these two gestures. That is, the interval of time where the PES actively requires nasality precedes the interval of time where the PES actively requires lip rounding. There is no requirement that every point of time during which there is nasality precede every point of time during which there is lip rounding.

This situation arises because intervals can be properly included in larger intervals:


(3.159)



In this diagram, both A and B represent intervals during which it is true that a significant degree of nasality exists. A is a subinterval of B. A might represent the *target* interval of the nasality gesture, the interval that is the interpretation of the onset's root node, the interval during which the phonology-phonetics interface actually cares about nasality. The non-A parts of B would then represent “unnecessary” nasality, nasality not required by the phonology-phonetics interface but nonetheless present in the phonetic event. This can be seen as the result of the process of interpolation discussed by Keating. For convenience, I will call these larger intervals that properly include the target the **extended interval** of a gesture, though it is surely an idealization to use such a discrete mechanism to represent an essentially gradient phenomenon.

Returning to our example of the syllable *nu*, it is easy to see that parts of gestures “associated” with a particular node in a PS might temporally overlap parts of gestures associated with another node, even though the targets of the two nodes are in a strict relation of temporal precedence. In the following diagram, where left-to-right orientation represents the time axis, the interval A might represent the onset's nasality target, B “unnecessary” interpolated nasality (i.e., nasality's extended interval), C the nucleus' roundness target interval, and D interpolated roundness. The arrow represents the only temporal precedence relation required by the phonology-phonetics interface.

(3.160)




Essentially the same thing is happening in vowel coarticulation. An anticipatory coarticulation between V_1 and V_2 will result if the extended interval of V_2 extends

backwards, entirely overlapping the target interval of V_2 's onset and overlapping with the extended interval of V_1 . It is probably exactly this kind of overlap, and a linguistic community's increasing tolerance for it and increasing expectation of it over time, that eventually leads to the phonologization of umlaut and other vowel harmony rules.⁵¹

One interesting possibility is that a gesture's extended interval may be completely overlapped by the extended intervals of other gestures, which can give rise to a phenomenon Browman and Goldstein (1989, 1990) called "acoustic hiding". For example, in a fluent pronunciation of the phrase *perfect memory*, the final *t* of *perfect* will often appear to be deleted. But, using X-ray tracking of lead pellets attached inside a speaker's mouth, Browman and Goldstein found that the speaker actually did perform an alveolar closure gesture, but this gesture was completely overlapped by the preceding velar closure and the following labial closure, as in (3.161), making it acoustically imperceptible and making the entire utterance sound as if the segment *t* had been deleted.

(3.161)



The present approach to the problem of temporal order in phonetic interpretation differs from that of Coleman (1992), though both share a commitment to a declarative mapping between phonological structure and an interval-based phonetic structure. Coleman's approach focuses on the boundaries of phonetic intervals, while the approach of this section focuses on the "centres". Coleman has principles to accomplish the proper alignment and precedence of interval endpoints; the only ordering that the phonology-phonetics interface cares about in the present framework is that between the centres of extended intervals, that is, their targets. I believe the present approach is superior. While articulatory targets clearly play a role in phonetics, it is not clear that boundaries can be said to be real at any level.

3.5.3 Phonetic Event Structures and phonetic events

Phonetic event types, represented formally by Phonetic Event Structures, are abstract entities. Perhaps the most appropriate way to think about the relationship between

⁵¹This presupposes, correctly I believe, that languages are to a certain extent able to manipulate the extent of their extended intervals. For example, it may be possible to constrain an onset's extended interval to strictly precede its nucleus' target. The question of how much power the phonetic component of a language has to refer to non-target intervals, how exactly this is accomplished, and whether there are any universals of such reference is beyond the scope of this discussion.

the two is to see phonetic event types as prototypes that actual phonetic events in the world may instantiate more or less exactly. Actual phonetic events are subject to two kinds of soft constraints. Put another way, there are two different measures of goodness that the producer of a phonetic event will try to maximize.

The first measure I shall call **goodness-of-fit**. This determines how closely the actual phonetic event adheres to the prototype of the PES it is meant to instantiate. The second I shall call, for lack of a better term, **articulatory goodness**. This measures how preferred a phonetic event is as an articulatory event and involves many of the considerations that have generally been termed “ease of articulation”.

The two measures are often in conflict with each other. For example, a phonetic event that involves an abrupt and difficult-to-coordinate transition from one gesture to another might score high on goodness-of-fit to the PES it is supposed to instantiate, but performs abysmally on the articulatory goodness measure. Both measures cannot be maximized at the same time. To determine the overall goodness of a phonetic event, different weights are given to the two measures depending on such factors as rate of speech and formality.

The goodness-of-fit measure may be concerned with some gestures more than with others. For example, the articulation of an unstressed vowel may be quite far away from its prototype without having very much effect on the goodness-of-fit measure, while a stressed vowel articulated in the same way could have a large negative impact. The inherent content of gestures can be a factor in such differences—one phoneme may be allowed considerably more freedom in its articulation than another phoneme. Some of this variability, such as that of velar site, is undoubtedly due to quantal effects (Stevens 1989), though it may perhaps be possible for a language to somewhat arbitrarily rank gestures in terms of how exactly they must be articulated. Capturing the differences in allowed variability would require at least three parameters (one each for articulator, site, and degree) to control the weighting of various aspects of a gesture for the purposes of determining goodness-of-fit. We can call these the **fussiness** parameters, and for the moment can assume that they are properties of the target interval (along with strictly gestural properties) and have their values determined by the phonology-phonetics interface (since they seem to require access to information contained in PSs).

3.5.4 The mapping principles

Though we may eventually want to make the phonology-phonetics interface more intricate, the basics are simple.

Each constriction node in a PS is interpreted as a target interval in a PES. When two nodes are joined by an arc, the arc is interpreted as a temporal relation (precedence or overlap) between the targets corresponding to the nodes.⁵² Let $\mathcal{P}(x)$ stand for the

⁵²I am not certain what would be the best way to implement the interpretation of sorts with phonetic effect, such as `line-2`.

PES interval that is the phonetic interpretation of node x . If our friend, the syllable *nu*, had c_7 , v_7 , and o_7 as the onset root node, nucleus root node, and onset-licence arc respectively, the phonetic interpretation would have: an interval $\mathcal{P}(c_7)$, the target for a nasal gesture, an interval $\mathcal{P}(v_7)$, the target for a lip rounding gesture, and a temporal precedence relation between the two induced by the onset-licence o_7 , that is $\mathcal{P}(c_7) \prec \mathcal{P}(v_7)$. The government types **onset-licence**, **coda-licence**, and **o-r** will be interpreted with the dependent's target temporally preceding the head's target; **secondary** will be interpreted with the dependent and the head overlapping; all others will have the head precede the dependent.

A PES is a phonetic interpretation of a PS if there is a one-to-one relationship between target intervals in the PES and constriction nodes in the PS, where the PES interval is an interpretation of the PS node,⁵³ and if the target intervals stand in the appropriate temporal relations demanded by the PS arcs.

Besides containing these straightforward mapping principles between nodes and intervals, the phonology-phonetics interface also decides the fussiness of each target interval. Though it is possible that some of this is done by language-particular principles, we can assume that the interface also contains a number of universal constraints on the relative fussiness of various intervals. One such universal might be:

- (3.162) For any two nuclear nodes, v_i and v_j , if v_i is **line-1** and v_j is not,

$$\text{fussiness of } \mathcal{P}(v_i) \geq \text{fussiness of } \mathcal{P}(v_j)$$

3.5.5 Phonology vs. phonetics: a cautionary tale for the border patrol

In section 3.5.1, we noted that Russian x is palatalized before i and uvularized before a . The first instinct of most phonologists would be to explain both assimilations by means of phonological rules. If both processes can be “generalized” into a single phonological rule, so much the better. Unfortunately, Keating (1988b) noted a significant phonetic difference between palatalization and uvularization, a difference that has phonological consequences. Palatalization of x to ζ is clearly a phonological spreading rule. It results in a new articulatory target in the Phonetic Event Structure. Uvularization of x to ʁ is not the result of a phonological rule. No new target is created. Whatever degree of backness the fricative may have is the result of interpolation.

⁵³Since phonology and phonetics may presumably use different vocabularies to talk about articulatory gestures, there will probably have to be a set of mapping principles that would relate, e.g., [a:Dor, s:Vel, d:0] to whatever the representation of a velar stop constriction gesture is in the vocabulary of PESs.

This one-to-one property of the relation may have to be weakened somewhat to take into account the effect of articulator-based defaults (see section 3.5.6).

Trying to explain a phenomenon away as a “phonetic effect” seems like a slippery and highly suspect manoeuvre, usually for good reason. But, as Keating’s discussion shows, it is a fact that some phenomena *are* phonetic effects, and phonological and phonetic theory will need a way to come to terms with that fact. Keating’s work, and the sort of considerations discussed earlier in this section, suggest starting points that we may be able to develop into a set of criteria for determining when a phenomenon is phonological and when it is phonetic. I cannot offer such a set of criteria here. But I can present another example of a phonetic phenomenon that has been mistaken for a phonological one. The discussion should offer some clues about the sorts of characteristics we should expect a phonetic phenomenon to have. It should also reinforce the point that greater sensitivity to the difference between phonetic and phonological phenomena is needed and that a proper division of labour can considerably reduce the power that a theory of phonology needs to have.

Since Poser (1982) first presented an autosegmental account of Chumash sibilant harmony, it has stood as one of the most recalcitrant obstacles to simplifying autosegmental theory. Against attempts to make all harmonic processes into feature-adding processes, Chumash has been used to prove the need for feature-changing processes as well. Against attempts to give constrained definitions of locality, Chumash has been used as an example of action-at-a-distance.

Phonologists have managed to analyze almost every harmony process as feature adding: the harmonic feature is spread onto a position that until then had been unspecified for that feature. But in Chumash (and a similar process in Navajo), the harmonic feature is spread onto an already specified position, deleting whatever value of the feature was there beforehand. Though Chumash is the only language Lieber (1987) is aware of that requires such power, she still believes the evidence warrants giving Chumash its own category in a typology of possible harmony processes.

Chumash has also frustrated attempts to formulate a definition of phonological locality. Phonologists believe almost as an article of faith that processes should be allowed to affect only positions that are adjacent to each other at some level of representation. Chumash, however, apparently allows a sibilant to affect another sibilant to its left across an arbitrarily large distance, skipping over all vowels and irrelevant consonants in between. While accounts could be devised that make the sibilants adjacent at some level (e.g., Shaw 1991), they are not as emotionally satisfying as straightforward definitions of locality.⁵⁴

As commonly presented in the autosegmental literature, all sibilants in a Chumash word agree with the rightmost sibilant in their place of articulation or in one of the

⁵⁴The model developed in this dissertation also conflicts with the usual analysis of Chumash on both counts, which means I have a sort of vested interest in showing Chumash sibilant harmony not to be phonological. The present framework cannot allow two sibilants to be local to each other, and thus share features, without also affecting all the other positions in the PS connecting the two sibilants. Nor can it allow any feature changing whatsoever, let alone feature-changing harmony.

Though Poser (1982) gives very little indication that sibilant harmony in Chumash is anything other than an automatic exceptionless rule, the situation is the same as in Navajo. Harrington's 1928 study of sibilants in the Ventureño dialect, published posthumously as Harrington (1974), is worth quoting at length:

But in actual practice the raising or lowering [to *s* or *š*] is largely only partial and frequently does not occur at all. Intermediate sounds between *s* and *š*, here written *ś*, arise by such imperfect assimilation or by a lowering of sounds before *t*, *l*, *n*... The assimilation is moreover less thorough with some speakers than with others. Especially in slow speech and when detached words are furnished it is apt to be absent.

The assimilation is as a rule retrogressive. Progressive assimilation is rare and never extends far. The probable reason for this backward direction is that the phonetically strongest sibilants of Chumashan are the final sibilants....

It is interesting in the light of general phonetics that *š* is much more thorough and far-reaching in its working of assimilation than is *s*. Just as in language growth in general it is supposed that *s* more often becomes *š* than vice versa, just as a drunken man may allow his *s*'s to lapse into *š*'s but does not *s*-ize his *š*'s, so also here in Chumashan it seems that *š* has more power to pull *s* down than *s* to raise *š* up.

It should be noted that the harmony rarely extends further back than through a single word and that the article *si-*, when it has this form, seems especially resistant to assimilation.


In this passage, almost every characteristic of Poser's rule is brought into doubt. If the rule created new segments like any other assimilation rule, we should expect their new identities to be clearly the same as the trigger; instead we often find articulations intermediate between *s* and *š*, often enough that Harrington felt the need to devise the symbols *ś* and *ç* to transcribe them. Formally, there is symmetry between the *s*→*š* process and the *š*→*s* process; in reality, one direction is preferred to the other. The stated domain of the rule is the word; but there are frequent cases where the domain is smaller, and some cases where the domain is larger. The stated direction of the rule is right-to-left; but there is a limited tendency for left-to-right assimilation as well. Instead of applying wherever its structural description is met, like other phonological rules (e.g., Hungarian vowel harmony), it occurs more often in fast speech and can be suppressed entirely in careful speech.

In short, sibilant harmony has all the characteristics of a phonetic effect of fast speech and none of the characteristics of a rule of the lexical phonology. It has more to do with the reason I can't say "She sells sea shells by the seashore" quickly than with the reason I can't say "cat+z".

Applying the ideas of Keating's discussion of Russian /x/ to Chumash, sibilant harmony resembles *ia* much more than it resembles *açi*. If there really had been

phonological assimilation of a feature, as Poser's and Shaw's rules demand, we should expect a series of identical sibilant targets, as illustrated below. Such a series would not pose a great amount of articulatory difficulty. Identically pronounced sibilants would be optimum for both the goodness-of-fit measure and the articulatory goodness measure.

(3.167) Phonetic event predicted by the feature-changing analysis




However, this is not what actually happens. Instead we find the “harmonized” sibilants are pronounced with a point of articulation between the original value and that of the trigger. A typical phonetic event instantiating sequence of Chumash sibilants might appear somewhat as in the following diagram.

(3.168) Typical phonetic event



If all the sibilants did indeed share the same constriction features, a phonetic event like (3.168) would be less optimum by both measures: their goodness-of-fit would be much worse and they would be harder to pronounce than a series of identical targets. One measure of goodness is usually sacrificed in order to optimize the other. It is hard to imagine speakers consistently sacrificing both at the same time for no apparent reason. This is a strong indication that the original sites of the harmonized segments have not been replaced with those of the trigger. Instead they keep their original PS specifications. Although the targets corresponding to the original features are not as clearly realized as they would be in a non-harmonic context, they continue to contribute to the overall contour of the phonetic event. The situation can be diagrammed as:

(3.169) Phonetic event predicted by the “phonetic effect” analysis



A phonetic event trying to instantiate a series of non-identical sibilant targets faces a quandary. If it hugs close to the targets, it scores large penalties on the measure of articulatory goodness. If it tries for ease of articulation, keeping the tongue blade at more or less the same site throughout the event, it loses on the goodness-of-fit measure.⁵⁵ Most actual events will compromise between the two pressures, resulting in sibilants that are to some degree intermediate between their targets and that of the final sibilant.⁵⁶ Some preference is usually given to one of the two measures of goodness depending on factors like speech rate and formality. In slow or highly formal contexts, where the pressure to minimize physical difficulty is not as strong, goodness-of-fit will be preferred, and the tongue blade will hug close to its targets. In faster speech or less formal contexts, the preference may be for articulatory goodness, which will result in a higher degree of surface assimilation. This is exactly the behaviour Harrington describes.

In short, Chumash does not force phonology to allow feature changing harmonies, because no features are in fact changed. Harrington’s phonetic description is consistent with a situation where all sibilants keep their original underlying values through all levels of phonology and phonetics, but it is not at all consistent with the claim that sibilants receive new articulatory targets as the result of a phonological feature-changing rule.

The most immediate moral of the story is that analyzing Chumash sibilant harmony as a phonetic effect not only considerably simplifies phonological theory by reducing the power needed by phonological rules, it also provides a superior explanation of the behaviour of Chumash sibilants than an account that rushes to capture the variation with a feature-changing phonological rule. A moral with broader implications is that

⁵⁵In languages like Chumash and Navajo, the goodness-of-fit principles are less fussy about the exact constriction site of sibilants than they are in languages like English (which are themselves relatively lax compared to many other segment types). As Harrington’s discussion clearly suggests, the goodness-of-fit principles in Chumash become more fussy the closer to the end of the word one is, and are fussier about *š* than *s*.

⁵⁶The existence of intermediate values is particularly troublesome for an analysis that uses a feature-changing phonological rule. There is no reason in the world that we should find sibilants straying back towards original values that have supposedly been obliterated by the phonology.

phonologists need to be pickier about what they accept as the data to be explained by the phonological component of a grammar. It is essential to distinguish between phonetic phenomena and phonological phenomena. The bad news for those of us who prefer to work in armchairs is that deciding between the two often involves gathering and interpreting instrumental data. For the purposes of distinguishing the two, generalizations based on broad transcriptions cannot be accepted at face value. This makes it much more difficult to use published grammars as sources of data, especially since few descriptive grammar writers equal Harrington in providing the level of phonetic detail necessary to determine whether a phenomenon is phonological or phonetic.

3.5.6 Defaults

The first major distinction to be drawn has already been touched on in chapter 1, that between redundancy rules and default rules.

Redundancy rules are just like any other phonological constraint. They correlate pieces of phonological structure with pieces of other phonological structure. They essentially say, “Any PS that has X in it must also have Y in it.” As such, they judge the well-formedness of PSs. As might be expected, there is no clear boundary between redundancy rules and other types of constraints.

Default rules are not strictly part of phonology, but of the phonology-phonetics interface. The absence of a default specification is incapable of making a well-formed PS ill-formed.⁵⁷ The function of default rules is to determine the proper phonetic interpretation of a well-formed PS that for some reason fails to contain all the information necessary for phonetic interpretation. For example, default rules determine the interpretation of **unspecified nodes**. Redundancy rules must operate monotonically, default rules may operate non-monotonically.

A second major distinction is between defaults that are node-based and defaults that are articulator-based. **Node-based defaults** fill in missing pieces of information on a node in a PS, for example, an unspecified degree feature or an unspecified articulator feature. Node-based defaults cannot apply if the PS has no node for them to apply to. **Articulator-based defaults**, on the other hand, determine what kinds of gestures should be performed with certain articulators when the PS says nothing about them. These typically correspond to what would have to be represented with entire gesture nodes in a PS, except that the PS has no corresponding node. For example, an articulator-based default might stipulate that, in the absence of information to the contrary in the PS, the tongue root should be retracted or the velum should be raised.

Basically, the difference between the two types of defaults can be seen as the ways the phonology-phonetics interface has of answering two different questions. Node-based defaults answer the question: “What am I supposed to do with this gesture node in the PS that isn’t fully specified?” Articulator-based defaults answer the question: “What

⁵⁷Though, as will be seen shortly, we may want to give the *presence* of a default specification the ability to make a PS ill-formed.

am I supposed to be doing now with this articulator that the PS isn't giving me any information about?" Node-based defaults are driven by the logical needs of the formal system to specify gestures completely. Articulator-based defaults are driven by the physiological need to decide what a piece of anatomy should be doing at certain points in time.⁵⁸

I assume the non-monotonic interface can somehow order node-based default rules in terms of their strength, roughly in accordance with the Elsewhere Principle. I shall not formalize any particular scheme here, but there are several default logics or autoepistemic logics that can accomplish the task.⁵⁹ This will involve some complication of the phonology-phonetics mapping principles of section 3.5.4. It will no longer be the case that a target interval of a PES will have a certain property (say a site) only if that property corresponds directly to some piece of the PS; the property may also be the result of a default rule, or rather the strongest default rule applying to the node in question. It will most likely be the case that the phonology-phonetics interface does not determine phonetic properties directly off of the PS, but off of an **extension** (in the sense of default logic) derivable from the PS by non-monotonic inference rules using the defaults.

The greatest difference between node-based and articulator-based default rules may turn out to be the role they play in the phonology itself. Articulator-based defaults are utterly irrelevant to phonology and can have no effect whatsoever on the legality of a PS. Node-based defaults, on the other hand, while they cannot affect the legality of a PS directly, may be able to have an impact on phonology.⁶⁰

It is possible that we might want to incorporate into the framework some way of preventing ternary power in our feature specifications. That is, we may not be able to draw a three-way distinction between, say, a node that is specified [s:Vel], a node specified for some site other than [s:Vel], and an **unspecified** node that will be interpreted as [s:Vel] by a default rule.⁶¹ We could prevent the situation on a

⁵⁸It is possible that articulator-based defaults are all universal. For example, it is perfectly reasonable to believe that the universal thing to do with your velum is to keep it raised unless told otherwise.

⁵⁹For an overview of the various versions of non-monotonic logic that can be bought off the shelf, see Brewka (1991). A more in-depth look at a class of versions based on work begun by Reiter (1980) is given in Besnard (1989).

⁶⁰The two types of defaults may also behave somewhat differently in phonetics. For example, it is a plausible hypothesis worth exploring further that node-based defaults result in gestures that are universally more fussy (or at least as fussy) as the gestures of articulator-based defaults, according to the concept of fussiness of section 3.5.3.

⁶¹See, for example, Lightner (1963), Stanley (1967), and Archangeli (1988), for discussions on the formal undesirability of allowing ternary power. Archangeli and Pulleyblank (1992), however, argue that ternary power may not be so bad after all, and give many analyses that are considerably simplified by exploiting the three-difference between plus specification, a minus specification, and no specification at all. In this section, I leave aside the question of whether we *should* prevent ternary power and concentrate instead on whether it could in principle be prevented if we decided to.

case-by-case basis by banning the appearance in a PS of what would be the default specification; for example, the grammar of Hungarian would have a constraint banning an overt Velar site on vowels:

$$(3.170) \quad \sim \exists v \ v \xrightarrow{s} Vel$$

Languages where the default vowel was *i* would have a similar constraint against overt Palatal sites. This strategy would prevent the ternary use of features in individual cases, but it could not prevent it in the general case. It would remain a mystery why grammar after grammar, having the power to use features ternarily, conspired not to.

A more satisfactory solution is to let some information about defaults into the phonology. We can do this with a three-place predicate:

$$(3.171) \quad default(x, g, y)$$

where x is a node, g is an arc-sort from the set of gesture arc-sorts (**articulator**, **site**, **degree**), and y is a gestural atom (e.g., Pal, Apc, 1) appropriate to the arc-sort. I shall usually use the more iconic notation:

$$(3.172) \quad default(x \xrightarrow{g} y)$$

Default principles may now be stated directly in the phonology. For example, a language with *i* as the default vowel might contain the description:⁶²

$$(3.173) \quad \forall v \ nucleus(v) \rightarrow default(v \xrightarrow{s} Pal)$$

Our original remark about defaults not being part of phonology is still true in a limited sense: the *statement* of default rules is part of phonology (on a par with any other constraint), the *application* of default rules to fill in default values is not.

Using the predicate *default*, we can now ban ternary use of features with the following (universal) constraint:

⁶²It should be noted that this default statement holds of *all* nuclear positions, even those with other overt sites like [s:Vel] and those specified **null**. For the default logic of the phonology-phonetics mapping, overt gestural features and null sorts are “stronger” than default specifications. The same prioritized default logic will choose between competing defaults if two are applicable to a given segment. Because the *default* predicate is manipulable by the description language, it is possible for two different default statements to hold of the same node. It is also possible for two nodes of the same sort but in different environments to have different applicable default rules. For example, Mohawk verbs use three different default vowels (Michelson 1989): a within verb-stems, prothetic *i* initially to satisfy minimal word requirements, and *e* elsewhere.

$$(3.174) \quad \forall x, g, y \text{ default}(x \xrightarrow{g} y) \rightarrow \sim x \xrightarrow{g} y$$

This prevents any node in a PS from having an overt specification for what would be its default value.

If a particular feature is predictable for a node, it is generally possible to express the predictability using either a redundancy constraint or a default rule. The proper balance between the two strategies requires more investigation. For the purposes of this dissertation, however, I shall give preference to redundancy constraints over default rules. That is, I shall tentatively assume that any specification is actually present in a PS unless there is evidence that some constraint needs that specification to be absent in order to work properly.

Chapter 4

Locality: harmonies and reduplication

This chapter explores some of the implications and applications of the concept of local domain introduced in the last chapter. It is shown how two phenomena that mainstream phonology has treated as processes *par excellence*, spreading and reduplication, can be dealt with in a “static” system of constraint satisfaction.

Section 4.1 generalizes the notion of locality. The first version of locality, introduced in the last chapter, involves only the sharing of primitive gestural features (site, articulator, degree) between two nodes that are directly connected by a local-domain-creating arc. Now, the notion will be extended to the “sharing” of complex nodes, such as secondary articulations and onsets. A definition will be developed of what it means for, e.g., a nuclear-licence arc to be a local-domain-creator for onsets. The presence of an onset for one of the two nuclei joined by such an arc will require the presence of an onset on the other nucleus. Furthermore, these two onsets will themselves form a local domain. This weakens the claim made in the last chapter that two nodes could form a local domain only if they were connected by a local-domain-creating arc: now two nodes can also be local if, in a sense, they inherit their locality properties from their respective parents, who in turn form a local domain with each other.

The next three sections are reasonably detailed discussions of actual vowel harmony systems. The first is the frontness harmony of Hungarian, an example of a **symmetric** vowel harmony system, that is, a system where the frontness of every vowel in the word is determined by that of the stem. An interesting wrinkle is the existence of neutral or transparent vowels, that is, the behaviour of some front vowels as if they were not present for the purposes of spreading frontness. Kalenjin is a classic example of an **asymmetric** system, where one of the two values for the spreading feature is “dominant”, and its presence in any morpheme of the word will cause its spread to

all other vowels.¹ The third language, the Pasiego dialect of Spanish, involves two different harmony processes, one spreading [ATR] and one spreading vowel height. The height harmony component, as it has been analyzed until now, also involves neutrality of [+low] vowels, which behave as if they were not present for the purposes of spreading the feature [high]. The height harmony component of Pasiego, like Chumash sibilant harmony, has also been used to argue for the necessity of feature-changing harmonies, a non-monotonic mechanism that could not be integrated into the present framework.

Section 4.5 looks at reduplication. Various processes of reduplication are shown to follow from the revised definition of locality (where children can recursively inherit locality from their parents). Systems that copy entire prosodic constituents arise from arcs that are local-domain-creators for (among other things) prosodic government types like *trochee*. Systems that seem to disregard prosodic constituency and operate on simple string adjacency involve arcs that are local-domain-creators for the composed government type *n-o*, the relation between a nucleus and the following onset. We shall look at cases that truly seem to favour a static constraint-satisfaction explanation of the identity between parts of a reduplicated form. The most obvious procedural account of these cases would involve copying the base, making changes to the duplicate (changes whose environment can only be satisfied after the copy has already been made), then copying the changes back to the original, a course of events that procedural analyses of reduplication have not been able to handle. But this kind of situation is a natural result of a model where reduplication is only a passive requirement that certain parts of a word be “the same as” certain other parts, without regard to how that sameness might have been created. Finally, we shall see that the model developed in this chapter suggests an interesting explanation for why reduplication seems to be so intimately connected with prosody.


4.1 Recursive locality

In the discussion of locality in chapter 3, we only saw examples of the sharing of gesture features (articulators, sites, and degrees). This is easily handled. The PS objects that specify gestures are *atoms*: Pal, Vel, Dor, 2, etc. These are denoted by constants of the description language. The effects of sharing such features is straightforward.

The picture is not as clear, however, when we are dealing with *complex* objects. What would happen if, for example, we wanted to make a particular government relation a local domain creator for secondary articulations? Would we end up with


¹“Symmetry” in this usage refers to the behaviour of the plus and minus values of the harmonic feature, not to the behaviour of stems as opposed to affixes or to directionality of spreading. Hungarian is symmetric because both [+back] and [-back] behave equally in spreading from the stem to the suffixes. Kalenjin is asymmetric because [+ATR] and [-ATR] behave quite differently. Asymmetric systems are often called **dominant** or **dominant-recessive** systems, again referring to the behaviour of the two feature values.

a structure like the following, where the entire secondary articulation node is shared between nuclei?




We shall see below in the analysis of Kalenjin vowel harmony an example where this kind of structure-sharing would not work. We shall want some way of making the presence of a secondary articulation on one nucleus require the presence of a secondary articulation on another nucleus, and we shall want a way of requiring these two secondary articulations to be similar in much the way simple root node harmonies require root nodes to be similar, but there will also be a need to allow the secondary articulations to be somewhat different.

There is another case where simple sharing is inadequate for dealing with complex nodes. We shall eventually want the system of locality constraints we develop for harmonies to be extendible to reduplication, and reduplication offers many instances where “original” and “copy” can be somewhat different. Consider, for example, the reduplication of *trabaho*. There are languages that would reduplicate this word using a light-syllable template, giving the result *tra-trabaho*. It seems the most straightforward way to deal with this would be to make the nuclear-licence arc joining the prefix and the base a local domain creator for gesture features (hence, the copy of the nucleus) and also for onsets:



The reduplication of the *r* of the branching onset (the onset's release) comes for free by the fact that the entire onset node is shared, complete with all its dependents.


But there are other languages where the *r* will not appear in the copy, resulting in *ta-trabaho*. This cannot coherently be represented as the result of a single onset node linked to two different nuclei. Simply to represent the phonetic reality, we would need a structure like:



This fails to do what we want it to do: derive the similarity of c_1 and c_2 by the mechanisms of locality. We could simply decree c_1 and c_2 to be a local domain for gesture features but not for releases. But this move would violate the single most constraining property we have proposed for local domains: that they can only be created between positions that are joined by a local-domain-creating government arc. c_1 and c_2 are not joined by any arc. We need a way of allowing c_1 and c_2 to form a local domain that is not so unconstrained that it allows any two randomly chosen nodes to do so as well.


I propose a slight weakening of the strict definition of locality given earlier to allow children to inherit locality properties from their parents under certain circumstances. I call this **recursive locality**. Thus, in (4.4), mother v_1 and daughter v_2 can form a local domain under the strict version of locality discussed so far. But the aunt c_1 and niece c_2 can also form a local domain (recursively) by inheriting locality from their respective parents.

(4.4)



I will use the term **communal list** to refer to the set of government types that a local domain is local for. For example, *ta-trabaho* might result from the communal list $\{a, s, d, o\}$ shared between the first two nuclei:

(4.5)




If two nodes are in a local domain, they share outright any atomic children (i.e., gesture features) whose governments appear in their communal list. v_1 and v_2 thus have identical articulators, sites and degrees.

Complex children are not shared outright, instead they inherit the communal list from their parents. If one node in local domain dominates a complex child through an arc whose sort appears in the communal list, the other must dominate a complex child through the same sort of arc, and the two complex children will themselves form a local domain for the communal list of their parents. In (4.5), c_1 and c_2 form a local domain for the communal list $\{a, s, d, o\}$. They thus share all their articulators, sites, and degrees.²

²This has the somewhat strange consequence that onsets are in a local domain for onsets

The sharing of atomic nodes (and the mutually implied existence of complex nodes) required by locality applies only to government types in the communal list. The communal list in (4.5) does not contain the release relation, r . So the existence of a release specification on c_2 does not require a similar release on c_1 . c_1 may be forced to have a release by other well-formedness constraints of the language, but the two release nodes will not—and cannot—form a local domain. c_1 's release will be filled in by default principles as a simple unspecified release, giving *ta-trabaho*. Languages that reduplicate to *tra-trabaho* are the same in all respects, except their communal lists are $\{a, s, d, o, r\}$, requiring the release nodes as well to form a local domain, as in:



The formal definition of recursive locality has a basis step saying that two nodes are local whenever they are joined by a local domain creating arc. Furthermore, two nodes can be local *only* if they are joined by a local domain creator, or if their parents were local.

$$(4.7) \quad \forall x, y, S \text{ local}(x, y, S) \leftrightarrow (\exists g \text{ arc}(x, g, y) \wedge \text{local-domain-creator}(g, S)) \vee (\exists p_x, p_y, S2 \ p_x \xrightarrow{S2} y \wedge p_y \xrightarrow{S2} y \wedge \text{local}(p_x, p_y, S) \wedge \text{local}(p_x, p_y, S2))$$

In terms somewhat closer to English: two nodes x and y can be in a local domain for the arc-sort S if and only if they are joined by an arc that is a local domain creator

arcs. If onsets were allowed to onset-license other onsets, c_1 and c_2 would indeed each have to have an onset dependent if the other did, and these onsets would in turn form a local domain. This is why I refer to this definition as recursive locality. Fortunately, independent appropriateness constraints will usually limit the possible depth of recursion.

for S or if they both are children (via some arc-sort $S2$) of parents that form a local domain for both S and $S2$.

The effects of locality are handled by a revision of the Spreading Constraint of (3.148). We could specify the effects of locality on gestures and non-gestures separately, but this is somewhat complicated. I will simplify the definition by requiring *all* children, atomic or complex, by the appropriate government to be local, with the proviso that any atom is local to itself:³

$$(4.8) \quad \forall x, y, s \text{ local}(x, y, s) \rightarrow \\ \forall z_x \ x \xrightarrow{s} z_x \rightarrow \exists z_y \ y \xrightarrow{s} z_y \wedge \\ \forall s2 \ \text{local}(x, y, s2) \rightarrow \text{local}(z_x, z_y, s2)$$

$$(4.9) \quad \forall x, s \ \text{atom}(x) \rightarrow \text{local}(x, x, s)$$

local is of course a reflexive relation for the nodes involved:⁴

$$(4.10) \quad \forall x, y, s \ \text{local}(x, y, s) \leftrightarrow \text{local}(y, x, s)$$

4.2 Symmetric vowel harmony: Hungarian

The frontness/backness harmony of Hungarian has long been a testing ground for accounts of vowel representations and harmony processes (Vago 1980, Ringen 1977, 1988, Booij 1984, Goldsmith 1985, Demirdache 1988). The intriguing thing about Hungarian is not the fact vowels in a word almost always agree in frontness, but that some of the front vowels often behave, for the purposes of the harmony, as if they simply weren't there. An *i* can occur in the middle of a word that otherwise consists of back vowels. To complicate matters, there are other occurrences of the same front vowels that *do* trigger the harmony and require that all other vowels in the word also be front. The challenge for an account of Hungarian vowel harmony is to find a way to represent both those cases where front vowels are neutral (or "transparent", as others have often called them), i.e., how a harmony process can skip over their frontness specifications as if they weren't there, and those cases where the front vowels are *not* neutral, where their frontness specifications are not ignored but enter into the harmony. An ideal account would be able to avoid diacritic features that mark some words as exceptions.

³which should technically be added as another clause of (4.7).

⁴Also a complication for (4.7).

4.2.1 Data

Ringen (1988) presents the vowel inventory of standard (Budapest) Hungarian in the following table, using the letters of Hungarian orthography and their IPA values.

		Front				Back		
	Short	Long	Short	Long	Long	Short	Long	
High	i [i]	í [i:]	ü [y]	ű [y:]		u [u]	ú [u:]	
Mid		é [e:]	ö [œ]	ő [œ:]		o [o]	ó [o:]	
Low	e [ɛ]				á [a:]	a [ɔ]		
		Unrounded	Round		Unrounded	Round		

Other writers disagree slightly over the exact phonetic values of the letters. It appears that orthographic *e* varies dialectally between [ɛ] and [æ], with a similar variation in height for orthographic *a*. As well, *a* might not be phonetically as rounded as IPA cardinal vowel 6—writers agree that whatever phonetic rounding it may have is not phonologically relevant.

Vowel harmony affects what vowels can co-occur in a stem, but its effects are most spectacular in suffixes. Most suffixes in Hungarian have two forms, one that occurs in front-vowel words and one in back-vowel words. For example, the dative suffix alternates between *-nek* and *-nak*.⁵

		ház-nak	‘house (dat.)’
		város-nak	‘city (dat.)’
Front		űr-nak	‘gap (dat.)’
		öröm-nak	‘joy (dat.)’

The three front vowels *i*, *í*, and *é* can occur in back-vowel words without disturbing their status as back-vowel words—the dative suffix will remain in the back form *-nak*:

		radír-nak	‘eraser (dat.)’
		kavics-nak	‘pebble (dat.)’
		tányér-nak	‘plate (dat.)’

For this reason, *i*, *í*, and *é* have been called “neutral” or “transparent” vowels. A word containing only neutral vowels will generally act as a front-vowel word (4.14), but about fifty act as back-vowel words (4.15).

		víz-nak	‘water (dat.)’
		fillér-nak	‘penny (dat.)’

⁵Data are from Ringen (1988) unless otherwise noted.

- (4.15) híd-nak ‘bridge (dat.)’
cél-nak ‘goal (dat.)’

As well, it is possible for a stem to begin with a front neutral vowel, continue with back vowels, and behave as a back-vowel word for the purposes of harmony.

The one class of stems that is exempt from the requirement that all vowels have the same frontness consists of loanwords. In loanwords, any combination of front and back vowels is possible in the stem. The frontness of any suffix vowel will depend on the frontness of the last stem vowel.

- (4.16) büro-nak ‘bureau (dat.)’
sofőr-nek ‘chauffeur (dat.)’

4.2.2 Analysis

Every nuclear position in a Hungarian word will have exactly one secondary articulation node, whose site specification will be Lip, Pal, or Pha. In addition, the nuclear position itself, that is, its root node, may have a site specification of Pal and may have a degree of 2. Together with reasonable assumptions about the default articulators for each site and the defaults of an unspecified root node, these simple principles are enough to characterize the Hungarian vowel system (ignoring length distinctions). Each possible combination of the allowed substructures results in a legal Hungarian vowel with the appropriate phonological properties, and each Hungarian vowel can be represented by one of the possible combinations. The representations I assume for the vowels are in table 4.1.

The fact that *i* and *é* can result from two different PSs will be significant.

The details of restricting segmental PSs to only those in table 4.1 can be handled by the following constraints:⁶

- (4.17) All vowels must have at least one secondary articulation, whose site will be from the set {Pha, Pal, Lip}.
- $$\forall v \text{ nucleus}(v) \rightarrow \exists \text{sec } v \xrightarrow{2} s \wedge (\text{sec} \xrightarrow{s} \text{Pha} \vee \text{sec} \xrightarrow{s} \text{Pal} \vee \text{sec} \xrightarrow{s} \text{Lip})$$

- (4.18) There will be at most one secondary articulation.
- $$\text{unique}(\text{secondary})$$

⁶*unique* is a sort over arc sorts. Most other government types also belong to it, e.g., onset-licence, articulator. It allows a concise expression of the uniqueness constraint for all its members:

$$\forall g, x, y, z (\text{unique}(g) \wedge x \xrightarrow{g} y \wedge x \xrightarrow{g} z) \rightarrow y = z$$

<p><i>/u/</i></p> <p>v \ o Lip</p>	<p><i>/u/</i></p> <p>v \n Pal o Lip</p>	<p><i>/o/</i></p> <p>v /\n 2 o Lip</p>	<p><i>/o/</i></p> <p>v / \n 2 Pal o Lip</p>
<p><i>/i/</i></p> <p>v \ o Pal</p>	<p><i>/i/</i></p> <p>v \n Pal o Pal</p>	<p><i>/e/</i></p> <p>v /\n 2 o Pal</p>	<p><i>/e/</i></p> <p>v / \n 2 Pal o Pal</p>
		<p><i>/a/</i></p> <p>v /\n 2 o Pha</p>	<p><i>/e/</i></p> <p>v / \n 2 Pal o Pha</p>

Table 4.1: Representations of Hungarian vowels

(4.19) The only possible root node site is Pal.

$$\forall v, Site \text{ nucleus}(v) \wedge v \xrightarrow{s} Site \rightarrow Site = Pal$$

The defaults and redundancy rules I am assuming are as follows. The sites [s:Lip], [s:Pal], and [s:Pha] will redundantly have the articulators [a:Lab], [a:Dor], and [a:Rad] respectively. If the segment has no [s:Pal] specification, the root node's default will be [a:Dor,s:Pha] if a [2:[s:Pha]] is present (judging from Ringen's phonetic value of [a:] for *á*), and otherwise [a:Dor,s:Vel], that is, a mid or high back vowel. There is a redundancy rule requiring [d:2] in the environment of [2:[s:Pha]], otherwise the default rule filling in [d:1] will apply in the phonology-phonetics mapping.


Since all and only **constriction** nuclei in Hungarian participate in harmony, we can assume that the spreading operates along line 0 of metrical structure, where nulls are not allowed to be **line-0**. The government relation between positions on line 0 is **10g**. We want to say that in native words, **10g** arcs are local domain creators for the feature **site**.⁷

(4.20) Hungarian Palatal harmony

$$\forall g \text{ 10g}(g) \wedge \text{within-native-word}(g) \rightarrow \text{local-domain-creator}(g, \text{site})$$


For example, in *öröm-nek* the line 0 nuclei are connected to each other by government arcs that create local domains for site, so they are all required to share their site features. (For readability, I suppress onset specifications in the following diagrams and label nuclear position with the orthographic spelling of the “syllable”.)

(4.21)




⁷ *within-native-word*(*g*) will be true if it is true that $\exists W \text{ within}(g, W) \wedge W \xrightarrow{\text{level}} \text{word} \wedge W \xrightarrow{\text{class}} \text{native}$. In the discussions of vowel harmony in this chapter, I shall often use the predicates *within* and *within-word*, which will be defined in chapter 5.

For further readability, from now on I shall also suppress the nuclear spine that line 0 is built on, absorbing any consonants dominated by null nuclei into the label of the preceding nucleus. *öröm-nek* would now look like:




Constraint (4.19) guarantees that the only site specification that could ever be in a position to be shared among the root nodes of nuclei is [s:Pal]. The creation of local domains by 10g means that either every nuclear root node in a word will have [s:Pal] or none of them will. The first possibility results in front-vowel words like *öröm-nek* in (4.22), the second in back-vowel words like *város-nak* in (4.23):




Here, none of the nuclei have an [s:Pal]. The phonology-phonetics mapping will thus give them their default phonetic interpretations, [s:Vel] or [s:Pha], resulting in non-front vowels.

The possibility that an [s:Pal] specification could be either at the root node of the nucleus or at the secondary articulation node is what allows the neutral behaviour of the front vowels. In *radír-nak*, the *í* has its palatal site specified at the secondary articulation node rather than the root node.



10g creates a local domain only for **site**, not for **secondary**, and creates that local domain only between the two nodes that it links. Within this local domain, there is simply no [s:Pal] available for spreading. The [s:Pal] hiding downstairs in the secondary articulation is no more available for spreading than secondary [s:Pha]s and [s:Lip]s are. Since the other vowels in (4.24) have no tongue body specifications, they will receive their defaults. This is how a front vowel can occur in what is otherwise a back-vowel word. It also explains why only unrounded front vowels can be neutral. Only unrounded vowels can have their [s:Pal]s hiding downstairs in the basement. A front round vowel like *ü* needs a secondary articulation of [s:Lip]. Since there can be only one secondary articulation, which is already taken, *ü*'s [s:Pal] can only be on the root node, making it available for spreading and creating a front-vowel word.

There is no reason to expect that only vowels in the middle of a word can exploit the possibility of having [2:s:Pal] rather than [s:Pal]. It should be possible for the initial vowels of words to use this structure as well, resulting in stems that begin with (or consist entirely of) front vowels but nonetheless harmonize like back-vowel stems. And this is exactly what we find with words like *híd-nak*.



Finally, loanwords are exempt from the *local-domain-creator(10g,site)* requirement of native words. As a morphological requirement, however, their suffixes must still be linked to the stem with a domain-creating 10g. Though there may be other local domains in a loanword—nothing bars the possibility—the only local domain that is guaranteed to exist is the one between the suffix and the last vowel of the stem, resulting in the observation that the frontness of a loanword's suffix depends on that of the last vowel.

4.3 Asymmetric vowel harmony: Kalenjin

Kalenjin is the name of a family of closely related Nilo-Saharan languages, often also known as Southern Nilotic. The most detailed phonetic study of various Kalenjin dialects is Tucker (1964). Hall et al. (1974) discuss the vowel harmony process of Kalenjin in some depth. Kalenjin vowel harmony has also made cameo appearances in more recent autosegmental and metrical literature, e.g., Halle and Vergnaud (1981), Lieber (1987), Hammond (1988).

4.3.1 Data


Like many others in Africa, Kalenjin languages are characterized by two series of vowels. In traditional terminology, one set is “close” or “tense”, the other “open” or “lax”. More recent researchers, like Hall et al. (1974), have labelled the appropriate feature Advanced Tongue Root.

(4.26)	[+ATR]	[-ATR]	
	i	u	I U
	e	o	E O
	a ₊		A ₋


With a handful of exceptions to be discussed below, all vowels in a word must belong to the same class. While in symmetric harmony systems like Hungarian, which class the word’s vowels belonged to was determined by the class of the stem, in an asymmetric system like Kalenjin, either stem or affixes can determine the class of the word as a whole. Morphemes belong to two classes: **dominant** (or “underlyingly” [+ATR]) and **recessive** (or “underlyingly” [-ATR]). If there is a dominant morpheme anywhere in the word, all vowels in the word are [+ATR], as in (4.27). But if the word consists solely of recessive morphemes, all vowels are [-ATR], as in (4.28). Data are from Hall et al. (1974), and represent the Elgeyo dialect of the Nandi-Kipsigis-Elgeyo branch of Kalenjin.

(4.27)	<i>Dominant stem keer:</i>		
	kI-	A-keer-In	→ kiageerin
	Distant Past-I-	see -2sg object	‘I saw you’
	<i>Dominant suffix e:</i>		
	kI-	A-kEr -e	-∅ → kiagere
	Distant Past-I-	shut-Non-Completive-3sg object	‘I was shutting it’
(4.28)	<i>All morphemes recessive</i>		
	kI-	A-pAr-In	→ kIAbArIn
	Distant Past-I-	kill -2sg object	‘I killed you’


So far, we have a system that could easily be analyzed as having only [+ATR] underlyingly specified. All [-ATR] vowels are underlyingly unspecified. [+ATR] can spread bidirectionally from any (dominant) stem or affix that contains it onto all unspecified vowels in the word, as in (4.29). Near the end of the derivation, [-ATR] will be filled in by default on all vowels that do not yet have an [ATR] specification (this would only have the chance to apply in words where all morphemes are recessive, that is, unspecified for ATR).



This would indeed be all there was to say about Kalenjin vowel harmony, except for the existence of **opaque** morphemes.⁸ Hall et al. have identified three morphemes which always appear with [-ATR] vowels, resist being made [+ATR] by a dominant morpheme, and shield any recessive morphemes lying beyond it from the effects of a dominant morpheme. For example, the negative prefix *mA* remains [-ATR], even in a word with the dominant stem *keer*, and shields the recent past morpheme to its left so that it too will surface as [-ATR]:



This might be taken as evidence that we also need an underlying [-ATR] for the opaque morphemes, giving ternary power to ATR specifications. The [+ATR] of a dominant morpheme could spread bidirectionally until it gets blocked by a prelinked [-ATR] on the vowel of an opaque morpheme:



Then normal [-ATR] default insertion would apply to the recent past morpheme *kA*.

We need not admit the possibility of ternary specification in the framework we have been developing. Opacity to spreading is not a property of nodes or their specifications, per se, but of the government relations that join nodes. It would be possible to get the effect of an opaque morpheme simply by decreeing the nuclear government arc that joins it to the rest of the word to be not *local*.


But before we rush head-long into an analysis based upon the description and comments so far, we should take a moment out for some phonetic honesty. The characterization of the two classes of vowels as [+ATR] and [-ATR] was not entirely accurate. For mid and high vowels, it is fairly clear that the relevant feature is ATR. The Nandi-Kipsigis-Elegyo versions of the [+ATR] mid and high vowels have been described variously by Tucker (1964) and Hall et al. (1974) as “tense”, “close”, “slightly breathy”, “hard”, “markedly bright”. The [-ATR] mid and highs are “open”, “lax”, “creaky”. The *a* of the [-ATR] set is transcribed by Tucker as [a], and described as “mid-way between Cardinal Vowel 4 [a] and Cardinal Vowel 5 [ɑ]”. It certainly seems

⁸and except for a phonetic complication to be discussed below.

appropriate to represent this vowel as what I have called a bare RTR specification: [a:Rad, s:Pha, d:R].

The *a* of the [+ATR] set, however, is markedly unlike anything ATR. Tucker usually transcribes it as [æ], sometimes [a]. It “varies in value between Cardinal Vowel 5 [a] when unstressed, and Cardinal Vowel 6 [ɔ] when stressed” (Tucker 1964:452). Both Tucker and Hall et al. remark that it is often indistinguishable from lax /O/. This would seem to be the sort of sound that the segmental model of Chapter 2 would have to represent as distinctly RTR:

(4.32)



Indeed, Tucker admits, “It has no accompanying hollowness of voice, nor any acoustic claim to be regarded as a ‘Close’ vowel.”

For the remainder of this section, I will write the [-ATR] set’s [a] as /A/, and the “[+ATR]” set’s [æ]~[a] as /a/.

If we were developing a framework that could use non-monotonic rewrite rules, this kind of phonetic glitch would be no problem. We could simply spread [+ATR] onto the appropriate *as*, and fix the result later with some low-level phonetic clean-up rule to change the [+ATR] to [-ATR], being careful to change some other features too so that /a/ and /A/ remained phonetically distinct. It would also pose no problem if we were willing to assume that the phonology-phonetics mapping had the perverse ability to systematically interpret a constriction specification as its exact opposite. Since neither of these is true of the present enterprise, we must come up with an adequate account that respects phonetic reality.

4.3.2 Analysis

As has been suggested in the foregoing informal discussion, it is useful to assume, at least for mid and high vowels, that [-ATR] is the unmarked feature and hence unspecified. We can distinguish between the mid and high vowels by their site (Palatal or Velar), their degree (1/high or 2/mid), and the presence or absence of a [+ATR] secondary articulation:

$$(4.33) \quad \begin{array}{cccc} /i/ & /I/ & /u/ & /U/ \\ \begin{bmatrix} s & \text{Pal} \\ d & 1 \\ 2 & [d:A] \end{bmatrix} & \begin{bmatrix} s & \text{Pal} \\ d & 1 \end{bmatrix} & \begin{bmatrix} s & \text{Vel} \\ d & 1 \\ 2 & [d:A] \end{bmatrix} & \begin{bmatrix} s & \text{Vel} \\ d & 1 \end{bmatrix} \\ /e/ & /E/ & /o/ & /O/ \\ \begin{bmatrix} s & \text{Pal} \\ d & 2 \\ 2 & [d:A] \end{bmatrix} & \begin{bmatrix} s & \text{Pal} \\ d & 2 \end{bmatrix} & \begin{bmatrix} s & \text{Vel} \\ d & 2 \\ 2 & [d:A] \end{bmatrix} & \begin{bmatrix} s & \text{Vel} \\ d & 2 \end{bmatrix} \end{array}$$

These are the fullest possible representation of these vowels. They include some specifications that we shall later decide are more appropriate as defaults.

The [-ATR] set low vowel, /A/ or [a], can be represented as a bare RTR specification, [a:Rad, s:Pha, d:R]:

$$(4.34) \quad \text{recessive /A/} = [a] \begin{bmatrix} a & \text{Rad} \\ s & \text{Pha} \\ d & \text{R} \end{bmatrix}$$

As discussed earlier, the most phonetically faithful representation for “[+ATR]” /a/ would be the same representation we use for [a] in other languages:

$$(4.35) \quad \text{dominant /a/} = [a] \begin{bmatrix} a & \text{Dor} \\ s & \text{Pha} \\ 2 & \begin{bmatrix} a & \text{Rad} \\ s & \text{Pha} \\ d & \text{R} \end{bmatrix} \end{bmatrix}$$

that is, with a primary constriction between the tongue body and the pharynx and a secondary constriction between the tongue root and the pharynx.

The unmarked low vowel /A/ must have the ability to alternate between /A/ in recessive environments and /a/ in dominant environments. So the minimal description for “underlying” /A/, that is, the description of a low vowel in the lexical constraint of a recessive morpheme, must be general enough that it can be satisfied by both (4.34) and (4.35). We can see that the only things that the two representations have in common is the [s:Pha] feature on the root node. This leaves us with the following simple representation for /A/:

$$(4.36) \quad \begin{array}{l} v \\ s| \\ \text{Pha} \end{array}$$

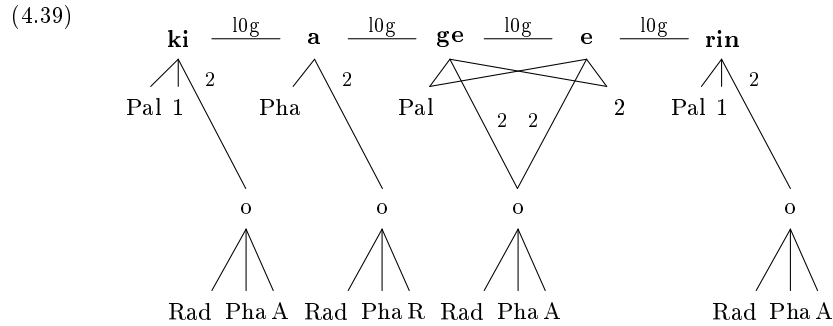
We can fill this out by some redundancy rules. Basically, the articulator should be Dorsal if there is a secondary articulation and Radical if there is none. We can

further generalize this rule to cover the Dorsal articulators of mid and high vowels as well: a vocalic root node should be [a:Rad] if [s:Pha] and there is no secondary articulation, otherwise [a:Dor]. The statement of this is simplified considerably if the [a:Rad] component is a redundancy rule and the [a:Dor] is a default:

$$(4.37) \quad \forall v (v \xrightarrow{s} Pha \wedge \sim \exists sec \ v \xrightarrow{2} sec) \rightarrow v \xrightarrow{a} Rad$$

$$(4.38) \quad \forall v \text{ default}(v \xrightarrow{a} Dor)$$

As in Hungarian, we can assume that it is line-0 government arcs (10g) that are the local-domain-creators. We must decide what features they create local domains for: we must figure out exactly what it is that is spreading, that is, what features are being shared by all vowels in a dominant environment, and hence what features are missing from vowels in a recessive environment. Consider the following overly complete representation for the dominant word *kiageerin*:



The most obvious difference between dominant and recessive vowels is the presence of an ATR secondary articulation on dominant mid and high vowels and its absence from recessive high and mid vowels. A secondary articulation is also present in dominant /a/ but absent from recessive /A/. Clearly, secondary articulation governments will play a role in the local domains, i.e.:

$$(4.40) \quad \text{local-domain-creator}(10g, \text{secondary})$$

This is in fact sufficient to derive the spreading effects. Recall that under the recursive definition of locality, complex nodes like secondary articulations need not be shared outright. The only requirement is that the presence of a secondary articulation on one node of a local domain should imply the presence of secondary articulations on the other nodes. This is enough to distinguish our two vowel sets: the [-ATR] vowel

set has no secondary articulations, the “[+ATR]” set has. If the communal list of 10g is only {2}, how the secondary articulations of the “[+ATR]” set are filled out with gesture specifications is unaffected by locality. /a/ is free to use a different degree on its secondary articulation than the mid and high vowels.⁹

All we need now is a redundancy rule to fill in the relevant gestural features on secondary articulations: [a:Rad], [s:Pha], and, if the root node is not [s:Pha], [d:A]:

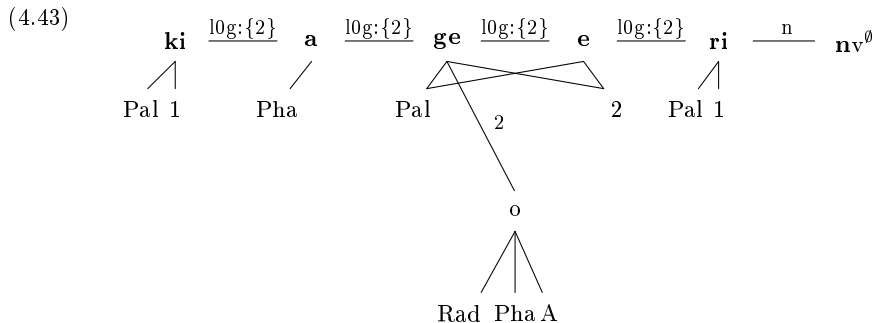
$$(4.41) \quad \forall v, sec \quad v \xrightarrow{2} sec \rightarrow \\ sec \xrightarrow{a} Rad \wedge sec \xrightarrow{s} Pha \wedge (v \xrightarrow{s} Pha \vee_{ex} sec \xrightarrow{d} A)$$

and a default rule to interpret all other radical articulator nodes as [d:R]:

$$(4.42) \quad \forall x \quad x \xrightarrow{a} Rad \rightarrow default(x \xrightarrow{d} R)$$

This last default rule also applies to the root node radical gesture when /A/ is in a recessive environment.

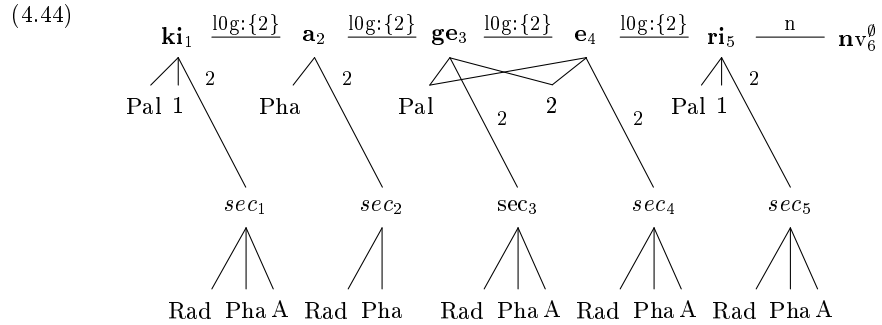
Let us look in more detail at how these constraints apply to Kalenjin words. In *kiageerin*, for example, the dominant ATR specification is a property of the verb stem. Let us assume that (at least) the first vowel of the verb stem is lexically constrained to have an ATR secondary articulation. The parts of the PS that the lexical constraints of the stem and affixes are interested in can be diagrammed as follows:¹⁰



The more widely applicable constraints of Kalenjin require that any PS containing the pieces in (4.43) must also contain a number of other pieces and that a number of properties will be true of the PS. A more complete diagram of *kiageerin*'s PS is:


⁹Of course, recursive locality will also require all the secondary articulations to form a local domain for 2. But since secondary articulations dominating secondary articulations are ruled out by an independent universal principle, this requirement will always be satisfied vacuously.

¹⁰Although it is tempting to think of the diagram in (4.43) as the “underlying representation” of the word, it is simply an aid to comprehension. It is not a legal PS of Kalenjin and should not be taken as having any theoretical status.



The final v_6 is null, and therefore not a line 0 position. It does not participate in the harmony. The other vowels have constriction nodes, are line 0 positions, and are joined by 10g arcs. 10g arcs are local domain creators with the communal list {2}. Therefore, v_1 and v_2 are in a local domain for the communal list, as are v_2 and v_3 , v_3 and v_4 , v_4 and v_5 . The communal list is in this case simply {2}. If any root node has a dependent governed through a 2 arc, then all of them must. The lexical constraint of the verb stem indeed requires the existence of s_3 joined by a 2 arc to v_3 , so all other root nodes will also have a secondary articulation. The redundancy rules in (4.41) will require the presence of the correct sites and degrees on the s_i s.

In an all-recessive word, no vowel will have a secondary articulation. In *ki-A-bAr-In*, the minimal requirements of the lexical constraints of the morphemes can be diagrammed as:



None of the nodes are specified with secondary articulations, so the sharing requirement of locality is satisfied vacuously. The node-based default rules will interpret unspecified [a:Rad] nodes as [d:R]s and the articulator-based default rules will result in [d:R] tongue root gestures on all other vowels.

As mentioned in section 4.3.1, the three “opaque” morphemes can be handled by lexically stipulating their opacity. Simply declaring the arc g joining, say, the negative morpheme to the rest of the verb to be $\sim local-domain-creator(g, secondary)$ is not quite enough. If the locality requirement were simply that *all* 10g arcs in a word were local domain creators, this would result in a contradiction that could not be satisfied by any PS. A trickier definition of the locality requirement would be necessary, something along the lines of: for every morpheme m , either m has the semantics of ‘negative’

(*mA-*), the semantics of ‘perfectivizer’ (*kA-*), the semantics of ‘reflexive’ (*kEE-*), or a phonology such that every *lOg* between its top and bottom is a local domain creator for {2}, as is *lOg* arc leading from its bottom if this is not also the bottom of a phonological word.

This is certainly possible, but it is a little complex. I suspect it is more likely that Kalenjin uses a simpler characterization of local domain creators (i.e., all *lOgs* between the top and bottom of a phonological word) and that the three morphemes in question are not part of the same phonological word as the verb stem, at least in the narrow sense in which English *good* and *-ness* are not part of the same phonological word. The semantics of the morphemes does not strongly suggest otherwise: they have the kinds of meanings that cross-linguistically are not infrequently represented by phonologically independent particles, modals, or pronouns. The theory of prosodic constituency and its possible mismatches with morphosyntactic constituency, as presented in section 5.4, will provide a mechanism by which these three morphemes can force a prosodic boundary. Of course, much more would need to be known about Kalenjin phonology and morphology in order to adequately evaluate this hypothesis.

4.4 Pasiego

The Pasiego dialect of Montañes Spanish is an especially interesting example for this chapter, containing not one but two harmony processes—a height harmony similar to Hungarian’s front harmony and an ATR harmony similar to Kalenjin. Like Hungarian frontness harmony, the height harmony component of Pasiego has an additional complication in that low vowels are apparently transparent to it—they seem not to participate in it, but they do not block it.

Since its first description by Penny (1969a,b), Pasiego has played a large role in autosegmental theory. McCarthy (1984) analyzed the height component as a bivalent feature-changing harmony, that is, when the [high] spreads it destroys whatever former [high] values lie in its path. Since in this analysis both values, [+high] and [–high], are required to spread, Pasiego has also been used as evidence for bivalent features against privacy, and for contrastive specification against radical underspecification (e.g., Steriade 1987). Vago (1988), however, has shown how a framework that has only one feature value in underlying representation and spreads only that value can account for Pasiego height harmony and in fact does so more satisfactorily than the contrastively specified feature-changing analysis.

4.4.1 Data

Tense harmony

As in Kalenjin, there are two sets of vowels in Pasiego, which we may for convenience call the tense set and the lax set:

(4.46)	<i>tense</i>	<i>lax</i>	
	i u	I U	
	e o	O	
	a	A	

As Kalenjin might have led us to suspect, languages do not necessarily cooperate by having straightforward phonetic correlates for their harmonic classes. “Lax” /A/, for example, is described by Penny (1969a: 49) as very palatal, close to close [e], and is probably best represented the way we have been representing ATR /æ/, that is [a:Dor, s:Pha] with a secondary ATR articulation. Another wrinkle is the absence of /ε/ in the lax set. /e/ appears in words of either set.

As in Kalenjin, all vowels in a word must be drawn from the same set. In the natural order of things, vowels come from the tense set. All words which involve lax vowels contain the masculine singular suffix *-U*. As the following data from McCarthy (1984:293) show, we find morphophonemic alternations, for example, between masculine plural and masculine singular nouns (4.47), between feminine nouns and their masculine (often diminutive) forms (4.48), and between feminine adjectives, masculine mass adjectives, and masculine count adjectives (4.49).

(4.47)	<i>Tense words</i>		<i>Lax words</i>	
	soldáus	‘soldiers’	sOldÁU	‘soldier’
	kastáñus	‘chestnut trees’	kAstÁñU	‘chestnut tree’
	simpátikus	‘congenial (pl.)’	sImpÁtIkU	‘congenial (sg.)’

(4.48)	pitrína	‘waistband’	pItrínU	‘waistband (dim.)’
	trípa	‘belly’	trÍpU	‘belly (child’s)’
	gulundrína	‘swallow (fem.)’	gUIUndrínU	‘swallow (masc.)’

(4.49)	fem.sg.	sg. mass	sg. count	
	mála	málu	mÁIU	‘evil’
	límpja	límpju	lÍmpjU	‘clean’
	súθja	súθju	sÚθjU	‘dirty’

The neutrality of /e/ is shown in the following words, where it is the only vowel that does not alternate between lax and tense:

(4.50)	<i>Tense words</i>		<i>Lax words</i>	
	ermáñus	‘brothers’	ermÁñU	‘brother’
	peñáskus	‘cliffs’	peñÁskU	‘cliff’
	komfesonárjus	‘confessionals’	kOmfesOnÁrjU	‘confessional’
	kampečáñus	‘noble (pl.)’	kAmpečÁñU	‘noble (sg.)’

Height harmony

All non-low vowels in a word¹¹ have the same height as the vowel of the stressed syllable. (The data below are taken from Vago 1988.)

This harmonic principle operates both distributionally and morphophonemically. Distributionally, it acts as a morpheme-structure constraint, requiring all non-low vowels in a word to be either all high or all mid.

- (4.51) *All high and tense* *All mid and tense*
 biniθír ‘to bless’ xeléča ‘fern’
 čipúduš ‘hunchbacks’ belórta ‘hay-rake’
 lubúkus ‘young wolves’ destorθér

All high and lax
 piθígU ‘pinch’
 kÚntÍntU ‘happy’ (count)
 mInÚdU ‘small’ (count)

It also plays a more active role, creating morphophonemic alternations.

- (4.52) *All mid*
 bebér ‘drink’ koxér ‘take’ infinitive
 bebémus koxémus 1pl present indicative
 beberé koxeré 1sg future
- All high*
 bibí:s kuxí:s 2pl present indicative
 bibía kuxía 1sg imperfect indicative
 bibírí:s kuxírí:s 2sg conditional

The harmony also affects proclitics and the definite article:

- (4.53) el pélu ‘the hair’ (mass)
 Il kÚrdÍrU ‘the lamb’
 i mi díxu ‘he said to me’
 me lo kompró ‘he bought it for me’

Low vowels are neutral to height harmony. They seem not to participate in it, but they do not block it either. In *legatérna* ‘lizard’, the stressed *é* can affect the initial *e* across the intervening *a*.

¹¹Actually, the domain of spreading is somewhat larger, including some clitics on verbs and the definite article on nouns.

(4.54)	legatérna	‘lizard’
	IskAlAmbrÚkU	‘dog-rose’
	se kasó	‘he got married’
	Il mAdÍrU	‘the log’

An interesting situation arises if the stressed vowel of a word is low. In this case, stems divide between high and mid:

(4.55)			subjunctive	cf. infinitive
	bebámus	‘drink’	1pl present	bebér
	koxámus	‘take’	1pl present	koxér
	sintáís	‘feel’	2pl present	sintír

Many of these stems like ‘feel’ that surface as high in the absence of any triggering stressed high vowel nevertheless have mid-vowel alternants, as in the 1pl present indicative *sentémus*. This has been taken as evidence by researchers like McCarthy that both values of [high] must be specified underlyingly, and height harmony must be able to destroy these underlying values.¹²

There are some aspects of the height harmony system that I shall not be dealing with. A complete analysis would be possible, but peripheral to the main points of this section. I shall not discuss the behaviour of final vowels, which are limited to the set *e, u, U, a*, and often are not harmonic with a preceding stressed vowel. McCarthy argues that harmony applies to final vowels anyway, but they are reduced by a late rule and so may become non-harmonic. For the purposes of the present section, we may assume that height harmony only occurs between the stressed vowel and the vowels to its *left*, though a complete analysis would want to capture some of the interesting regularities in the behaviour of final vowels as well. Another aspect I shall not deal with is the behaviour of glides as harmonic triggers.¹³

4.4.2 Feature-changing and feature-adding analyses

This section reviews some of the arguments put forward for analyzing the height harmony of Pasiego as feature-changing or feature-adding, for the most part following the

¹²Nouns show a similar contrast in “underlying” height: *pigáa* ‘magpie’ versus *ontárga* ‘lard’.

¹³In words like *bibjéndu* and *miludjó*, [+high] harmony is triggered by the glide rather than the apparent nucleus, which may disharmonically be a mid vowel. This could be easily handled if we assumed that these glide–vowel sequences are in fact light diphthongs and are represented as argued for previously, that is, the high “glide” portion is specified on the nucleus’ root node and the mid (or low) portion is specified on a dependent release node. In this type of representation, it is the high vowel that lives on the nuclear tier and is available for harmony, while the non-high vowel is hidden downstairs and has no requirement to be harmonic. Obviously, a much more in-depth analysis would be required to determine if this kind of light diphthong representation is consistent with other aspects of Pasiego phonology and morphology.

work of McCarthy (1984) and Vago (1988).

McCarthy argues that non-low vowels in Pasiego are underlyingly marked either as [+high] or [-high]. Low vowels are unspecified for the feature [high] and are specifically marked as not belonging to the class of segments that can bear [high]. Thus the stems of the verbs *beber* and *sintir* are underlyingly [-high] and [+high] respectively. The underlying values surface when the stressed vowel of the word is an *a*, which triggers no harmony:

$$(4.56) \quad \begin{array}{cccccccc} b & e & b & \acute{a} & m & u & s & & s & i & n & t & \acute{a} & i & s \\ & & | & & & & & & & & | & & & & & \\ & & [-\text{high}] & & & & & & & & [+high] & & & & & \end{array}$$

When the stems are followed by a non-low stressed vowel, however, these underlying specifications must be deleted and replaced by those of the stressed vowel, as in *bibir* and *sentemus*.

$$(4.57) \quad \begin{array}{cccccccc} b & i & b & i & i & r & & s & e & n & t & e & m & u & s \\ & \models & & \backslash & / & & & \models & & & | & & & & & \\ & [-\text{high}] & & [+high] & & & & [+high] & & & [-high] & & & & & \end{array}$$

For McCarthy, height harmony is in fact a type of deletion rule:

$$(4.58) \quad \text{McCarthy's height harmony (mirror image)} \\ \begin{array}{ccc} [\text{high}] \rightarrow \emptyset \% & \text{---} & [\text{high}] \\ & | & | \\ & [-\text{str}] & [+str] \end{array}$$

Spreading from the stressed vowel then follows in a separate stage.

Vago (1988) takes another approach, trying to reconcile the data of Pasiego with the framework of Radical Underspecification. Vago argues that while there is ample evidence for raising mid vowels, there is no similarly strong evidence for lowering high vowels. From this, he proposes that only [+high] is active in the phonology of Pasiego, and that [-high] is filled in as a default. As Vago points out, there is plenty of evidence for “raising contexts”. Verbs like *beber* that emerge with mid vowels ordinarily (when the stressed vowel is anything but a high vowel) emerge with high vowels with perfect regularity before a number of verbal suffixes: *-ís* ‘2pl present indicative’, *-í* ‘imperfect’, *-í* ‘perfect’, *-ís* ‘2pl future’, *-í* ‘conditional’, and also in past participles, e.g., *bibú*, *kumú*.

In contrast, there is no similarly strong evidence for lowering high vowels before stressed low-vowelled suffixes. For example, Vago presents the following paradigms for the “underlyingly” high vowel verb *sintir*:¹⁴


¹⁴Some of these forms are reconstructed by Vago. *sintirémus*, for example, is not attested in Penny’s descriptions, though *sintiré* and other crucial forms like it are.

(4.59)

<i>Present indicative</i>	<i>Present subjunctive</i>	<i>Future</i>
síntu	sínta	sintiré
síntes	síntas	sintirás
sínte	sínta	sintirá
sentémus	sintámus	sintirémus
sintí:s	sintáis	sintirís
sínten	síntan	sintirán

The crucial forms here are the 1pl present indicative and the 1sg and 1pl future. While *sentémus* shows the lowering predicted by the feature-changing harmony rule, *sintiré* spectacularly and regularly fails to. Indeed, the 1pl present indicative of the small class of *-ir* verbs is the *only* place in the language where any such lowering occurs, and even in this context it is highly variable. Some other *-ir* verbs fail to show this lowering, e.g., *iskupémus* ‘we spit’, others do so optionally, e.g., *iémus / eémus* ‘we say’, and still others show variation in the stressed vowel itself, e.g., *salémus / salímus* ‘we leave’. More examples of disharmonies of this type from the nominal derivation system can be found in McCarthy (1984:297–8). From these facts, Vago concludes that the *sentémus* forms are aberrations and do not result from a regular height harmony rule.¹⁵

The resulting pattern—where stressed high vowels in suffixes can raise mid vowels in stems, but where stressed mid vowels in suffixes cannot lower high vowels in stems—is exactly what we would expect from a feature-adding harmony system where only [+high] was specified and [-high] was the default. This is what Vago proposes. His harmony process is the result of the following cyclic spreading rule:

(4.60) *Vago’s H-spread (mirror image)*

For example, to derive the contrast between the infinitive *koxér* ‘to cook’ and the 1sg conditional *kuxiría*, Vago has the [+high] of the suffix spread onto the underlyingly unspecified vowel of the stem (and infinitive suffix):

¹⁵He writes: “There are thus reasons to consider the 1pl present indicative forms of *-ir* verbs not be representative of the regular patterns of height harmony. There might even be a functional reason why the stem vowel is not the expected /i/: Penny (1969a:123) conjectures that since *sintímus* is the 1pl perfect inflection, mid vocalism in the 1pl indicative avoids homonymy. In any event, the number of verbs that assimilate in height to *-émus* is highly limited: the great majority of verbs belong to the *-ár* class (Penny 1969a), where the present indicative forms have /á/. In brief, the set of verbs in which lowering applies in the context of *-émus* can simply be memorized.” (Vago 1988:353)

$$(4.63) \quad \begin{array}{cc} /i/ & /I/ \\ \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Pal} \\ \text{d: 1} \end{array} \right] & \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Pal} \\ \text{d: 1} \\ \text{2: } \left[\begin{array}{l} \text{a: Rad} \\ \text{s: Pha} \\ \text{d: R} \end{array} \right] \end{array} \right] \end{array} \quad \begin{array}{cc} /u/ & /U/ \\ \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Vel} \\ \text{d: 1} \end{array} \right] & \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Vel} \\ \text{d: 1} \\ \text{2: } \left[\begin{array}{l} \text{a: Rad} \\ \text{s: Pha} \\ \text{d: R} \end{array} \right] \end{array} \right] \end{array} \\ \\ \begin{array}{cc} /e/ & /e/ \\ \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Pal} \\ \text{d: 2} \end{array} \right] & \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Pal} \\ \text{d: 2} \\ \text{2: } \left[\begin{array}{l} \text{a: Rad} \\ \text{s: Pha} \\ \text{d: A} \end{array} \right] \end{array} \right] \end{array} & \begin{array}{cc} /o/ & /O/ \\ \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Vel} \\ \text{d: 2} \end{array} \right] & \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Vel} \\ \text{d: 2} \\ \text{2: } \left[\begin{array}{l} \text{a: Rad} \\ \text{s: Pha} \\ \text{d: R} \end{array} \right] \end{array} \right] \end{array} \end{array}$$

These are the *maximal* representations. Clearly not all parts will be specified by the parts of lexical constraints that are interested in segmental content. [a:Dor] specifications on the root node can safely be made the subject of either a default rule or a redundancy rule, as can [a:Rad] and [s:Pha] specifications on secondary articulation nodes. For concreteness and the simplification of other default rules, I shall assume this is a default rule. Following the general pattern of other Spanish dialects, I shall assume Palatal is the default site, though nothing in the analysis depends on this:

$$(4.64) \quad \forall v \text{ default}(v \xrightarrow{\text{a}} \text{Dor})$$

$$(4.65) \quad \forall v \text{ default}(v \xrightarrow{\text{s}} \text{Pal})$$

As Vago's analysis suggests, we shall want to make [-high], that is [d:2], the default for root nodes, so [d:1] is the only value that constraints need to specify.

$$(4.66) \quad \forall v \text{ default}(v \xrightarrow{\text{d}} 2)$$

The [d:A] or [d:R] specifications on secondary nodes are predictable from the root nodes: [d:A] in the case of /e/ (that is, if the root node is *unspecified*) and [d:R] otherwise:

$$(4.67) \quad \forall v, \text{sec } v \xrightarrow{2} \text{sec} \wedge \text{unspecified}(v) \rightarrow \text{default}(\text{sec} \xrightarrow{\text{d}} \text{A})$$

$$(4.68) \quad \forall v, sec \ v \xrightarrow{2} sec \rightarrow default(sec \xrightarrow{d} R)$$

The ATR gestures that occur in the phonetic interpretation of non-low vowels without secondary articulations are the result of an articulator-based default rule.¹⁷

The two low vowels will be distinguished by the usual $\text{æ} \sim \text{a}$ contrast, represented maximally as:¹⁸

$$(4.69) \quad \begin{array}{l} [\text{æ}] = \text{“lax” } /A/ \\ \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Pha} \\ 2: \left[\begin{array}{l} \text{a: Rad} \\ \text{s: Pha} \\ \text{d: A} \end{array} \right] \end{array} \right] \end{array} \quad \begin{array}{l} [\text{a}] = \text{“tense” } /a/ \\ \left[\begin{array}{l} \text{a: Dor} \\ \text{s: Pha} \\ 2: \left[\begin{array}{l} \text{a: Rad} \\ \text{s: Pha} \\ \text{d: R} \end{array} \right] \end{array} \right] \end{array}$$

As with the mid and high vowels, /A/’s membership in the “lax” set will be marked by the simple presence of its secondary articulation node, and /a/’s membership in the “tense” set by the absence of a secondary node:

$$(4.70) \quad \begin{array}{l} \text{“tense” } /a/ \\ \left[\begin{array}{l} \text{s: Pha} \\ 2: [] \end{array} \right] \end{array} \quad \begin{array}{l} \text{“lax” } /A/ \\ \left[\begin{array}{l} \text{s: Pha} \\ 2: [] \end{array} \right] \end{array}$$

/a/ will be phonetically realized as RTR by an articulator-based default rule. /A/’s ATR will be filled in by a node-based default rule:

$$(4.71) \quad \forall v, sec \ v \xrightarrow{2} sec \wedge v \xrightarrow{s} Pha \rightarrow default(sec \xrightarrow{d} A)$$

Removing all the specifications that are result of default rules, we obtain the following stripped-down PS representations of Pasiego vowels:

¹⁷In fact, it will be crucial that degrees on secondary nodes are unspecified in PSs and filled in by default. Height harmony will require **degree** to be in the communal list. Tense harmony will require **secondary** to be in the communal list. Under the recursive locality proposal of section 4.1, this would require all secondary articulation nodes to share any overt degree features as well, which could result in clashes between the [d:A] of /e/ and the [d:R]s of the rest of the lax set. This is solved by making secondary node degrees a matter for default rules, and marking the difference between lax and tense vowel sets simply by the presence or absence of a secondary articulation node, as in (4.72).

¹⁸Actually, the phonetic value of the “tense” set /a/ is rather more central than we would expect from the representation of [of an articulator-based default rule, we might expect somewhat more variation or deviation than we would if were the result of a node-based default rule or a PS specification.

$$\begin{array}{cccc}
 (4.72) & /i/ & /I/ & /u/ & /U/ \\
 & [d: 1] & \begin{bmatrix} d: 1 \\ 2: [] \end{bmatrix} & \begin{bmatrix} s: Vel \\ d: 1 \end{bmatrix} & \begin{bmatrix} s: Vel \\ d: 1 \\ 2: [] \end{bmatrix} \\
 & /e/ & /e/ & /o/ & /O/ \\
 & [] & [2: []] & [s: Vel] & \begin{bmatrix} s: Vel \\ 2: [] \end{bmatrix} \\
 & & /a/ & /A/ & \\
 & & [s: Pha] & \begin{bmatrix} s: Pha \\ 2: [] \end{bmatrix} &
 \end{array}$$

Vago required the completely underspecified segment to be /a/, a choice for which there is little theory-external evidence. With these segmental representations, Pasiego again resembles other dialects of Spanish in having /e/ as the completely underspecified vowel.

Tense harmony

With these representations, we can analyze tense harmony as the result of sharing secondary articulation nodes, in the loose sense developed in section 4.1. That is, the existence of a secondary articulation on one vowel requires the existence of a secondary articulation on all the others, and all of these secondary articulations will form a local domain for the communal list they inherit from their parents. Since it is secondary articulation nodes that are shared, this communal list will include at least **secondary**. (As will be seen shortly, it will also include **degree**.)

The sharing can be enforced by a constraint requiring any metrical **line-0-government** within a word to be a local domain creator for the communal list $\{\mathbf{degree}, \mathbf{secondary}\}$.¹⁹

$$\begin{array}{l}
 (4.73) \text{ Pasiego tense harmony} \\
 \forall g \text{ line-0-government}(g) \wedge \text{within-word}(g) \rightarrow \\
 \text{local-domain-creator}(g, \mathbf{secondary})
 \end{array}$$

The locality requirements implied by constraint (4.73) can be satisfied either by a PS which has a secondary articulation node on every vowel, as in (4.74), or, vacuously, by a PS with no secondary articulation nodes at all, as in (4.75).

¹⁹Again, we postpone to chapter 5 the question of how exactly to determine if the 10g is within a word—merely using a (so far) undefined predicate *within-word* to indicate the requirement. Again, it should be noted that this is an oversimplification. The most likely domain of vowel harmony in Pasiego is the clitic group (cf. Nespov and Vogel 1986) rather than the phonological word.

$$(4.74) \quad \begin{array}{cccc} \mathbf{sIm} & \underline{10g:\{2\}} & \mathbf{pA} & \underline{10g:\{2\}} & \mathbf{tI} & \underline{10g:\{2\}} & \mathbf{kU} \\ | & & | & & | & & | \\ 2 & & 2 & & 2 & & 2 \\ \circ & & \circ & & \circ & & \circ \end{array}$$

$$(4.75) \quad \mathbf{sim} \underline{10g:\{2\}} \mathbf{pa} \underline{10g:\{2\}} \mathbf{ti} \underline{10g:\{2\}} \mathbf{kus}$$

A secondary articulation is introduced into a PS by the masculine singular count suffix $-U$, which we may assume has the following as part of its lexical constraint:

$$(4.76) \quad \text{Part of lexical constraint for } -U \\ \dots \exists v_u, sec_u \dots v_u \xrightarrow{s} Vel \wedge v_u \xrightarrow{d} 1 \wedge v_u \xrightarrow{2} sec_u$$

If a word contains this suffix, this lexical constraint will require a secondary articulation node, sec_u , on the final vowel, which in conjunction with the harmonic constraint (4.73) and the definition of locality will require all the other vowels in the word to have secondary articulation nodes as well, resulting in an all-“lax”-vowel word, as in (4.77). Otherwise, no vowel will have a secondary articulation and the result will be an all-“tense”-vowel word, as in (4.78).

(4.77) *kOmfesOnÁrjU* ‘confessional’


$$\begin{array}{cccccc} \mathbf{kOm} & \underline{10g:\{2\}} & \mathbf{fe} & \underline{10g:\{2\}} & \mathbf{sO} & \underline{10g:\{2\}} & \mathbf{nA} & \underline{10g:\{2\}} & \mathbf{rjU} \\ \begin{array}{c} s \\ \swarrow \\ \text{Vel} \end{array} & | & | & \begin{array}{c} s \\ \swarrow \\ \text{Vel} \end{array} & | & \begin{array}{c} s \\ \swarrow \\ \text{Pha} \end{array} & | & \begin{array}{c} s \\ \swarrow \\ \text{Vel} \end{array} & \begin{array}{c} d \\ \searrow \\ 1 \end{array} \\ \circ & & \circ & \circ & & \circ & & \circ & \end{array}$$

(4.78) *komfesonárjus* ‘confessionals’


$$\begin{array}{cccccc} \mathbf{kOm} & \underline{10g:\{2\}} & \mathbf{fe} & \underline{10g:\{2\}} & \mathbf{sO} & \underline{10g:\{2\}} & \mathbf{nA} & \underline{10g:\{2\}} & \mathbf{rjU} \\ \begin{array}{c} s \\ \swarrow \\ \text{Vel} \end{array} & | & | & \begin{array}{c} s \\ \swarrow \\ \text{Vel} \end{array} & | & \begin{array}{c} s \\ \swarrow \\ \text{Pha} \end{array} & | & \begin{array}{c} s \\ \swarrow \\ \text{Vel} \end{array} & \begin{array}{c} d \\ \searrow \\ 1 \end{array} \\ \circ & & \circ & \circ & & \circ & & \circ & \end{array}$$

Height harmony

Recall Vago’s proposed height harmony rule (ignoring its “mirror image” aspect), repeated here for convenience:




Embedded in the framework of autosegmental phonology, this rule is an instruction for the grammar (or whatever) to go out and actively build some more structure (in this case, an association line) whenever it finds a situation that looks like the structural description. But it can also be interpreted in a more passive way that is more appropriate for a constraint-based framework. From this point of view, the rule simply requires that whenever you have two vowels that are adjacent at some level (for Vago, on the dorsal tier), and the first vowel is not stressed and the second vowel is [+high], then the first vowel must also be [+high] or the entire structure will not be a legal representation in Pasiego. We can translate this almost directly into our constraint-based framework, interpreting [+high] as [d:1] and [-stress], for the purposes of illustration, as *not* being a line 2 head. Our framework does not have a dorsal tier available for determining adjacency, so we shall have to make do with our faithful standby, **line-0-government**.




For any two vowel positions adjacent on line 0, if the first is not a main stress (i.e., **line-2**) and the second has [d:1], then the line 0 government that joins them is a local domain creator for **degree**. This has the effect of “spreading” [d:1] leftwards onto unstressed vowels.

Consider the difference between 1pl present indicative *koxémus* and 1sg conditional *kuxiría* ‘take’. We may assume that the stem ‘take’ makes no demands on PSs in terms of vowel height. In the further absence of any demands imposed by a suffix, all vowels will be interpreted with the default height, [d:2] or mid. When neither stem nor suffix requires a [d:1], the result is *koxémus*:




In *kuxiría*, on the other hand, although the stem still makes no demands on vowel height, the conditional suffix *-í* requires presence of a [d:1] on its vowel. Because of constraint (4.80), the 10g arc joining *-í* and the preceding vowel will be a local domain creator for *degree*. So the preceding vowel will also have a [d:1], which will in turn require another local domain with the vowel to *its* left, and so on. The result is:

(4.82) *kuxiría*



Notice that the formulation of the harmonic constraint in (4.80) is permissive enough to allow forms like *sintiré*, where the mid vowel of the suffix does *not* create a local domain with the vowel to its left. The vowel to the left keeps the [d:1] required of it by the lexical constraint of the verb ‘feel’, but the [d:1] does not spread rightward to the future suffix *-é*, which is realized as a mid vowel by default.²⁰

(4.83) *sintiré*



Lastly, we need to consider the neutrality of low vowels. Vago tried to account for this by having rule (4.60) scan for adjacency along the dorsal tier, and stipulating that low vowels have no dorsal node. Since the present framework has no independent tiers for individual features or class nodes, we can only refer to adjacency along the nuclear spine or one of its subsets (a metrical line). The low vowels that do not block harmony and appear not to participate in it certainly occur on the nuclear spine, and there is

²⁰More needs to be said about how lexical constraints for *-ir* verbs require [d:1] in all forms except the 1pl present indicative, how some of them fail to require it in the 1pl present indicative (e.g., *sentémus*), and how any degree features required by the stem are spread to the infinitive suffix. Though analyses are easily imaginable, deciding on the right ones would require a long discussion that would add little to the main points of this section.

overwhelming evidence from the stress system that we cannot exclude low vowels from any of the metrical lines just to make height harmony work out properly. Thus, there is no way in the present framework for height harmony to skip over intervening low vowels as if they were not there.

Fortunately, there is no need to. Low vowels do in fact participate in Pasiego height harmony. This statement might sound ridiculous at first, but it is so only if we assume that what is spreading is the feature [+high], which manifestly does not dock onto low vowels in any principled way. In the present analysis, however, what is being shared among vowels is not [+high], but [d:1]. [d:1] does not characterize vowel height per se, rather it characterizes the aperture between whatever articulator and site happen to belong to the node it is on. For root nodes of low vowels, this is not the tongue body and the roof of the mouth, but the tongue body and the pharynx.

I have not until now said much about the degree specifications of dorso-pharyngeal gestures, because until now it did not seem like a very answerable question or one that had any phonological relevance. Clearly the tongue body cannot perform a complete closure [d:0] at the pharynx. This leaves the formal possibilities of [a:Dor, s:Pha, d:1] and [a:Dor, s:Pha, d:2] that are not ruled out physiologically. I know of no language that contrasts the two gestures phonemically.


If we wish to take seriously the possibility that Pasiego low vowels do in fact participate in degree-harmony, there are two possibilities for reconciling this with the observation that degree never plays a distinctive role in dorso-pharyngeal gestures in other languages.

1. It is always a formal possibility for languages to make distinctive use of the [d:1]/[d:2] distinction in dorso-pharyngeal gestures. But the physiological control of such gestures is not precise enough for them to be reliably distinguished, and languages would generally find it pointless to avail themselves of the formal possibility. In Pasiego, however, the cues for distinguishing [d:1] from [d:2] do not rely solely on the dorso-pharyngeal gesture, but are spread out over an entire harmonic domain that typically includes vowels where the contrast is easy to hear.
2. It is not a formal possibility for languages to contrast two degrees of [a:Dor, s:Pha] gestures. Only one value is available. In Pasiego, and perhaps universally, this is [d:1]. A low vowel's root node may have an overt [d:1] spread onto it by height harmony in "high-vowel" words, otherwise (in "mid-vowel" words) it will be interpreted as [d:1] by default principles.

Choosing between these two possibilities would require good instrumental phonetic data from Pasiego speakers and a thorough analysis of the implications of each for the rest of the language and the rest of the theory. In the absence of both, I shall restrict myself to pointing out that both possibilities will result in exactly the same PSs and conclude that these PSs are plausible representations for the relevant Pasiego words, regardless of how we may eventually decide to handle their phonetic interpretation.


Compare *legatérna* ‘lizard’ and *IskAlAmbrÚkU* ‘dog-rose’. In *legatérna*, the *a* in the second syllable does not belong to a local domain for **degree** and has no **degree** feature on its root node’s dorso-pharyngeal gesture.

(4.84) *legatérna* ‘lizard’



In *IskAlAmbrÚkU*, on the other hand, because the stressed vowel is high, that is [d:1], there must be a chain of local domains for **degree** to its left. The *as* that occur between the stressed syllable and the first syllable *do* belong to these local domains and *do* bear [d:1] features on their root nodes. (The secondary articulations from tense harmony are irrelevant for the present point).

(4.85) *IskAlAmbrÚkU* ‘dog-rose’



4.4.4 Summary: Harmonic neutrality and transparency

A position may easily be neutral to harmony and opaque to that harmony, that is, block any further spreading of the harmonic feature. In the Pasiego height harmony constraint in (4.80), for instance, stressed vowels were opaque to height harmony. A [d:1] could not spread from a vowel leftwards onto a stressed vowel because this environment does not meet the criteria for setting up a local domain. Neutrality with opacity is straightforward.

More interesting is neutrality with transparency. That is, a vowel seems not to participate in the harmony, yet does not block the spread of the harmonic feature. This type of phenomenon poses a challenge for the present framework, since the transparent position must clearly belong to a local domain for the harmonic feature, yet shows none of the effects of belonging to the local domain. We have seen three different ways how

this effect can arise, one involving root nodes only, one secondary articulations only, and one the interaction between the two. All involve a situation where the phonology-phonetics interface gives two different structures the same phonetic interpretation, or at least apparently the same interpretation.

1) The type of transparency involving only secondary articulations is illustrated by the transparency of Pasiego /e/ to tense harmony. In these cases, the communal list includes *secondary*. What is phonologically relevant is not the content of the secondary articulations, but merely their presence or absence. The apparent transparency arises because the phonology-phonetics interface can fill in a completely underspecified secondary articulation in such a way that the result is indistinguishable from a segment that had no secondary articulation in the first place. In Pasiego, the apparent neutrality results from two different structures receiving the same phonetic interpretation: an unspecified bare root node [], and an unspecified root node plus an unspecified secondary articulation [2:[]]. A language may use this mechanism to get transparency for at most one feature. If, for example, a language phonologically marked both ATR and roundness on secondary articulations, a completely unspecified secondary node could be realized as only one of these and only that feature could show this type of transparency.

2) A similar type of transparency involves both the root node and a secondary articulation, and is illustrated by the neutrality of Hungarian front vowels. In this type, the local domain does not involve the sharing of a secondary articulation but of one of the root node's gestural features (in Hungarian, *site*). The transparency results from the same kind of phonetic double-interpretation as in the last type. In Hungarian, two different structures can receive the same phonetic interpretation: a root node specified for site [s:Pal], and an unspecified root node with a secondary articulation specified for site [2:[s:Pal]].

3) The third type of transparency involves only root nodes. It is illustrated by the neutrality of Pasiego low vowels to height harmony. Here, we are not dealing with the same interpretation for a bare root node and a root node plus secondary articulation, but with the same interpretation for two different root nodes. In Pasiego, these root nodes are [a:Dor, s:Pha] and [a:Dor, s:Pha, d:1]. Especially if we assume an anti-ternarity principle of the kind discussed in section 3.5.6, we should expect this kind of transparency to be extremely limited. Of the segmental gestures dealt with in this dissertation, only the dorso-pharyngeal gestures of low vowels have the necessary properties. (It may also turn out to be a possibility for other gesture-types that have not been dealt with in depth, e.g., nasality, laryngeal gestures.)

4.5 Reduplication

What all analyses of reduplication have in common, whether they are framed in representation-based or rule-based morphological or phonological models, is the operation of copying. In earlier rule-based generative phonology, this copying was done

piece by piece, as in the following Aspects-style Tagalog rule proposed by Carrier (1979), where an extra index in the structural change portion indicates a piece of the structural description to be copied:

$$(4.86) \quad \begin{array}{cccccc} \# & C & (C) & V & X & \\ 1 & 2 & 3 & 4 & 5 & \rightarrow 1\ 2\ 4\ 2\ 3\ 4\ 5 \end{array}$$

A later rule-based account of reduplication, Steriade (1988), performs a single copy operation on the entire word once at the beginning of the series of processes that constitute a reduplication. Steriade argues that this operation copies everything in the word, including its prosodic structure. After this, a battery of adjustment rules apply (deletions, insertions, etc.) in order to make the duplicate conform to a number of prosodic conditions that together act as a template for the reduplication.

A representation-based model of reduplication, along the lines of Marantz (1982), assigned most of the distinctive properties of a particular reduplication to a piece of representation, an underspecified template. For Marantz (1982), this template was a sequence of CV skeletal slots. For later researchers, like McCarthy and Prince (1986), this template consisted of prosodic constituents. No matter how much work could be assigned to these templates and to preferably universal association conventions, there was still the inescapable need for a special operation to copy the segmental melody before it could be reassociated to the template.

In more recent hybrid models, like Spring (1990a) or McCarthy and Prince (1990b), copying is again a basic operation irreducible to other principles of grammar. Indeed, Spring (1990a) argues that templates and restrictions on prosodic bases are independent mechanisms of grammar that can be found in other morphological phenomena, just as they can optionally be found in reduplication. The only defining characteristic of reduplication, Spring argues, is this copy operation.

The problem with all these approaches is the obligatory copy operation, which is like nothing else in phonology. It is an extra mechanism that has had to be added to phonological theory for the sole purpose of dealing with reduplication. Yet the behaviour of reduplication is not quite as unique as we might expect it to be if it had its very own grammatical mechanism. There are several cases that could be analyzed using either the special copy operation of reduplication or using the more mundane operations of autosegmental spreading. The possibility of doing reduplication by autosegmental spreading is by no means limited to copying a single C or V. With the device of separate C and V planes that McCarthy (1989a) argues can exist even within a single morpheme, even phenomena that seem uncontroversially to be reduplications, like the CVV reduplication of Tagalog (e.g., *mag-linis* → *mag-lii-linis*), can be accomplished by autosegmental spreading without crossing association lines. Any decision between the two analyses would be essentially arbitrary. This should cause us to question whether spreading and reduplication are really entirely independent operations after all.

The model of reduplication proposed in this section treats reduplication and vowel harmony as manifestations of exactly the same grammatical principles, albeit on a slightly different scale. Just as vowel harmonies were analyzed using local-domain-creating government arcs with communal lists that contained government types like **site**, **degree**, and **secondary**, reduplications involve communal lists that include inter-segmental government types, like **onset-licence** or **moraic-trochee**.

This section is not intended to be a complete treatment of reduplication. It will almost certainly turn out that the principles discussed here will not be sufficient to handle all cases of reduplication in the world's languages. The theory of locality may even require some quite radical revisions. The purpose of this section is simply to demonstrate the *prima facie* plausibility of the claim that a fully developed constraint-based model will be able to deal with reduplication at least as effectively as a rule- or representation-based model, if not more so.

Section 4.5.1 illustrates the basic application of the principles of generalized locality to reduplication involving prosodically simple duplicates (e.g., single CVs or moraic trochees). Section 4.5.2 demonstrates how the more string-like reduplications that fail to respect the prosodic structure of the original can be handled. Section 4.5.3 deals with the distinction between bases and templates that has been made within Prosodic Morphology and shows how the same effects are achieved in the present framework. Section 4.5.4 discusses a type of reduplication that it would appear constraint-based models are actually better at than either rule- or representation-based models—cases where changes that are made to the reduplicant seem to be “copied back” to the original. Finally, section 4.5.5 deals with an interesting answer suggested by the model developed here to the question of why reduplication should care about prosody in the first place.

For consistency, I shall refer throughout to the two parts of a reduplicated structure as the **original** and the **reduplicant**.

4.5.1 The basic mechanism


We have already seen in section 4.1 the basics of the analysis of reduplication in terms of recursively defined locality.

A simple CV reduplication that does not involve complex onsets can be handled in the same way as a complete vowel harmony, except with the addition of the **onset-licence** government type to the communal list of the local-domain-creating arc.


For example, consider the following CV reduplication from Paamese, an Oceanic language of Vanuatu (Crowley 1982):

(4.87)	sitali	si-sitali	‘emerge’
	mesai	me-mesai	‘sick’
	kaa	ka-kaa	‘fly’
	suai	su-suai	‘disappear’

This is the result of a prefix that is joined to the beginning of the base with a **nuclear-licence** arc that is a local domain creator for **onset-licence** as well as for the gestural features relevant in Paamese. The communal list for the local domain created is {**site**, **degree**, **secondary**, **onset-licence**}. The morpheme for the prefix imposes no other phonological conditions on it. Looking only at the local-domain-creating arc joining the prefix and base, the situation is:



Of course, this is not a legal PS of any language. Two positions joined by a local domain creator must share any atomic children they have and pass on the property of locality to any non-atomic children they have. In order to satisfy the locality constraints, the PS for *si-sitali* would have to look like:




v_1 and v_2 are in a local domain for the communal list {**s**, **d**, **2**, **o**}. Thus they share any atomic children they have. For the sake of illustration I assume that *i* is fully specified as [s:Pal, d:1], so both these specifications are shared in (4.89). Neither nucleus has a secondary articulation, so they vacuously satisfy the locality requirements resulting from the presence of **secondary** in the communal list. They do have onsets, however. In order to satisfy locality requirements, the onsets c_1 and c_2 must inherit the communal list {**s**, **d**, **2**, **o**} from their parents, and must in turn share any atomic children they have, i.e., [s:Ant] and [d:1]. (They have no non-atomic children—and indeed could not possibly have onsets of their own—so the locality requirements for those members of the communal list are satisfied vacuously.)


Reduplication of larger prosodic constituents works the same way, with the addition of prosodic government types to the communal list. Consider the following moraic trochee (CVCV) reduplication, also from Paamese:

- (4.90) hiteali **hite**-hiteali 'laugh'
 hotiini **hoti**-hotiini 'find'
 hulai **hula**-hulai 'spray'
 saani **saa**-saani 'give'

Here, the prefix is attached to the base with a line-1 government (11g) that is a local domain creator for sites, degrees, secondary articulations, and onsets (as we just saw for simple CV reduplication), and the moraic trochee government type as well. Again, the prefix's morpheme imposes no further phonological requirements on it. The basic situation is:



Again, in order for this to be a legal PS, a number of other nodes, arcs, and local domains must also be present. The PS that satisfies all the locality requirements would look like the following. (In order to avoid an illegible tangle of crossing lines, feature specifications or abbreviations thereof have been drawn under each node. It should be understood that the constraints require these to be identical for each pair of nodes in a local domain.) The four local domains involved are indicated by dashed lines:



Just as in the CV case, v_1 and v_3 are in a local domain for the communal list $\{s, d, 2, o, \mu t\}$, as are c_1 and c_3 , having inherited locality from their parents. But the communal list now also includes μt , which means that any children v_1 and v_3 have via moraic trochee arcs must also inherit their locality properties. So v_2 and v_4 also

form a local domain for the communal list $\{s, d, 2, o, \mu t\}$, which they in turn pass on to their onset children, c_2 and c_4 .

Notice that in order to move from CV reduplication to CVCV reduplication, we only had to add one statement—*local-domain-creator*($g, \mu t$)—to the morphemic constraint of the prefix. There was no need for an entirely new type of copy operation to be added to the theory of grammar. So far, reduplication is just like vowel harmony, only more so.

4.5.2 “String” reduplication

Many reduplications where the reduplicant has the form of a moraic trochee cannot be obtained by putting the moraic trochee government type into the communal list of the local domain creator. These are cases where it really does appear that the melody of the original has been copied and then reassociated to a template without regard for prosody. For example, Agta marks plurality with a CVC reduplicative prefix, and the second C will be filled in the reduplicant regardless of its prosodic status in the original (data from Marantz 1982, 439):


(4.93)	bari	‘body’	bar-bari-k kid-in	‘my whole body’
	mag-saddu	‘leak (verb)’	mag-sadsaddu	‘leak in many places’
	na-wakay	‘lost’	na-wakwakay	‘many things lost’
	takki	‘leg’	taktakki	‘legs’

Marantz (1982) analyzes these by copying the entire melody of the original and associating it left to right to the template CVC:

(4.94)	C	V	C			C	V	C	V	C	
	w	a	k	a	y	+	w	a	k	a	y


Marantz uses cases like this as evidence that reduplication does not always copy a prosodic constituent. Simply reduplicating the first moraic trochee of *saddu* and *takki* will give the correct reduplicants: *sad* and *tak*. But this will not work for *bari* and *wakay*. If these words were reduplicated with μt in the communal list, as in (4.92), the results would be the incorrect **bari-bari** and **waka-wakay**.

Fortunately, there is a way to access c_2 in a $c_1.v_1.c_2.v_2$ structure without having to refer to its nucleus, v_2 . In all such structures, v_1 governs c_2 through the composed government relation **n·o**:



We have already seen this composed government relation at work, although in most diagrams so far it has been suppressed for readability. For example, in Hua, vowel features spread via this relation to a following onset that would otherwise be empty, as in /hu+e/ → [huβe]. Since this government relation behaves like other government relations in its ability to be a local domain creator, it is worth exploring the idea that it can also behave like other government relations in being able to belong to communal lists.

This turns out to be just what we need to account for cases like Agta. We can propose that the reduplicant prefix and the original base are linked by a *11g* that is a local domain creator for gestural features, for onset-licences, and for the composed relation *n·o*. That is, the communal list is {*s, d, 2, o, n·o*}. Notice that *μt* is not a member of the communal list. These requirements will give rise to a structure like:



As usual, *v*₁ and *v*₃ form a local domain for the communal list, as do their onset-licensed children, *c*₁ and *c*₃. Because *n·o* is in the communal list, the children of *v*₁ and *v*₃ via this government type, that is, *c*₂ and *c*₄, must also form a local domain and inherit the communal list from their parents. Thus, *c*₂ and *c*₄ share the gestural features for the segment *k*. Note that *v*₂ and *v*₄ do *not* form a local domain. They are not dependent on *v*₁ and *v*₃ through any of the government types in the communal list, so it is impossible for them to form a local domain. In the absence of any segmental requirement on *v*₂ imposed by locality or a morpheme, *v*₂ is null.

4.5.3 Bases and templates


Recent Prosodic Morphology analyses of reduplication (e.g., McCarthy and Prince 1990b, Spring 1990a) make a clear distinction between what we can call prosodic **bases** and prosodic **templates**. The template tells us about the prosodic properties of the reduplicant, but this does not exhaust the information we need to know in order to predict the final reduplicated form. There can also be prosodic conditions on the portion of the original that is visible to the copying operation in the first place, that is, conditions on the base.

Consider the following examples of Tagalog RA reduplication (Carrier-Duncan 1984):

- (4.97) nag+liinis nag+**lii**+liinis
 gupit+in **guu**+gupit+in
 nag+hintay nag+**hii**+hintay

For Marantz, the template would be CVV, which would explain why the vowel of the reduplicant is lengthened if it is not already long in the original. The theory of Prosodic Morphology, however, does not allow templates to refer to the syllabicity of segments. The Prosodic Morphology template for Tagalog RA reduplication would have to be a simple bimoraic syllable, $\sigma_{\mu\mu}$. The problem is that there is no way or restricting the second mora of this template to vowels. For *hintay*, we should expect *hin-hintay* to be possible.

The most straightforward solution within Prosodic Morphology is to restrict the portion of the original that is visible to the copy operation. For Tagalog RA, only a light syllable is visible.²¹ The *n* of *hintay* doesn't associate to the second mora of the template because it was never copied in the first place.




This distinction between template and base can be carried over into the model being developed here. The base, that is, the part of the original that is visible for “copying”, is determined by the contents of the communal list of the local domain creator. If there is no $\mu\mathfrak{t}$ in the communal list, no trochaic dependents of the original will be available to form local domains. In effect, trochaic dependents will not be “copied”. In addition, prosodic conditions may be imposed on the reduplicant affix that are unrelated to how much of the original is available for copying. These extra prosodic conditions on the affix have the effect of a template.

As in the case of Tagalog RA reduplication, the template may require the reduplicant to be prosodically larger than the part of the original that is visible to the reduplicant via the communal list. Tagalog requires the RA prefix to be a moraic trochee, yet it does not include the $\mu\mathfrak{t}$ government in the communal list of the **11g** arc that links the reduplicant and the original, meaning that the reduplicant cannot copy an entire moraic trochee from the original. The basic situation of the affixation is:


²¹In more hardcore Prosodic Morphology terminology, a light syllable is parsed out of the original by a prosodic circumscription operation and mapped onto the heavy syllable template.

(4.99)



In order to make this PS legal, several other things must be present. c_1 and c_3 must be in a local domain. v_2 must be present because of the prosodic template imposed on the prefix, but because μt is not in the communal list, it cannot form a local domain with v_4 . Yet somehow it must be filled. The only possibility being lengthening of the previous vowel:

(4.100)



4.5.4 “Copy-back” reduplication

The foregoing discussion has, I hope, made a convincing case that a framework relying only on the passive satisfaction of constraints (not necessarily this one) will be able to deal with reduplication. In this section, I would like to discuss a phenomenon that strongly suggests that a constraint-based model would not only be equal to, but actually better than, a rule- or a representation-based account. For want of a better term, I shall call this phenomenon “**copy-back**” reduplication.²²

The phenomenon in question is well-illustrated by the classic example from Bloomfield (1933) involving the interaction of reduplication and nasal assimilation in Tagalog. The prefix *pa*

- (4.101) *atip* ‘roofing’
paŋ-atip ‘that used for roofing; shingle’
pu:tul ‘a cut’
pa-mu:tul ‘that used for cutting’

²²Wilbur (1973) and others following her have used the term “rule overapplication”.

When prefixed to a reduplicated stem, however, this nasalization affects both the reduplicant and the original:

- (4.102) *pa-mu-pu:tuł
pa-mu-mu:tuł 'a cutting in quantity'

Other examples include (data from Mester 1986):

- (4.103) pulah pa-mu-mulah 'turning red'
ta:kot pa-na-na:kot 'frightening'
ka?ilaŋan pa-ŋa-ŋa?ilaŋan 'needing'

The problem is this. If *pa* correctly apply to the stem, but then reduplication would then somehow have to reach inside the prefix-stem complex and copy only the stem, whose exact identity has become somewhat problematic. On the other hand, if reduplication applies before *pa* prefixation, then reduplication can be defined straightforwardly, but the result after prefixation and nasal assimilation would be *pa* correct *pa*

have to be copied back to the original. (Hence the term “copy-back” reduplication for these cases.) Theories of reduplication have so far lacked any principled way of performing such a copy-back operation, so analyses of these cases have tended to concentrate on the first alternative, somehow giving reduplication a way of reaching inside the prefixed form.

Another example of what copy-back reduplication looks like may be found in Paamese, a language we have already seen in section 4.5.1. Consider the following:

- (4.104) muni munu-munu 'drink'
luhi luhu-luhu 'plant'
uhi uhu-uhu 'blow'

Crowley (1982: 48) tries to account for this with a rule $i \rightarrow u / u C _$, that will “back final *i* to *u* in disyllables when the preceding syllable has the vowel *u* and the form is reduplicated.”²³ Interestingly, there is another general *i*-backing process in Paamese that applies obligatorily when *i* precedes *u* across a morpheme boundary.²⁴ It would make sense if the backing that occurs in reduplications were a related process. But if this were the case, it would raise serious problems for both rule- and representation-based approaches. Simply applying the backing rule should result in *munu-muni*, *luhu-luhi*, and *uhu-uhi*. But the actually existing forms have the final *i* of the original backed as well. This poses an ordering paradox: the backing in the original could not

²³I know of no mechanism in generative grammar that would be able to restrict the application of this rule as written only to forms that are reduplications.

²⁴In the general process, the two vowels cannot be separated by a consonant.


have taken place before reduplication, because the environment for the backing rule is not yet met, and it could not have taken place after reduplication, because no known mechanism of reduplication has a way for changes made to the reduplicant to be copied back to the original.

Much of the problem posed by the ordering paradox is simply an artefact of assuming that ordering exists in the first place. But a significant part of the problem is the result of the mechanism that existing accounts of reduplication use to explain the similarity of the reduplicant to the original, namely the copy operation. The dependence in these accounts is strictly one-way: the form of the reduplicant depends on the form of the original because it was a copy of the original at some stage of the derivation. There is no way for the form of the original to depend on the form of the reduplicant, even though the most natural analyses of cases like Paamese and Tagalog is to say that changes that are made to the reduplicant are “copied back” to the original.

In a constraint-based approach of the type proposed here, the dependence between reduplicant and original is two-way. The reduplicant must be like the original in the properties implied by the communal list, but equally the original must be like the reduplicant. If the reduplicant has some additional requirements imposed on it, say as the result of a first person prefix, then there are only two alternatives: i) the effects of the extra requirements must be invisible to the original because of the make-up of the communal list, or ii) the effects of the extra requirements must be shared by the original.

Consider the PS for the Kihehe copy-back reduplication *kw-íta-kw-íta* ‘to pour a bit’, where the prefix *ku-*, because it has merged with the initial syllable of the reduplicant, is copied back to the original.²⁵ Just the reduplication, without the prefix *ku-*, would involve the PS:

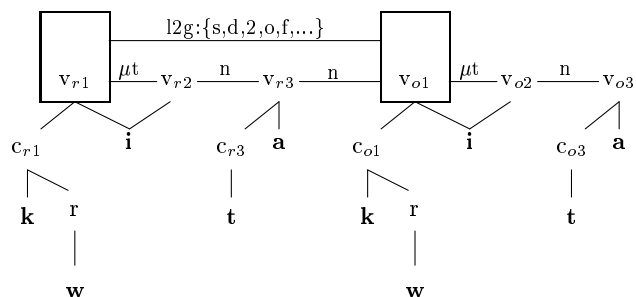
(4.105)



Because onset-licences are crucially part of the communal list, c_{r1} and c_{o1} must necessarily form a local domain. The prefix imposes an extra condition on c_{r1} , namely that it have the gestural specifications for a $k(w)$. Because the two onsets are in a local domain, c_{o1} must also have the gestural specifications for a kw , resulting in the PS:

²⁵Compare what happens with a consonant initial stem: *kú-haata-haáta* ‘to start fermenting’.

(4.106)



Although no entirely satisfactory analysis for cases like this have been offered within rule-based or representation-based approaches to reduplication, it is not inconceivable that some extra devices could be added to the models in order to mechanically generate the correct forms. But it would remain the case that such cases are serious problems for the underlying assumption that the similarity of original and reduplicant is a one-way relation brought about by a copy operation that was introduced into the theory of grammar specifically to deal with reduplication. On the other hand, a constraint-based model where the dependence is two-way not only handles the data easily and naturally, but is the only model that would lead one to predict that cases like Paamese, Tagalog, and Kihehe should exist.

4.5.5 Why does reduplication care about prosody?

Since the work of Shaw (1985), the crucially prosodic nature of reduplication has been recognized, but there has not been a great deal of progress in explaining it. It is entirely conceivable that languages would have been put together in such a way so that they could have built reduplicative templates out of prosodically arbitrary sequences of CVs, similar to the ones initially proposed by Marantz (1982).²⁶ Neither rule-based nor representation-based approaches have offered a convincing explanation for the intimate relation between prosody and reduplication.

In a rule-based framework, the relation is especially mysterious. In a model like that of Steriade (1988), the prosodic properties of reduplication are essentially the result of rule conspiracies. The battery of clean-up rules that applies to reduce the complete copy of the original into its final form in effect conspires to produce a final

²⁶The usual meta-theoretical objection to arbitrary CV templates, namely that they involve the ability of phonology to “count to three”, does not hold as much water as it may first appear. Under most theories, the lexicon is full of prosodically arbitrary strings of segments, which inspire no great fear of the number three. Under at least some of these theories, it is possible for arbitrarily many of these segment positions to be underspecified for their melodic content. If “content” morphemes can look like this, there is little a priori reason why reduplicative morphemes should not be able to as well, even in a theory that dislikes counting to three.

form that has certain prosodic properties. Since rule conspiracies were first discussed by Kisseberth (1970), rule-based frameworks have not been able to account convincingly for their nature and behaviour. Steriade (1988) comes a step closer to dealing with the problem, by having her characterization of a reduplication process explicitly state what the prosodic goal of the rule conspiracy is, but the ineffable link between the goal and the conspiracy remains unexplained and, so far, unenforceable. I see no way a rule-based framework could prohibit a rule battery that did not engage in a prosodically guided conspiracy, but instead performed a number of unrelated and prosodically arbitrary operations on the reduplicant: say, deleting the first onset, lengthening the first nucleus, substituting *oi* for the final nucleus, simplifying the final onset to a single consonant, and deleting everything in between. In rule-based approaches, the prosodic nature of reduplication remains a stipulation.

Representation-based accounts fare better, but still cannot fully explain the relationship. McCarthy and Prince (1986: 6), for example, make the claim:

(4.107) The fact that the templates are bounded by a language's prosody follows from their being literally built from that prosody.

This goes part of the way, but in and of itself it cannot account entirely for the behaviour of reduplication. Specifically, it would still allow reduplicants to indulge in unprosodic behaviour in their parts that were not subject to templatic conditions.

We have a good example of what languages should be able to look like if McCarthy and Prince's claim in (4.107) were the full story about the relation between prosody and reduplication. Yawelmani templatic morphology has been convincingly argued by Archangeli (1991) to involve templates that are literally constructed from prosody, as required by (4.107). In Yawelmani, a prosodic template, say a moraic trochee, can be imposed on a root consisting of the segments *cupn* and the suffix *hin*, resulting in the first part of the word having the form *coo*. After this, according to Archangeli, the rest of the word is syllabified according to the general syllabification principles of the language (complete with such processes as epenthesis), giving *coopunhun*.²⁷

If the prosodicness of reduplication really did follow entirely from the fact that templates are constructed out of prosodic constituents, as McCarthy and Prince claim, then why do we not find reduplications that behave similarly to Yawelmani templatic morphology? If a reduplicative template consisted of an iamb, we might expect it to be possible for the copied melody of the reduplicant to associate to this iamb right-to-left and then have the rest of the reduplicant to the left get syllabified according to the general principles of the language. This does not happen. If a reduplicative template

²⁷The lowering of *u* to *o* in the moraic trochee is a regular process of Yawelmani, as is the harmonic behaviour of the suffix vowel and the epenthetic vowel between *p* and *n*.

In chapter 6, we shall see a similar situation in Moroccan Arabic, where derivational morphemes impose prosodic conditions on the first part of the word and leave the rest of the word to fend for itself, with the distribution of **null** and **constriction** nuclei being determined by the general principles of the language.

consists of an iamb, then the only parts of the copy that survive are those that fit into the iamb. The other segments of the melody, it is usually assumed, are deleted by the process of stray erasure. While bringing in some vague notion of “recoverability” may make it easier to see why stray erasure is *tolerated* in reduplication (but not templatic morphology), it does nothing to explain why it is *required*. Clearly there is more to the question than can be explained by McCarthy and Prince’s claim.


The model of reduplication developed in this section suggests an intriguing possibility for an explanation, which in fact follows as a sort of theorem from what we have already discussed. The reason why reduplication seems so strongly connected to prosody is that putting prosodic government types in the communal list is the only mechanism available for copying more of the original than a simple CV or CVC.

A simple CV reduplication can be accomplished by a communal list that includes nothing more than gestural governments and the onset-licence government. As we saw with Agta, also including the composed government $n\text{-}o$ in the communal list will also give us the effect of a CVC reduplication that does not respect the prosodic constituency of the original. These cases have seemed to support the relation between reduplication and prosody for, although they need not involve anything specifically prosodic in the current framework, they could possibly be analyzed using prosodic templates like σ , σ_c , or $\sigma_{\mu\mu}$. In order to copy any portion of the original larger than this, though, the participation of prosodic government types is absolutely essential.

Consider the possibilities that are open for a framework like the present one for accomplishing a moraic trochaic reduplication, such as *hite-hiteali* diagrammed in (4.92). Somehow the second nuclei of the reduplicant and the original (the two *es*) need to form a local domain with each other. Because there is no local-domain-creating government arc linking them directly, they must inherit their locality from their parents. In order to inherit locality, they must be connected to their parents by a government type that appears in the communal list. Now the only classes of government types that could possibly relate two nuclei (and hence could be used to bequeath locality from one to another) are: i) nuclear-licence or n , ii) metrical line governments, e.g., $10g$, $11g$, and iii) prosodic government types (e.g., μt , σt , i). The first two of these could not possibly be used in a communal list without resulting in incoherent structures.

Let us consider what would happen if the nuclear-licence government type ever appeared in a communal list of a reduplication. The general structure that would be intended for such a reduplication would be something like the following:

(4.108)



The reduplicant’s v_{r1} and the original’s v_{o1} would be connected by some level of met-

rical government arc, say a 12g, that was a local domain creator. By general principles of phonology and morphology demanding an unbroken nuclear spine, the end of the reduplicant, v_{r5} , and the beginning of the original, v_{o1} , would also have to be connected by a nuclear-licence. This will turn out to be important.

At first glance, this might appear to be a perfectly reasonable structure. It certainly gives us the right local domains for the nuclei we are interested in. Since **n** is in the communal list, v_{r2} and v_{o2} form a local domain, inheriting locality from their parents. Likewise v_{r3} and v_{o3} form a local domain, as do v_{r4} and v_{o4} and v_{r5} and v_{o5} . This is all very well so far, but now we run into an unintended consequence of including the nuclear-licence type in the communal list. v_{r5} and v_{o5} are in a local domain for a communal list that includes **n**, hence if one of them has a child via a nuclear-licence, the other must as well, and these two children must in turn form a local domain. v_{r5} does in fact have a child through a nuclear-licence, that is, v_{o1} . There is no way for v_{r5} to satisfy the requirement that it too have a nuclear-licensed child. It cannot double back in the PS, dominating v_{o1} through a nuclear-licence, without violating a number of independent constraints on the well-formedness of PSs. Even if v_{o5} *did* have a nuclear child of its own, this child would now be in a local domain with v_{o1} , would in turn have to have an nuclear child, and the problem would be passed on in infinite regression. No finite PS could possibly satisfy the locality requirements imposed by including a nuclear-licence in the communal list. Similar arguments apply to the metrical lines, 10g, 11g, etc.

The only option that remains for a language that wants to reduplicate more than a CV, CVV or CVC is to include prosodic government types in the communal list. Any reduplication of this scale will involve a prosodic government, and all the prosodic constraints that that involvement implies. This is one of the main reasons why reduplication seems inextricably connected to prosody.

Chapter 5


Constraint-based morphology

5.1 Preliminaries

5.1.1 Sign structure

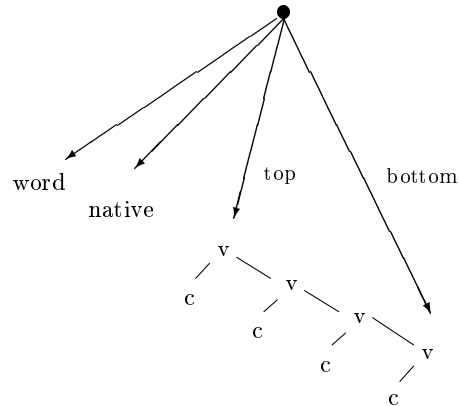
Recall from section 3.1 the tripartite structure it was assumed that signs possess. A sign (a linguistic object at some level) contains phonological, syntactic, and semantic information:

(5.1)



The phonology branch has a structure like:

(5.2)



Each of these three branches has well-formedness conditions of its own (e.g., syllable structure constraints in phonology, binding principles in syntax). As well, there are constraints on well-formed combinations of branches. Many of these constraints deal with things that have often been lumped under the rubric “morphology”.

In section 3.1.1 we considered the possibility that a complete picture of linguistic structure would require more branches than these three, for example a branch for pragmatic information or one for morphological information. While admitting the general possibility that a separate morphological representation (such as that proposed for Autolexical Syntax in Sadock 1990) may turn out to be necessary, in this chapter I would like to explore how much we can do with only these three branches. That is, I shall try to see how much morphology can be done without a separate morphological “component” or “module” of grammar.

It should be noted that this is not as radical claim as it might first appear, or as radical as claims made by some other researchers. Lieber (1992), for example, argues that all word formation can be done by devices already available in the syntax of sentences, and specifically that there should be no grammatical principle that can refer exclusively to things lower than the X^0 , or zero-bar, level of a syntactic tree. This is not the claim that is being made here. It will be quite possible for constraints to refer (even exclusively) to syntactic elements under the zero-bar level just as it would be possible to refer to maximal projections or the X^2 level.

The hypothesis being explored here is that there is no need for an independent level of *representation* devoted entirely to morphological information and to which purely morphological constraints apply. As a corollary (of this hypothesis and of the entire constraint-based enterprise), we will be working with the assumption that there is no hermetically sealed portion of the linguistic system, either in virtual space or in virtual time, that deals only with morphology. This disagrees with many recent pictures of

the lexicon.

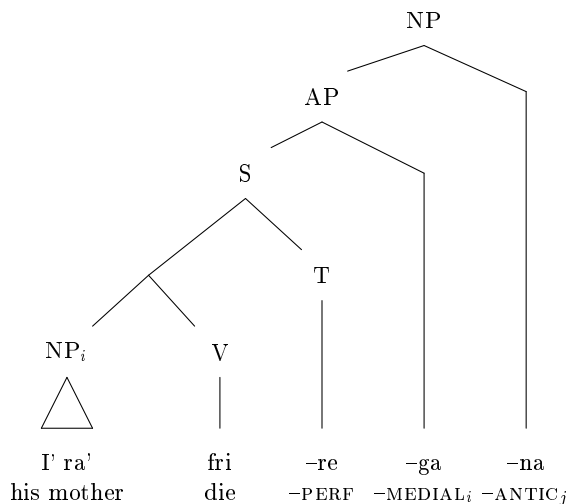
The preceding chapters have given us a reasonable idea of the sort of things that serve as phonological representations. Since morphology is all about the mapping between syntax, semantics, and phonology, we should have some idea of syntactic and semantic representations look like. Unfortunately, I have nothing to say about the nature of semantic representations. This makes it all the more necessary to at least sketch some of my assumptions about the nature of syntactic representation.

The terminal nodes of syntax

Neither syntacticians nor phonologists have always been completely explicit about the entities that are assumed to live on the leaves of syntactic trees. Possibly much of the haziness can be traced back to the tree formalism's early history in (automata-oriented) formal language theory, where the terminal nodes were logically defined to be symbols drawn from a finite alphabet. The abbreviatory devices that have made changes in linguistic theory more palatable to consumers have looked similar enough to these early tree diagrams that the issue of terminal nodes was seldom seen as a problem.


In early transformational grammar, terminal nodes tended to be “morphemes”, in somewhat the conception inherited from the Item-and-Arrangement stream of American structuralism. In an Aspects framework, words could be split up and spread out over several levels of syntactic structure, as in the following diagram of a Hua sentence taken from Haiman (1980):

(5.3)



Following the lexicalist hypothesis advocated in Chomsky (1970), the overt material living on terminal nodes (as opposed to traces and so on) was usually seen somewhat hazily as “words”. While entire theories of morphology and the lexicon were devised to build morphological words, presumably for eventual insertion as terminal nodes in a syntactic tree, it was not always clear exactly what these morphological words were or how they could be distinguished from other suspiciously word-like things, such as clitic groups and phrasal idioms. Even in current GB analyses, the fact that one morphological formative (say a tense marker) usually seems to occur in more or less the same phonological word as some other morphological formative (say a verb stem) is taken as incontrovertible evidence that some sort of movement (like Head Raising) has occurred in order to put the two morphological formatives on the same terminal node. The tree diagram of the Hua sentence would now have to be redrawn as:


(5.4)



This assumption—let us for convenience call it the “one word, one node” assumption—is somewhat stronger than any version of the Lexicalist Hypothesis, and seldom is it clearly articulated as an assumption. Most unification-based approaches to syntax have taken the “one word, one node” assumption to extremes, many of them denying the possibility of any empty categories. In GPSG, for example, if a sentence has three “words”, that sentence’s syntactic tree has exactly three terminal nodes.

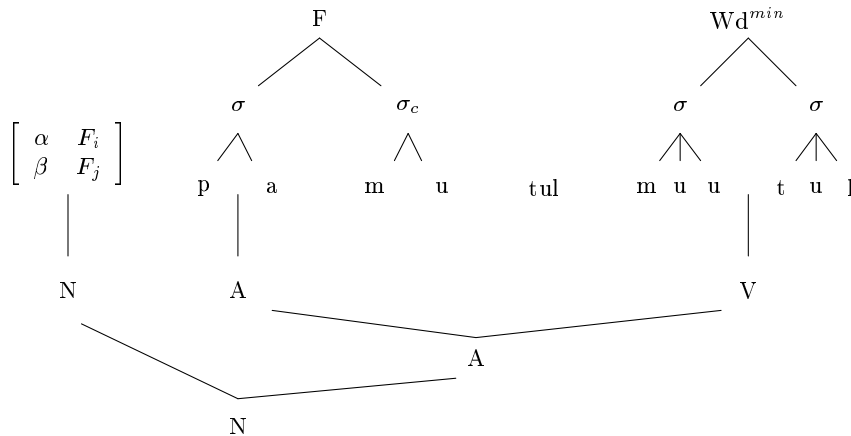
A simple equation of terminal nodes with X^0 s is not likely to find many friends among people who are interested in the internal syntax of words. Generative morphologists have found much of interest going on beneath the X^0 level and have taken to drawing trees with even lower bar levels, like X^{-1} and X^{-2} :

(5.5)




This raises a whole new set of questions about the nature of terminal nodes. What is it, for example, that is attached to the X^{-2} ? Unfortunately, most generative phonologists have simply assumed without comment that these terminal nodes consist of phonological representations of one sort or another. This has led to some curious diagrams, like the following from Lieber (1992:185).

(5.6)




The line joining the (morpho)syntactic node V to what we can only presume is supposed to be its daughter mysteriously terminates somewhere in the empty space between two syllables. Clearly Lieber wants to get across the idea that there are simultaneous morphological and prosodic parsings of the same phonological string. Unfortunately this idea lost its coherence when all the strings in phonology became extinct and were

succeeded by large, intricate autosegmental diagrams. We might rescue the idea by saying that the morphological tree is parsing only a central string-like piece of the phonological representation, such as the skeletal tier. V would now not have a single terminal node, but would be an n -ary branching structure dominating all the relevant skeletal slots:



This was essentially the choice taken by many researchers in the early 1980s (e.g., Stemberger 1981, and Odden and Odden 1985), who treated morpheme nodes essentially as autosegments.

But what then are we to do with templatic morphology systems? Lieber gives the following diagram of an intermediate stage in the derivation of Arabic *katab*, morphologically relevant features notated in matrices but, tellingly, no morphosyntactic tree:



There is no way that morphosyntactic nodes like V^{-2} could hook up to the appropriate skeletal slots, since only one of the three relevant morphemes here contains any skeletal slots at all.¹

¹In fact, there are many appealing aspects to Lieber's idea of morphology and prosody providing parallel bracketings of the same string. I shall return to the idea of parallel bracketing in section 5.4 and propose a way of getting the effects of parallel bracketing that I hope avoids the difficulties faced by Lieber's version. The nuclear spine of morphologically complex PSs will stand in for the "string" that is to be parsed.

The problem here lies in trying to forcibly integrate into a syntactic representation information that is not syntactic but phonological. Syntax and phonology are two different domains of linguistic structure, and their representations should not be grafted together lightly. We have a clearer idea now of what we *cannot* take to be syntactic terminal nodes—morphological words vaguely defined, pieces of phonological representation—but we are not much further towards deciding what they *should* be.

Assumptions about syntactic representation

The kinds of work that a syntactic representation does in a linguistic structure concern things like: head-argument relations, co-indexing, hierarchical relations, scope relations. We will run into fewer problems if we make our syntactic structures as pure and abstract a representation of these properties as possible.

To do this we will clearly need various syntactic features; possibilities include $[\pm N]$, $[\pm V]$, $[\pm \text{anaphoric}]$, $[\pm \text{pronominal}]$, $[\text{singular}]$, $[\text{neuter}]$, $[\text{2nd person}]$, $[\text{perfective}]$ —though the exact inventory, language-specifically or universally, is irrelevant for present purposes. These features can serve as terminals in syntactic structures the same way that phonological atoms like *Pa1*, *Rad*, and *1* serve as terminal nodes in PSs.

We will also need argument structures (however they are to be represented) which may be fully or partially discharged at various levels of the tree.² These argument structures form most of the syntactic content of members of lexical categories.³

It has become a common practice in recent GB analyses to posit all sorts of non-overt abstract operators (e.g., *FOCUS*, *GENERIC*, *WH*) that live in various functional positions like *COMP* or *Determiner*. All functional elements in a syntactic representation can be treated this way. The difference between a preposition and an abstract preposition-like operator assigning case can lie in the syntax-semantics-phonology mapping—the preposition overtly requires a certain stretch of the related PS to be such-and-such a way, the abstract operator behaves like a zero-morpheme (as defined below)—not necessarily in the nature of their terminal nodes.

A complete theory would need some method of representing coindexation, though

²See Pollard and Sag (1987, forthcoming) for specifics on how argument structures are represented and manipulated in HPSG using the same fundamental formalism as the one used here.

³I leave it as an open question whether or not members of lexical categories in syntactic representation need some kind of unique identifier (say, the integer 2092 for *eat*). It is entirely possible that the *V* node for *eat* consists entirely of its categorial features $[\pm V, \pm N]$ and its argument structure, and that it is the semantic branch of the sign (in conjunction with the syntax branch) that determines the phonological form /iit/. If this turns out to be the case, it is a natural way of expressing the intuition that large parts of syntax can be done without any reference at all to the precise identity of the lexical items involved. In frameworks that depend on virtual-time derivations, this intuition is often captured using very late (possibly post-syntactic) lexical insertion (cf. Anderson 1992).

exactly what method is chosen will be irrelevant to the main points of this discussion.⁴ The formalism developed so far suggests different possibilities for structure-sharing. It may even be possible to come up with a generalized representation for structure-sharing that unifies bounding principles in syntax and the adjacency required for local domains in phonology. I leave open the question of whether syntactic chains (cf. Chomsky 1992) will need some sort of independent representation beyond structure sharing in the overall syntactic tree. It should be pointed out that the formalism is rich enough to support several different types of coindexation, such as the often-used distinction between subscripting and superscripting.

In our discussion of phonology, in order to represent asymmetric relations between a head and a dependent we had to choose between using a constituent-based approach (where X governs Y through the mediation of a higher node \bar{X}) or a dependency-based approach (where X directly dominates Y). For phonology, we chose a dependency-based approach. For syntax, I will tentatively assume a constituency-based approach. Most of the reason for this choice is simply the familiarity of the notation. Where a dependency-based approach in phonology resulted in diagrams that more closely resembled those that phonologists have been used to working with,⁵ most work in formal syntax has relied almost exclusively on constituency-based notations. It may or may not be possible to recast syntax into a dependency-based approach and such a move may or may not result in better (or worse) theoretical or empirical coverage, but these questions are beyond the scope of this work. The focus of the present discussion is morphology, seen as the general principles of relating syntactic and phonological structures, and for these purposes any plausible syntactic representation will suffice.

In syntax there is often a need for a head and its projections to have identical specifications for a certain class of features. This is something that comes for free in a dependency-based approach (the head and its projection are exactly the same object — of course they have the same features); in a constituent-based approach, some extra work is required. This extra work is done in generative morphology by some version of “percolation” (e.g., Selkirk 1982, Lieber 1989) and in GPSG and HPSG by the Head Feature Convention. I will not deal with the mechanisms here, except to point out that it is a simple matter to implement percolation or the Head Feature Convention in a manner analogous to local domain creation in phonology. The relations between a head and its projections are local domain creators for head features, and will be required to share those by a principle analogous to the one in phonology. For example, if a head is feminine and singular, and gender and number are head features, we can say that gender and number are in the “communal list” of the head-projection relation, so the head’s projections will also have those features:

⁴Again see Pollard and Sag (1987, forthcoming) for the formal representation of coindexation in HPSG.

⁵Even in areas like syllable and foot structure where constituency-based notations have been common, there have still been strong dependency-based competitors and the resulting diagrams are not too alien.

$$(5.9) \quad \begin{array}{c} \text{YP} \\ \left[\begin{array}{l} \text{gend: fem} \\ \text{num: sg} \end{array} \right] \\ \{ \text{gend, num} \} \\ \text{X}^0 \\ \left[\begin{array}{l} \text{gend: fem} \\ \text{num: sg} \end{array} \right] \end{array}$$

Certain non-head daughters might also be required to be in a local domain with the head's projection for some features. This would be the basic mechanism for head-dependent agreement.

Selectional restrictions of members of lexical and functional categories can be encoded directly as constraints on the properties of these elements' sisters in the syntactic tree.

In summary, I shall be assuming a syntactic representation that consists solely of syntactic objects and contains no pieces of phonological representations, which are kept strictly segregated in the PSs. The kinds of syntactic objects that live in syntactic trees are syntactic features, abstract operators (possibly indistinguishable from features), argument structures, and possibly indices. A typical structure for a case-marked noun phrase might be the following (where in order to simplify the diagram and avoid a tangle of crossing lines, many features are represented as labels on the node rather than dependents):

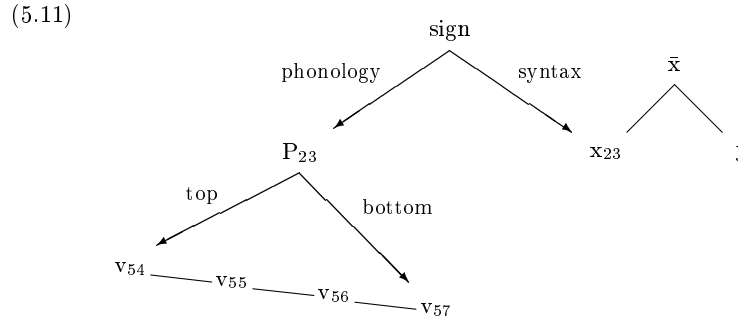
$$(5.10) \quad \begin{array}{c} \text{KP} \\ \left[\text{case:dat, num:sg} \right] \\ \\ \begin{array}{cc} \text{K} & \text{DP} \\ \left[\text{case:dat} \right] & \left[\text{case:dat, num:sg} \right] \\ \\ \begin{array}{cc} \text{D} & \text{NP} \\ \left[\text{num:sg} \right] & \left[\text{case:dat, num:sg} \right] \\ \\ \text{N} \\ \left[\text{case:dat, num:sg} \right] \end{array} \end{array} \end{array}$$

Node labels like D and DP can be seen as abbreviating feature complexes like [cat:Det, bar:0] and [cat:Det, bar:2]. In principle, this structure is independent of the lexical items that the syntax organizes. That is, this same structure could be associated in signs with a number of different semantic structures and phonological structures.

Notational devices

It will be convenient in what follows to be able to refer to the PS associated with a syntactic node by a sign, if any such PS exists, without having to refer explicitly to the intervening structure that joins them. To do this easily, I introduce the description language function $Ph(X)$ whose value is the phonology node associated with the syntactic node X in a sign if such a phonology node exists, and otherwise has no value, or rather has as its value the null object \perp . I further introduce the functions $Ph_t(X)$ and $Ph_b(X)$ whose values are the top and bottom dependents respectively of $Ph(X)$.

Consider the following sign structure:



Here, the value of $Ph(x_{23})$ is p_{23} . $Ph_t(x_{23})$ has the value v_{54} , and $Ph_b(x_{23})$ the value v_{57} . Assuming that there is no sign with y as the tail of its **syntax** node, then the value of $Ph(y)$ is the null object \perp . The diagram in (5.11) also illustrates the fact that the syntactic node pointed to by a sign can, and usually is, part of a larger syntactic tree.


The identity of $Ph(X)$ can be constrained within the description language by:⁶

$$(5.12) \quad \forall x \text{ syntax-node}(x) \rightarrow (\forall s \text{ } s \xrightarrow{\text{syntax}} x \rightarrow (\forall p \text{ } s \xrightarrow{\text{phon}} p \rightarrow Ph(x) = p)) \wedge (\sim \exists s \text{ } s \xrightarrow{\text{syntax}} x \rightarrow Ph(x) = \perp)$$

I shall occasionally speak loosely of the “top” or “bottom” of a syntactic node. This should be understood to mean the top or bottom of the PS that is associated with that syntactic node in a sign, i.e., $Ph_t(X)$ or $Ph_b(X)$.

We sometimes need to express the fact that a certain **n** arc is “within” a certain PS, that is, it is part of the stretch of the nuclear spine that lies between the top and the bottom of the PS. We shall use the predicate *within* to do this.

⁶This definition presupposes that any syntactic node can be the tail of only one syntax arc. Further, we should stipulate that the null object \perp participates in no graph structure; for example, the term $arc(\perp, \text{site}, \text{Pal})$ will always receive a false truth value.



The definition of *within* is straightforward:⁷

$$(5.14) \quad \forall g, Ph \quad within(g, Ph) \leftrightarrow \exists x, y, t, b \\ arc(x, g, y) \wedge Ph \xrightarrow{top} t \wedge Ph \xrightarrow{bottom} b \wedge \\ dominates(t, x) \wedge dominates(y, b)$$

With this, we can now define the *within-word* predicate that was used in the analyses of vowel harmony in chapter 4. *within-word*(*g*) is true if and only if *g* is *within* some phonology node that has the prosodic level of “word” (see section 5.4):

$$(5.15) \quad \forall g \quad within\text{-}word(g) \leftrightarrow \exists Ph \quad within(g, Ph) \wedge Ph \xrightarrow{level} word$$


5.2 General properties of concatenation

Let us first consider the general case of what linguistic structures look like when two morphemes are concatenated. We shall assume that the two syntactic nodes involved in a concatenation are in a sister relationship under a mother node that is a projection of one of the two (the head):


$$(5.16) \quad \begin{array}{c} X^{n+1} \\ / \backslash \\ X^n \quad Y \end{array}$$

The phonological effects of concatenation generally look like the following, where each of two adjacent stretches of a PS satisfies the phonological conditions imposed with respect to one of the syntactic nodes:

⁷ Assuming a recursive version of *dominates*: $dominates(x, y) \leftrightarrow x = y \vee \exists z (x \rightarrow z \wedge dominates(z, y))$.

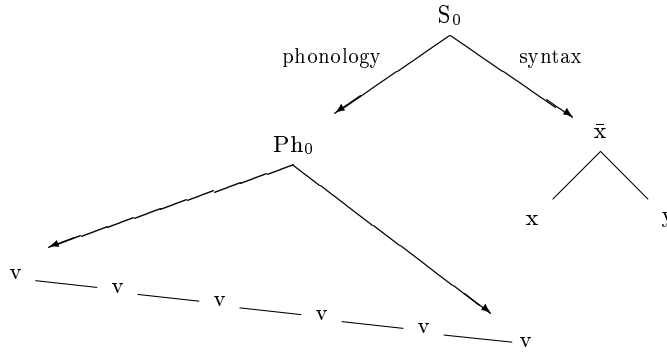


The simplest and most direct way of dealing with the mutual dependence of the semantic, syntactic, and phonological information involved in satisfying a single morpheme is to assume that these three types of information form the branches of a sign. Recall that signs can point to pieces of larger phonological and syntactic structures, so we can have a large number of signs all pointing to the same three structures (phonological, syntactic, and semantic), relating them piece to piece. For the head of (5.16), I propose a sign S_1 of which X is the syntax branch and the left part of the PS is the phonology branch. This is only one of the signs that will be relating the two structures.




Similarly, the Y node in the syntactic tree and the right part the PS are associated in a sign, S_2 . In addition, we will need a sign that encodes the relation between the entire syntactic construction and the entire phonological construction. In (5.19), S_0 forms the sign of which the mother node X is the syntax branch and the entire PS is the phonology branch:

(5.19)



Putting all these sign relations together, we have the following:

(5.20)



This is a rather complicated diagram, so this type of structure will usually be simplified in what follows by suppressing the tree structure of the signs and indicating that a syntactic node and a (sub-)PS are associated in a sign using the $Ph(X)$ notation.

A situation where the first “morpheme” is consonant-final and the second one vowel-initial would be captured by this structure. It can also be appropriate for situations where a vowel-final and a vowel-initial “morpheme” come together and one of the vowels is “lost” or there is a coalescence. The relevant term in the concatenation constraint would look like $Ph_b(x) = Ph_t(y)$.

In general, however, it is not always guaranteed that the edge-most overt segment of your sister is going to be either a consonant or a vowel. While we want the ability to require either that $Ph_b(x) = Ph_t(y)$ or that $Ph_b(x) \xrightarrow{n} Ph_t(y)$, we shall also want to be able to write more general constraints that will be satisfied by either situation. To describe this kind of situation, we shall use a new three-place predicate analogous to *arc*, namely *r-arc*(x, g, y), where x and y are nodes and g is a government arc. We shall abbreviate *r-arc*(x, g, y) as $x \xrightarrow{S_r} y$, where g is of sort S . This predicate behaves just like *arc* except that it is also reflexive. $x \xrightarrow{S_r} y$ is true either if $x \xrightarrow{S} y$ is true (if there is an arc of sort S between x and y) or if x and y are the same node:

$$(5.23) \quad \forall x, y, S \quad x \xrightarrow{S_r} y \leftrightarrow x \xrightarrow{S} y \vee x = y$$

So, if we are given two syntactic sisters, x and y , whose phonological correlates are in a pure concatenative relation with x 's to the left of y 's, there are the following three possibilities for their alignment with respect to each other:

$$(5.24) \quad \begin{array}{l} \text{a) } Ph_b(x) = Ph_t(y) \\ \text{b) } Ph_b(x) \xrightarrow{n} Ph_t(y) \\ \text{c) } Ph_b(x) \xrightarrow{n_r} Ph_t(y) \end{array}$$


Below, we shall formulate the Sister Alignment Principle, which will require for any two syntactic sisters that both have phonological correlates (as opposed to being zero-morphemes), that the top or bottom of the first stand in a \xrightarrow{n} relation with the top or bottom of the second. This is the most permissive situation. Individual morphemes may be more restrictive, requiring a particular pair of sisters to use either the \xrightarrow{n} option or the = option.⁸

In order to fully specify the concatenative behaviour of a morpheme, we also need to say *which* sister it is whose associated PS will be on the left and which will be on the right. In other words, we need to say whether the construction will be head-initial or

⁸In the absence of a more restrictive morphemic requirement, a PS with either $Ph_b(x) = Ph_t(y)$ or $Ph_b(x) \xrightarrow{n} Ph_t(y)$ will satisfy the $\xrightarrow{n_r}$ condition imposed by the Sister Alignment Principle. Of course, one of the two versions may not be able to satisfy other constraints imposed on the forms. In a case where both versions can satisfy the Sister Alignment Principle and other constraints, the general preference principles of the non-monotonic component of the grammar will choose the smaller PS (the one where $x = y$) over the larger.

head-final. The difference lies in whether the two phonological positions being aligned are the bottom of the head and the top of the complement (head-initial) or the bottom of the complement and the top of the head (head-final).

For example, consider a language where tense is marked by a suffix on the verb stem and assume that tense is the syntactic head in a construction under the zero-bar level:



The lexical constraint for a future suffix (say, $-su$) might look like:

(5.26) Hypothetical future suffix $-su$

$$\begin{array}{ll}
 \forall x \text{ semantics}(x) \approx \text{“future”} \rightarrow \exists y, c_1, v_1 & \\
 x \xrightarrow{\text{cat}} \text{tense} \wedge x \xrightarrow{\text{bar}} -1 \wedge x \xrightarrow{\text{tense}} \text{future} \wedge & \text{major category} \\
 \text{complement}(x, y) \wedge y \xrightarrow{\text{cat}} \text{verb} \wedge & \text{selectional restrictions} \\
 Ph_t(x) = Ph_b(x) = v_1 \wedge c_1 \cdot v_1 \wedge & \text{phonology of /su/} \\
 c_1 \approx \mathbf{s} \wedge v_1 \approx \mathbf{u} \wedge & \\
 Ph_b(y) \xrightarrow{n} Ph_t(x) & \text{concatenation: head-final}
 \end{array}$$

For any syntactic node whose semantic correlate is “future”, this morpheme constrains the syntactic category of the node, the category of its syntactic sister (a selectional restriction), what the phonology branch will look like, and how the phonology branch of the sign and that of its syntactic sister will align (in this case, head-finally, that is $Ph_b(y) \xrightarrow{n} Ph_t(x)$).

While it is possible for each morpheme to impose its own requirement for directionality, it is not necessary. It is also possible for more general constraints to determine head-directionality across wide swaths of the morphology. For example, we might want a more general constraint in some language that says that all case markers are suffixes:


$$(5.27) \quad \forall x, y \quad x \xrightarrow{\text{cat}} \text{case} \wedge \text{complement}(x, y) \rightarrow Ph_b(y) \xrightarrow{n_r} Ph_t(x)$$

or that all verbal operators are suffixes:

$$(5.28) \quad \forall x, y \text{ complement}(x, y) \wedge y \xrightarrow{\text{cat}} +V \rightarrow Ph_b(y) \xrightarrow{n} Ph_t(x)$$

Finally it should be noted that the examples we have been looking at so far are the ideal of perfect concatenation. In accounting for these, we have been constraining the location of the top of one sister with respect to the bottom of the other. There is no reason that this should be the only way to do alignment. Formally, there is nothing to stop a language from trying to align, for example, the top of one sister with the *top* of the other. And indeed templatic morphologies provide perfect examples of this.

Consider briefly the Classical Arabic *kutib*, the passive first binyan form of the root *KTB* ‘write’. As usually analyzed since McCarthy (1979), the morpheme for the verb root contributes the sequence of consonants *KTB* and the passive morpheme contributes the sequence of vowels *UI*. (I leave aside the role of the CV or prosodic template itself here.) In this case, it is the tops of the two morphemes that are aligned with each other in order to produce the “concatenation”:



The lexical constraint of the passive might look something like:

(5.30) Classical Arabic passive morpheme (tentative)

$$\begin{array}{ll}
 \forall x \text{ semantics}(x) \approx \text{“passive”} \rightarrow \exists y, v_2 & \\
 x \rightarrow \text{passive} \wedge & \text{category} \\
 \text{complement}(x, y) \wedge y \xrightarrow{\text{cat}} \text{verb} \wedge & \text{selectional restrictions} \\
 Ph_t(x) \approx \mathbf{u} \wedge v_2 \approx \mathbf{i} \wedge Ph_t(x) \xrightarrow{\text{lig}} v_2 & \text{phonology of U-A} \\
 Ph_t(x) \xrightarrow{n} Ph_t(y) & \text{align the sisters' tops}
 \end{array}$$

All of the above considerations lead us to believe that the most general formulation of the Sister Alignment Principle requires only that the top or bottom of one sister


stand in a \xrightarrow{n}_r relation with the top or bottom of its sister. Taking into account the possibility that one of the sisters may act like a zero-morpheme, making sister alignment irrelevant, we can formulate the Sister Alignment Principle as:

(5.31) Sister Alignment Principle


$$\begin{aligned} \forall x, y \text{ sister}(x, y) \wedge Ph(x) \neq \perp \wedge Ph(y) \neq \perp \rightarrow \exists v_x, v_y \\ (x = Ph_t(x) \vee x = Ph_b(x)) \wedge (v_y = Ph_t(y) \vee v_y = Ph_b(y)) \wedge \\ v_x \xrightarrow{n}_r v_y \vee v_y \xrightarrow{n}_r v_x \end{aligned}$$

5.2.2 The Mother's Border Principle

We have now developed a way of making sure that the two PSs corresponding to syntactic sisters are aligned correctly with respect to each other, as in the following diagram:




We still have to deal with the PS corresponding to the mother node. We want the phonology branch of the mother node to encompass all the material of the PSs associated with its daughters, but no more, as in (5.21), repeated here:



Specifically we don't want a situation like:

(5.34)




Ideally, we would like the PS of the mother to be completely predictable from the PSs of the daughters, at least as far as the boundaries (top and bottom) are concerned. The most obvious way of doing this is to constrain the mother's top to coincide with the top of one of its daughters and the mother's bottom to coincide with the bottom of its other daughter. To do this we can use the following constraint:

(5.35) The Mother's Border Principle

$$\forall m, h, c \quad m \xrightarrow{\text{head}} h \wedge m \xrightarrow{\text{comp}} c \rightarrow \\ (Ph_t(m) = Ph_t(h) \vee Ph_t(m) = Ph_t(c)) \wedge \\ (Ph_b(m) = Ph_b(h) \vee Ph_b(m) = Ph_b(c))$$

Note that the constraint as formulated does not do exactly what was promised. That is, it does not require that the mother take its top and bottom from *different* daughters—it is quite possible for the mother to use the same daughter for determining its top and bottom. This would allow structures like:

(5.36)




There is stray phonological material here that belongs to one of the daughters but not to its mother.

We could outlaw the situation in (5.36) by making the Mother's Border Principle more complex so as to require the non-equality of the daughter that supplies the top

and the daughter that supplies the bottom. But this may not be the best approach, especially if we want the Mother's Border Principle to apply to non-concatenative morphology as well. For the structure of Classical Arabic *kutib* in (5.29), a stronger and more complex Mother's Border Principle could still be satisfied, since the mother's top can be supplied by the passive morpheme head and the bottom by the verb root complement, even if the top of the passive morpheme happens to coincide with the top of the verb root. What would not be possible would be a situation where one daughter was entirely contained within the other daughter with no borders coinciding:

(5.37)



This might be the situation in other non-concatenative morphology and for some infixes.⁹ If structures like (5.37) are in fact needed for some morphologies, we would not want the Mother's Border Principle to rule them out.¹⁰

In order to rule out situations like (5.36), we can instead use a more intuitively direct constraint that requires anything that is in a daughter's PS also to be in its mother's PS. For "anything", we can substitute *nuclear-licence* arcs without loss of generality:

(5.38) No Orphan Principle:


$$\forall m, d, g \text{ mother}(m, d) \wedge \text{nuclear-licence}(g) \wedge \text{within}(g, \text{Ph}(d)) \rightarrow \text{within}(g, \text{Ph}(m))$$

We can now rule out the situation where we have orphaned phonological material:

⁹Such infixes are again made possible by the general typology of sister alignment discussed in section 5.2.1. If a pure concatenation allows any one of $=$, \xrightarrow{n} , or \xrightarrow{n}_r between the two edges and templatic morphology allows the two edges to be of the same type (e.g., both tops), we might expect a combination of the two possibilities, for example, having two tops connected by $\xrightarrow{n}: \text{Ph}_t(x) \xrightarrow{n} \text{Ph}_t(y)$.

¹⁰We shall see another example in section 5.3.2, where the simplest analysis of zero morphemes will require the more permissive Mother's Border Principle in (5.35).

(5.39)



This structure would be ruled out because g_5 , which is within the PS of the syntactic daughter d (i.e., it lies on the nuclear spine between $Ph_t(d)$ and $Ph_b(d)$), is not also within the PS of the mother.

5.3 Special cases of “concatenation”


We look now at some types of morphology that have proved troublesome for other frameworks and show how they fit more or less naturally into the theory of sister alignment developed in the previous section. In section 5.3.1, portmanteau morphs and other forms of suppletion are discussed. In section 5.3.2, I look at zero morphemes. Finally, I consider the question of circumfixes, which are the hardest phenomenon to integrate into the current framework. I shall first discuss what changes would have to be made to the concatenation principles proposed in the last section in order to account for circumfixes, and shall then discuss whether we really need circumfixes after all—that is, does the one-to-many mapping relation hold between syntax and phonology (as is usually assumed) or between semantics and syntax?

5.3.1 Suppletion and portmanteau morphs

We have seen that the general principles of morpheme combination do not require strict linear ordering between the PSs of the two morphemes. Templatic morphologies offer examples where the phonologies of two morphemes can be more or less coextensive. Suppletive forms are perhaps the most extreme example.

It is reasonable to assume that the English verb stem *go* will have a constraint saying that its phonology is usually /gow/, except in the past tense, where it is /wɛnt/. Likewise, the past tense morpheme will usually require a coronal stop at the end of the word, except for when combined with certain verb stems—when combined with the stem *go*, the phonology of the past tense is /wɛnt/. The overall structure for this situation might be:

(5.40)



The PSs of the two syntactic sisters are coextensive, and the mother is coextensive with both of them.

This might not always be the best way to handle suppletion. For example, it might be argued that the English irregular past tense would be better handled by having a suppletive allomorph of the verb stem when the verb stands in syntactic construction with a past tense marker, but that the past tense marker has no phonological reflex at all (that is, it would be a zero morpheme in the sense described in section 5.3.2).¹¹ Some cases, though, will probably have a considerably simpler analysis if a structure like (5.40) is posited, especially those involving sequences of affixes (e.g., person and number markers that are often separate but sometimes merged).

5.3.2 Zero morphemes

Zero morphemes have long been a source of contention among phonologists and morphologists, and there is little consensus on how they should be handled, or even if they exist. I shall argue that many of the problems with zero morphemes stem not from the nature of zero morphemes, but from faulty assumptions about the nature of phonology, morphology, and syntax.

Recall the quotation from Hoeksema (1985:18) in chapter 3:


¹¹This argument depends on the prevalent assumption that the best grammar is the one that takes the least ink to write. The grammar that spells out the irregular phonology of only the verb stem while taking the suffix to be a zero morpheme could admit exactly the same set of linguistic structures as the grammar that spells out the irregular phonology for both sisters, i.e., uses the portmanteau analysis of (5.40). If we accept the assumption that grammars should be minimally redundant, then we should clearly prefer the zero morpheme analysis.

For example, the postulation of zero morphemes makes necessary certain arbitrary decisions about e.g. their position in the word: are they prefixes, or suffixes, or perhaps even infixes? Such questions are impossible to answer.

In the present framework, it does not make sense to ask if any, let alone a zero morpheme, *is* a prefix or a suffix, even a morpheme whose phonological behaviour is purely concatenative. Such questions arise from the confusion discussed in section 5.1.1, namely that the terminal nodes of morphosyntactic structure are made up of phonological stuff, or put more bluntly, that morphemes *are* phonological stuff. The result has been an assumption that every terminal node in a morphosyntactic tree—every tense marker, every person or number marker—by the very fact of being a terminal node must consist of some piece of phonological stuff, even if that piece is "zero".

In the present framework, a morpheme is not a piece of phonological representation, nor is it a rule that inserts such a piece at the appropriate point in time. A morpheme is simply a constraint on the relationship between a syntactic structure and a phonological structure, two structures that are otherwise independent of each other. Although it is perhaps the usual situation for a terminal node in a syntactic tree to have a constraint requiring there to be a corresponding swath of phonology (that is, requiring there to be a sign with that node as the syntax branch and a phonology branch that meets certain criteria), this is not necessary. There may be more than one constraint making demands on that node, there may be none at all. It is perfectly possible for there to be a syntactic terminal node which is not the syntax branch of a sign and which has no corresponding PS. This is what is commonly referred to as a "zero morpheme". I shall continue to use this term, although it is not entirely accurate within the present framework: a morpheme is a constraint on the phonology-syntax relationship; a "zero morpheme" is the absence of any such relationship at all.

The overall structure caused by a syntactic node with no phonological correlate (a zero morpheme) is the following:



This might, for instance, illustrate the situation in most of the non-past tense of the English verbal paradigm. Note that the head of the syntactic construction is simply a syntactic node, it does not serve as the syntax branch of any sign. Consequently,

the only two signs involved are those for the complement and those for the mother node. By the Mother's Border Principle, the phonological tops and bottoms of these two signs coincide.

We should take some time to verify that this situation does indeed satisfy the alignment constraints discussed in the last section. First of all, there are no sister alignment constraints to satisfy, as there are no constraints at all imposed on the sign- and PS-associations of the syntactic head. The No Orphan Principle is satisfied for the syntactic complement (every *n* arc in the complement's PS is also in the mother's PS) and is also satisfied vacuously for the head. The Mother's Border Principle of (5.35) is clearly satisfied: the top of the mother coincides with the top of one of its children ($Ph_t(c)$) and the bottom of the mother corresponds to the bottom of one of its children ($Ph_b(c)$)—the same child supplying both top and bottom.¹²

5.3.3 Infixes

I mentioned briefly in section 5.2.2 the interaction of the Mother's Border Principle with infixes. In this section, we take a closer look at what infixes in general look like with respect to their “concatenation”. We shall look at two types of infixation: 1) where the infix splits the initial onset and nucleus of the stem, and 2) where the infix occurs after the initial mora. The examples given will be those discussed by Anderson (1992:206–10).


The first type of infixation is easy to deal with in terms of the alignment principles discussed so far. This is the type familiar from Tagalog ($-um-$ + *bili* → *bumili*) or Chamorro:

- (5.42) *tristi* ‘sad’
 trumisti ‘becomes sad’
 trinisti ‘sadness’

Like the passive morpheme of Classical Arabic, these infixes are simply cases where the PS of the two syntactic sisters are aligned by their tops (rather than by the bottom of one and the top of the other, as in standard concatenation):

¹²Indeed, this is the only possible assignment of top and bottom for the mother in a binary-branching syntactic tree with one of the children being a zero morpheme. Trying to access the phonological top or bottom of a zero-morpheme syntactic node will result in the null object \perp . Since the null object is not equal even to itself, $Ph_t(m) = Ph_t(h)$ cannot be true. Thus the only way for $(Ph_t(m) = Ph_t(h) \vee Ph_t(m) = Ph_t(c))$ to be true is for $Ph_t(m) = Ph_t(c)$ to be true. Similarly, $Ph_b(m) = Ph_b(c)$ will have to be true. Hence the structure in (5.41) is the only way to assign the top and the bottom of the mother node without violating the Mother's Border Principle.

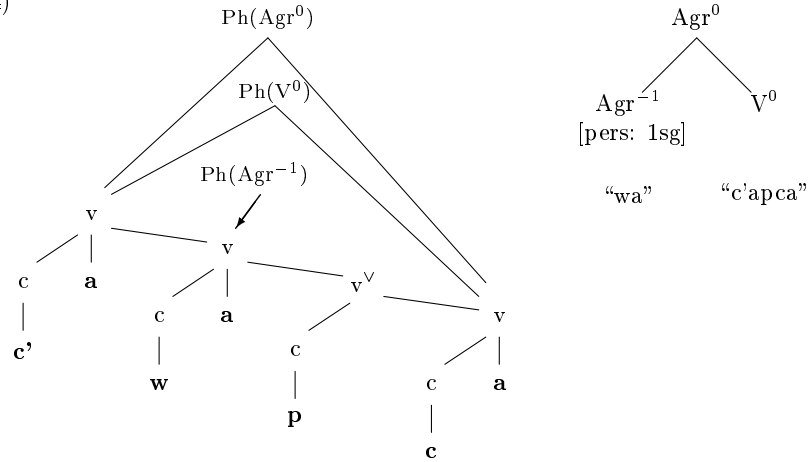
(5.43)



The lexical constraints of the verb stems in these languages will have to be slightly more complicated. The Chamorro lexical constraint for ‘sad’ cannot require that the first onset *tr* be a dependent of the “first” nucleus *i*. While this would work for the simple form *tristi*, it would be violated by the infixed forms (where the nucleus licensing *tr* and the nucleus specified for *i* are not the same). Rather, the lexical constraint will have to use two different existentially quantified variables for the nucleus licensing *tr* and the nucleus specified for *i*, in a manner reminiscent of our treatment of tone in chapter 3. These two variables may often end up referring to the same nucleus, but they need not. Despite the more marked fashion of constraining the gestural content of the verb stems, the alignment of the stem and the infix is straightforward.

In the second kind of infixation, the infix does not split the first consonant and the first vowel, rather it occurs after the first nuclear position (or mora). In Dakota, for example, the first person infix (/wa/, /ma/) occurs after the first vowel but before any coda consonant that a traditional framework would analyze as belonging to the same syllable: *čawapca* ‘I stab’, *?maktomi* ‘I am Iktomi’. This type of infixation is not quite as straightforward, but it is still easily handled within the system of alignment developed in section 5.2. Indeed, the system predicts that cases like this should exist. We have seen that the edges of two syntactic sisters can be aligned according to one of the three relations: =, \xrightarrow{n} , or \xrightarrow{r} . We have also seen that the two edges that enter into the = relation may be either the top of one sister and the bottom of the other, or both tops, or both bottoms. Putting these possibilities together, we should expect to find situations where two tops or two bottoms are related by a relation other than =. The second type of infixation is an example of this:

(5.44)



In Dakota, the *top* of the first person morpheme (*wa* or *ma*) is aligned with the top of the stem (e.g., *ćapca*), and the two are required to stand in the $\overset{n}{\rightarrow}$ relation. In (5.44), it is the case that $Ph_t(\acute{c}apca) \overset{n}{\rightarrow} Ph_t(wa)$.¹³

5.3.4 Circumfixes


Circumfixes involve a one-to-two mapping between branches of a linguistic sign. In this section, we will look at two possible assumptions about which two branches this one-to-two mapping occurs between. First, we shall explore the possibility that it occurs in the syntax-to-phonology mapping, that is, that one node in a syntactic tree corresponds to two different pieces of a PSs. This assumption will prove problematic for the set of hypotheses that we have been developing in this chapter. We shall look at the kinds of changes that would have to be made in order to accommodate the one-syntax-to-two-phonology analysis of circumfixation, and we shall do so at some length because the discussion offers an excellent illustration of how restrictive our current set of hypotheses is. After this, we shall look briefly a second possibility, namely that the one-to-two mapping takes place between the semantics branch and the syntax branch rather than between the syntax and phonology branches.

¹³Infixes that occur at the end of their bases are handled in a similar manner, using the **bottom** pointer instead of the **top** pointer. There also appear exist infixes that are inserted before or after the *stressed* nucleus of their base. These could be handled by doing sister alignment using the primary stress pointer of the phonology node that I tentatively suggested earlier. This would require a slight complication of the Sister Alignment Principle as formulated in (5.31).

Syntax 1, phonology 2

Let us consider the possible ways that one syntactic node could correspond to two discrete chunks of a PS. We will consider two possibilities: one involving a single sign, the other involving two signs.

First, if the circumfix corresponded to a single sign, the general structure that would be involved when it concatenates with its stem would look something like:



In this diagram, B is the piece of the PS corresponding to the stem. A and C are the pieces of the PS that correspond to the prefix and suffix parts of the circumfix respectively. The boundaries of the mother node correspond to those of the circumfix.


(Actually, we do find situations similar to the one diagrammed in (5.45) with infixation. Compare the diagram in (5.44) of the Dakota first person infix construction. From a slightly skewed perspective, we could think of Dakota not as infixing a person marker into a verb root but as circumfixing a verb root around a person marker.)

But there are some problems with this proposal as a general explanation of circumfixation. We need some way of aligning the sisters, that is, the stem and the circumfix. There should be some constraint relating the top or the bottom of the circumfix with the top or the bottom of the stem. But with circumfixation, there is no general relation we could rely on. Clearly we cannot align the top of one with the bottom of the other, as we did with normal concatenation, since the boundaries of the circumfix lie further out than the boundaries of the stem. Suppose we aligned the top of the stem with the top of the circumfix. We have tried to reduce the possible alignment relations to the three: =, \xrightarrow{n} , and $\xrightarrow{n,r}$. If we keep to this hypothesis, then the stem that is top-aligned with a circumfix could be separated from the top of the circumfix by at most one **nuclear-licence** arc — that is, the prefix portion of a circumfix would be universally limited to at most one mora.¹⁴ Even if this does turn out to be a universal limitation


¹⁴Actually, since the bottom of the stem could also be aligned with the bottom of the

on circumfixes, which does not seem likely, or if we weaken the hypothesis by allowing other alignment relations, then the single-sign circumfix proposal still suffers from the fact that it says nothing about the alignment of the *other* half of the circumfix, at the other end of the word.

A second possibility is if the prefix and the suffix portions of the circumfix are the phonology branches of two different signs, both of which share the same syntactic node for their syntax branch. The signs involved in the circumfix would look like:



Abbreviating the signs using bracketed subscripts, as usual, the structure of the whole construction looks like:




There are problems with this proposal too. The largest problem is that a single node is now the syntax branch of two different signs. We have been assuming so far circumfix, the more accurate prediction is that in any circumfix either the prefix portion or the suffix portion is restricted to at most one mora.

that any syntactic node can be the syntax branch of only one sign. Our definition of the functions Ph , Ph_t , and Ph_b depended on this assumption. If X is the syntax branch of two signs, then $Ph(X)$ is no longer uniquely defined. This in turn poses problems for the Mother's Boundary Principle as it was formulated. At the very least, the statement of the Mother's Boundary Principle would have to be considerably complicated in order to allow for the possibility of multiple phonology branches associated with a syntactic node.¹⁵

In sum, there are at least two possibilities for handling circumfixation as a mapping from one syntactic node to two pieces of a phonological representation. Both are formally workable, with some degree of additional complexity, and the first is probably what underlies some cases of circumfixation, but neither is very desirable as a general explanation for all circumfixes.

Semantics 1, syntax 2

The second general way of handling circumfixation is to locate the one-to-two mapping in the relationship between the semantic structure and the syntax. That is, simply because a circumfix has a single, unanalyzable semantic meaning, it does not necessarily follow that it corresponds to a single syntactic node. It may quite possibly correspond to two different syntactic nodes, each of which might in isolation have a semantic value unrelated to that of the circumfix. A single semantic node can be associated with more than one syntactic node:



Put slightly differently, circumfixes can be seen as phrasal idioms (e.g., *kick the bucket*) on a somewhat smaller scale. Just as one "content" atom in the semantics might be constrained to be associated with a quite complex hierarchical structure of syntactic lexical and functional categories, we might expect some less "content"-ful semantic atoms to correspond to complex structures of functional categories.

¹⁵I can see no principled way of restricting the number of signs to two.

In this sense, the relation between semantics and syntax is similar to the relation between syntax and phonology. A single syntactic node is usually constrained to be associated with a quite complex PS consisting of several different nodes organized in a specific hierarchical way. The relation between semantics and syntax allows for the same possibility, although perhaps less spectacularly.

5.3.5 Ablaut and allomorph selection

One possibility that should be pointed out is that a syntactic head, besides imposing phonological conditions on the PS associated with it, can also impose conditions on the PS of a sister that it selects. This can give rise to a couple of different effects that are difficult to give a unified treatment to within a representation-based approach. If the conditions imposed by the head are consistent with the all (perhaps underspecified) allomorphs of the sister, then the effect will be one of ablaut. If on the other hand the conditions are incompatible with some of the allomorphs of the sister, the effect will be one of phonologically conditioned allomorph selection.

There are several examples of the second type of effect (see for example the discussion of Carstairs 1988). A well-known example is the allomorphy of the ergative suffix in the Australian language Dyirbal, where the allomorph $-ŋgu$ is chosen if the PS of the sister consists of a vowel-final disyllable and $-gi$ otherwise (Dixon 1972). In this case, which allomorph of the head is chosen depends on which phonological conditions on the sister are satisfied.

In the first type of effect, the head seems to add phonological material to its sister. For example, many Fore morphemes can induce a high tone on the first or second syllable immediately following the morpheme (Scott 1978). Many cases of umlaut would also be most appropriately in this way. In a representation-based account, this might be captured by proposing that the head's morpheme consists of a contiguous stretch of segmental material, plus some floating sub-segmental material that docks somewhere on the sister.


We will see a good example of this kind of condition imposition on sisters in our discussion of Moroccan Arabic templatic morphology. The primary phonological effect of the passive participle is the addition of the prefix $m-$, but a secondary effect is (pre-theoretically speaking) the imposition of the template CCuUC on the verb root. It would not be clear what course to take if we were forced to choose between addition of structure (ablaut) and selection between structures as the correct method to deal with case; fortunately, a constraint-based approach does not force us to make this choice.

5.4 Prosody and morphology

5.4.1 The prosodic hierarchy

Building on work such as Selkirk (1984), Nespor and Vogel (1986) argued in favour of the idea that phonological representations contained a very rich superstructure of prosodic information. Different types of prosodic constituents were organized into a hierarchy. Each constituent at some level of this hierarchy consisted of (i.e., dominated) one or more constituents of the next lower level of the hierarchy. Nespor and Vogel's proposed levels were: syllable, foot, phonological word, clitic group, phonological phrase, intonational phrase, and phonological utterance:


(5.49)



Nespor and Vogel tried to justify each of these constituent types by finding phonological rules that applied within a constituent of that type but not between constituents.

Nespor and Vogel marked headship using the *s* and *w* (strong and weak) subscripts of metrical phonology. This type of almost diacritical marking was necessary because, for them, prosodic constituents were *n*-ary branching, that is, not restricted to two members. As well, they assumed the Strict Layer Hypothesis of Selkirk (1984), requiring each constituent to be composed entirely of constituents of the immediately lower type on the hierarchy. Skipping levels, as in (5.50), was prohibited:

(5.50)



The further up in the hierarchy one gets, constituent construction is determined more by syntactic and semantic factors and less by phonological ones. For example, Nespov and Vogel propose the following principles for the definition of the phonological phrase, making heavy reference to aspects of syntactic structure:

(5.51) Phonological Phrase Formation (Nespov and Vogel 1986: 168)

- I. *φ domain*: The domain of ϕ consists of a C which contains a lexical head (X) and all C s on its nonrecursive side up to the C that contains another head outside of the maximal projection of X.
- II. *φ construction*: Join into an n -ary branching ϕ all C s included in a string delimited by the definition of the domain of ϕ .
- III. *φ relative phenomenon*: In languages whose syntactic trees are right branching, the rightmost node of ϕ is labelled s [strong]; in languages whose syntactic trees are left branching, the leftmost node of ϕ is labelled s . All sister nodes of s are labelled w [weak].


Many more examples of the relation between syntactic and prosodic structure are provided in the volume edited by Inkelas and Zec (1990).

5.4.2 Representing the prosodic hierarchy

We can integrate higher levels of the prosodic hierarchy into the present framework by treating them as special kinds of PSs, that is, as phonology nodes that are not necessarily the phonology branches of signs. (In this sense they can be like the phonological equivalents of zero-morphemes. Where zero-morphemes were syntactic nodes that could have been pointed to by the syntax branch of a sign but happened not to be and hence had no associated phonology, these phonology nodes can happen not to be pointed to by the phonology branch of a sign and hence have no associated syntax.) In section 3.1.1, it was mentioned that the phonology node was probably the most appropriate place to represent certain kinds of information, among them the prosodic status of the PS, e.g., phonological word or phrase. This idea can be extended to all

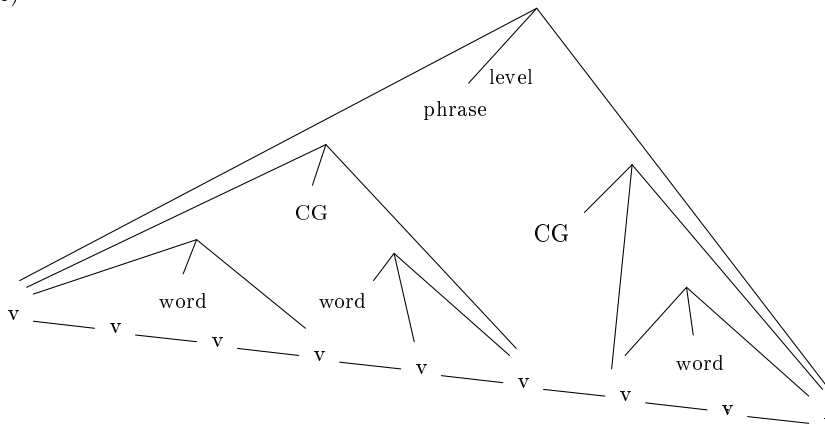
levels of the prosodic hierarchy that we wish to adopt.¹⁶ A prosodic constituent can be represented as a phonology node marked for its prosodic level and pointing to its boundaries, the top and bottom nuclei (and perhaps the primary stress):

(5.52)



Several different levels of prosodic constituents can be represented at the same time in the same way:


(5.53)



In a sense, this is similar to the kind of representation we considered and rejected for metrical structure, where feet and superfeet were independent nodes:

¹⁶There has been some debate about whether all the levels proposed by Nespors and Vogel (1986) are in fact needed. Since the question does not directly affect the general points being made here, I shall not address the question of what levels of the hierarchy are actually needed and shall continue to assume the levels of Nespors and Vogel.

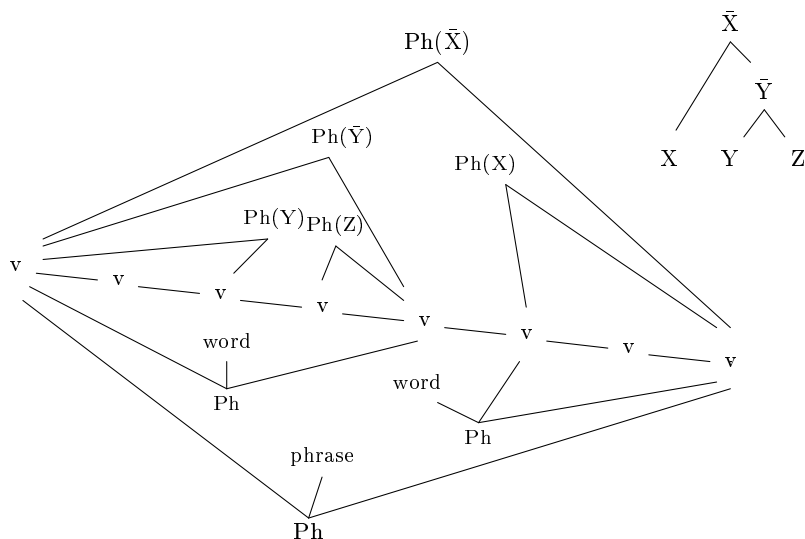
(5.54)



There is an important respect in which (5.53) differs from (5.54). In (5.54), we need new sorts of metrical nodes that are entirely different from other nodes and can dominate other metrical nodes. For example, a superfoot node can dominate a foot node. The result is a large number of non-gestural nodes engaging in complex hierarchical dependency relations among themselves. In (5.53), the nodes are still non-gestural, but they are not a new type introduced simply for the sake of representing constituents of the prosodic hierarchy. Rather, they are **phonology** nodes, which we have already seen at work in signs, and have the usual properties of **phonology** nodes, except for not necessarily being part of a sign. These phonology nodes do not dominate other non-gestural nodes. Instead, like all phonology nodes, they dominate only positions on the nuclear spine via top and bottom arcs. Representing the fact that constituents of one type are made up of constituents of another type, e.g., a phonological phrase is made up of one or more clitic groups, is not done by having the node for the phonological phrase dominate the nodes for the clitic groups. Rather the boundaries of a prosodic constituent are aligned with respect to the boundaries of the next lower constituents on the hierarchy, in much the same way that the Mother's Boundary Principle aligned the phonological boundaries of a syntactic mother node with the phonological boundaries of its syntactic daughters. Nespors and Vogel's diagrams give the impression of an almost syntax-like constituent-structure diagram. The impression of the representations being proposed here is more like a set of nested bracketings of the same string.

In fact, there can be a fairly intuitive correspondence between the collection of the various phonology nodes pointing to the nuclear spine and the various kinds of bracketing used in other work on morphology. In particular, we can see morphological principles and prosodic principles as in effect giving parallel bracketings or parsings of the nuclear spine:


(5.55)



This situation is very reminiscent of the kinds of parallel bracketings proposed by Lieber (e.g., 1992), as illustrated in (5.6). One of the main differences is that the dominating nodes that do the bracketing are not qualitatively different from each other. The “morphological bracketers” are not an altogether different beast from the “prosodic bracketers”; a morphological bracketer is a **phonology** node like the prosodic bracketer, the main difference being that it is also part of a higher-level sign structure that associates it with a piece of syntax. Prosodic bracketers are labelled with prosodic-level information (e.g., clitic group, intonational phrase); morphological bracketers need not be. The visual separation of morphological bracketers and prosodic bracketers on different sides of the nuclear spine in (5.55) may be misleading—there is no clear-cut distinction between morphological and prosodic bracketers. The same phonology node may be both associated in a sign with a syntactic node *and* labelled for prosodic level (for example, in the prototypical cases where a syntactic atom does indeed correspond to a phonological word).

Consideration of (5.54), the rejected possibility for representing metrical structure, brings out a significant difference in the present framework between the treatment of metrical structure and the rest of the prosodic hierarchy. While higher levels of the prosodic hierarchy are treated so that constituency is represented with edge-pointers (and headship possibly with a head-pointer), for metrical structure constituency is represented with inter-node government relations (i.e., **trochee** and **iamb**) and headship using sorts (i.e., **11g**, etc.). Since Nespor and Vogel (1986), other researchers, such as Inkelas (1989), have also come to the conclusion that metrical structure is really


(5.57)




5.4.3 Prosodic edge requirements

We have seen that morphological principles and prosodic principles can result in parallel parsings of the nuclear spine of PS. While the two parsings are in principle independent of each other, it is possible (and quite common) for a morpheme to require its PS to be in certain relationships with constituents of the prosodic hierarchy. This can be done in a couple of ways: i) directly, by requiring the PS associated with a syntactic node to bear a prosodic level specification, as in (5.58), or ii) indirectly, by requiring the top or bottom of a syntactic node's associated PS to also be the top or bottom of a prosodic constituent (which need not be the same **phonology** node as the syntactic node's associated PS), as in (5.59):

(5.58)




(5.59)




The indirect way is frequently used with higher levels of the prosodic hierarchy. Much of the literature in prosodic phonology is devoted to the question of the correct indirect alignment principles relating syntactic constituents of various levels to the boundaries of prosodic constituents. We have already Nespore and Vogel's phonological phrase formation rule. (See also Selkirk 1986 and the papers in Inkelas and Zec 1990.)

Clitics are a special case of these general properties. There is no guarantee that the prosodic structure required by a morphemic constraint (or more general syntax-prosody mapping principles) will coincide with the morphosyntactic boundaries of the PS in the morpheme's sign. Or, if the boundaries do coincide, there is no guarantee that the prosodic levels will be commensurate. For example, if the mother node of a clitic corresponds to a single intonation group, it need not be the case that the clitic and its sister will both correspond to prosodic constituents of the next level down, a phonological phrase, as in (5.60). Instead, the sister may not correspond to any single prosodic constituent, and the clitic may demand to be part of the leftmost constituent of a level quite low in the hierarchy, as in (5.61):

(5.60)



(5.61)




Let us consider a hypothetical example where a determiner cliticizes onto the rightmost phonological word of the NP with which it is associated. For concreteness, let us assume that the phonology of the determiner consists of just one nucleus specified for /ti/, and that the structure of the determiner construction is as proposed in much recent work in GB:

(5.62)

		DP
		\bar{D}
	NP	D
		\bar{N}
CP		N

Recall that the linear order inescapably suggested in the drawing of this syntactic tree is irrelevant, as it is in any syntactic structure. For syntactic structure, only hierarchical relations are important. (For convenience, I have drawn the diagram in a way that reflects the linear ordering that will be imposed by the syntax-phonology mapping.) Let us assume that the D and its sister NP stand in a typical sister alignment relation, that is, $Ph_b(NP) \overset{n}{\rightarrow} Ph_t(D)$. Leaving aside the contribution of prosodic constituents, the phonological and syntactic structure of the construction is so far unexceptional:

(5.63)




A clitic effect will be the result of extra prosodic alignment conditions on the determiner's PS. We can require the bottom of $Ph(D)$ to coincide with the bottom of a phonological word constituent:

$$(5.64) \exists W W \xrightarrow{\text{bottom}} Ph_b(D) \wedge W \xrightarrow{\text{level}} \text{word}$$

This could allow a situation where the top of the determiner's PS is also the top of the phonological word required by (5.64):

(5.65)

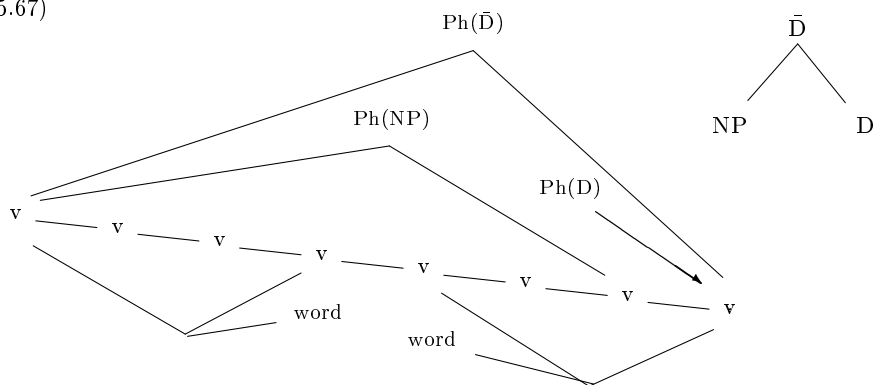


In order to avoid this, we would have to further ban the top of the clitic from being top of a phonological word:

$$(5.66) \sim \exists W W \xrightarrow{\text{top}} Ph_t(D) \wedge W \xrightarrow{\text{level}} \text{word}$$

This would achieve the desired effect of making the clitic part of the final phonological word of the NP:

(5.67)



Another more straightforward way of achieving a clitic effect would be simply to require that the bottom of the NP and the top of the determiner are within the same phonological word, that is, that the *within-word* predicate is true of the arc joining the two positions.¹⁸

(5.68) $\exists gr\text{-}arc(Ph_b(np), g, Ph_t(d)) \wedge nuclear\text{-}licence(g) \wedge within\text{-}word(g)$

Putting these pieces together, the full(er) morpheme for the determiner might look something like:

(5.69) Morpheme for hypothetical clitic determiner

$\forall d \xrightarrow{cat} det \wedge semantics(d) \approx \text{“definite”} \rightarrow$	
$\exists np \text{ complement}(d, np) \wedge np \xrightarrow{cat} NP \wedge$	selection restrictions
$Ph_t(d) = Ph_b(d) \wedge \exists c \text{ } Ph_t(d) \xrightarrow{o} c \wedge$	
$Ph_t(d) \approx \mathbf{i} \wedge c \approx \mathbf{t} \wedge$	segmental phonology
$\exists g \text{ } r\text{-}arc(Ph_b(np), g, Ph_t(d)) \wedge$	
$nuclear\text{-}licence(g) \wedge$	sister alignment
$within\text{-}word(g)$	prosodic “clitic hood” condition

A more compact treatment of clitics could be achieved within a language if prosodic word boundaries were determined by more general principles, in much the same way that phrase boundaries are determined in the accounts of Nespov and Vogel (1986: 168) or Selkirk (1986). For example, a general constraint might be that word boundaries

¹⁸Recall that *r-arc* is the unabbreviated form of \rightarrow_r , as *arc* is the unabbreviated form of \rightarrow , which allows the government arc to be named and referred to.

occur (only) at the left edge of the PSs associated with lexical categories. Then the morphemes for individual clitics would not have to specify where word boundaries do *not* occur, they could simply rely on the general principles to specify where they *do* occur. Of course, languages can still achieve clitic effects the more complicated way, without having such general principles.


A foretaste of Nisgha

As a somewhat more extended example of the prosodic alignment of clitics, let us consider the determiner *-t* in Nisgha.¹⁹ This determiner has an added wrinkle compared to most clitics. It attaches phonologically not to the left edge of the NP that belongs to it, but to the right edge of the preceding phonological word, regardless of what the syntactic category of that word might be. A typical transitive sentence will look schematically like:

$$(5.70) \quad \begin{array}{ccccccc} V & -t & & NP_{Ag} & -t & & NP_{Pa} \\ \left[\quad \right] & & & \left[\quad \right] & & & \left[\quad \right] \\ \left[\quad \right] & & & \left[\quad \right] & & & \left[\quad \right] \end{array} \begin{array}{l} \text{syntactic constituency} \\ \text{phonological constituency} \end{array}$$

Anderson (1992) discusses a similar case of clitic placement in the Wakashan language, *K^wak^w'ala*.

For argument's sake, let us assume an analysis of Nisgha's VSO word order in which the Verb "raises" to Tense, resulting in a syntactic structure like:




Let us work our way up the syntactic tree to see what effects each level has on an associated PS.²⁰ Starting at the object DP, we have the usual morphosyntactic

¹⁹Tarpent (1987) glosses this morpheme as "non-determinate". Its exact semantics need not concern us here.

²⁰This is for the purposes of exposition only and should not be understood as an actual derivational process of the language.


relationship between mothers and daughters, the only difference being that there is an extra prosodic requirement that the nucleus associated with the determiner also be the bottom of a prosodic word. Looking only at this narrow window of the overall sentential structure, we cannot yet see how this constraint will be satisfied.

(5.72)



The next syntactic level up is the V^1 or \bar{V} level. We may safely assume that the co-indexed trace left behind by the verb when it raised acts like a zero-morpheme,²¹ resulting in the following structure.

(5.73)

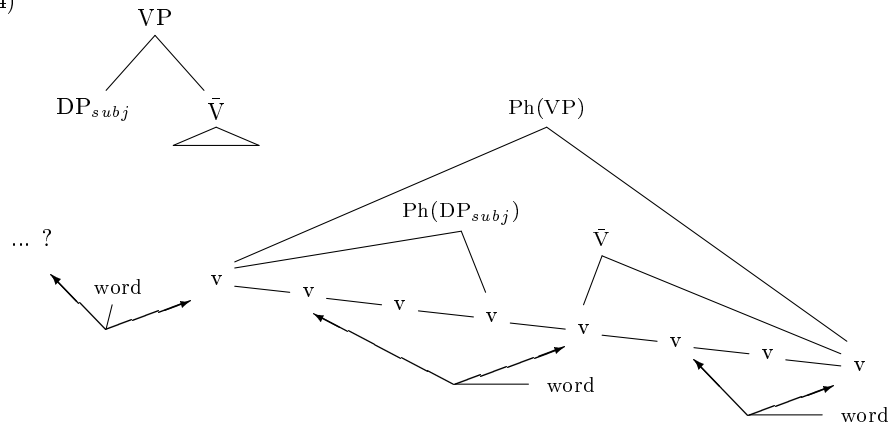


We are still not sure how the prosodic condition on the determiner is going to be satisfied. The subject DP has a morphosyntactic structure exactly like the object

²¹A more monostratal, temporally neutral way of characterizing this trace would be as the V^0 position that is bound by the T^0 higher in the tree (by virtue of the T^0 containing a V^0) and standing in the right structural relationship to allow (or force) such a binding.

DP diagrammed in (5.72). Moving up to the VP level, the \bar{V} is constrained to align with the specifier DP according to the relation $Ph_b(DP_{subj}) \xrightarrow{n} Ph_t(V^1)$, that is, the specifier precedes \bar{V} .²²

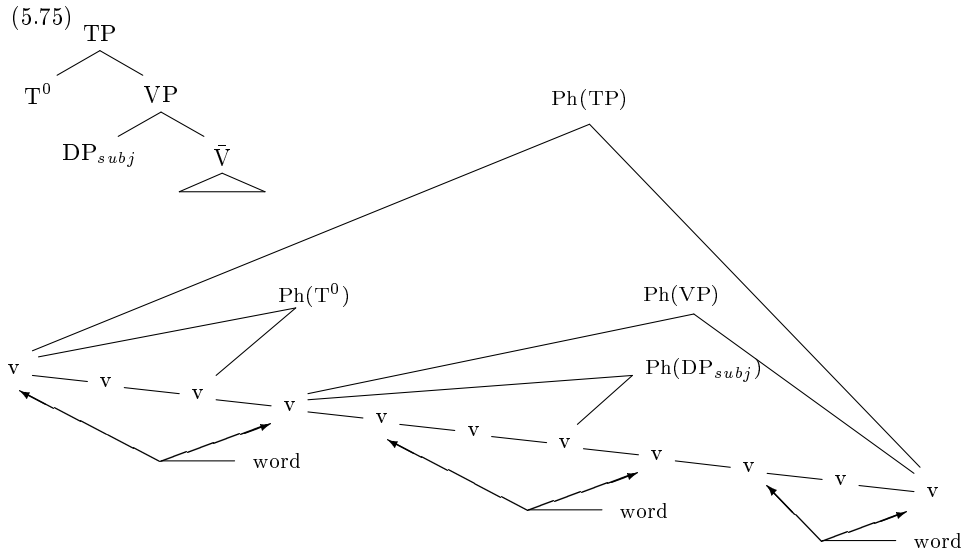
(5.74)



Now we can see how the prosodic requirement of DP_{subj} 's determiner is satisfied: by integrating the determiner into the same phonological word as the preceding noun, although the preceding noun in fact belongs to an altogether different DP. Now it is DP_{subj} 's determiner that has a prosodic requirement that we cannot yet see fulfilled within the window we are looking at.

Moving up to the TP level, and ignoring the internal structure of T^0 and those parts of the internal structure of VP that we have already looked at, we have the following situation:

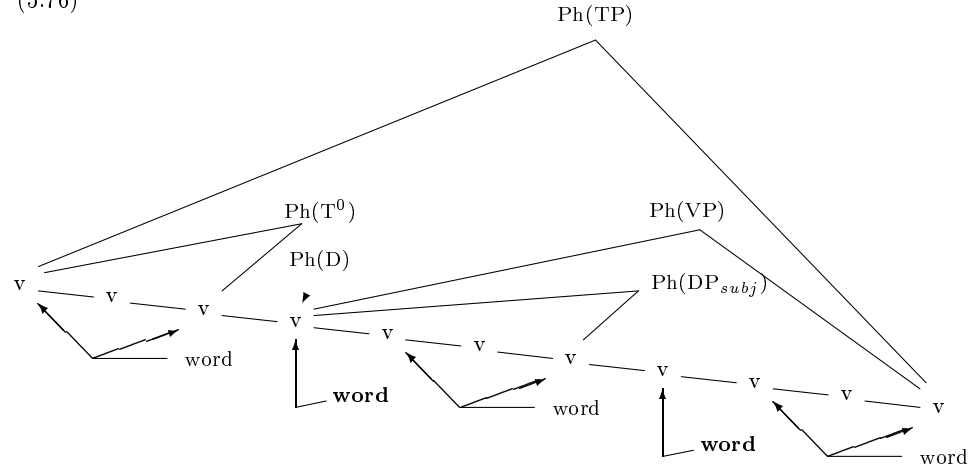
²²Justification for the reflexive aspect of this relation will be given in section 5.6.



Just as DP_{obj} 's determiner was attached to the noun preceding it, DP_{subj} 's determiner is prosodically attached to the phonological word preceding *it*—regardless of the fact that the preceding phonological word otherwise corresponds to a syntactic verb complex.

Before concluding this discussion, we should spend some time considering why the prosodic requirement imposed on the determiner—namely, that it be at the bottom of a phonological word—could not be satisfied by making the determiner into its own phonological word, as illustrated below, rather than merging it with the preceding one.

(5.76)



Why can't the constraint that the bottom of the determiner be the bottom of a prosodic word be satisfied just as easily by a PS where it is the top of the prosodic word as well? The answer is that nothing stops it. (5.76) does satisfy all the constraints. We must now resort to our nonmonotonic principle for selecting the smallest PS from among the candidates that all obey a given set of constraints. Both (5.75) and (5.76) satisfy all the morphosyntactic and prosodic constraints; but of the two, (5.75) is the more compact, containing fewer nodes than (5.76), so, as all else is equal, it will be the PS that is paired with this particular syntactic structure.²³

The Nisgha determiner *-t* is just one example, though perhaps a more spectacular example than most, of how morphosyntactic and prosodic parses of the nuclear spine, even though they are mutually constraining, need not coincide.

5.5 Morphological overdetermination

Over and over in languages, we find situations where a single piece of a PS marks more than one morpheme. We can call this phenomenon **morphological overdetermination**. This is in effect what is going on with portmanteau morphs, but there is another interesting class of cases where two phonologically similar morphs that can occur independently “merge” with each other when adjacent, resulting in a single piece of phonology that is the exponent of both morphemes. Stemberger (1981) has discussed examples of this phenomenon, under the rubric of “morphological haplology”, from

²³Or, at least, the overall sign containing (5.75) rather than (5.76) will be the one that is used for phonetic interpretation.

several different language, including English, Swedish, Spanish, Classical Greek, Sanskrit, French, Mandarin Chinese, and Manchu. As an example of the sort of behaviour we are dealing with, consider Stemberger's example of the Swedish present tense suffix *-r*, which "fails" to appear after stems that already end in an *r*:

- (5.77) *bygg-* 'build' pres. *bygger*
rör- 'move' pres. *rör* (**rörer*)

These situations have often caused problems for theories of phonology and morphology, especially those extreme representation-based approaches that implicitly subscribe to the Physical Integrity of Morphemes assumption discussed in chapter 1. Because the Physical Integrity of Morphemes assumption rules out a priori the most obvious explanation, that the piece of the PS actually does "belong" to more than one morpheme, generative morphology has tended not to view these situations as having anything in common beyond their problematicity for the standard theory. Different kinds of patches were invented for different counterexamples, with varying degrees of success.

It is the purpose of this section and the next to point out some extreme cases of morphological overdetermination where the usual patching techniques cannot work, where one must either complicate accounts in baroque ways or else abandon the Physical Integrity of Morphemes assumption. A constraint-based approach makes the second choice obvious and natural. While it is hard to imagine within a representation-based account how two discrete chunks of phonological representation could be taken out of different storage cells in the lexicon and somehow end up fused together during concatenation, it is perfectly normal for the same piece of a representation to satisfy more than one constraint. We have already seen several examples that exploited the distinction between representations and descriptions that enabled the same node of a PS to "belong" to more than one morpheme, most notably the discussion of tone in chapter 3.

A good example of the most extreme version of morphological overdetermination comes from what are known as cumulative morphs in Athapaskan (Kari 1989).²⁴ In the position class of the Navajo verb that Young and Morgan (1987: 80-82) number VIa, there are more than a dozen morphemes that all have the consonant *d* as their primary phonological reflex. Since they inhabit the same position class, we should expect the morphemes to be mutually exclusive. This is not the case. Instead, it seems that an arbitrary number of them can contribute to the semantics of a single verb, yet in the phonological form of the verb there is only a single *d*.

One response to this state of affairs would be to abandon any attempt at giving these verbs an even remotely compositional morphological analysis. This is essentially the line taken by Kari (1989: 441), who posits a single morpheme *d* with a very abstract meaning. This is not surprising, because the alternatives in either a representation-based or rule-based morphology are not attractive. If four of these *d*-morphemes are

²⁴I thank Jim Kari for bringing this phenomenon to my attention.

at work in a verb, a representation-based analysis would have to concatenate together four different *ds* and then bring in special clean-up rules to delete three of them.

A rule-based account fares no better. It could, like the representation-based account, have four morphemic rules/processes inserting a *d*, and again use a clean-up rule to delete three of them.²⁵ A second alternative for the rule-based account is to claim that the second, third, and fourth insertion rules are somehow “blocked” if they would insert a *d* where one is already present. Note, though, that this notion of “blocking” is very different from the one that is usually at work in morphology. Elsewhere in morphology, if a morphological insertion rule is blocked by a phonological condition, there is no corresponding word at all. What does *not* happen is for the semantic content of the morpheme to get marked by the form without the material that would have been inserted by the rule. For example, *insane* does not meet the prosodic requirements for adding the English comparative suffix *-er*. So there is no *insaner*. But neither is *insane* the comparative form of *insane*. If the insertion rule cannot add *-er*, then there is no morphological comparative at all. But the kind of “blocking” where *insane* could be the comparative form of *insane* is just the kind of blocking that a rule-based account of Navajo would need in order not to avoid adding extra *ds* to a word while still marking the semantics of the morpheme.

A constraint-based approach can handle this kind of morphological overdetermination with ease. If several nodes of a syntactic tree all have morphemes making demands that a certain consonant node of the PS be a *d*, the same *d* can satisfy all the morphemes simultaneously.

This is a phenomenon that both representation- and rule-based approaches to morphology have extreme difficulty dealing with. But not only are these cases easy to deal with in a constraint-based framework, a constraint-based framework is the only one that actually leads us to expect that they should exist.

Not all cases of morphological overdetermination are as clear-cut as these, but it would seem that this is essentially what is going on with a wide range phenomena. Tonal morphemes, other morphemes that seem to consist of “floating autosegments” (e.g., initial mutation in Celtic), morphemes that consist only of an ablaut process, templatic morphemes—these are all cases where more than one morpheme makes a claim on the same piece of a PS. These less dramatic examples all involve somewhat different types of information belonging to the same piece of a PS (e.g., a tone and a melody or a melody and a prosodic constituent), so representation-based approaches have been able to make a fair degree of progress by assuming that the morphemes really involved different pieces of the PS after all. By segregating the types of information referred to by the different morphemes onto different tiers or, if all else failed, onto separate morphemic planes, autosegmental phonology went a long way toward being able to maintain the Physical Integrity of Morphemes assumption discussed in chapter 1. But the Physical Integrity assumption simply cannot be maintained for

²⁵If this solution is less of an embarrassment to the rule-based framework, it is only because it makes less pretense of avoiding complex and arbitrary rule batteries in the first place.

cases like Navajo. The Navajo *d* morphemes cannot be segregated because they refer to exactly the same phonological features (those that make up the *d*). In order even to generate the correct surface forms, a representation-based account needs a whole new grammatical mechanism or else needs to rely on suspicious multiple deletion rules, which still amounts to an implicit claim that Navajo *d* morphemes are qualitatively different from consonant mutations or tonal morphemes. A constraint-based approach lets us see Navajo as differing from all the other overdetermination phenomena only in degree, not in kind.

Another example of morphological overdetermination may be the English plural possessive, commonly represented in the orthography as *-s'*. The plural alone requires a PS that ends in an *-s*. The possessive alone requires a PS that ends in an *-s*. For a word that is both plural and possessive, we might expect two *ss* in a row. Instead, it seems that exactly the same *-s* in the PS satisfies the requirement of both morphemes. We can say that the *-s* is morphologically overdetermined.²⁶

Many deletion rules in the autosegmental literature are necessary only when one tries to force a language into the Physical Integrity of Morphemes mould. They do not involve real disappearances of phonological information at all. All the phonological information still exists in the PS—it is simply not distributed in the way demanded by the Physical Integrity assumption. These spurious deletion rules are not arguments against a monostratal and monotonic approach to grammar, but arguments against the underlying assumptions that forced the deletion analysis in the first place.

In the next section we shall see a dramatic and extremely productive example of morphological overdetermination in Nisgha, where a single segment can satisfy the requirements of a number of separate morphemes with widely differing categories.

5.6 Case study: Nisgha

We have already looked briefly at Nisgha. We saw that determiners cliticized onto the end of the phonological word preceding the noun phrase, regardless of the syntactic category of that word. Nisgha determiner clitics turn out to be even more interesting than this.

Determiner clitics show the kind of morphological overdetermination we have been talking about. Specifically, a single segment in a PS may satisfy the requirements of the determiner clitic, as well as the requirements of one, two, or three other morphemes at the same time. This gives the appearance that some sort of process has merged up to four independent morphemes into a single segment. Tarpent (1987) analyzes the behaviour as involving simple concatenation of all morphs, then deletion of all but one of them.

²⁶See Stemberger (1981) for a more detailed discussion of the merging behaviour of English suffixes with the shape *-s*, including the plural, possessive, third person singular, and adjectival nominalizing (e.g., *linguistic-s*) markers.

As a simple example, consider the following sentences, taken from Tarpent (1987: 866), starting with the simplest version.²⁷

- (5.78) gin'am-i -t loo-y' She gave it to me
 give -CTL-3sg to -me

Here, neither the subject nor the direct object is overt. I shall use boldface for important pieces of a sentence that I am trying to call attention to. In this sentence, the important piece is the third person singular suffix *-t* on the verb, which will mysteriously vanish in many future sentences. The morpheme *-(t)i-* glossed 'CTL' is what Tarpent (1987) refers to as the "control" morpheme, which she does not intend to be understood in the GB sense of PRO-control. It appears on the verb in what Tarpent refers to as "predicate-focussed" clauses and Belvin (1990) refers to as "independent" clauses. It does not appear in "regular" (Tarpent) or "dependent" (Belvin) clauses, which involve a modal or aspectual auxiliary, a subordinator, or a fronted question-word. Its presence seems to mark what in GB terms would be best analyzed as the verb undergoing head-movement into Comp (as opposed to "dependent" sentences, where it would only be able to raise as far as Infl).

Compare (5.78) to the following, where the subject is made overt:

- (5.79) gin'am-i =t hanak' loo-y'
 give -CTL=DET woman to -me
 The woman gave it to me

Here we meet again our old friend, the determiner clitic = with the following noun, but has phonologically cliticized onto the end of the preceding word. However, in the process, it seems to have obliterated the third person suffix *-t* that was there in (5.78).

One conceivable explanation for what is going on would be that it is not phonological at all, but syntactic. Perhaps the 3sg suffix *-t* is really a clitic pronoun that

²⁷Nisgha sentences will be given in a transcription that follows as closely as possible the Nisgha practical orthography. See Tarpent (1987) for a discussion of the relation between the practical orthography and a phonemic transcription. The voicing distinction represented in the orthography is not phonologically relevant. An underlined velar letter represents the corresponding uvular. An apostrophe after a consonant indicates a (pre-)glottalized segment and is elsewhere a glottal stop. A short *i* or *a* often represents a phonological schwa, which in the environment of a uvular or laryngeal is very *a*-like and elsewhere is very *i*-like. The greatest deviation here from the practical orthography will be in the use of the symbol /t/ for the voiceless lateral fricative, represented in the orthography by the digraph *hl*.

All example sentences will be taken either from Tarpent (1987) or from the *Nisgha Phrase Dictionary* (cited as NPD) published by the Bilingual-Bicultural Centre of School District no. 92 (Nisgha).

serves as the argument of the verb, along the lines of Jelinek (1984),²⁸ and is therefore mutually exclusive with an overt noun phrase. While such an analysis could be made to mechanically produce the right facts in just this case, there are severe problems with it. First, this is not the only obliteration going on. This kind of merger is a consistent phenomenon that occurs throughout the entire zone between the last consonant of the verb stem and the determiner clitic of the following noun phrase, affecting a number of morphemes whose only point in common is their phonological shape: a simple (i.e., non-glottalized) coronal consonant. The 3sg suffix *-t* will be obliterated just as surely if you try to suffix it to a verb stem that ends in

clitic =

pronoun for exactly those verbs that happen not to end in *t*, *s*, or

problem with the analysis is that *-t* and an overt noun phrase are *not* mutually exclusive. When the *-t* and the determiner =

the assertive evidential clitic =*a'a*, both survive:

- (5.80) gin'am-i -t =a'a =t hanak' loo-y'
 give -CTL-3sg=assert=DET woman to -me
 The woman *did* give it to me!

A more dramatic example involves several coronal consonants, only one of which appears to survive to the surface. The verb stem *naks-* 'marry' ends in an *s*. Tarpent (1987) argues that "determinate" nouns (i.e., proper names and deictic pronouns) take the determiner =*t* (analogous to the = nouns, e.g., 'woman' above). In addition, she proposes that a "determinate connective" =*s* appears between a verb and a determinate noun phrase when the latter is in the subject position of a transitive clause. (We could also see this =*s* as the ergative case marker for determinate nouns.) For motivation of all these morphemes, and examples of their use in isolation, see Tarpent (1987). When all these are combined with the 3sg suffix *-t*, only one of the *ss* survives. Following Tarpent's practice, I place square brackets around all but one of the segments to show that they are not phonetically realized. Of course, the decision of which *s* it is that surfaces and which is deleted is essentially arbitrary.

- (5.81) ṭa naks -[t] =[s] =[t] Peter → ṭa naks Peter
 nowmarried-3sg=CONN=DET Peter Peter is married now.

More details

I shall now show in somewhat more detail that the merger facts really are as I just sketched them. Readers who are willing to accept my characterization at face value may skip ahead to the analysis in section 5.6.1.

²⁸This possibility is, by the way, rejected in Jelinek's own analysis of Nisgaha (Jelinek 1986).

Let us first review the morphological *dramatis personae*:²⁹

- (5.82)
- | | | |
|----|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -t | 3SG | agreement marker suffixed to the verb. |
| =s | ERG | case marker for determinate nominals (proper names and deictic pronouns) in transitive subject position, cliticizes onto the immediately preceding constituent (i.e., the verb complex). |
| =t | DET | determiner for determinate nominals, cliticizes onto the immediately preceding constituent, whatever its category. |
| = | DET | determiner for non-determinate nominals (everything that isn't a proper name or a deictic pronoun), cliticizes onto the immediately preceding constituent, whatever its category. |

Appearances will also be made by verb stems (the most interesting ones being those that end in *s*, *t*, or being interestingly coronal).

In a full independent transitive sentence with an overt subject, the order of morphemes between the verb stem and the subject nominal will be:

- (5.83) stem - CTL - 3sg = (evidentials) = ERG = DET noun

The third person singular suffix *-t* is part of the following paradigm of person markers (Tarpent 1987: 612):

(5.84)

	SG	PL
1	-y'	-m'
2	-n	-sim'
3	-t	-t/-diit

Besides marking agreement on verbs, these suffixes also mark possession in nouns and form independent person pronouns when suffixed to the base *n'i-* and indirect personal pronouns when suffixed to the base *loo-* (compare the *looy'* of sentences (5.78ff)).³⁰

The determiner *=t* occurs with determinate nominals, i.e., proper names and deictic pronouns. When these occur in the subject position of a transitive clause, the *=t* will

²⁹The affix glossed here as ERG is the one that was named “determinate connective” by Tarpent and glossed as CONN in (5.81).

³⁰*-t* can also often be used to mark third person plural. The other 3pl suffix, *-diit* is peculiar both in its distribution and phonology. Phonologically, for example, it never cooccurs with the “control” or complementizer morpheme discussed in the previous section, nor does it participate in the merger phenomena under consideration. Its behaviour clearly deserves closer examination. See Tarpent (1987: 616–9) for some discussion.

- (5.89) gin'am-i -t =gat =ł hanak' loo-n
 give -CTL-3sg=EVID=DET woman to -you

gin'amitga

I hear the woman gave it to you.

- (5.90) gin'am-i -t =gat =s =t Mary loo-n
 give -CTL-3sg=EVID=ERG=DET to -you

gin'amitgas Mary loon

I hear Mary gave it to you.

Before closing the descriptive part of this section, one wrinkle in the regularity of coalescence should be mentioned. This will receive only a tentative analysis later. The one environment where coalescence fails to occur is before the object of an independent transitive clause (what Tarpent (1987) called a “predicate-focussed clause”, that is, a main clause in which there are no syntactic oddities, such as aspectual auxiliaries or fronted constituents. Compare (5.91), where the overt nominal ‘woman’ acts as the subject of the independent clause and its determiner causes coalescence as expected, with (5.92), where there is no overt subject, ‘woman’ acts as the object and its determiner fails to coalesce with the 3sg *-t* of the verb. (Tarpent 1987: 868)³²

- (5.91) wilaay-i -t =ł hanak' wilaayi
 know CTL-3sg=DET woman The woman knows it/him/her.

- (5.92) wilaay-i -t =ł hanak' wilaayit
 know CTL-3sg=DET woman S/he knows the woman.

Compare this difference in coalescence patterns between subject and object interpretations with the absence of a difference in dependent clauses, for example, clauses introduced by the complementizer *wil*:

- (5.93) wil -t wilaax-t =ł hanak' wilt wilaax
 that-DET know -3sg=DET woman ...as the woman knows it/him/her.

³²There is an additional difference between the two sentences. As in other Nisgha clauses, there is a difference in the level of sentential stress between subjects and objects, with subjects receiving primary stress and objects secondary stress. Tarpent gives the stress patterns of these two sentences as: *wiláayi* ‘S/he knows the woman.’ This strongly suggests that the verb and the subject belong to the same constituent at some level of the prosodic hierarchy, perhaps the phonological or the intonational phrase, while the verb and the object belong to two different constituents.

(5.94) wil -t wilaax-t =t hanak' wilaax
 that-DET know -3sg=DET woman ..as s/he knows the woman.

Aside from the absence of the “control” morpheme *-i-* indicating independent clauses, we notice here identical coalescence patterns are seen in both subject and object readings.³³

5.6.1 Analysis

Intuitively, though somewhat inaccurately, we might look at this behaviour involving the obliteration of coronals as a “competition” between morphemes for the right to be expressed in a single slot. The basic generalization is that in a competition between a morpheme with the shape of *t* and one with the shape of *s*, the *s* wins. In a competition between *t* and

Let us assume that complete closure [d:0] is the default degree feature for Nisgħa consonants, and that fricative [d:1] is the only degree that is phonologically specified. (This is supported by other stop/fricative alternations in Nisgħa, including some involving velars.) I shall not address the exact representation of the lateral fricative instead abbreviating its relevant specifications as [Lat]. The only requirements for the present analysis are that incompatible with that of *s*.³⁵ *t*, *s* and

(5.95) /t/ /s/ /ʃ/
 $\begin{array}{c} \text{c} \\ \text{a} / \end{array}$ $\begin{array}{c} \text{c} \\ \text{a} \wedge \text{d} \end{array}$ $\begin{array}{c} \text{c} \\ \text{a} \wedge \end{array}$
 Lam Lam 1 Lam [Lat]

The basic idea behind a constraint-based analysis of Nisgħa morphological overdetermination is that any morpheme that requires a *t* in its PS will be equally satisfied by a PS

³³The only remaining difference disambiguating between the two readings is the difference in stress pattern mentioned in the previous footnote.

For a brief discussion of the alternation of the verb stem between *wilaax* and *wilaay*, see section on the sorts *consonantal* and *vocalic*.

³⁴One of the reasons why this characterization is inaccurate is that there is clearly no sense in which there is a single “slot” for the competing morphemes. Both morphemes surface intact if they are separated by some other morpheme (e.g., by an evidential post-clitic) or if one of them involves an incompatible non-coronal consonant.


³⁵For this second requirement, I shall simply assume a constraint against the cooccurrence of the [Lat] specifications and [d:1] on the same root node. If have to occur on a secondary articulation.

The most appropriate characterization of the coronal series seems to involve the Laminal articulator. Site is not constant—/s/ especially showing variation between [s] and [ʃ]. (Cf. a former spelling of the name, *Nishga*, which does accurately reflect how the word is often pronounced.)

that contains s . Every gestural feature demanded for the t is present in the s , though there is also an extra [d:1], presumably required by some other morpheme.³⁶ Consider two syntactic nodes, x and y , whose morphemes both end up imposing conditions on the same consonant position.

$$(5.96) \quad \begin{array}{l} \forall x \dots \exists c_7 \text{ Ph}_b(x) \xrightarrow{o} c_7 \wedge c_7 \xrightarrow{a} \text{Lam} \quad \text{i.e., } c_7 \approx \mathbf{t} \\ \forall y \dots \exists c_{13} \text{ Ph}_t(y) \xrightarrow{o} c_{13} \wedge c_{13} \xrightarrow{a} \text{Lam} \wedge c_{13} \xrightarrow{d} 1 \quad \text{i.e., } c_{13} \approx \mathbf{s} \end{array}$$

$$(5.97) \quad \begin{array}{l} \llbracket c_7 \rrbracket = c_1 \\ \llbracket c_{13} \rrbracket = c_1 \end{array}$$



From this point of view, we can see that the apparent deletion of t is really no deletion at all. Both morphemes continue to be “present”, they are simply satisfied by exactly the same node of the PS.


We now need to show how more than one morpheme comes to impose its requirements on the same position in the first place. This is a natural result of the morphological framework we have been developing so far, specifically of the possibility of accomplishing sister alignment with the \xrightarrow{n}_r relation.

Let us assume that in the sentence

an IP, where the verb has raised to Infl. For convenience we shall ignore the multitude of intermediate functional categories that Infl has recently split into in GB literature, as all the ones not discussed here will act as zero morphemes and have no effect on the ultimate shape of the PS. Let us assume that the 3sg suffix $-t$ is an Agr^0 living under Infl, that the “determinate marker” $=t$ for proper names is a determiner that projects


³⁶The extra [d:1] will not appear if no morpheme demands it, by the general principles that will select the smaller of two candidate PSs that both meet all the constraints.

a DP, and that the “determinate connective” =s is an ergative case marker (K) that projects a KP. The phrase *naks Peter* will thus have the following structure with an null direct object:




For convenience, syntactic terminal nodes have been labelled with the phonological “strings” that they correspond to (or at least correspond to in ideal non-deleting environments). This notation should not be given any theoretical status.

Let us first consider the complex structure under I^0 and the PS that it will be associated with. The verb node V will likely be the subject of a lexical constraint like the following (where I ignore the syntactic properties of the root and concentrate on the phonological).



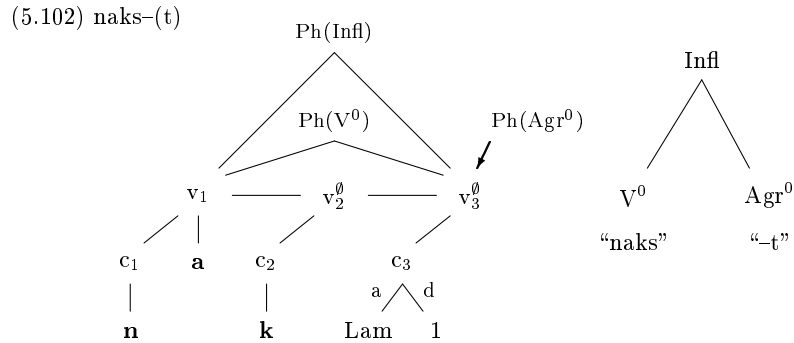
In other words, any structure that involves the verb ‘marry’ will have to look something like:




The Agr^0 will be the subject of a morphemic constraint like:

- (5.101) Morphemic constraint for 3sg $-t$
- | | |
|----------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| $\forall x \ x \xrightarrow{\text{cat}} \text{Agr}^0 \wedge x \xrightarrow{\text{pers}} 3 \wedge x \xrightarrow{\text{num}} \text{sg} \rightarrow$ | |
| $\exists V \ V \xrightarrow{\text{cat}} \text{verb} \wedge \text{complement}(x, V) \wedge$ | selectional restriction |
| $\text{Ph}_t(x) = \text{Ph}_b(x) \wedge \text{null}(\text{Ph}_t(x)) \wedge$ | a single null nucleus |
| $\exists c_{11} \ \text{Ph}_t(x) \xrightarrow{o} c_{11} \wedge c_{11} \xrightarrow{a} \text{Lam}$ | null nucleus' onset is /t/ |
| $\text{Ph}_b(V) \xrightarrow{n} \text{Ph}_t(x)$ | sister alignment |

The sister alignment portion of $-t$'s constraint requires that the bottom of the verb's PS and the top of its PS stand in the potentially reflexive \xrightarrow{n} relation. If possible, the bottom of the verb and the top of $-t$ will be exactly the same nucleus. In *naks* it is possible, resulting in the morphosyntactic structure:



There is a larger PS that also meets all the constraints:



Here, the $Ph_b(V) \xrightarrow{n} Ph_t(x)$ option of $Ph_b(V) \xrightarrow{nr} Ph_t(x)$ has been exercised. But because (5.102) is smaller than (5.103), it is the candidate solution that will be chosen. Other times, it will be impossible for the bottom nucleus of the verb and the top nucleus of $-t$ to coincide. This will be the case if the final onset of the verb is not an anterior coronal but, say, a k . Then there is no way that a PS like (5.102) can satisfy both constraints, because the verb root and the suffix would be making conflicting demands on the segmental content of c_3 . In this case, a PS like (5.103) would be the only candidate solution, and thus would be the correct solution.


Between the case marker $=s$ and the determiner $=t$ we find a similar conflation of consonant positions. The determiner $=t$ would have a morphemic constraint like:

(5.104) Morphemic constraint for determiner $=t$

$$\begin{array}{ll}
 \forall d \ d \xrightarrow{cat} det \ \wedge \ d \rightarrow \text{“determinate”} \rightarrow & \\
 \exists np \ np \xrightarrow{cat} noun \ \wedge \ np \xrightarrow{bar} 2 \ \wedge & \text{selectional restriction} \\
 \text{complement}(d, np) \ \wedge & \text{single null nucleus} \\
 Ph_t(d) = Ph_b(d) \ \wedge \ null(Ph_t(d)) \ \wedge & \text{null nucleus' onset is /t/} \\
 \exists c_{21} \ Ph_t(d) \xrightarrow{o} c_{21} \ \wedge \ c_{21} \xrightarrow{a} Lam \ \wedge & \text{leftward cliticizing} \\
 \exists W \ W \xrightarrow{level} word \ \wedge \ W \xrightarrow{bottom} Ph_t(d) \ \wedge & \text{sister alignment} \\
 Ph_b(d) \xrightarrow{n} Ph_t(np) &
 \end{array}$$

The determiner and its noun phrase are sister aligned with an n arc (absolutely, with no possibility of reflexivity), giving the following structure that in more fortunate circumstances (say in direct object position) would result in t Peter:

(5.105) t Peter




At the KP level, the DP t Peter, gets prefixed with the case marker $=s$, whose morphemic constraint looks like:

(5.106) Morphemic constraint for ergative “connective” =s

$$\begin{aligned}
 \forall k \quad k \xrightarrow{\text{cat}} \text{case} \wedge k \xrightarrow{\text{case}} \text{ergative} &\rightarrow \\
 \exists dp \quad dp \xrightarrow{\text{cat}} \text{det} \wedge dp \xrightarrow{\text{bar}} 2 \wedge & \\
 \quad \text{complement}(k, dp) \wedge & \text{selectional restriction} \\
 Ph_t(k) = Ph_b(k) \wedge \text{null}(Ph_t(k)) \wedge & \text{PS has single null nucleus} \\
 \exists c_{31} \quad Ph_t(k) \xrightarrow{o} c_{31} \wedge & \\
 \quad c_{31} \xrightarrow{a} \text{Lam} \wedge c_{31} \xrightarrow{d} 1 \wedge & \text{null nucleus' onset is s} \\
 \exists W \quad W \xrightarrow{\text{level}} \text{word} \wedge W \xrightarrow{\text{bottom}} Ph_t(k) \wedge & \text{leftward cliticizing} \\
 Ph_b(k) \xrightarrow{n_r} Ph_t(dp) & \text{sister alignment}
 \end{aligned}$$


The sister alignment relation between the case marker and its determiner phrase is the potentially reflexive $\xrightarrow{n_r}$. In this case, the consonantal demands of the case marker =s and the left edge of the DP (i.e., the determiner =t) are compatible, so it is possible for the alignment to use the $Ph_b(k) = Ph_t(dp)$ option, resulting in the following morphosyntactic structure:

(5.107) s Peter



We can assume that the verb trace and the null object in (5.98) act as zero morphemes and have no effect on the associated PS. Therefore, the PS associated with the KP in (5.107) is also the PS associated with the entire VP:

(5.108) s Peter



Finally, we have to align the sisters under IP: I^0 whose phonology is so far *naks* and VP whose phonology is so far *s Peter*. This can be done by a simple directionality constraint that will order any I^0 and VP:³⁷

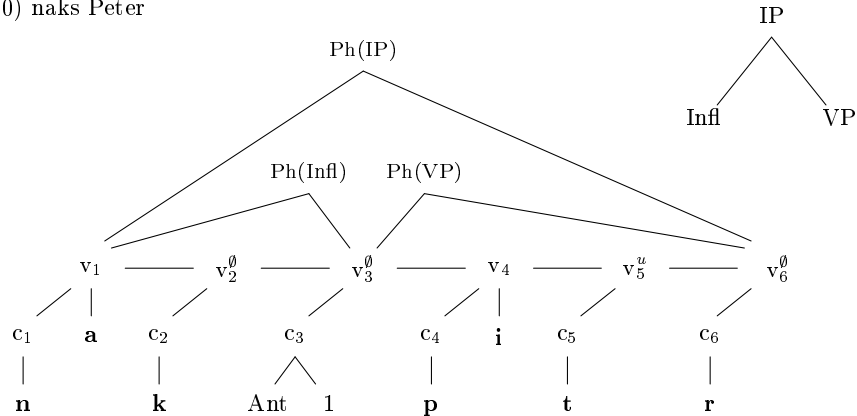
(5.109) Directionality of \bar{I} branching

$$\forall i, vp \quad i \xrightarrow{\text{cat}} \text{infl} \wedge vp \xrightarrow{\text{cat}} \text{verb} \wedge vp \xrightarrow{\text{bar}} 2 \rightarrow \\ Ph_b(i) \xrightarrow{n} Ph_t(vp)$$

Exploiting the $Ph_b(i) = Ph_t(vp)$ option, the bottom of Infl and the top of the VP can coincide. The consonantal requirements imposed by the right edge of Infl and the left edge of VP are compatible, indeed identical, so the same consonant can satisfy the constraints for both syntactic constituents.


³⁷Since Nisgħa is largely right-branching, this may be a special case of a much more general constraint. Individual categories could impose further requirements consistent with this general constraint. For example, we have seen that determiners rule out the reflexive option of \xrightarrow{n} , allowing only \xrightarrow{r} .

(5.110) naks Peter



Note that in the final PS, a single consonant, c_3 , satisfies the requirements of four different morphemes:

(5.111)



This has been accomplished without deleting (or failing to insert) the phonological content of any of the morphemes.

I turn now briefly to the question of why coalescence fails to occur between a plain main clause verb and the determiner of a following object, in the absence of a subject, as discussed at the end of the descriptive section. Intuitively, it would seem that the object is somehow “too far away” from the verb for coalescence to occur. Somehow this intuition will have to be formalized, and whatever the ultimate formalization is will have to indicate why in subordinate clauses, the object is clearly *not* “too far away” from the verb, and coalescence does take place. A good answer would take much more research into the syntax and phonology of Nisgha. In the meantime, I have only the following tentative suggestion for the property that a full formalization might hinge on.

What all the sentences where coalescence fails to occur as expected have in common is the presence of the morpheme $-(t)i-$ that Tarpenet glosses as “control”. As mentioned earlier, this morpheme is present in all main or independent clauses where nothing out of the ordinary is going on. Some of the places it fails to occur in include subordinate clauses introduced by a complementizer, in main clauses that contain an auxiliary verb, and in interrogative main clauses. It is not unreasonable to suppose that the “control” morpheme is in fact a complementizer and lives in the C^0 position of the relevant clauses. Since the verb stem occurs to the left of the control morpheme, it would follow that the verb has undergone long-distance head-movement to C^0 ,³⁸ bypassing Agr^0 on the way up, which remains in its expected base position to the right of the “control” morpheme. In clauses without the “control” morpheme, V^0 would undergo only normal head-to-head movement to Agr^0 and would not end up in C^0 . It is possible that there is some property of the trace of V^0 left behind by normal head-to-head movement versus long-distance head movement that the prosodic principles of the language can refer to and use to constrain the possible boundaries of prosodic constituents. Specifically, the trace of a V^0 that has undergone long-distance head movement might require an edge of a prosodic phrase; this would have the effect of putting the end of the agreement marker and the beginning of the object’s determiner into two separate phrases, thus preventing coalescence between them.

To sum up the discussion of Nisgha, we have seen how up to four different morphemic constraints can be simultaneously be satisfied by the same piece of a PS. If we assumed the Physical Integrity of Morphemes hypothesis, we would have no choice but to admit that all but one of the “morphemes” (that is, all but one of the pieces of phonological structure) had been deleted or else, by some as yet mysterious mechanism, had failed to be inserted in the first place. But when we accept that morphemes are constraints or descriptions that demand that certain properties be true of a PS (without regard to how many other morphemes may be making demands on the same piece of a PS), the behaviour of Nisgha coronal consonants, and phenomena like it in other languages, is easily explainable. Furthermore, while rule- and representation-based approaches to morphology would predict situations like this to be quite rare (a prediction belied by the number of examples found in even Stemberger’s (1981) short overview of mostly Indo-European languages), a constraint-based framework would lead us to expect that such things should happen with some degree of frequency in the languages of the world.

5.7 Summary and implications

We have seen that the principles of morpheme combination are subject to a set of simple constraints, such as the Sister Alignment Principle and the Mother’s Boundary

³⁸The long-distance movement might be forced by selectional properties of the “control” morpheme that demand a V^0 sister. Cf. Ouhalla (1991) on morphological selection.

Principle. These principles leave open a few options about how exactly they will be satisfied (e.g., by a top pointer or a bottom pointer). Each possible choice will result in PSs that exemplify traditional morphological processes like simple concatenation, infixation, templatic morphology, etc. Syntactic trees under the level of X^0 look much like work done in the tradition of Selkirk (1982), though because the syntax-phonology mapping principles only check for correct alignment (and do not have to actively create the correct alignment through operations like concatenation), unlike in Selkirk (1982) there is no need to suppose that templatic morphology is handled by a completely different component of grammar than the principles that deal with prototypical concatenation.

We have also seen how the passive checking nature of constraint-based grammars can explain in a natural way the kinds of morphological overdetermination that procedural frameworks need deletion rules (or some additional mechanism) for. Even spectacular examples of morphological overdetermination, like the merger of post-verbal coronals in Nisgha, follows with perfect regularity from the same principles that are at work in more prototypical cases of concatenation.

The proposals discussed so far have some serious ramifications for various versions of the Lexicalist Hypothesis, the hypothesis that syntax cannot refer to the internal structure of words. If one looks closely at how a Lexicalist Hypothesis could possibly apply in this framework, one will notice an even more disturbing point: there are no words.

The absence of words from the framework is not as great a shortcoming as it might first appear. Researchers have tried for years, without an overwhelming amount of success, to give theoretical content to the pre-theoretical notion of “word”. One of the most detailed discussions of the question is given by di Sciullo and Williams (1987), who distinguish between several concepts that usually coincide with the pre-theoretical notion: phonological wordhood, syntactic atomicity, being listed in the lexicon (because of some unpredictable, non-compositional property). Clearly these several dimensions do not always coincide. Certainly there are large numbers of prototypical words, things that are X^0 nodes syntactically, are prosodic words phonologically, and must be listed in the lexicon. But there are a perhaps even greater number of cases where these disparate properties do not single out the same entities (clitics and phrasal idioms are two of the most obvious exceptions). Much ink has been spilt by taking one of these properties to be the real, fundamental definition of a morphological word and then trying to squeeze into the mould all the other things that have been taken pre-theoretically to be words.

In the current proposal, I do not take any of these properties to be the uniquely defining characteristic of a “word”. Rather, I treat them as independently defined (and independently interesting) aspects of linguistic structure. There are phonological words, that is, a level of the prosodic hierarchy whose constituents are of a size that is usually comparable to pre-theoretically defined words. There are also syntactic constituents of the zero-bar level. These are indeed interesting entities, though I see little to be gained by insisting on referring to them as “words” rather than as X^0

constituents. There are also constraints that “list” the unpredictable properties of some linguistic objects. Again, I see no advantage in trying to single out a subset of these constraints as word constraints, especially since there will be little that formally distinguishes them from other constraints.

So, in a sense, the entire issue the Lexicalist Hypotheses are based on is sidestepped by denying the existence of the entities they try to predicate unique properties of. In fact, the sidestepping is more radical than this. Besides denying the existence of the entities, the framework developed so far also denies the existence of the unique properties that so fascinate proponents of the Lexicalist Hypotheses. Let us consider a couple of representative characterizations of what the content of a Lexicalist Hypothesis is:

Syntactic rules are not allowed to refer to, and hence cannot directly modify, the internal morphological structures of words. (Lapointe 1980: 222)


The content of [the Lexicalist Hypothesis] is that transformations should only be permitted to operate on syntactic constituents and to insert or delete named items (like prepositions). This means that they can't be used to insert, delete, permute or substitute parts of words. (Spencer 1991: 72-3)

The things that a Lexicalist Hypothesis wants to prevent from happening to a “word” are things that a monostratal and monotonic approach to language does not allow to happen to *anything*, word or not. The inability to be inserted, deleted, permuted, or substituted does not set apart any special class of linguistic objects at all, let alone one that corresponds to the pre-theoretical notion of word.

It should be noted that this rejection of the Lexicalist Hypothesis does not imply a free-for-all, where syntactic constituents of the zero-bar level or lower can appear anywhere in a syntactic tree they choose. Nor does it involve the strong claim of Lieber (1992) that no syntactic principle can refer exclusively to constituents under the zero-bar level. I see no problem with a constraint predicating interesting properties of the X^{-1} level, any more than there should be problems predicating interesting things of the X^2 level or of nouns. It is simply a claim that all the effects that a Lexicalist Hypothesis is interesting in capturing can be captured just as effectively without having to posit and define a special type of linguistic object corresponding to the pre-theoretical notion of word.

Let us consider a more concrete example. The generalizations that various versions of the Lexicalist Hypothesis are intended to explain involve questions like: why can't you topicalize a tense suffix by moving it to the front of a sentence? The way that this question is phrased presupposes that the best way to handle topicalization is as a literally real movement of part of the sentence. But this is not the only way to deal with the long distance dependency between the topicalized constituent and the gap in the rest of the sentence that it is logically connected with. Even within GB, there is a trend toward seeing “movement” simply as a metaphor for understanding the

connection between a topicalized constituent and its trace and for describing the kinds of S-structure or LF binding relations that are and are not legal. We can extend this treatment to the suffix topicalization case, and allow or disallow these structures using only constraints on representations. A topicalized suffix structure might look like:



Leaving aside purely phonological constraints that might prevent this syntactic tree from being part of a legal sign (for example, a requirement that the PS corresponding to T^{-1} be in a dependent position in a phonological clitic group whose primary stress is to its left), there are representational syntactic constraints that might rule out (5.112) without relying on constraints on some kind of movement transformation. For example, it is widely accepted that any constituent that is in the specifier of a maximal projection (X^2) must itself be a maximal projection— T^{-1} is not. T^{-1} may be unable to bind its trace underneath T^0 ; perhaps T^0 acts in this language as a barrier in a sense related to that used by Chomsky (1986). (Compare Lieber's (1992) discussion of binding under the X^0 level.) Or perhaps the selectional restrictions of T^{-1} (e.g., it must have a V^0 complement) are not satisfied properly in (5.112).

For the present discussion, it is irrelevant which one or more of these constraints prohibits a structure like (5.112). The main point is that it is very likely that the generalization against topicalizing suffixes (and generalizations like it) can be expressed using only constraints on representations.³⁹ These types of facts do not force us to accept movement transformations as a literal fact, or to formulate conditions on the operation of such a transformation, or to define a linguistic entity corresponding to the pre-theoretical notion of word to be used in the definition of movement constraints.

³⁹Baker's (1988) analysis of incorporation and similar phenomena is a good example of the amount of work that can be done by constraints on representations (e.g., possible binding configurations) and of the kind of constraints that will be needed.

The theoretical work of the Lexicalist Hypotheses⁴⁰ can be done using the mechanisms and entities we already have.

⁴⁰as opposed to the sociological work of allowing one to reject Generative Semantics out of hand

272

9

Chapter 6

Templatic morphology: Moroccan Arabic

In this chapter, I look in more detail at one example of a templatic morphological system, that of Moroccan Arabic (MA). I shall show how the properties of MA fit into the general theory of morpheme combination developed in the last chapter. The analysis also provides an extended example of how the system of `null` nodes proposed in section 3.3.1 works in a real language. Moroccan Arabic is particularly interesting for a discussion of templatic morphology. Short vowels of Classical Arabic have historically been deleted in MA, leaving behind complex patterns of consonant clusters. Very few previous analyses have managed to give a convincing account even of the syllabification behaviour of consonant clusters, and none that I am aware of have integrated such an account with an analysis of the templatic morphology system.

6.1 Moroccan Arabic

This section outlines some of the phenomena of Moroccan Arabic that will form the basis of the rest of the chapter.

6.1.1 Segments

The following table shows the basic consonant inventory of MA (cf. Harris 1942, Harrell 1962, Heath 1987).

(6.1)		labial	dental	palato alveolar	velar	uvular	pharyngeal	laryngeal
	stop		t ṭ d ḏ		k g	q		ʔ
	fricative	f	s ṣ z ẓ	š ž		x ǧ	ħ ʕ	h
	nasal	m	n					
	lateral		l ḷ					
	flap-trill		r ṛ					
	semivowel	w		y				

There are two entries for almost every line in the dental¹ column. The sounds marked with the underdot are usually referred to as **emphatic** or **pharyngealized**, the rest as **plain**. An emphatic consonant's pharyngealization is for the most part realized on the neighbouring vowel. As Heath points out, *tī* is usually heard [te]. The exact set of emphatic consonants is a matter of some disagreement. Harrell, for instance, includes emphatic labials. Since the question is not central to the following discussion, I shall simply assume that some subset of MA consonants can be specified [2:[s:Pha, d:R]]. This specification also covers the uvulars and, of course, the pharyngeals.

There are three **full** or, in Harrell's terminology, **stable**, vowels in MA: *a*, *i*, and *u*. Most of these are descended from CA long vowels. Near plain consonants, these are realized as [æ], [i], and [u]. Near emphatic consonants, the vowels are pharyngealized too and are realized as [ɑ], [e], and [o]. See Chapter 2 on the representation of MA low vowels.

In addition there are some **reduced**, **short**, or **unstable** vowels. The most common is the schwa, which for typographical convenience I shall follow Harrell in transcribing *e*. Kaye (1990) argues that the MA schwa is Government Phonology's cold vowel: high, back, unrounded, lax [ɨ]. This is not an inaccurate characterization, though the realization of *e* is highly variable, depending largely on the nature of the neighbouring consonants, hence Harrell's description of it as unstable. Near pharyngealized consonants, *e* is usually quite pharyngeal. Harrell transcribes it as *ǣ* in the environment of pharyngeals and uvulars. Harris transcribes his schwa phoneme as [ʷ] alternations as the full vowels *i* and *u*.

There are neighbourhoods of some MA words that are affected by labialization. This can be realized as a secondary articulation or release on a velar or uvular consonant [k^w, g^w, q^w, x^w, ǧ^w] or by a labialized schwa, transcribed *ũ* by Heath. Compare the verb *tqũb* 'puncture' with its simple verbal noun *t^wqib* and its instantiating verbal noun *t^weqba* (Heath 1987:81). Heath has long discussions about the underlying repre-

¹Heath (1987) labels this series alveolar. Harrell describes it as apico-dental. I do not think it inaccurate to characterize this series as anterior laminal, that is, in the system of chapter 2, [a:Lam, s:Ant].

sentation of this labialization. I shall not discuss the problem here, partly because it is peripheral to the concerns of this chapter and partly because the “problems” Heath finds are to a large extent a by-product of proceduralist assumptions.

6.1.2 Clusters, epenthesis, and syncope

Clusters of two consonants are possible in all positions of the MA word.² Consider the initial clusters in the following perfective verbs:

- (6.2) kteb ‘he wrote’
 šreb ‘he drank’
 mdeǧ ‘he chewed’
 lbes ‘he wore, he put on’
 qbel ‘he accepted’

While it is conceivable that *šr* and perhaps even *kt* and *qb* constitute branching onsets, this is not possible with *mǧ* and *lb*. Syllables in other languages consistently obey what has been called the **Sonority Sequencing Constraint** or Generalization (cf. Steriade 1982, Selkirk 1984): segments decrease in sonority the further they are from the nuclear peak of the syllable. This means no consonant in a branching onset can have a greater sonority than the consonant that follows it.³ If *mǧ* and *lb* were onsets in MA, this would require abandoning a proposed universal principle that has served well in language after language.

It is also possible to have a cluster of three consonants if the first two are geminate, e.g., *qeddmet* ‘she presented’, *t-ṭbib* ‘the doctor’, where the definite article prefix *l-* has assimilated to the first coronal of the stem. This largely occurs in morphologically restricted environments, since most gemination in MA arises from morphological processes.⁴ Even more so than with two-consonant clusters, most of these three-consonant clusters cannot be seen as onsets, even when word-initial. This would require MA, alone among the world’s languages, to allow geminates in its onsets—several phonological models on the market today do not even have a way of representing such a situation.

²although word-final clusters are morphologically restricted to, e.g., some kinds of nouns, first person singular verb forms.

³Languages usually also demand a certain sonority “distance” between elements of a branching onset. Thus, from *to* to *is* is a large enough increase in sonority for *to* to count as a legal onset in Greek, but not in English.

⁴The *l-* of *l-ṭbib*, for example, is the result of the templatic constraints imposed by the causative morpheme, to be discussed later.

In other environments, cliticization can create three-consonant sequences without geminates. For example, one type of negation of a first person singular perfective verb adds *l-* to a verb already ending in a cluster because of the first person singular suffix *-ni*: *l-ḥad* ‘I found’, *l-ḥadni* ‘I didn’t find.’

So the first challenge for a model of MA phonology is to provide a syllabic structure for such clusters that avoids the problems discussed above. There should be independent evidence for any proposal — saving the Sonority Sequencing Constraint should not be its only motivation.

One possible clue, and another fact to be explained, is that these clusters are not always clusters. The examples in (6.2) are only one inflected form of the verb, the third person singular masculine perfective. Each verb has other forms in which the initial consonants in (6.2) are broken up by schwas or full vowels. For example, the third person *plural* perfective, besides adding the suffix *-u*, re-positions the schwa in each word:

(6.3)	Third person perfective		
	<i> masc. singular</i>	<i> plural</i>	
	kteb	ketbu	‘write’
	šreb	šerbu	‘drink’
	mḍeġ	medġu	‘chew’
	lbes	lebsu	‘wear’
	qbel	qeblu	‘accept’

This kind of alternation is not restricted to verbs. We also find it in the difference between plain nouns and those with possessive suffixes:

(6.4)	<i>ktef</i>	‘shoulder’	<i>ketfi</i>	‘my shoulder’
	<i>šġūr</i>	‘childhood’	<i>šġūri</i>	‘my childhood’

It is reasonable to conclude that this is a general process of the language and not a peculiarity of particular morphological form.

The basic generalization to be had is that (with the exception of certain morphological environments) there can be at most one consonant at the end of a word and at most two in a row medially. It would seem as though we could propose that schwas do not exist underlyingly and are inserted by a rule like:

$$(6.5) \quad CC \rightarrow CeC / _ \left\{ \begin{array}{c} C \\ \# \end{array} \right\}$$

However, not all occurrences of schwa are predictable by a rule like this. Nouns can show both the CCEC pattern predicted by the rule, as in *sqef* ‘roof’, and a CeCC pattern, as in *bent* ‘girl, daughter’. There are some nouns that can occur in both patterns in different dialects, e.g., ‘gravy, sauce’ which can be either *mṛeq* or *meṛq*. While speakers may have a preference for one of these variants, they accept the other as “still Moroccan”. This strongly suggests that the positioning of the schwa cannot be reduced to a predictable effect of some property of the surrounding consonants, but must be lexically specified. Something more than a simple epenthesis rule is needed in order to account for the distribution of reduced vowels in MA.

6.1.3 Templatic morphology

The kind of templatic morphology that operates in MA can be illustrated by the following sets of words (Harrell 1962:23).

- (6.6)
- | | |
|--------|------------------|
| kteb | ‘he wrote’ |
| iketbu | ‘they write’ |
| kateb | ‘having written’ |
| mektub | ‘written’ |
| ktab | ‘book’ |
| ktub | ‘books’ |
| mketba | ‘writing desk’ |
- (6.7)
- | | |
|--------|----------------------------|
| qtel | ‘he killed’ |
| qtal | ‘carnage, slaughter’ |
| qtil | ‘(action of) killing’ |
| qtila | ‘murder, assassination’ |
| qettal | ‘killer, murderer; deadly’ |

Semantically, each of the words in (6.6) deals in some way with the concept of writing. Phonologically, they all have in common the occurrence of the consonants *k*, *t*, and *b*, in that order. Similarly, all the words in (6.7) involve the concept of killing and contain the consonants *q*–*t*–*l*. This characteristic sequence of consonants for each verb is usually called the verb **root**.

There is a traditional typology that divides MA verb roots along two major axes. The first distinction is between **triliteral** and **quadriliteral** roots, that is, whether the root (as shown in a “neutral” form like the third person singular masculine perfective) contains three or four full segments (consonants or vowels). The second is between **strong**, **hollow**, **weak**, and **irregular** roots.

In a **strong** root, each of the three or four full segments of the root is a consonant. The *ktb* of (6.6) and *qtl* of (6.7) are strong trilaterals. The third person singular masculine perfective (3sm.pf) verb forms derived from these roots are *kteb* and *qtel*.⁵ A root may contain a full vowel instead of a consonant in any position but the first. A **weak** root has a vowel in place of a final consonant, e.g., *šra* ‘buy’, *šeqša* ‘inquire’. A **hollow** root has a vowel in place of the second consonant of a triliteral, e.g., *bas* ‘kiss’, or the second or third consonant of a quadriliteral, e.g., *ayen* ‘wait’, *ħmar* ‘redden’. Note that the presence of the full vowel in the inflected form of a weak or hollow root affects the positioning of reduced vowels; this will be discussed below.

⁵For the rest of this section, roots will be cited in their 3sm.pf forms.

Closely related to pure strong trilaterals are **geminate** or **doubled** trilaterals, where the second and third consonants are identical. This results in a different position for the reduced vowel in their citation forms, e.g., ‘he smelt’, ‘he poured’.

Although the vowel of a weak or hollow root will generally appear as a in the 3sm.pf form, in some other inflected forms the vowel can be different. Which full vowel is used is consistent across the forms of any one stem, but lexically idiosyncratic. For example, the weak trilaterals *šri* ‘buy’, *ḥba* ‘crawl’, and *bda* ‘begin’ have the imperatives *šri*, *ḥbu*, and *bda*, and the first person singular imperfectives *ne-šri*, *ne-ḥbu*, and *ne-bda*.

There are also a small number of **irregular** roots that do not fit into any of these categories, e.g., *ža* ‘come’, or fail to act like members of the category they appear to belong to, e.g., *xda* ‘take’ has *xud* as its imperative, where we would expect *xdu*, *xdi*, or *xda* if it were a trilateral weak root.

Templates are usually described informally using a notation that mixes upper and lower case letters, as in the trilateral template, *meFuL*. The lower case letters appear in the positions indicated in the resulting form. The upper-case consonants *F*, *ḥ*, and *L* are replaced by the three consonants of the root.⁶ Thus, *ktb* ‘write’ in the *meFuL* template form would be *mektub*.

Some of the characteristic templates of various morphological forms will now be examined. The forms that will be dealt with in the analysis section of this chapter are the base forms (e.g., 3sm.pf), the causative, the reciprocal, and the active and passive participles.

Base forms

The essential features of the base forms have been outlined above in the discussion of the various root types. In traditional grammars, this form is called measure 1, or binyan 1. The following table summarizes the 3sm.pf forms of the various trilateral root types in this measure.

(6.8)	root type	template	3sm.pf
	strong	FḥeL	<i>kteb</i>
	geminate	Feḥ _i L _i	<i>ḥell</i>
	hollow	FVL	<i>ba</i>
	weak	FḥV	<i>šra</i>

When the verb carries inflectional person-number suffixes, schwas are repositioned according to the principles sketched in 6.1.2 and analyzed below in 6.3. Also recall that some inflected forms of hollow and weak measure 1 verbs use the verb’s characteristic vowel rather than *a*. The first person singular imperfective of *ba*, for example, is *n-bi*.

⁶The “variables”, *ḥ*, *ḥ*, and *L*, are usually chosen because *ktb* was a verb root meaning ‘do’ in Classical Arabic. Classical Arabic speakers showed great foresight and consideration in selecting a ‘do’ verb root whose First and Last consonants would be mnemonic for English-speaking linguists.

Causatives

These forms belong to the class traditionally called measure 2, or binyan 2. A triliteral root that is an intransitive verb or an adjective in its base form will be a transitive verb with a causative or factitive meaning in measure 2. This is a highly productive method of deriving causatives from triliteral roots in MA.⁷

The characteristic mark of this measure is gemination of the middle consonant of the root:

(6.9)		3sm.pf	template	causative
	strong	<i>wqef</i>	FeʕʕeL	<i>weqqef</i>
	geminate	<i>šemm</i>	Feʕiʕi;eLi	<i>šemmem</i>
	hollow	<i>faq</i>	FeGGeL	<i>feyyeq</i>
	weak	<i>bka</i>	FeʕʕV	<i>bekki</i>

For hollow verbs, which have no middle consonant, the geminate is of the glide corresponding to the root's characteristic vowel (i.e., the vowel that appears in the imperative, imperfective, etc.). *i* corresponds to *y* and *u* to *w*. If the characteristic vowel is *a*, the causative will have *y* or *w* unpredictably.⁸

(6.10)	gloss	base 3ms.pf	base imperative	causative
	'wake up'	<i>faq</i>	<i>fiq</i>	<i>feyyeq</i>
	'get up'	<i>naḍ</i>	<i>nuḍ</i>	<i>newweḍ</i>
	'fear'	<i>xaf</i>	<i>xaf</i>	<i>xewwef</i>
	'appear'	<i>ban</i>	<i>ban</i>	<i>beyyen</i>

Reciprocals

These forms belong to the class traditionally known as measure 3, or binyan 3. The most obvious characteristic of this measure is the presence of a full *a* between the first and second consonants of the root.

While there are some words that instantiate the bare template *FaeL*, they tend to be semantically idiosyncratic and the pattern by itself is not productive. The most common use of the *FaeL* template is in conjunction with the prefix *t-* to form a reciprocal verb.

(6.11)	gloss	3sm.pf base	reciprocal
	strong	'kill' <i>qtel</i>	<i>t-qatel</i>
	geminate	'smell' <i>šemm</i>	<i>t-šamm</i>
	hollow	'fear' <i>xaf</i>	<i>t-xawef</i>
	weak	'buy' <i>šra</i>	<i>t-šara</i>

⁷A quadriliteral root can have no measure 2 form. Their causatives must be formed by syntactic periphrasis.

⁸Many varieties of MA will fully vocalize the schwa before a geminate glide, e.g., instead of *weqqef*. See Harrell (1962:30).

As with the causatives, the reciprocals of hollow verbs fill in the ζ position of the template with the glide corresponding to their characteristic vowel. *ba* ‘sell’ (imperative *bi*) has the reciprocal *t-baye*, *bas* ‘kiss’ (imperative *bus*) has *t-bawes*, and *xaf* ‘fear’ (imperative *xaf*, causative *xewwef*) has *xawef*.

Active participle

Like the reciprocal, the templatic pattern for the active participle of measure 1 verbs is *FaeL*, with no prefix.

(6.12)	3sm.pf	participle	gloss
	<i>k</i> teb	<i>kateb</i>	‘having written’
	<i>ḥ</i> ell	<i>ḥall</i>	‘having opened’
	<i>b</i> a	<i>baye</i>	‘having sold’
	<i>š</i> ra	<i>šari</i>	‘having bought’

Passive participle

The templatic pattern for the passive participle is *meFuL*.

(6.13)	3sm.pf	participle	gloss
	<i>k</i> teb	<i>mektub</i>	‘(having been) written’
	<i>ḥ</i> ell	<i>mehlul</i>	‘(having been) opened’
	<i>b</i> a	<i>mebyu</i>	‘(having been) sold’
	<i>š</i> ra	<i>mešri</i>	‘(having been) bought’

The ζ position for hollow verbs is filled by a glide in the same way described for other patterns. Note that the position of the schwa is the prefix *me-* is predictable from the principles of section 6.1.2.

Some other derivational templatic patterns exist. These tend, on the whole, to be less productive and more lexically idiosyncratic. The patterns given above are a representative sample of the types of alternations that are found in MA verbs and verb-derived words. Analyzing them will allow us to see the essential scaffolding of the MA templatic morphology system.

6.2 Government Phonology analyses of MA


The work of Government Phonologists on Moroccan Arabic gives one of the most interesting accounts yet of the schwa/zero alternations that characterize the language. Three articles deal with some depth with MA, applying the concepts of GP at various stages of its development: Kaye, Echchadli, and El Ayachi (1986), Kaye (1990a), and Kaye (1990b). Kaye (1990b) is the most consistent with the principles now generally accepted in Government Phonology.

A favourite slogan of GP is “Languages do not wear their syllable structure on their sleeves.” It is an error to assume that exhaustively dividing up a string of surface segments into the least offensive groupings will yield the syllables that are phonologically relevant. The problem is especially acute in MA, where naively grouping the segments of the verb *ibes* ‘wear, put on’ into a single syllable yields a typologically grotesque constituent that seriously compromises several universal principles of syllabification that have been shown to hold in almost every other language.


GP argues that the syllables that are phonologically relevant in MA are not immediately obvious from the apparent surface segmental string. Much of the work leading up to the present theory of GP pointed to the usefulness of having an **appendix** at the end of a word where a consonant could be prosodically licensed without being syllabified into the preceding syllable (Charette 1984, Kaye and Lowenstamm 1984). It was soon realized that this appendix need not be a novel type of prosodic constituent, rather it could be a normal syllable whose nucleus happened to be empty. This idea has become more widely accepted. McCarthy and Prince (1990a), for example, propose a syllable with no nucleus (σ^-) at both the right and left edges of the verb’s prosodic template. MA provides the strongest evidence yet that these kinds of syllables with empty nuclei can also occur in the middle of words as well as at their edges.

Kaye proposed that in every consonant cluster that does not obey the sonority sequencing generalizations,⁹ the consonants are not in fact adjacent but are “separated” by an empty nuclear position. Instead of trying to squeeze *ibes* into a single syllable:

(6.14)



GP proposes that each consonant is followed by a nuclear position, two of which are phonetically unrealized:





⁹Actually, the generalizations do not involve sonority per se, but a slightly different GP implementation of the concept involving the abstract phonological property of **charm**. See Kaye, Lowenstamm, and Vergnaud (1990b) for discussion of how the charm values of consonants affect syllabification.

Although GP does not recognize the existence of any higher node σ , pre-theoretically we can say that the GP representation of *ibes* contains three syllables.

The central explanatory idea GP brings to schwa/zero alternation in MA is that schwas and empty nuclei are fundamentally the same thing in slightly different environments. An empty nucleus is not in fact empty, rather it is filled by an entity that GP refers to as v^0 or the **cold vowel**. In some contexts the cold vowel is phonetically realized, in some it is not. The cold vowel has none of the four active or “hot” features ([–back], [–high], [+round], [+ATR]), so when it is phonetically realized, the result is a back high unrounded lax vowel, usually transcribed [ɨ].

Whether or not a cold vowel is phonetically realized depends on another property of phonological representations known as **proper government**. A phonetically realized nucleus may properly govern the nucleus to its left. A nucleus that is properly governed by the nucleus to its right will be phonetically unrealized. (A parameter setting will let some languages additionally allow a word-final cold vowel to be phonetically unrealized. Standard French and MA both have this parameter setting.)

Recall from section 3.3.1 the example of schwa/zero alternation from French. Charette (1988) accounted for the difference between *d t∅ l d∅mander* and *d∅ t l∅ dmander* as a difference in which nuclei governed which others:




In (6.16), the \tilde{a} properly governs the first nucleus of *demander*, making it phonetically unrealized. Similarly, the nucleus of *le* properly governs the nucleus of *te*, allowing it to be phonetically unrealized. The first nucleus of *demander* could not properly govern the nucleus of *le*, since the first nucleus of *demander* is phonetically unrealized and hence does not have the qualifications to be a proper governor; the nucleus of *le*, then, must be phonetically realized as schwa.

Kaye (1990b) states the principles controlling this effect as:¹⁰


- (6.18) *Empty Category Principle* (Kaye 1990b: 314)
- i. A licensed empty nucleus has no phonetic realization.
 - ii. An empty nucleus is licensed if (a) it is properly governed or (b) if it is domain final in languages which license domain-final empty nuclei.
- (6.19) A nuclear position α properly governs a nuclear projection β iff
- a. α is adjacent to β on its projection
 - b. α is not itself licensed
 - c. No governing domain separates α from β .

The GP treatment of MA is similar to the treatment of French just seen, except that there is no optionality in which nuclei are proper governors and which are not. In *kteb* 'he wrote', there are three empty nuclei, or rather three nuclei filled with the cold vowel v^0 , one after each of the consonants:



Since MA has the parameter setting that licensed word final empty nuclei to be phonetically unrealized, v_3^0 will not be pronounced. It is now no longer capable of properly governing v_2^0 , so v_2^0 cannot be licensed and must be phonetically realized. v_1^0 now properly governs v_2^0 , which will be licensed and phonetically unrealized.

Rather than having a drastically different syllable structure, *ketbu* is almost identical to *kteb*, except for the presence of the suffix *-u* and the distribution of proper government. The first two nuclei remain filled with the cold vowel v^0 .



¹⁰Requirement (6.19c) is intended to prevent three surface consonants in a row, which could arise if empty nuclei with codas were allowed to be unrealized [CØC][CV] (where proper government would have to cross the coda licensing governing domain) or if a nucleus with a branching onset could properly govern [CØ][CCV] (where proper government would have to cross the constituent government domain within the branching onset).

This time the third nucleus is phonetically realized and *can* properly govern v_2^0 , which will be phonetically unrealized and unable to properly govern v_1^0 , which will be unlicensed and pronounced.

Perhaps the weakest aspect of GP approaches to Semitic languages is unclear nature of its proposals on templates and template satisfaction, though it should be said that templatic morphology has not been the focus of GP research on MA. Kaye (1990b) gives the MA active participle template as:

$$(6.22) \quad \begin{array}{cccccc} \text{O} & \text{N} & & \text{O} & \text{N} & \text{O} & \text{N} \\ | & | & & | & | & | & | \\ \text{x} & \text{x} & \text{x} & \text{x} & \text{x} & \text{x} & \text{x} \\ & | & & & & & \\ & \text{a} & & & & & \end{array}$$

Roots are represented with the rather sketchily defined notation of vertical bars, e.g., |ktb|. The segments of the root are somehow inserted into the template, though exactly how this operation is to be performed is never made clear.¹¹ Both the template and the root are problematic. Kaye's concept of the template is subject to the same kinds of criticisms Prosodic Morphology makes against templates constructed from CV skeletal: the structure in (6.22) is prosodically arbitrary. There are no suggestions as to why we could not have templates even more complex than (6.22), say one requiring five heavy syllables in a row with the second and fourth onsets preassociated to *f* and *q*. There are also problems with the concept of the root. It is telling that the abbreviatory notation of |ktb| is never spelt out in full. I believe that it could not be. A root could not be a sequence of segments, because segments have no independent status in GP. The phonological content of a segment is represented with one or more privative elements like I^0 , A^+ , h^0 . In order to be defined as a segment, these elements need to be associated to a skeletal slot. But the roots themselves cannot contain skeletal slots — skeletal slots are already indispensable parts of the templates, and there is no non-ad-hoc way in GP in which two sequences of skeletal slots could become superimposed on each other.

These difficulties result from accepting the Physical Integrity of Morphemes assumption (that two different phonological objects cannot occupy the same space at the same time) and from GP's few remaining proceduralist assumptions (here, that

¹¹The most complete discussion of the question is footnote 10 of Kaye (1990a):

A word should be said concerning the association of radical segments with the positions of the template in question. I follow here the traditional autosegmental left-to-right approach. Further I follow the view that only certain positions are accessible to segments of a given morpheme. I assume that the causative template contains the same five positions as its non-causative counterpart. In addition, it contains a single position, that which follows the first nucleus[,] that is proper to the causative form. Finally, this position has no segmental content.

verb forms have to be derived in virtual time by actively linking up the two representations that form the morphemes). It will be shown in the next section how adopting a constraint-based approach to morphology will allow both the root morpheme and the “template” morpheme to refer to same skeletal positions of a PS without any ad hoc devices. It will also be seen that “template” morphemes will not need to be given elaborate structures like (6.22); they will consist of just one or two simple prosodic conditions.

Though nominal morphology will not be the primary focus of the next section, I should spend some time discussing cluster-final nouns such as *raʒl*, since they are the area in which my proposals and those of Kaye (1990a) are most divergent.

Kaye (1990a) discusses the forms of nouns at some length.¹² The generalization he is most interested in capturing is this: There is a fairly clear difference in the types of consonants found in nouns of the form CCeC and those of the form CeCC:

(6.23)	<i>bent</i>	‘girl’		<i>nmer</i>	‘tiger’
	<i>kelb</i>	‘dog’		<i>qfel</i>	‘lock’
	<i>merd</i>	‘sickness’		<i>msen</i>	‘stone’
	<i>melk</i>	‘angel’		<i>sbe</i>	‘lion’


The CeCC forms are problematic for a GP analysis that adopts the strong “one nucleus per consonant” position. *bent* would have to have the structure:

(6.24)		N		N		N
	x	x	x	x	x	x
	b	v ₁ ⁰	n	v ₂ ⁰	t	v ₃ ⁰

v₃⁰ would be licensed to be phonetically realized by the end-of-the-word parameter, but contrary to the definition in (6.19b) it would also have to properly govern v₂⁰.

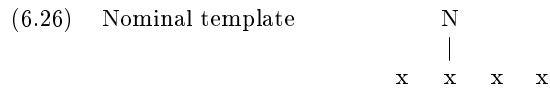
In order to avoid this problem, Kaye (1990a) proposes another structure for words like *bent*, noting that each of them ends in a consonant cluster that is a prime candidate for being a coda-onset sequence:

¹²Though published in 1990, Kaye (1990a) was written some years earlier and assumes an earlier version of GP that is in many respects incompatible with current proposals. For example, it assumes that word-final consonants can simply be left stranded at the end of the word and need not be followed by an empty nucleus.




Indeed there does seem to be a generalization that whenever the two final consonants meet the sonority criteria for standing in a coda licensing relation, they do so, but when they cannot, as with the *m* and *r* of *nmer*, they are separated by a schwa.

Kaye (1990a) tries to account for this by proposing a four-slot nominal template:



Whenever the consonants associated to the second and third slots can stand in a coda licensing relation, they do so. If they cannot without violating the sonority principles of GP, epenthesis must take place, inserting an extra nuclear position between the two:



In the more developed (and more principled) version of GP presented in Kaye (1990b), this kind of epenthesis operation is highly suspicious. As well, there are some empirical problems with the analysis. As mentioned in the last section, the possibility of both *meṛq* and *mṛeq* for ‘sauce, gravy’ makes it seem highly unlikely that any deterministic derivation will be able to produce the difference between CeCC and CCeC patterns based solely on the properties of the consonants. As well, there are exceptions to the sonority principles in apparent coda-onset sequences, both in underived nouns (e.g., *raʒl* ‘man’) and in a large class of morphologically derived deverbal nouns, known as *masdars*:

(6.28)

verb		masdar	
<i>leb</i>	‘play’	<i>leb</i>	‘playing’
<i>qfez</i>	‘jump’	<i>qefz</i>	‘jumping’
<i>wzen</i>	‘weigh’	<i>wez n</i>	‘weighing’

Kaye (1990a) tries to explain masdars by pointing out that they started out life as verbs, where it was established that there was no coda-licensing relation between the second and third consonants:

$$(6.29) \quad \begin{array}{ccccccc} & & & & \text{N} & & \text{N} \\ & & & & | & & | \\ [N[V & x & x & x & x & x & v]N] \\ & & & & \underbrace{\quad \quad \quad} & & \underbrace{\quad \quad \quad} \\ & & & & \text{no government} & & \end{array}$$

To derive the masdar, the verbal template is “adapted” to the nominal template (again a critical vagueness in defining the mechanism of template satisfaction), resulting in the loss of a skeletal slot:

$$(6.30) \quad \begin{array}{ccccccc} & & & & \text{N} & & \text{N} \\ & & & & | & & | \\ [N[V & x & x & x & x & v]N] \\ & & & & | & & | \\ & & & & \text{w} & & \text{z} & & \text{n} \end{array}$$

This representation is not compatible with the current tenets of GP. Indeed, I see no alternative to giving *wezn* the same structure that was considered and rejected for *bent* in (6.24). Some weakening of the Empty Category principle and definition of proper government is going to be needed in order to allow structures like (6.24), the kind of weakening I tried to accomplish with the idea of **extra-nulls** in section 3.3.1. Given that a structure like (6.24) is going to be needed anyway in order to represent masdars, much of the motivation disappears for wanting to represent *bent* differently, with the *n* in a coda position under a branching rhyme, as in (6.25).

Clearly, the framework of this dissertation does not permit a branching rhyme in any event, but I believe that even within a framework that *does* permit them, there are good reasons for not using them for MA nouns like *bent*. One of the strongest reasons is that these nouns commonly have broken plurals and broken diminutives in which the apparent coda consonant appears as an onset:

(6.31)	<i>singular</i>	<i>plural</i>	<i>diminutive</i>	<i>gloss</i>
	bent	bnat	bnita	‘girl’
	meṛḍ	mṛaḍ		‘sickness’
	kelb	klab		‘dog’
	selk	sluk		‘wire’

It is this kind of vowel/zero alternation in *kteb~ketbu* that led us to propose the presence of empty nuclei after the *k* and the *t*. It seems natural to extend this reasoning

and posit an intervening nucleus between the *n* and *t* of *bent*, which is sometimes phonetically unrealized and other times is filled by a full vowel.¹³

Within the present framework, accepting that nouns like *bent* have a PS with three empty nuclei, as in (6.24), the question raised by Kaye's analysis becomes: is anything gained by assuming a coda licensing relation between the last two consonants? I believe that the advantages of restricting coda licensing to geminates (as proposed in the next section) far outweigh the advantage of expressing the tenuous generalizations about sonority sequences in nouns.


6.3 The cold nucleus system

The analysis of MA involves two fairly independent sets of principles, one which I shall call the **cold nucleus system** and the other the **prosodic system**. The cold nucleus system consists of those principles that determine the distribution of empty positions throughout an MA word (**simple-null**, **extra-null**, and **unspecified** positions). I refer to it as the cold nucleus system because it involves principles that determine the phonetic realization or non-realization of nuclei that in GP analyses are filled by the cold vowel, v^0 . The prosodic system controls the nature and distribution of prosodic structure (metrical line sorts and foot governments), and allows morphemes to require their PSs to satisfy certain prosodic conditions. We shall see how the behaviour of MA words can be explained as the interaction of these two sets of principles. Although all parts of a word must of course respect both sets of principles, it will generally be the case that the effects of the prosodic system are most visible in the first part of the word (the part that is subjected to the "template") while the effects of the cold nucleus system are more visible nearer the end of the word, which are usually left unclaimed by any template. In this respect, my analysis is very much like Archangeli's (1991) analysis of templatic morphology in Yawelmani, where words were the result of prosodic templates imposed at the beginning of the word with general principles of syllabification applying to the rest of the word.

As argued for in Government Phonology analyses, I shall assume that apparent consonant clusters in MA are really separated by empty nuclei. Thus, *kteb* would have the following structure:


¹³The Projection Principle of GP ("Governing relations established at the level of lexical representation are maintained at all levels of representation" Kaye 1990a: 138) also recommends the empty nucleus analysis. Although some way might be contrived of denying that the and of (6.24) are in a coda licensing relation at the level of lexical representation (and thus illegally cease to be in the broken plural), this kind of escape hatch would seriously compromise the empirical content of the Projection Principle.

(6.32)




In fact, I push the GP claim one step further in extending it to geminates, which Kaye (1990a,b) still holds to have a traditional coda-onset structure. Where Kaye would argue for a structure of (6.33) for *weqqef* (where \Leftarrow represents the coda-licensing government relation between a coda consonant and the following onset), I propose (6.34).

(6.33)




(6.34)



Accepting that all onsets in MA belong to their own nuclei, the question becomes, what *sorts* of nuclei. MA uses the four basic nuclear sorts made available by universal grammar, as discussed into section 3.3.1. The four sorts can be grouped into two supersorts, an arrangement which may be diagrammed as follows.

(6.35)




PS diagrams will use the usual abbreviations for these sorts:

(6.36)	null	\emptyset
	simple-null	s \emptyset
	extra-null	e \emptyset
	constriction	c
	specified	sp
	unspecified	u

Recall that **constriction** nodes are those that represent an articulatory constriction gesture and must be interpreted as such by the phonetic component. Within the supersort of **constriction**, there are two basic sorts, **specified** and **unspecified**. **Specified** nodes are those who have their gestural features specified phonologically, within the PS. An **unspecified** node, on the other hand, cannot have any gestural features in a PS—it shares with **null** nodes the inability to be the source of an **articulator**, **site**, or **degree** arc. An unspecified node receives the default values of these features only during phonetic interpretation.

As partially evidenced by their unified behaviour with respect to gestural features, **unspecified** and **null** can be seen as forming another supersort, one which cross-cuts the supersorts diagrammed in (6.35). This supersort will be called **cold**. These are roughly those nuclei that orthodox Government Phonology would argue were filled with the cold vowel v^0 , whether pronounced or not—though I also include here those empty nuclei that dominate onsets which GP would treat as codas:


(6.37)



Null nodes cannot be phonetically interpreted as articulatory gestures, default or not. Their only contribution to the phonetic realization of the PS will be in the temporal ordering of the other nodes they stand in government relations with. As pointed


out in the discussion of null positions in section 3.3.1, **simple-nulls** tend to be severely restricted in their distribution, but have relatively few constraints on the kinds of onsets they can support. **Extra-nulls**, on the other hand, tend to be free of most of the distributional restrictions on **simple-nulls**, but will often be restricted in the kinds of onsets they can support (for example, the English constraint that onsets of extra-nulls be anterior coronals, that is, [s:Ant]) and/or in the type of morphological constituents they can occur in. For MA, I propose that there can be two conditions that allow an **extra-null** nucleus. The first allows an **extra-null** if its onset is coda-licensed; together with the local-domain-creating properties of the **coda-licence** government, this will result in **extra-nulls** that dominate the first member of a geminate consonant:

(6.38)



The second possible condition allows an **extra-null** nucleus as the **bottom** of a PS associated with a Noun in the syntactic tree. This is to allow *masdars* (e.g., *wezn* ‘weighing’) and underived nominals like *bent* ‘girl’ or *kelb* ‘dog’:

(6.39)



These nouns need an **extra-null** at their bottom in order to avoid violating the **Adjacent Nulls Constraint**.¹⁴

Ignoring the role of extra-nulls in nouns, these principles can be formalized as:

¹⁴Another intriguing possibility is that these nouns do not involve an extra-null at all, but that the avoidance of the **Adjacent Nulls Constraint** comes from the final two consonants not belonging to the same phonological word. We saw in the last chapter the possibility of

- (6.40) The onset of an extra-null nucleus must be coda-licensed. Whenever there is a nucleus of sort **extra-null**, there will be a **coda-licence** arc to its onset from the following onset.

$$\forall v_1 \text{ extra-null}(v_1) \rightarrow \exists c_1, v_2, c_2 \ c_1.v_1.c_2.v_2 \wedge c_2 \xrightarrow{\text{cl}} c_1$$

- (6.41) A coda-licence government arc creates a local domain for all segmental specifications.

$$\begin{aligned} \forall g \text{ coda-licence}(g) \rightarrow \\ & \text{local-domain-creator}(g, \text{articulator}) \wedge \\ & \text{local-domain-creator}(g, \text{site}) \wedge \\ & \text{local-domain-creator}(g, \text{degree}) \wedge \\ & \text{local-domain-creator}(g, \text{secondary}) \wedge \end{aligned}$$

Simple-nulls are of course subject to the Adjacent Nulls Constraint formulated in section 3.3.1, which we repeat here:

- (6.42) *Adjacent Nulls Constraint*

$$\sim \exists v_1, v_2 \ v_1 \xrightarrow{n} v_2 \wedge \text{simple-null}(v_1) \wedge \text{simple-null}(v_2)$$

or more fully:

$$(6.43) \quad \sim \exists v_1, g, v_2 \ \text{arc}(v_1, g, v_2) \wedge \text{simple-null}(v_1) \wedge \text{simple-null}(v_2) \wedge \text{within-word}(g)$$


For MA, we can formulate a similar constraint against adjacent **unspecified** nuclei.¹⁵

a single “word” containing a phonological word boundary (e.g., English . It may be that, even in the absence of any obvious triggering affixes, MA systematically allow the bottom nucleus of a noun to begin a new phonological word (forcing the preceding nucleus to end a phonological word). This would be a somewhat different type of mismatch between morphosyntactic and prosodic parsings of the nuclear spine than we have yet seen, one in which a phonological word PS is *smaller* than the PS of a syntactic atom. Hewitt (to appear) argues for an analysis of ternary stress in Alutiiq that uses prosodic words that are similarly smaller than the PS of the corresponding syntactic atom. Clearly there would need to be more examination of the ramifications of this proposal before we could adopt it.

¹⁵This constraint is for the most part meant to force a nucleus to be null whenever possible. Presumably, this work could be done by the nonmonotonic portion of the grammar that select for each form only one PS from the set of PSs that each satisfy all the constraints, of the selection metric counted a PS with an **unspecified** position as “larger” than an otherwise identical PS with a **null**. In the interests of expressing as much as possible using hard constraints, I will continue to assume the constraint in (6.44). An advantageous side-effect of (6.44) is to systematically ban long (bimoraic) schwas without complicating the prosodic system.

- (6.44) *Adjacent Unspecifieds Constraint*
 $\sim \exists v_1, v_2 \ v_1 \xrightarrow{n} v_2 \wedge \text{unspecified}(v_1) \wedge \text{unspecified}(v_2)$


As an example of how these basic parts of the cold nucleus system work, consider *kteb* ‘he wrote’. The shape demanded by the lexical constraint of the verb root can be diagrammed by:




The verb root is the only lexical constraint that troubles itself over the shape of *kteb*. The 3sm.pf form being morphologically “unmarked”, *kteb* gets no full vowels or templates from any inflectional or derivational morpheme; in other words, any inflectional or derivational features with which the verb stands in construction in the syntactic tree act as zero-morphemes in the sense of chapter 5. Because no other morpheme contributes a full vowel to the form, each of the nuclei will be **cold**. The question is: which **cold** nuclei will be **simple-null** and which will be **unspecified** (and surface with a default schwa)? Assuming that the final nucleus must be null, by a constraint discussed below, there are four possible structures:

- (6.46)
- | | | | | |
|----|----------------|----------------|----------------|--------------|
| a. | kv^\emptyset | tv^c | bv^\emptyset | <i>kteb</i> |
| b. | kv^c | tv^\emptyset | bv^\emptyset | <i>ketb</i> |
| c. | kv^\emptyset | tv^\emptyset | bv^\emptyset | <i>ktb</i> |
| d. | kv^c | tv^c | bv^\emptyset | <i>keteb</i> |

v_2 and v_3 of (6.46b) violate the Adjacent Nulls Constraint. (6.46c) violates the Adjacent Nulls Constraint twice over. (6.46d) violates the Adjacent Unspecifieds Constraint of (6.44). Only (6.46a) satisfies all the constraints, resulting in the PS:



The third person plural, *ketbuu* ‘they wrote’,¹⁶ is also obtained by these principles analogous to Government Phonology. We may assume that the morpheme of the 3pl suffix *-uu* requires its top to be joined to the bottom of the verb root, in the relation $Ph_b(ktb) \xrightarrow{n} Ph_t(uu)$:



v_1 will be **unspecified**, v_2 **simple-null**, resulting in *ketbuu*.¹⁷

The last remaining question is what sorts of nuclei are allowed to occur in word-edge positions, that is, as the governees of **top** and **bottom** arcs. The **top** arc apparently imposes no special constraints on its governee. The **bottom** arc on the other hand does not allow an **unspecified** position as its governee. A **constriction** node is permissible, resulting in a final full vowel as in *ketbuu*. A simple-null is permissible, resulting in an apparent consonant-final word, as in *kteb*. What is not permitted is an **unspecified** position, a constriction node with no gestural specifications, that would receive a default schwa phonetically. *ketbe* is not a legal MA word.

This is clearly a language-particular constraint on word-final nuclei. Yawelmani is good example of what MA would be like if it allowed unspecified final nuclei. Many word-final Yawelmani suffixes lexically demand a simple-null. The lexical constraint of aorist *-hin*, for example, requires that the nucleus dominating the *n* be a simple-null. But there are some suffixes that do not make this stipulation, and here we find the variation we would expect. The consequent gerundial suffix, *-mi* \sim *-im*, for example, makes no claims on the sort of the *m*'s nucleus. The final nucleus will be a simple-null when permitted by Yawelmani's alternation conditions, which are almost identical to MA's. In *walxo-m*, the final nucleus can be a simple-null. In other cases, like *oglin-mi*,

¹⁶Since I adopt the Government Phonology argument that full vowels in MA are bimoraic, I shall write all full vowels as long from now on.

¹⁷The reason the reverse assignment is not possible, i.e., v_1 as **simple-null** and v_2 as **unspecified**, resulting in *ktbu*, is a prosodic requirement on the relation between v_2 and v_3 that will be discussed in the next section.

where a final simple-null is not possible, the final nucleus will be **unspecified** and will phonetically receive Yawelmani's default vowel, *i*. (Cf. Newman 1948:134–5.)

6.4 The prosodic system

All MA nuclei of whatever sort are **line-0** positions. All nuclei are thus available to participate in prosodic structure, but they are not required to. Unlike many languages, there is no Exhaustivity Condition in MA (cf. Halle and Vergnaud 1987). There are no dire consequences awaiting a MA nucleus that is neither a head nor a dependent of a foot government arc.

The final nucleus of a word is prosodically atypical. This nucleus can enter into prosodic relations without many of the restrictions that apply to other nuclei. For example, a null nucleus may only be a line 1 position at the end of a word. Similarly, this is the only position where the second consonant in a moraic trochee of the form [CeC] need not be the first half of a geminate.

6.4.1 Moraic trochees

A nucleus may be **line-1** if it is the head of a moraic trochee, the fundamental constituent of the MA prosodic system. In more traditional metrical terms, a MA moraic trochee is obligatorily branching, i.e., must contain two morae.

$$(6.49) \quad \begin{array}{c} v^{l1} \\ \quad \swarrow \mu^t \\ \quad \quad v^{l0} \end{array}$$

(A nucleus may also be **line-1** at the end of a word without having to branch.) The obligatory branchingness of the moraic trochee is the result of the following constraint:¹⁸

$$(6.50) \quad \begin{array}{l} \text{Moraic trochees are obligatorily branching or are word-final. Non-word-} \\ \text{final line 1 positions must govern another nucleus in a moraic trochee.} \\ \forall v_1 \text{ line-1}(v_1) \leftrightarrow \\ \text{word-bottom}(v_1) \vee \exists v_2 \ v_1 \xrightarrow{n} v_2 \wedge v_1 \xrightarrow{t} v_2 \end{array}$$

The conjunction of $v_1 \xrightarrow{n} v_2$ (or more precisely $v_1 \xrightarrow{\text{log}} v_2$) and $v_1 \xrightarrow{t} v_2$ will usually be abbreviated as $v_1 \xrightarrow{\mu^t} v_2$.

¹⁸We can define a nucleus as *word-bottom* if it is pointed to by the **bottom** arc of a phonology node that is word-level and is also associated in a sign with some syntax node. (We need this last requirement in order to rule out the possibility of (spurious) non-morphosyntactically defined prosodic constituents licensing illicit prosodic properties.) The definition might look like: $\text{word-bottom}(v) \leftrightarrow \exists x \text{ Ph}_b(x) = v \wedge \text{Ph}(x) \xrightarrow{\text{level}} \text{word}$.


These moraic trochees are involved in two types of structures, traditionally considered as heavy syllables: i) a long vowel, which in MA is any full vowel, *a*, *i*, or *u*; ii) a schwa followed by the first member of a geminate consonant. Thus, the dependent in a moraic trochee must be of the sort **rhyme-dependent**, introduced in section 3.3.1, which requires either the nucleus or its onset to be null. Furthermore, except at the end of a word, the dependent must be either **specified** (if it is the second half of a long vowel) or **extra-null** (in the case of a geminate); more concisely, it must *not* be **simple-null**.

- (6.51) The dependent of a moraic trochee government must be of the sort **rhyme-dependent** and can only be **simple-null** at the end of a word.


$$\forall v_1, v_2 \quad v_1 \xrightarrow{\mu^t} v_2 \rightarrow \text{rhyme-dependent}(v_2) \wedge (\text{simple-null}(v_2) \rightarrow \text{word-bottom}(v_2))$$

A long vowel and a geminate consonant involve the following structures:

- (6.52) Long vowel



- (6.53) Geminate consonant



The geminate consonant structure follows from the requirement that the dependent of a moraic trochee be **extra-null** and from constraints (6.40) and (6.41) of the cold nucleus system, requiring a local domain creating coda-licence in the neighbourhood of an **extra-null**.

We should also express the generalization noted in Government Phonology analyses of MA, that the full vowels, *aa*, *ii*, and *uu*, are necessarily bimoraic:

- (6.54) Any **specified** nucleus must be either the head or dependent of a moraic trochee.

$$\forall v \text{ specified}(v) \rightarrow \exists u \quad u \xrightarrow{\mu^t} v \vee v \xrightarrow{\mu^t} u$$

We prohibit vocalic diphthongs by making moraic trochee governments local-domain-creators for vocalic features:

$$(6.55) \quad \forall g \text{ moraic-trochee}(g) \rightarrow \text{local-domain-creator}(g, \text{site})$$

Note that this last constraint is also valid for those moraic trochees that involve geminate consonants. Since neither the **unspecified** nucleus of the schwa nor the following **extra-null** nucleus governing the first half of the geminate has any gestural features at all, they vacuously share all their site specifications and thus respect the Spreading Constraint.

6.4.2 Iambs

Iambic feet are involved in the prosody of MA. A **line-1** position, as well as obligatorily standing in a trochaic relation with the immediately following nucleus, may also stand in an iambic relation with the preceding nucleus, resulting in the structure proposed in section 3.3.2:

$$(6.56) \quad v_1^{l0} \text{---} i \text{---} \boxed{v_2} \text{---} \mu t \text{---} v_3^{l0}$$

ll


The strongest evidence for iambs in MA comes from the templatic behaviour of the vast majority of broken plurals.¹⁹ Consider the prosodic properties of the beginnings of these representative broken plurals (from Heath 1987: 103):

(6.57)	<i>qelb</i>	<i>qluub(aa)</i>	'heart'
	<i>sqef</i>	<i>squufaa</i>	'roof'
	<i>ṭiiṛ</i>	<i>ṭyuuṛ</i>	'bird'
	<i>žbel</i>	<i>žbaal</i>	'mountain'
	<i>feddaan</i>	<i>fdaaden</i>	'field'
	<i>buun</i>	<i>bwaan</i>	'coupon'
	<i>saaruut</i>	<i>swaaret</i>	'key'

Each of these broken plurals can be analyzed as beginning with an iamb, e.g., [qØluu]b(a):


¹⁹Arabic plurals are generally divided into **sound** plurals and **broken** plurals, with many nouns having (at least) one plural of each type. Sound plurals are formed by adding a suffix, but otherwise involve no change from the form of the singular, except perhaps schwa/zero alternations as determined by the cold nucleus system. Broken plurals involve templatic principles and are often quite different from the forms of the corresponding singulars.

(6.58)



Following McCarthy and Prince (1990b), who argue convincingly for an iambic broken plural template in Classical Arabic, we can propose an iambic prosodic condition for these plurals in MA. The effects of this iambic template are especially visible in the plural of *saaruut*, *swaaret*:

(6.59)



The singular *saaruut* already begins with a moraic trochee. In the plural, we find an extra templatic default consonant *w* whose sole purpose is to allow the form to begin with an iambic government relation.

One peculiarity of MA iambs is that the left-hand element must be null:

(6.60) Obligatory Iambic Weakening

$$\forall v_1, v_2 \quad v_1 \xrightarrow{i} v_2 \rightarrow null(v_1)$$

Cross-linguistically there is a strong tendency for vowels in this prosodic position to be reduced or “deleted” altogether, a phenomenon often referred to as Iambic Weakening and discussed in works on metrical phonology such as Hayes (1991). In MA, this tendency has simply been made obligatory.

6.4.3 Syllabic trochees

Constraint (6.50) on the obligatory branchingness of moraic trochees left an escape hatch for the end-of-word position. While it is generally true that a line-1 position must govern a dependent, and furthermore that a line-1 position cannot be null,

I am not strongly committed to the reality of this pointer. It does considerably simplify the presentation of the analysis, though I believe a more longwinded analysis could be made to work without it. Just as the function $Ph_t(x)$ relates a syntactic node x to the nucleus pointed to by the top pointer of its associated PS, we shall use the function $Ph_\eta(x)$ to relate x to the tail of its PS's η arc.

I propose that the lexical constraint of an MA verb will involve one of the two special predicates, $root_3$ or $root_4$. These are three- and four-place predicates respectively whose arguments refer to nodes (and will usually be satisfied by onset nodes). The lexical constraint can specify the segmental content of each of these consonant nodes using the predicate *full-value*. Using the $\overset{f}{\approx}$ abbreviation for the *full-value* specifications, we can write the phonological part of *KTB*'s lexical constraint as:

$$(6.64) \quad \text{Lexically specified phonology of the strong root KTB 'write'}$$

$$root_3(C_1, C_2, C_3) \wedge C_1 \overset{f}{\approx} k \wedge C_2 \overset{f}{\approx} t \wedge C_3 \overset{f}{\approx} b$$

I use capital C s for the argument variables, because lower-case c s might misleadingly suggest that the three consonants must belong to successive nuclei. (In the case of weak and hollow verbs, the three “consonants” need not even be onsets.)

The three arguments of the *root* predicate are related to the pointers in (6.63) fairly straightforwardly. The top nucleus onset-licenses C_1 . The bottom nucleus onset-licenses (or is) C_3 . And the nucleus that is the tail of the η pointer onset-licenses (or is) C_2 .

(6.65) Any syntactic node that is a V^0 has an associated PS of which the *root* predicate is true, and the arguments of the *root* predicate are aligned with the top, η , and bottom pointers of the PS.

$$\forall x \ x \xrightarrow{\text{cat}} \text{verb} \rightarrow \exists C_1, C_2, C_3$$

$$root_3(C_1, C_2, C_3) \wedge$$

$$Ph_t(x) \xrightarrow{o} C_1 \wedge$$

$$Ph_\eta(x) \xrightarrow{o}_r C_2 \wedge$$

$$Ph_b(x) \xrightarrow{o}_r C_3$$

This constraint requires that the $root_3$ predicate be true of *all* verbs, even quadrilaterals. This turns out to be a real generalization of MA. All the constraints that apply to trilaterals systematically apply to quadrilaterals as well. Indeed, the major difference seems to be the addition of an extra consonant position in the onset of the nucleus immediately following the top nucleus:²⁰

²⁰Some templates will not be able to apply to quadrilaterals because their requirement for a moraic trochee will conflict with the root's requirement for the identity of C_2 , which will typically not be able to be integrated into a moraic trochee.


$$(6.66) \quad \text{Quadrilateral roots} \\ \forall C_1, C_2, C_3, C_4 \text{ root}_4(C_1, C_2, C_3, C_4) \rightarrow \text{root}_3(C_1, C_3, C_4) \wedge \\ \exists v_1, v_2 \ C_1.v_1.C_2.v_2$$

Besides using *full-value* to specify the features of consonant positions, lexical constraints also determine if the verb is weak or hollow. I propose that this is done with the predicate *weak*, which may apply to the \mathfrak{f} nucleus (for a hollow verb) or the **bottom** nucleus (for a weak verb). The weak argument of root_3 will be given vocalic specifications. For example:

$$(6.67) \quad \text{Lexical constraint of the hollow root NuḌ ‘get up’} \\ \forall x \ x \xrightarrow{\text{cat}} \text{verb} \wedge \text{semantics}(x) \approx \text{“get up”} \rightarrow \\ \text{root}_3(C_1, C_2, C_3) \wedge C_1 \overset{f}{\approx} \text{n} \wedge C_2 \overset{f}{\approx} \text{u} \wedge C_3 \overset{f}{\approx} \text{ḍ} \wedge \text{weak}(\text{Ph}_9(x))$$


Although this chapter will not be primarily concerned with the principles that determine the form of hollow and weak verbs, I shall spend a moment showing that they are at least consistent with the root constraints proposed so far. Consider the hollow root NuḌ ‘get up’, which has the 3sm.pf *naaḍ*, the imperfective stem and imperative *nuuḍ*, and a causative *newwēḍ*. In constraint (6.65), the \mathfrak{f} and bottom nuclei are related to C_2 and C_3 by the potentially reflexive $\overset{0}{\rightarrow}_r$ relation rather than the absolute $\overset{0}{\rightarrow}$ relation. This property together with the conditionality of the root’s *full-value* specifications explains the behaviour of the vowel specified for the middle position of a hollow root. In the imperfective/imperative form *nuuḍ*, the \mathfrak{f} pointer points to v_2 , and C_2 coincides with v_2 , exploiting the possibility that $\text{Ph}_9(x) \overset{0}{\rightarrow}_r C_2$ can be satisfied when $\text{Ph}_9(x) = C_2$:

$$(6.68) \quad C_1 = c_1 \\ C_2 = v_2 \\ C_3 = c_3$$




In the causative *newwēḍ*, C_2 coincides with c_3 , exploiting the possibility that $\text{Ph}_9(x) \overset{0}{\rightarrow}_r C_2$ can be satisfied when $\text{Ph}_9(x) \overset{0}{\rightarrow} C_2$:

- (6.69) $C_1 = c_1$
 $C_2 = c_3$
 $C_3 = c_4$



Even 3sm.pf *naad*, where no *u* appears at all, is no problem. The PS itself is as in *nuud*, except that the full vowel specification is the *a* demanded by the 3sm.pf morpheme. But C_2 , the position that is required to have a full-value specifications for *u*, can coincide with an onset position, c_2 . Since c_2 is null, it is not a full position, and there is no need for the full-value specifications to appear in the PS. The lexical constraint is again satisfied.

- (6.70) $C_1 = c_1$
 $C_2 = c_2$
 $C_3 = c_3$




There is one more interesting property of roots in MA. Consider the relation between the nuclei of C_2 and C_3 in the three forms 3sm.pf *k_teb*, 3p.pf *ketbuu*, and the passive participle *mektuub*. Particularly, consider the possible prosodic constituencies of the portion of each word from *t* to the end:


- (6.71) k [teb] moraic trochee
 ke [tbuu] iamb
 mek [tuub] syllabic trochee

In each case, there is a prosodic government between the nucleus of C_2 and the nucleus of C_3 . This relation is an *i* in *ketbuu*. It is a *τ*, as well as a *l0g*, in *kteb*, that is, a moraic trochee arc. It is a *τ* and a *l1g* in *mektuub*:


(6.72) *kteb*



(6.73) *ketbuu*



(6.74) *mektuub*



We never find the C_2 and C_3 split between two different feet.

We can express the constraint that the nuclei of C_2 and C_3 stand in a prosodic relation, that is, that they be joined by an arc belonging to the supersort *foot* (consisting of *iamb* and *trochee*), as follows, using the \mathfrak{f} and *bottom* pointers:

$$(6.75) \quad \forall x \quad x \xrightarrow{\text{cat}} \text{verb} \rightarrow Ph_9(x) \xrightarrow{\text{foot}} Ph_t(x)$$

This constraint may have to be weakened in order not to apply to hollow and weak verb roots.

It is this property that is responsible for the form of the 3pl form of *kteb* being *ketbuu* rather than *ktebuu*. The nuclei governing the *t* and the *b* must be in a prosodic government relation. Given the presence of the final inflectional *-uu*, the only relation this could be is an iambic one. By constraint (6.60), the leftmost nucleus in an iamb must be null. Hence, *ke[tØbuu]* rather than *ktebuu*.

6.6 Moroccan Arabic “templates”

We are now ready to see how the “templates” of Moroccan Arabic, together with the constraints governing verb roots, the cold nucleus system, and the prosodic system, result in the ultimate form of verbs. Because so many properties of words are already handled by these other systems of constraints, there is no need to specify them again in the lexical constraints of “templatic” morphemes such as the causative. Our templates will not look like CeCCeC (cf. McCarthy 1979, Heath 1987), with the need for language-specific association conventions and extra epenthesis rules for fine-tuning. More in the spirit of McCarthy and Prince (1986, 1990), the templates will simply be a prosodic category (or rather, a prosodic requirement). There is still no need for any special association conventions; the proper association of the template’s prosodic constituent and the root’s melody is handled by the same sister alignment principles that are responsible for normal concatenation. In the case of derivational templates, it is usually the tops of the two morphemes that are aligned: $Ph_t(\text{root}) = Ph_t(\text{template})$.

6.6.1 Base forms

In the base form of verbs (or measure 1), no derivational morphemes have been applied, or more precisely, any “derivational” node in the syntactic tree behaves as a zero-morpheme. There are no external prosodic conditions applied to the root. The only things determining the pattern of Cs and Vs are the cold nucleus system and the presence or absence of third person inflectional suffixes (*-uu* ‘3pl’, *-t* ‘3sg feminine’).

First and second person inflections on perfective verbs behave as if they did not belong to the same phonological word as the root:

$$(6.76) \quad \begin{array}{llll} ktebt & \text{'1sg'} & ktebnaa & \text{'1pl'} \\ ktebtü & \text{'1pl'} & ktebtüw & \text{'2pl'} \end{array}$$

These all seem to be composed of the 3sm.pf base *kteb* as an independent word, plus the suffix. Specifically, the consonants of the suffixes have no effect on the distribution of schwas and nulls in the verb root. Given the 1sg suffix *-t*, we might expect a 1sg form of *ketbet*, a form we do indeed find for the 3sg feminine, whose suffix is also *-t*. The difference between the 1sg and the 3sg feminine follows naturally if the 3sg

feminine suffix belongs to the same phonological word as the root, affecting its nuclei by the cold nucleus system, while the 1sg is not in the same phonological word. We can capture this behaviour by a constraint requiring the top of a first or second person perfective suffix also to be the top of a phonological word.


6.6.2 Active participles

Deriving the active participle (e.g., *kaateb* for *kateb* in fact involved no prosodic conditions at all that have not already been stated as part of the prosodic system. All the morpheme for the active participle needs to specify is that the top nucleus be specified for the vowel a (or just for the site [s:Pha]). The rest, including the length of the vowel, is taken care of by principles we have already seen.

For concreteness, let us assume the following lexical entry for the active participle, abbreviating the syntax:


$$\begin{array}{ll}
 (6.77) \quad \forall x \ x \approx \text{“active participle”} \rightarrow & \\
 \quad \exists v \ \text{complement}(x, v) \wedge v \xrightarrow{\text{cat}} \text{verb} \wedge & \text{selectional restriction} \\
 \quad \text{Ph}_t(x) \xrightarrow{s} \text{Pha} \wedge & \text{vowel specification} \\
 \quad \text{Ph}_t(x) = \text{Ph}_t(v) & \text{sister alignment}
 \end{array}$$

This results in a morphosyntactic structure like:



The prosodic system requires that, since the top of the PS is a **specified** nucleus, it must be part of a moraic trochee, hence the second nucleus must also be a specified nucleus with a pharyngeal site. The PS for *kaateb* is:

(6.79)



As the dependent of a **trochee** arc, v_2 is of the sort **rhyme-dependent**, so c_2 must be **rhyme-dependent** as well. This means either v_2 or c_2 must be **null**. Since v_2 is already required to be **specified**, as the second half of the long *a*, this leaves c_2 to be **null**. This null consonant position could not possibly be specified for the *t* features demanded by the root's C_2 , so C_2 must coincide with c_3 . The fact that v_3 is **unspecified** (pronounced schwa) and v_4 is **null** (not pronounced) is determined by the normal operation of the cold nucleus system.

Interestingly, this “templatic” constraint has nothing prosodic at all about it. All it specifies is the articulatory content of a vowel, something McCarthy (1979) would think more appropriate for an inflectional morpheme marking voice. All the prosodic consequences, such as the fact that the resulting full vowel must be bimoraic, are stated in separate constraints. In this respect, MA has an easier time of it in its templatic morphology than Classical Arabic did. Since MA has reduced its historic short vowels to **unspecified** vowels, any vowel specification whatsoever automatically creates a moraic trochee. A Classical Arabic template, on the other hand, would have to explicitly state the trochaic requirement separately from the vowel quality requirement.

6.6.3 Reciprocals

The active participle is homophonous with measure 3 of the verb. The requirements for the reciprocal form are identical to those of the active participle, with the addition of a *t-* prefix. Some of the implications of this extra prefix will become more apparent in the discussion of the passive participle forms in section 6.6.5.

6.6.4 Causatives

In the causative template, we see an example of a prosodic requirement. *wqef*'s causative, *weqqef*, can be derived by two simple requirements on the first nucleus:


that it be a line 1 position (i.e., head of a prosodic foot) and, almost but not quite in direct contradiction, that it be *cold*. Again, the top of the causative's PS is aligned with the top of the verb root's PS.

$$(6.80) \quad \forall x \ x \approx \text{“causative”} \rightarrow$$

$\exists v \ v \xrightarrow{\text{cat}} \text{verb} \wedge \text{complement}(x, v) \wedge$	selectional restriction
$\text{line-1}(\text{Ph}_t(x)) \wedge \text{cold}(\text{Ph}_t(x)) \wedge$	the “template”
$\text{Ph}_t(x) = \text{Ph}_t(v)$	sister alignment

Since nulls can only be prosodic heads at the ends of words, the only way this top node could be both *line-1* and *cold* is for it to be *unspecified*. So we know that the first vowel will be a schwa. Being a prosodic head, the top nucleus must also engage in a trochaic relation.


Let us consider a false start — a PS where the C_2 q required by the verb root occurs in the PS's c_2 and the next non-null consonant is the C_3 f :



Several factors rule against a configuration like this. First of all, in the absence of full vowels, the surface form would have to be *weqef* or *weqfe*, both of which are banned by the *cold* nucleus system (*weqef* for violating the *Adjacent Unspecified Constraint*, *weqfe* for having an unspecified position word-finally). Furthermore, v_2 is the dependent in the trochaic relation required by the template, so cannot participate in any further prosodic relations — but as the nucleus pointed to by the \uparrow pointer, it must bear some kind of prosodic relation to the nucleus of f , the nucleus pointed to by the *bottom* pointer.

The only alternative is for C_2 's nucleus to be one step further away. We can draw the diagram so far as:


(6.82)



We can also give the appropriate empty sorts to the final two nuclei with our current knowledge.

The only question remaining is what v_2 and c_2 are. v_2 obviously cannot be a full vowel, or it would have to share its gesture features with v_1 , making v_1 a **specified** position against the requirements of the causative template. So v_2 must be **cold**. It cannot be **unspecified** without clashing with the unspecified v_1 , violating the Adjacent Unspecifieds Constraint. It cannot be **simple-null**, since it is only word-finally that a **simple-null** can be the dependent of a trochee. Fortunately, there is a remaining possibility. v_2 could be an extra-null. This would require c_2 to be the first half of a geminate, but nothing in any constraint forbids this. The only PS that meets all the constraints is therefore:

(6.83)



6.6.5 Passive participles

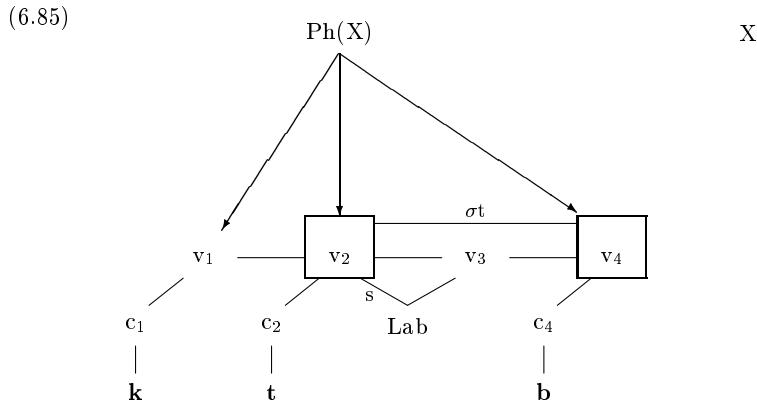
An analysis of the passive participle forms is one of the easiest places to see the difference between a constraint-based approach of the kind developed here and a more traditional representation-based account of templatic morphology. It also illustrates the \S pointer at work.

The passive participle of *kṭeb* is *mektuub* ‘(having been) written’. The expression of the passive participle consists of two parts, the prefix *m-* and the “templatic” vocalism *uu*. The conditions for each part in isolation are straightforward. Despite some initially apparent complications, their combination is just as straightforward.

The conditions on the vocalism of *-uu-* in the passive participle are extremely similar to those on the vocalism *-a-* in the active participle. Where the morpheme for the active participle required a certain vocalic quality on the nucleus pointed to by the **top** pointer, the passive participle morpheme uses the **ɿ** arc. Where the active participle morpheme contained the statement in (6.84a), the passive participle contains (6.84b).

- (6.84) a. Active participle vocalism: $Ph_t(x) \xrightarrow{s} Pha$
 b. Passive participle vocalism: $Ph_o(X) \xrightarrow{s} Lab$


We momentarily postpone discussion of the identity of X. Just as the single requirement for a Pharyngeal site interacted with other constraints to allow only the PS in (6.79) for *kaateb*, the requirement for a labial site in (6.84b) will result in the following PS:



Since a nucleus with a Labial site must be part of a moraic trochee, the additional nucleus v_3 is required. As required by constraint (6.75), the nuclei of *t* and *b*, v_2 and v_4 , must stand in a foot relation, a requirement satisfied here by a syllabic trochee relation.

The prefix component of passive participle marking is also straightforward. The PS of the syntactic head A^{-1} consists of a single (cold) nucleus with an onset specified for the consonant *m*. We do not find Nisgha-like merger in cases where the verb root also begins with *m*, so the bottom of the prefix and the top of the verb root must stand in a strict \xrightarrow{n} relation. The morphological structure would look like:

(6.86)



We can require this structure with a constraint like:

$$\begin{aligned}
 (6.87) \quad \forall a^{-1} \quad a^{-1} \text{ "passive participle"} \rightarrow & \\
 \exists V \text{ complement}(a^{-1}, V) \wedge V \xrightarrow{\text{cat}} \text{verb} \wedge & \text{selectional restriction} \\
 \text{Ph}_i(a^{-1}) = \text{Ph}_b(a^{-1}) \wedge & \\
 \exists c \text{ Ph}_b(a^{-1}) \xrightarrow{0} c \wedge c \approx \mathbf{m} & \text{content of prefix} \\
 \text{Ph}_b(a^{-1}) \xrightarrow{n} \text{Ph}_i(V) & \text{sister alignment}
 \end{aligned}$$

There remains the question of how to integrate the prefix requirement and the vocalism requirement into a single morpheme. We cannot include the *m* and the *u* in the same PS without losing the simple sister alignment properties that determine that *m* is to be a prefix. Fortunately, there is no need to. It is not necessary for every phonological property required by a morpheme to be within the PS that the morpheme associates in a sign with the syntactic node the morpheme is most intimately concerned with. Specifically, as was suggested in section 5.3.5, it is possible for the morpheme for a syntactic head to make demands on the PS of its sister. This is what is happening here. The passive participle morpheme, as well as demanding that A^{-1} 's PS have the shape of *m-*, also demands that A^{-1} 's sister have *u* at the nucleus pointed to by the sister's PS's \hat{q} pointer.

Putting these pieces together, the overall constraint for the passive participle morpheme would look like:

$$\begin{aligned}
(6.88) \quad \forall a^{-1} \quad a^{-1} \text{“passive participle”} \rightarrow & \\
\exists V \text{ complement}(a^{-1}, V) \wedge V \xrightarrow{\text{cat}} \text{verb} \wedge & \text{selectional restriction} \\
Ph_t(a^{-1}) = Ph_b(a^{-1}) \wedge & \\
\exists c \quad Ph_b(a^{-1}) \xrightarrow{o} c \wedge c \approx \mathbf{m} & \text{content of prefix} \\
Ph_b(a^{-1}) \xrightarrow{n} Ph_t(V) & \text{sister alignment} \\
Ph_\theta(V) \xrightarrow{s} Lab & \text{demand made on sister}
\end{aligned}$$

This property that two piece of phonological information can be the subjects of the same morpheme without necessarily belonging to the same PS allows analyses of templatic morphology without many of the complications that were needed in order to make a past analyses work. Under a representation-based approach, the assumption that the prefix *m-* and the vocalism *uu* are parts of the phonological content of the same morpheme is problematic. As part of the same morpheme, the two pieces should start out their derivational life in the same region of the representation. From a prefix location, the vocalism would then have to slide somehow over intervening material and dock onto the nucleus after the *t*, using some sort of (often allegedly universal) association convention. Allowing this kind of sliding is almost the sole motivation for such theoretical devices as the Morpheme Tier Hypothesis. A constraint-based approach allows the morpheme of the head to impose conditions on certain defined positions of its sister (such as the terminus of the \S pointer), without additional powerful devices that could compromise the application of locality principles in other areas of the grammar.

Chapter 7

Comparisons with other frameworks

The ideas in this dissertation have not developed in a vacuum. The framework argued for here has substantial similarities to several other frameworks and has been strongly influenced by many of them. This chapter discusses some of the most recent of these: Government Phonology, Harmonic Phonology, the Theory of Constraints and Repair Strategies, Autolexical Syntax, and the declarative phonology work of Bird, Scobbie, and others. While most of this dissertation was written before detailed accounts of Optimality Theory became available, there are also many similarities with (and differences from) this theory that will be discussed.

For want of a better plan of organization, the order I shall discuss these frameworks in follows three groups based on the consequences they propose for violations of constraints:

i) constraints may be violated by intermediate forms in the course of a derivation, but such violations immediately trigger repair operations. This description applies to the Theory of Constraints and Repair Strategies, Harmonic Phonology, and in many respects to Government Phonology, though the latter will be discussed first in its own section.

ii) constraints can be violated, but under the right circumstances these violations can be ignored. This description applies to Optimality Theory and Autolexical Syntax.

iii) constraints cannot be violated—the violation of any constraint results in an illegal representation. This has been the idea guiding the work of Bird (1990), Scobbie (1991), and the present dissertation.

Government Phonology is the theory that has most strongly influenced the choice of representations used in the present framework, so I shall begin with a discussion of GP.

7.1 Government Phonology

The framework presented in this dissertation is based in several respects on work from the research programme of Government Phonology. There are many clear similarities, which are perhaps most easily seen in the treatment of such phenomena as syncope and epenthesis in Moroccan Arabic. At the same time, there are distinct differences. Segmental representations are very unlike those proposed in Kaye, Lowenstamm, and Vergnaud (1985, 1990) and the notion of “charm” plays no role. Suprasegmental aspects of PSs, such as syllable structure, differ somewhat from the version of KLV (1990), though many aspects of the differences have been argued for independently within the Government Phonology tradition.¹ One of the greatest differences lies in the overall grammatical mechanisms within which the proposals on segmental and syllabic structure are embedded. Despite the stated goal of developing a “no rules” approach to phonology, work in GP still usually assumes a rich array of processual mechanisms (e.g., deletions, restructurings, quasi-cyclic concatenation) whose exact definition and behaviour is usually not made explicit. In this section, we shall look at some of the differences between the present framework and traditional GP, and attempt to motivate the choices made for the present framework.

7.1.1 Segmental structure

GP’s theory of segment-internal structure is based on privative **elements**. Primitive elements like I^0 , A^+ , h^0 , or R^0 are either present in or absent from the representation of a segment. They are pronounceable in isolation. To build more complex segments, elements can be combined by an asymmetric **fusion** operation, where one element (or cluster of elements, cf. Coleman 1990, Kaye 1990c) is the head and one element is an operator. Differences in headship can under some circumstances result in phonetically different segments, for example $A^+ \cdot I^0$ represents / ε / while $I^0 \cdot A^+$ represents / æ /. Kaye, Lowenstamm, and Vergnaud (1985) outline a calculus that can translate elemental representations of vowels into more familiar SPE-style feature matrices.²

Central to GP’s theory of segmental structure is the notion of **charm**, which can be regarded roughly as a formal representation of sonority. The +, −, or 0 charm value inherent in each element determines which other elements and element clusters it can combine with. The overall charm value of a segment is calculated from the charm values of its component elements, and determines which other segments it can govern or be governed by in inter-segmental government relations.

¹For example, strict CV alternation with no rhymes is proposed for Semitic languages by Lowenstamm and Guerssel (in preparation), Petros Banskira (1992), etc.

²Of course, this cannot be regarded as a rigorous semantics for the element formalism, partly because there is no straightforward way to extend it to the kinds of consonant representations proposed in KLV 1990, and partly because SPE feature matrices are themselves merely formal “syntactic” devices that will ultimately need to be semantically interpreted in terms of the properties of phonetic events.

The interaction of charm values provides something of a justification for classifying GP as a violation-repairing framework. Two elements with positive charm cannot legally fuse. If morphology attempts to create such illegal fusions, the attempt can either be blocked or it can succeed, with a repair operation applying to the resulting violation. For example, in Ola's (1992) analysis of Yoruba, if an ɪ^+ tries to spread onto a positively charmed $(\text{v}^0.\text{A}^+)^+$ complex, the spreading is blocked. In Kaye, Lowenstamm, and Vergnaud's (1985) analysis of Kpokolo, if an ɪ^+ tries to spread onto a positively charmed $(\text{v}^0.\text{A}^+)^+$ complex, the attempt will succeed, but a repair operation will switch the headship of the complex, resulting in a neutrally charmed $(\text{A}^+.\text{v}^0)^0$ that can be a legal host for ɪ^+ . (As in most of autosegmental phonology where constraints are not the central object of concern, this tension between constraints as blockers and constraints as repair-triggerers is not resolved in GP.)

The notion of charm plays no role in the framework presented in this dissertation, although there is a rough and inexact correlation between positive charm and gestures with constriction degree [d:2] or [d:A], between neutral charm and [d:1] or [d:R], and between negative charm and [d:0]. I have made no attempt to capture the relationship between sonority differentials and the ability to stand in government relations such as onset-licensing, coda-licensing, or foot relations, although such relationships clearly exist. If it turns out that these relationships cannot be handled easily within the system developed so far, it may be possible to integrate a sonority measure much like GP's charm, perhaps using the mechanism of sorts.

The different assumptions concerning charm are easy to reconcile compared to differences concerning the fundamental core of segmental representation—privative elements in GP, potentially multi-valued gestural features in the present framework.

The privativity of GP elements may seem conceptually more appealing than having multi-valued features, but privativity must be bought at a price that dilutes much of the conceptual attractiveness and often amounts to covert admissions of polyvalency. For example, GP must give grammars the ability to conflate two lines or tiers that are otherwise separate. Favourite choices for this treatment are the frontness line, on which the element I^0 lives, and the roundness line, on which the element U^0 lives. In languages with no front round vowels, these two lines will usually be conflated into a single front/round line, on which both I^0 and U^0 must learn to coexist. This has the same effect as positing a single feature which can have two values (I and U). In the present framework, these cases are usually dealt with by allowing two possible site values on the root node of a vowel, [s:Pal] and [s:Vel]. Considering the general combinatorial properties that should exist, the device of line conflation is drastically underused by the world's languages. We do not seem to find languages that conflate N^+ and U^0 or h^0 and H^- . Even assuming a constraint such as “Only elements of like charm values may have their lines conflated,” the conflations that are actually observed³ are far fewer than those that are possible. It is interesting to note that the

³which, in addition to I^0/U^0 , could plausibly include $\text{A}^+/\text{cations H}^-/\text{L}^-$.

plausible confluations are exactly those that simulate a binary feature.⁴

Another way in which GP duplicates the effects of binary features is by allowing mechanisms that in effect allow the grammar to spread the absence of a feature. A look at a GP analysis of Yoruba will illustrate this point. Within the framework of autosegmental phonology, Archangeli and Pulleyblank (1989) argue that the ATR harmony in Yoruba involves [-ATR] as the active or spreading feature. This would pose a problem for GP, which has nothing corresponding to a [-ATR], only an element \mathfrak{r}^+ corresponding to [+ATR]. It is easy to use spreading to account for the presence of a \mathfrak{r}^+ on a mid vowel triggered by the presence of a \mathfrak{r}^+ on the vowel to its right, as in (7.1a), but there is no immediately obvious way of accounting for the *absence* of a \mathfrak{r}^+ when the vowel to the right has no \mathfrak{r}^+ . A naive version of GP would predict the possibility of *ek*


(7.1)

a)	ε	k	ɔ		b)	e	k	ɔ
	N		N			N		N
	x	x	x			x	x	x
	\mathfrak{I}^0		\mathfrak{v}^0			\mathfrak{I}^0		\mathfrak{v}^0
	\mathfrak{v}^0		\mathfrak{U}^0			\mathfrak{v}^0		\mathfrak{U}^0
	\mathfrak{A}^+		\mathfrak{A}^+			\mathfrak{A}^+		\mathfrak{A}^+
						\mathfrak{r}^+		\mathfrak{v}^0

This problem is discussed by Ola (1992), who uses the notion of right-to-left inter-nucleus government to explain the distribution of [ATR]. Ola argues that the presence of a \mathfrak{r}^+ in a governed nucleus must be licensed by the presence of a \mathfrak{r}^+ in the governing nucleus. One of the things that can create this inter-nucleus government relation is the operation of the OCP on the \mathfrak{A}^+ tier. In *ek* creates a government domain in which the \mathfrak{r}^+ on the left nucleus must be licensed by a \mathfrak{r}^+ on the right nucleus. Failing this, the \mathfrak{r}^+ must delink:

⁴For example, $\mathfrak{I}^0/\mathfrak{U}^0$ simulates the contrast between [s:Pal] and [s:Vel], given [a:Dor]; \mathfrak{A}^+ / simulates the contrast between [d:R] and [d:A].

(7.2) e k ɔ ⇒ εkɔ



Destroying a [+ATR] in the absence of a licensing [+ATR] is essentially the same as spreading a [-ATR] in the presence of a licensing [-ATR].


While other nonmonotonic properties of GP will be discussed in more detail shortly, it is worth pointing out here the intimate link between nonmonotonicity and segmental proposals involving privativity. It turns out that a large proportion of the nonmonotonic operations in GP analyses, for example the deletion of ɪ^+ in Ola's analysis of Yoruba, are made necessary precisely in order to circumvent the effects of privativity. For those of us to whom monotonicity is at least as attractive a property as privativity, the trade-off is not worth it. Admitting that at least some specifications can behave in a multi-valued fashion would allow GP to go a long way towards entirely eliminating nonmonotonic operations like deletion.

It is likely that GP's element theory will be unable to represent a number of possible speech sounds in a straightforward and compositional way, e.g., retroflexes, clicks, linguolabials, or contrasts between the apical and laminal coronals found in many Australian languages. This reflects a deliberate decision to prefer undergeneration of possible sounds to overgeneration (Kaye, p.c.). While the segmental model presented in chapter 2 may err slightly in the opposite direction, towards overgeneration, I believe it is a reasonable compromise between the two extremes: it is capable of expressing most of the speech sounds found in human languages without predicting many non-existent sounds (given some basic constraints that express unavoidable facts about the anatomy of the vocal tract).

7.1.2 Suprasegmental structure

The suprasegmental structure argued for in this dissertation owes a great deal to work in Government Phonology and is, if anything, an extension of that work. The most obvious borrowing is the idea of empty nuclei. While GP takes empty nuclei to be filled by a special primitive element, the cold vowel v^0 , for reasons largely to with the segmental theory, in the present framework empty positions are truly empty in the sense of not containing any segmental content at all.⁵

The most obvious departure of this dissertation from the kind of suprasegmental structure used in most GP work is the strict CV structure. GP allows both of the representations in (7.3):



while the present framework neutralizes the difference between the two in favour of the analogue of (7.3b).

Some work in GP has already moved in the same direction. Lowenstamm and Guerssel (in preparation) propose a strict CV alternation to deal with Semitic languages, even in cases that meet most of the criteria to be analyzed using codas. Gibb (1992) argues for an analysis of Finnish where what appear to be the first half of geminate consonants are in fact onsets of empty nuclei rather than coda consonants.⁶ Gibb also argues that onset-to-onset government relations are possible and desirable things to have in GP theory. If such mechanisms are independently required within GP itself, there remains little justification for a structure like (7.3a) and a separate primitive government relation of coda-licensing, when all the work can be done by a structure like (7.3b), with coda-licensing that is a species of onset-to-onset government.

⁵In cases where where GP's cold vowel is not licensed to be phonetically null, it is pronounced as a high, back, unrounded, lax vowel. In the present framework, empty positions of the sort **unspecified** receive phonetic interpretation by means of default rules, and no epenthetic vowel is given for "free". Since GP too will need default insertion principles to handle those languages where the epenthetic vowel is not a high, back, unrounded vowel (but, say, a mid schwa as in French or a front /i/ as in Yawelmani), I do not take it as a weakness that the present framework also needs default principles to interpret an empty nucleus as a high, back, unrounded vowel.


⁶A move like this is essential within traditional GP, where the possibility of a geminate consonant occurring after a long vowel in Finnish would violate the important rhyme-binarity theorem (Charette 1988) if the first half of the geminate were analyzed as a coda consonant.

7.1.3 Monotonicity


Although it comes closer to a true “no rules” approach than most generative frameworks, GP still contains substantial nonmonotonic elements. Pieces of a phonological representation can often be destroyed in the course of a derivation through virtual time. Some of these deletions are the result of a commitment to privativity and the empirical need to avoid some of its consequences (as already discussed); others are the result of the fact that GP’s implicit theory of morphology remains a firmly representation-based approach. Two types of phonological elements that most frequently get deleted are skeletal positions and segmental elements.

As an example of skeletal slot deletion, consider Kaye’s (1990b) analysis of Yawelmani vowel shortening. In order to account for the alternation in length between the verbal forms *a:miltaw* and *amlal*, Kaye assumes that the vowel *a* is underlyingly long and that the *m* and the *l* are separated by an empty nucleus whose cold vowel alternates between a phonetic realization of [i] and no phonetic realization at all, in accordance with the principles of GP. The first parts of each word would have the following structures:

(7.4) *a:miltaw*



(7.5) *amlal*



In *amlal* in (7.5), the empty nucleus between *m* and *l* is properly governed by the *a* of the suffix *-al*, and so is licensed to be phonetically null. For the shortening of the first *a*, Kaye states the triggering context as:

- (7.6) A long vowel shortens when the following nucleus is a licensed empty nucleus.

Various mechanisms for performing the deletion in the triggering context are available to GP—we can assume for the moment a brute-force statement that branching rhymes can only be licensed by a following unlicensed nucleus. The exact mechanism is irrelevant. The point is that any analysis along these lines involves a nonmonotonic deletion operation. Even within GP, this is problematic: giving the grammar the power to delete skeletal slots and thereby change the governing relations within a word seriously vitiates the content of the Projection Principle.

The alternative is that there is no deletion at all. The short/long alternation needs to be seen as a true shortening process only if we take this small piece of the language as a textbook problem, out of the context of Yawelmani's rich templatic morphology. Specifically, it is only by ignoring Yawelmani's prosodic system that Kaye was forced to assume an underlying long vowel in the first place. In fact, there are no underlying long vowels in Yawelmani verb roots—rather, long vowels are created when an inherently lengthless vowel is associated to a prosodic template. The alternation between CVV and CVC reflects the two different ways of satisfying the requirements imposed by a moraic trochee template. It is possible to analyze the short/long alternation in Yawelmani without assuming that the vowel was at any time long and underwent a shortening process.⁷ No pieces of a representation need to be destroyed. While not all cases of apparent skeletal slot deletion in GP will be amenable to exactly this analysis, they will probably all be analyzable using the tools available to a constraint-based morphophonology. (If apparent deletions are not a unified phenomenon, there is no reason to expect them to share a unified explanation.)

Segmental elements also undergo deletion in many GP analyses. We have seen an example of this in section 7.1.1 with Ola's (1992) analysis of Yoruba [-ATR] harmony. As pointed out in that discussion, the need to indulge in many of these deletions stems entirely from a refusal to admit non-privative features. Not all cases of element deletion involve circumventing privativity, however. Kaye and Harris (19xxxx) analyze several cases of lenition as involving the loss of one or more segmental elements.

While Kaye and Harris' analysis of these lenitions is attractive, I do not believe the data themselves are firm enough to justify introducing potentially unconstrained nonmonotonic devices into the grammar. In general there needs to be much stronger justification for taking these lenitions to be real synchronic phenomena. Many are clearly historical changes (and there is no reason to assume a priori that historical

⁷One of the most commonly accepted justifications for assuming that these vowels are long at some intermediate level of the derivation is that the condition for vowel lowering can then be stated simply in terms of vowel length. But the generalization can also be stated in another way: vowels are non-high when they occur in a moraic trochee imposed by a template, and hence the non-height can be seen as one of the segmental requirements imposed by the constraints that make up the templates. This would apply to vowels in either a [CVV] or [CVC] moraic trochee, without any dependence at all on vowel length per se.

change takes place by the same mechanisms responsible for morphophonological alternations in a synchronic grammar). Many analyses involve comparison of dialects and seem to be in danger of succumbing to the all too common generative assumption that if some other dialect differs from my dialect, mine must be the underlying form that they derive theirs from by synchronic processes. Without this assumption, the status of many lenition “facts” becomes questionable, as does the need for the nonmonotonic machinery necessary to derive the “facts”. I expect that all demonstrably synchronic lenitions will be explainable within the present framework—using underspecification if it is a non-gradient morphophonemic alternation, using the dual-pressure phonetic model of section 3.5 if it involves free variation conditioned by factors such as speech rate and formality.

7.2 Harmonic Phonology

Harmonic Phonology is a framework that has been developed in works such as Goldsmith (1990), Bosch (1991), Wiltshire (1992), and Brentari (1990). Like other constraint-based frameworks, Harmonic Phonology argues for a set of well-formedness conditions, known as phonotactics. If a form violates one of these conditions, it is subject to repair operations. The term “harmonic” is loosely based on the connectionist Harmony Theory developed by Smolensky (e.g., 1986).


Perhaps one of the most obvious differences between Harmonic Phonology and the framework developed here is that the former explicitly recognizes multiple levels of phonological representations. The literature proposes three levels, or rather, three *types* of level:

- (7.7) **M** Morpheme level
 W Word level
 P Phonetic level

The M level is essentially the (underspecified) underlying representation of a morpheme. The W level corresponds roughly to the work that is usually done in the lexical component of Lexical Phonology. It also remains reasonably underspecified. The P level is the level of “systematic phonetics”, where all the underspecified features of the W level are filled in. Each of the three types of levels can have different types of constraints, or phonotactics, that apply to representations in that level, though as we shall see, the evidence for constraints applying at the M and P levels is far slimmer than the evidence for constraints at the W level.

Each of the three levels may contain a number of different representations that are derivationally related. The different representations are related to each other by intra-level rules that supposedly apply only to correct violations of phonotactics. A rule mapping a W-level representation to a W-level representation is symbolized by (W,W), and there are similarly (M,M) and (P,P) rules. The final representation of

one level is related to the initial representation of another level by inter-level rules, i.e., (M,W) and (W,P) rules, although, since the mapping between levels is supposedly non-directional, there could also be (W,M) and (P,W) rules. An overall derivation in Harmonic Phonology would look like:



The representation in Harmonic Phonology that is closest to the types of PSs used in this dissertation would be the final representation of the W level. It is probably no coincidence that this is the level of representation for which the strongest evidence exists in Harmonic Phonology. The final W-level representation is the one that satisfies all the phonotactic constraints imposed by the W level,⁸ and the vast majority of constraints that have been motivated within Harmonic Phonology are W-level constraints. It would seem that the main reason for proposing the existence of M level and P level in the first place is an uncritical acceptance of the details of earlier models of phonology that don't bear specifically on the points Harmonic Phonology is trying to make: M level being based on the unexamined assumption that morphemes must obviously have underlying representations, and P level on the assumption that there must be a level of representation where phonetic facts are expressed in a phonological vocabulary.

Despite the fact that these assumptions can (and, I believe, should) be absent from a coherent theory of phonology, there would be some motivation for believing in the existence of M level or P level in Harmonic Phonology if it could be shown that, like W level, they impose phonotactic constraints on their representations, and that these constraints could not possibly be imposed at W level. I know of no arguments in favour of M level.⁹ Wiltshire (1992) attempts to provide some evidence for P-level constraints, but this evidence is tenuous compared to that amassed in favour of W-level constraints. For example, in a discussion of Malayalam coda constraints that responds to issues raised by K.P. Mohanan (1986) and Tara Mohanan (1989), Wiltshire offers clear evidence for W-level licensing conditions on codas (such as phoneme distributions). But she also claims that there are no codas at all at P level, based mostly on

⁸Intermediate representations in the W level have no effect on the ultimate surface form and seem to serve no purpose but to be replaced because they violate constraints.

⁹Some proposals for such arguments might include the apparent morpheme structure constraints on the consonant sequences of Arabic verb roots or English *sCVC* words. For arguments against such generalizations as morpheme structure constraints and in favour of the position that all morpheme structure constraints that do in fact hold can be applied "at W level", see Paradis and Prunet (to appear).

external evidence, such as language games and where native speakers will pause when dictating words one “syllable” at a time.¹⁰

There is little in Harmonic Phonology’s ideas concerning representations and constraints on representations that is unique to or original to Harmonic Phonology. Representations are borrowed wholesale from more traditional types of phonology, with the exception of the idea of the appendix, which can be traced to Charette (1984) or perhaps further. The idea of licensing conditions bears strong similarities to that used in Itô (1986) and in more recent work by Itô and Mester. What *is* more unique to Harmonic Phonology is the idea that phonological rules apply only to repair violations of constraints and therefore that these rules do not need to have any context specified.

This being so, it is surprising that work in Harmonic Phonology has devoted so little attention to the mechanics of harmonic rule application. Most of the published work in Harmonic Phonology deals with what phonotactic constraints exist in particular languages or universally, to the exclusion of how exactly these phonotactics are brought to bear on representations. I know of no systematic discussion of how exactly harmonic rules are triggered, affect representations, and interact with each other. I have been unable to find answers to such basic questions as: What happens when a representation violates more than one constraint? Which violation is repaired first, or are both repaired simultaneously? (What if repairing one violation in the usual way would as a side-effect remove the second violation? Would this situation force a bleeding order between the two rules?) What if there are two possible repair strategies in a language’s inventory of rules that could repair the same constraint violation? Which applies? (Can the Elsewhere Principle somehow extended to context-free rules to determine their order of application intrinsically?) Or is the entire situation somehow systematically prevented from arising?

The Theory of Constraints and Repair Strategies, discussed in the next section, goes a long way toward answering many of these questions.

In sum, Harmonic Phonology has some interesting ideas that can provide useful insights when applied to phenomena on a small-scale, but is not one of the best suited theories on the market for the ultimate task of writing complete phonologies for individual languages (indeed it is difficult to contemplate what such a complete harmonic grammar would look like).

¹⁰In the present framework, the W-level constraints on codas translate into constraints on what onsets can be licensed by empty nuclei. While I do not consider performance behaviour such as language games and unnaturally slow and detached speech to be part of phonology proper, the failure of this behaviour to treat nucleus+coda sequences as belonging to a single constituent is not surprising, since in the present framework they *don’t* belong to a single constituent (at least any constituent smaller than a foot). This failure to treat as a constituent does not need to be explained by physically moving coda consonants out of one syllable and into the onset of the next syllable.

7.3 The Theory of Constraints and Repair Strategies

One research programme that has dealt with the issues surrounding repair operations is the Theory of Constraints and Repair Strategies (TCRS), which has been developed in work such as Paradis (1988a,b, 1990, 1993), Paradis and El Fenne (1993), and La Charité (1993).

Like Harmonic Phonology, TCRS recognizes roughly three types of “levels” of phonological representation: underlying forms stored in the list of the language’s morphemes (called DICT), levels relevant for lexical phonology, and levels relevant for post-lexical phonology. Also like Harmonic Phonology, TCRS holds that constraints on phonological form can be violated by non-final levels of representation. Violations trigger repair operations that, because of their external triggering, do not have to have any context specified. TCRS differs from Harmonic Phonology in the types of constraints that have been the focus of attention and in its more fully worked-out set of hypotheses on how repair operations work.

TCRS constraints can be of two types: principles and parameters. Principles are universal constraints that define what is possible or impossible in all languages (e.g., the OCP, the principle of prosodic licensing). Parameters are set by individual languages to permit or prohibit structures that are at least possible cross-linguistically. Parameters deal with phonological content or with phonological structure. Content constraints determine such things as which groups of features are allowed to cooccur, which segment sequences are allowed, which features can spread, and what the default segments are. Structure constraints determine things such as what configurations are allowed (e.g., diphthongs, complex onsets, geminates) and what are possible syllables and metrical structures. TCRS holds these parameters to be universal, that is, language will have a setting for each one. Some examples of content constraints in the language Fula from Paradis (1988b) are:

(7.9) The first part of a fused segment must be more sonorous than the second part: (on)

(7.10) velar: [+continuant] [+voiced] [-round]: (off)

(7.9) allows prenasalized segments and falling diphthongs like *aU* and *oI* while banning post-nasalized segments and rising diphthongs like *Ua* and *Io*. (7.10) prohibits the segment / γ /.

Constraints will specify which part of the content or structure is its **focus**. The focus of a constraint is one of the things that determines which repair strategy will be used to repair violations of it (cf. Paradis 1990). Constraints will also specify their **domain**, e.g., lexical, post-lexical, stratum 2, everywhere.

Some examples of configurational constraints are:

$$(7.11) \quad \begin{array}{ccc} \text{Nucleus} & (\text{on}) & \text{C} & (\text{off}) & \text{Onset} & (\text{off}) \\ \wedge & & \wedge & & \wedge & \\ \text{x} & \text{x} & \text{x} & \text{x} & \text{x} & \text{x} \end{array}$$

The focus of these three constraints is the syllable. The following constraints from Gere have different foci:

$$(7.12) \quad \begin{array}{ccc} & & \sigma_2 & & \\ & & | & & \\ \text{x} & (\text{on}) & \text{x} & (\text{off}) & \text{x} & (\text{off}) \\ \wedge & & \wedge & & \wedge & \\ \text{V} & \text{V} & \text{V} & \text{V} & \text{V} & \text{V} \\ \text{focus: skeletal} & & \text{focus: segmental} & & \text{focus: metrical} & \end{array}$$

TCRS proposes that repair strategies come from a universal inventory of simple context-free insertion and deletion rules. Repairs are subject to the Minimality Principle:

- (7.13) Minimality Principle: A repair must apply at the lowest phonological level at which the violated constraint it preserves refers.

The lowest level is determined according to the usual phonological hierarchy: metrical structure, syllables, skeleton, root nodes, non-terminal features, terminal features.

Repairs are the only type of purely phonological rule—all others must be morphologically conditioned. Repairs cannot create a constraint violation, unless forced to do so by a conflict between two or more constraints. So constraint violations can only occur in three circumstances: if created by a morphological process; if created by a repair strategy operating under a constraint conflict; or if the ill-formed structure was present in the underlying form of the morpheme in DICT.¹¹

Paradis (in press) argues for the need to allow morphemes to be underlyingly ill-formed. The ill-formedness must be corrected as soon as the form enters a stratum in which the relevant constraint applies. By the time this happens, the phonological and morphological context may be quite different for the same morpheme in two different words, causing two different repair strategies to apply, and it would seem as though the morpheme had two allomorphs.

Where workers in Harmonic Phonology spend most of their time addressing the question of what a possible constraint is and little on the question of how they apply, TCRS does the reverse. There is little idea given as to what a possible constraint is. Since constraints are supposed to be universal (i.e., universally available to languages as a parametric choice), this is to a large extent an empirical question. But implementing

¹¹Another possible source of constraint violation, loanwords, has received considerable attention from researchers in TCRS recently. (See, e.g., Paradis, Lebel, and LaCharité 1993.)

the idea may not be quite so straightforward. One might, for example, wonder if it makes much sense to say *[+high] and *[-high] are both universal parameters. Another problem may be the sheer number and complexity of the constraints. For example, Paradis (1988b) proposes a metrical-focus constraint for Fula banning the following configuration:

$$(7.14) \quad \begin{array}{c} \sigma_{2-3} \\ | \\ x \quad x \\ \vee \\ C \end{array}$$

That is, a consonant may not be doubly linked if the first skeletal slot it is linked to belongs to the second or third syllable of the word. A constraint this complicated raises serious questions about whether there are any reasonable limits at all on what a constraint can look like. It is of course always possible to claim that constraints like (7.14) are nevertheless made available as universal parameters. But if it turns out that we would need tens of thousands of similar “parameters” to handle phenomena in all the languages of the world, it is not clear that the task facing a language learner in determining all the settings would be any easier than if the learner had to induce all the active constraints from scratch.

One of the most positive aspects of TCRS is that it argues explicitly and tries to present evidence for a central hypothesis of generative grammar that has for several years usually been blindly assumed rather than argued for, namely, the hypothesis that the features that control allomorphy are the same features that define the content of segments.

Much allomorphy can be seen as additive. Each morpheme has certain minimum requirements that must be met by all representations that instantiate it. In addition to these minimum requirements, some forms require extra material to be present in some contexts. But not all cases of allomorphy can be handled perspicuously this way, the most obvious examples being cases of complete suppletion (e.g., *go~went*).¹² Somehow the grammar must contain information concerning the form of the allomorphs and their context, and we can use the term **diacritic** in a very general way to refer to this information.

There are essentially three ways of implementing these allomorphy diacritics. The first approach, common in much of generative phonology, assumes that the diacritic information about allomorphy is implemented by means of special diacritic features devoted specially to that information. The features are usually given mnemonically convenient names such as [+en], [- Rule 23], or [+defective]—though [+green], [Fred],

¹²These could be seen as additive in a trivial sense: the minimum requirements imposed on all forms of the lexeme GO is the empty set; the extra requirements imposed by individual forms are responsible for all the observed phonological properties.

and [-interesting generalization] would work just as well. Normal generative rewrite rules that are sensitive to the presence of these diacritic features in a lexical entry will spell out the abstract morpheme with the underlying representation of the correct allomorph, or will transform the representation of the “basic” allomorph into that of one of the other allomorphs, e.g.:

$$(7.15) \quad \emptyset \rightarrow en / \left[\begin{array}{l} +V \\ +en \end{array} \right] \text{ —————}$$

Clearly, this approach does not say anything particularly interesting about the nature of diacritics and suggests no ways in which the power of a system using diacritics might be restricted.

The second approach to the nature of these diacritics is that they simply state what the varying properties are and where they occur. This is essentially the approach being taken if someone proposes two separate lexical entries for a morpheme (i.e., two different underlying representations) with conditions on their use. It is also the approach of the present framework, where non-additive allomorphy is handled by constraints of the form:

$$(7.16) \quad \textit{if lexeme} = L \textit{ then} \\ \text{PS constraints common to all forms of L} \\ \textit{and} \\ \text{(context for allomorph 1 and PS constraints of allomorph 1)} \\ \textit{or} \\ \text{(context for allomorph 2 and PS constraints of allomorph 2)}$$

The shape of the allomorphs is controlled using exactly the same vocabulary as that used in controlling the shape of any morph of the language, and the description of the contexts also uses the same vocabulary. In this approach, allomorphy diacritics are composed of the same “stuff” (descriptions) as the rest of the grammar. As in Natural Generative Phonology, the descriptions that constitute the morphemes are all surface-true.

The third approach is similar in that it holds diacritics to be made of the same “stuff” as other parts of the grammar, but differs in exactly what that stuff is, namely, pieces of representation and rewrite rules. The features that ultimately control the distribution of allomorphs are the same features you would find anywhere in a phonological representation: [+lateral], [-back], and so on. In SPE-style phonology, these features can be used in an entirely arbitrary way to trigger or block the application of extrinsically ordered rewrite rules. A language that has no distinctive laterals might nonetheless use the feature [+lateral] as a diacritic, more or less along the lines of [+green] or [Fred], to prevent some words from undergoing a rule they would otherwise be expected to undergo (by putting [-lateral] in the structural description of the rule), and then delete the [+lateral] in an absolute neutralization operation.

TCRS is an example of the third type of approach, but has added the teeth necessary to make the approach interesting. Essentially diacritic features or feature combinations cannot trigger just any arbitrary phonological rewrite rule. The most the diacritic can do is to violate a constraint of the language, triggering a repair whose operation is largely beyond the control of the individual grammar. Paradis' (1993) proposal for the abstract segment / γ / in Fula is an example of how the diacritic mechanism of TCRS works.

Since the mid-1970s, when for a large number of North American phonologists the necessity of abstract (surface-untrue) representations became a background assumption rather than a position that needed defending against the proposals of Natural Generative Phonology, many of the original phenomena that were used to argue for abstractness have received interesting alternative explanations using technologies such as underspecification. TCRS's work in trying once again to justify abstract representations is exactly what is needed if the idea is to continue to be one of the assumptions of generative phonology. While TCRS may not succeed¹³ in establishing the need for surface-untrue underlying representations with derivations in virtual time, the effort is vital in clarifying precisely what the central theoretical differences are between TCRS and declarative approaches to phonology and in finding data that can bear on deciding between them.

7.4 Autolexical Syntax

Autolexical Syntax, discussed most fully in Sadock (1991), is a theory of parallel representations that bears close similarities to the framework proposed in this dissertation.

Sadock argues that linguistic objects consist of tuples of representations, one from each module, e.g., morphology, syntax, semantics, much like the sign structures used here and in frameworks such as Head-Driven Phrase Structure Grammar (à la Pollard and Sag 1987).¹⁴ Sadock also argues strongly that these representations are not organized in a hierarchical manner, or perhaps more accurately a manner based on virtual time, as they are in most versions of generative grammar, where the output of one module is fed into the next module. Rather, each type of representation has an existence that is independent of the other types of representations and is subject to the constraints of its own module.

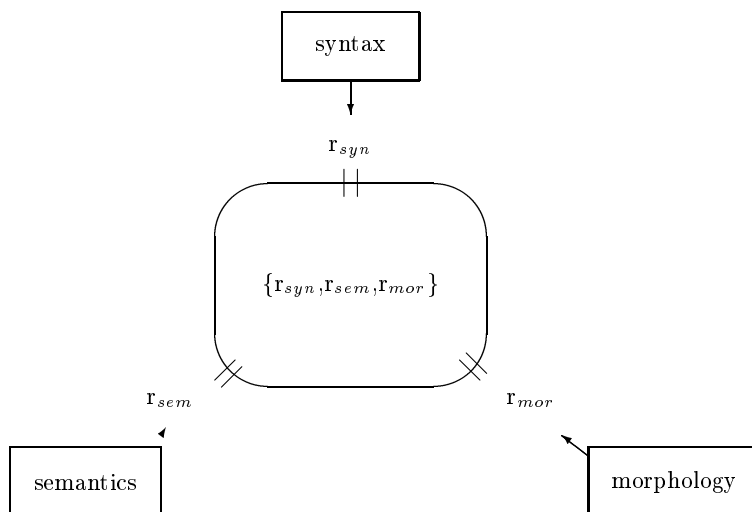
This argument for the independence of different types of representations is not a claim that any syntactic representation can be paired with any phonological representation, for example, pairing the syntactic structure [[Adj Adj N] [[V Adv]] with the phonology /ai heit kæts/. It is not enough for each representation to be acceptable to

¹³and, as one may surmise after the preceding six chapters, I suspect it will not

¹⁴Further research in the autolexical framework has suggested the addition of a number of new modules, for example, an illocutionary module, a pragmatic module, and separating morphology into a morphosyntactic module and a morphophonological module.

its own module. In the tuple of representations, $\{\Gamma_{syn}, \Gamma_{sem}, \Gamma_{mor}\}$, each r_i must conform to the constraints imposed by its module, but the entire tuple must also conform to a number of constraints that make up the **interface**. Sadock (1991: 20) diagrams the situation as:

(7.17)



Autolexical Syntax is exactly like the present framework in arguing that this interface is not a linear ordering of modules. A syntactic representation and a morphological representation are related in a logical way by constraints, not by the fact that the output of syntax was fed as input into the morphology.

As in the present framework, in Autolexical Syntax lexical entries play a central role in the interface between representations. Some typical English lexical entries look like:

(7.18) *Fido*:

syntax = N[2]
 semantics = $[Q[Q[-1]DEF][F[F[-1] ____] x]]$
 morphology = N[-1]

(7.19) *want*:

syntax = [SF5]
 semantics = O^{-2}
 morphology = V[-0]

The entry for *Fido* imposes requirements on the syntax, semantics, and morphology of any utterance it occurs in: it relates a noun phrase (an N with bar level 2) in the syntax to a one place function (under the scope of a definite operator) in the semantics, to a noun in the morphology. Similarly, the entry for *want* relates a verb in the syntax¹⁵ to a two place operator in the semantics to a verb stem in the morphology. One possibility is that a lexeme that may impose requirements on some modules but not on others, for example, it may impose requirements on syntax and semantics but no requirements on the morphology (I have argued that this is the situation underlying what others have referred to as “zero morphemes”).

Autolexical Syntax conceives of the operation of the lexicon slightly differently than the present framework does. In Autolexical Syntax, all terminal nodes in all representations in a tuple must be “lexicalized”, that is, each terminal node must be licensed by a lexical entry, which is conceived of as having the lexical entry attached to the terminal node. Any terminal nodes in other representations connected to the terminal node in question must be attached to the same lexical entry. The difference can be summed up as follows: in Autolexical Syntax, each terminal node must be satisfied (by a lexical entry); in the present framework, each lexical entry must be satisfied (the overwhelming majority will be satisfied vacuously).

While many of the central ideas of the theories are strikingly similar, there are of course some differences in detail. The most obvious is probably in the choice of what types of representation are held to exist. In addition to the syntactic, semantic, and morphological representations just discussed, Sadock and his co-workers have proposed modules to deal with phonology,¹⁶ illocutionary force, and some other pragmatic factors. While I remain agnostic about the map of linguistics in semantics and beyond, I have argued explicitly in this dissertation that there is no need for an independent level of representation for morphology.

Some of my reasons for not believing in a morphological component were outlined in the summary of chapter 5, where the existence of the very subject matter for morphology, morphological words, was questioned. There are fairly clear grounds for identifying phonological “words” (i.e., constituents of a certain level in the prosodic hierarchy) and for syntactic “words” (i.e., X^0 categories), but the evidence for the existence of some type of object that may be identified as a morphological “word” is slim. While such a creature may make some analyses marginally more convenient, I do not believe there is sufficient justification for introducing an entirely new type of entity into linguistic theory. Much of the work given to the morphology module in Autolexical Syntax could probably be reassigned to phonology, once a theory of the

¹⁵specifically a verb with the subcategorization feature [SF5], which, as in GPSG, controls which Immediate Dominance phrase structure rule holds for the verb, in this case the rule: $V[1] \rightarrow V[0, SF5] V[1, [to]]$.

¹⁶though a detailed proposal on the nature of the phonological module has not yet emerged.

phonological module is developed.¹⁷

A second major difference, is the one that justifies including Autolexical Syntax together with Optimality Theory as frameworks which can under certain circumstances ignore constraint violations. The constraints in question are certain principles governing the interface that are referred to as “default” interface principles in Sadock and Schiller (1993). Certain constraints will ordinarily govern how representations from different modules may be associated to each other (e.g., the hierarchical or linear order should be the same in both), but these defaults can be ignored when required by individual lexical entries. See Sadock and Schiller (1993) and the references therein for a fuller discussion of the default principles and some ways they can be violated. In the present framework, all interface constraints must be hard. The Sister Alignment Constraint of chapter 6, for example, must be obeyed by all representations in all languages. This is possible because the Sister Alignment Constraint is very general, it doesn’t demand much. Individual languages may impose more stringent conditions on sister alignment for some lexical items or constructions, but these conditions are consistent with the more general universal one, they do not override it.

7.5 Optimality theory

Optimality Theory, recently proposed in work such as Prince and Smolensky (1993) and McCarthy and Prince (1993a,b), is another framework that has the notion of constraints as one of its central ideas. It shares with the present framework a rejection of the derivational approach to building phonological representations (as still practised in Harmonic Phonology and TCRS), replacing it with the idea of a set of constraints choosing the well-formed representation from the infinite set of candidate representations.

A major difference between Optimality Theory and declarative approaches to phonology is that its constraints are not hard, i.e., they can under some circumstances be violated without causing the rejection of the representation as ill-formed. Each language takes the constraints given to it by Universal Grammar and ranks them in a hierarchy. Candidate representations are first evaluated by the highest-ranked constraint of the hierarchy. If all but one of the candidates violate this constraint, that one candidate is the well-formed representation. Otherwise, the set of candidates that violated the constraint the least (note that this does not necessarily mean they did not violate it at all) is passed on to the next constraint in the hierarchy, there to be judged again. The process is repeated down through the constraint hierarchy until one candidate remains (or until there are no more constraints, in which case all surviving

¹⁷For example, Sadock (1991) argues, and I agree, that there is no good reason for most cases of cliticization to be analyzed as forming a single constituent with their heads *in the syntax*. But it does not follow that the level at which they *do* form a single constituent must be morphology rather than phonology.

representations should be equally well-formed).

One of the claims of Optimality Theory is that every constraint is present in every language. For example, the universally supplied constraint ONSET is present in the constraint hierarchy even of a language that clearly allows syllables without onsets—it is just that in this language, ONSET is ranked so low on the hierarchy that higher constraints have already narrowed the candidate set down to one before it ever has a chance to apply. (McCarthy and Prince, and Prince and Smolensky, point to cases where low-ranked constraints can influence the well-formedness of two otherwise equally well-formed representations in some environments, even if the constraint never gets to apply in most words and is flagrantly violated by several words of the language.)

There are some problems with McCarthy and Prince's proposals. One problem is the lack of explicitness in the characterization of many of the constraints, subjecting the theory to the difficulties of interpretation discussed in section 1.5.1. Concerning the proposed constraint ALIGN:

$$(7.20) \]_{stem} =]_{\sigma}$$

McCarthy and Prince (1993: 36) admit: “To be fully accurate, the statement should explicitly mention the universal quantification over stem-edges and the existential quantification over syllable edges, ... but the concise statement in [(7.20)] is more memorable.” This is perfectly acceptable for ALIGN, where there is apt to be little confusion, but few of their constraint definitions are as concise or as memorable. Most are simply phrased in English, using a rich technical vocabulary that is seldom explicitly defined, and leaving out some crucial information about how the constraints are supposed to apply.

Perhaps the most obvious example of a crucial aspect of constraint application that does not get formalized is whether the constraint can be violated several times per form or only once.¹⁸ McCarthy and Prince argue that many constraints can be violated more than once. A word with two onset-less syllables violates ONSET more than does a word with only one onset-less syllable. A word with four epenthetic segments is a more serious violation of FILL than a word with only three. But there is no discussion of which constraints can have multiple violations and which cannot. Though it is not immediately apparent, their analysis of Axininca Campa requires at least one constraint that is crucially “all or nothing”, namely, SFX-TO-PRWD. SFX-TO-PRWD requires the left edge of a suffix to coincide with the right edge of a prosodic word.¹⁹ From their application of this constraint to the case where the suffix *-aanc^hi* to the stem *na-*, we can infer that SFX-TO-PRWD can only be violated once. In Axininca Campa, the minimal prosodic word is bimoraic. In the following two candidate forms,

¹⁸This is a different distinction than the binary/non-binary distinction made in Prince and Smolensky (1993: 68–73).

¹⁹It is difficult to determine if this constraint is supposed to apply to *all* suffixes, or only to some suffixes, and if so, which ones.

capital T and A represent epenthetic segments,] the right edge of a prosodic word, and | the left edge of the suffix:

- (7.21) a) naTA] T|aanc^hi
 b) na T|aanc^hi]

In order for McCarthy and Prince's analysis to go through, it is crucial that the violation in (a), where the left edge of the suffix follows the prosodic word boundary with one intervening segment, count as an exactly equal violation to the one in (b), where the suffix edge precedes the prosodic word boundary with five intervening segments.²⁰ If the two violations are exactly equal, SFX-TO-PRWD cannot decide between the two candidates and the decision is passed to the next constraint in the hierarchy, FILL, which will correctly choose (b), using its ability to count epenthetic segments, an ability that cannot be shared by SFX-TO-PRWD. Another implication also strongly suggests that McCarthy and Prince would not want to allow SFX-TO-PRWD to be violated more than once. In a case where both a three-segment suffix -xxx and a four-segment suffix -yyyy were subject to a counting version of SFX-TO-PRWD, candidate (a) must be chosen over candidate (b):

- (7.22) (a) PrWd-xxx-yyyy
 (b) PrWd-yyyy-xxx

since (a) incurs only three violations (suffix -yyyy is separated from the prosodic word by three segments), while (b) incurs four violations (suffix -xxx is separated from the prosodic word by four segments). In other words, we should expect to see morphological systems where the only factor determining the order of a number of affixes was their phonological length. I am aware of no such systems. In order to avoid having the theory make this prediction, SFX-TO-PRWD should be an "all or nothing" constraint.²¹

The problem is not the fact that SFX-TO-PRWD must be an all-or-nothing constraint, but the fact that this is not made clear in the presentation and that the "formalization" of the constraint makes no mention of this crucial aspect of its mode of application. There is no distinction made between proposed constraints that must be all-or-nothing and those that must be multiply-violable.

Leaving aside technicalities of formalism and turning to more substantive matters, there remain some significant differences between Optimality Theory and the framework developed here. One of the most profound of these lies in what the nature of morphemes is taken to be, and more generally in the role of morphology in the grammar. Optimality Theory stops short of being a completely constraint-based theory, incorporating the constraint hierarchy as just one stage of an otherwise serialist

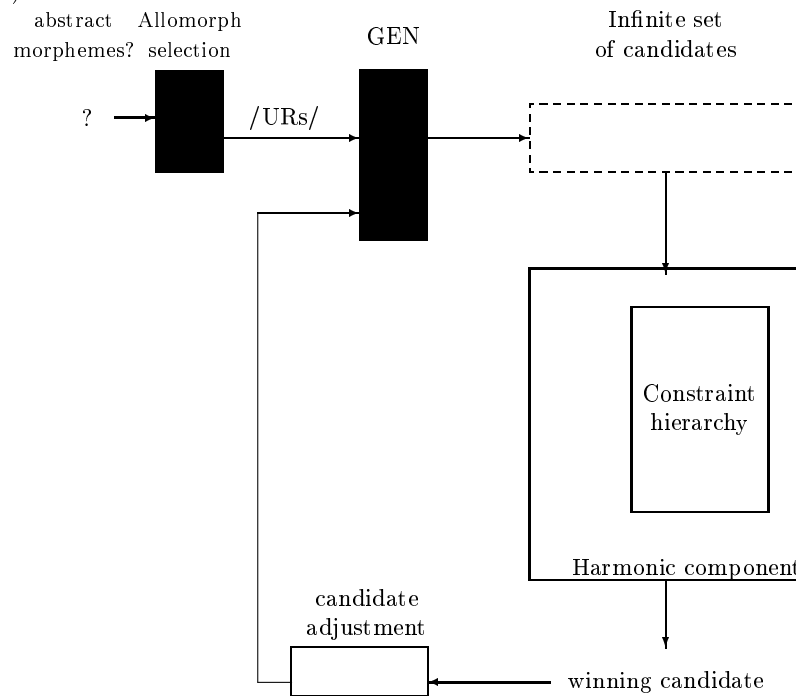
²⁰I thank Donca Steriade for pointing this out.

²¹Interestingly, a similar prediction would probably follow from the constraints LEFTMOSTNESS and RIGHTMOSTNESS, which McCarthy and Prince *do* want to be multiply-violable.

derivation. Morphemes are still seen as underlying representations that consist of a discrete piece of phonological stuff. These pieces are fed into a function called Gen (or Generate) that has as its output an infinite number of candidate representations based more or less on the morphemic input. The candidates produced are then fed into the harmonic component where they run the gauntlet of the constraint hierarchy; only one candidate emerges out the other end. But the output of the constraint hierarchy is the output of just one level of a multi-levelled phonology (analogous to the strata of Lexical Phonology). After various further adjustments are made to the successful candidate, such as stray erasure of unparsed segmental material, it then has the possibility of looping back to be fed again into Gen, this time together with some new affixes. The authors propose no restrictions on the number or nature of these levels.

The overall architecture of the theory can be diagrammed as:

(7.23)



McCarthy and Prince touch on some of the considerations of selecting the right version of the underlying representation, the right allomorph, to feed into Gen given various aspects of the morpheme's environment. The mechanism by which this would be done is not made clear, nor is it self-evident that it could be done without access

to information that may not be available until after the candidate set has been determined by Gen, or indeed until after the winning candidate has been chosen by the harmonic component. (Consider, for example, the discussion in McCarthy and Prince (1993: 110–112) on how the marked version of the Dyirbal ergative suffix – be prevented from occurring anywhere except on disyllabic vowel-final nouns, and the striking absence of discussion on the opposite problem: how the default ergative suffix –*gu* could be prevented from occurring on disyllabic vowel-final nouns.)

It is hard to tell exactly what unavoidable serialist aspects this architecture has retained from previous frameworks, largely because so much of the architecture is left undefined or only sketched. While the behaviour of the constraint hierarchy is clearly spelt out in Prince and Smolensky (1993) and McCarthy and Prince (1993), Gen essentially remains a black box of indeterminate internal structure and function. McCarthy and Prince list several desiderata that Gen would have in a full account, but it is not immediately apparent how these could be achieved. Some of these desiderata are themselves problematic, their primary motivation being to enforce a tacit version of what I have referred to as the “Physical Integrity of Morphemes” hypothesis. But, in a blow to the possibility of a restrictive overall architecture for Optimality Theory, McCarthy (1993) argues that language-particular rules must be able to influence the operation of Gen to induce it into producing candidates that it would not have produced in its normal mode of operation.²²

Most of the remnant serialism of Optimality Theory and most of the problems raised by Gen result from uncritically carrying over from earlier work a representation-based theory of the nature of morphemes. The architecture needs some way to turn

²²The desiderata for Gen’s operation include the following three properties:

- a. **Freedom of Analysis:** Any amount of structure may be posited.
- b. **Containment:** No element may be literally removed from the input form.
- c. **Consistency of exponence:** No changes in the exponence of a phonologically specified morpheme are permitted.

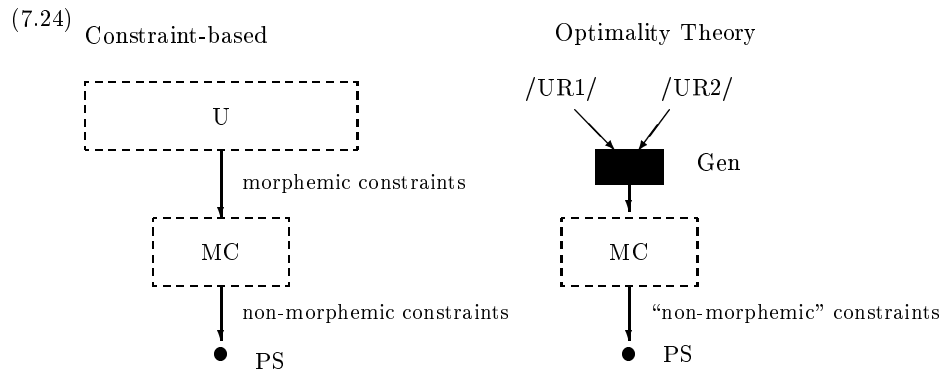
The problems with the Gen architecture discussed here and in the remainder of the section are not just abstract considerations of who has the most aesthetically pleasing theory. There are empirical consequences of the choices McCarthy and Prince have made. Specifically, the principle of Containment is phrased and interpreted in such a way as to disallow coalescence across morpheme boundaries. For example, when an *m* and a *p* become adjacent across a morpheme boundary in Axininca Campa, Gen is supposed to be unable to produce an analysis where their two Labial nodes have merged into one, which would have resulted in a legal coda-onset structure. Since this option is unavailable, epenthesis is forced.

If we restrict our attention to Axininca Campa, this is a good thing, since it is what Axininca Campa is empirically observed to do. Unfortunately, McCarthy and Prince have to use a “universal” principle to get this language-particular result, and in the process render the universal principle (and Gen as a whole) utterly incapable of dealing with the equally, if not more, numerous languages where such merger *does* take place. Gen’s normal method of operation would be unable to deal with the Nisgha coalescence phenomenon discussed in chapter 5, and, yet again, ad hoc language-particular rules would have to manipulate the candidate set directly.

discrete chunks of phonological material (assumed to be the underlying representations of morphemes) into an infinite set of candidate representations. This motivation for Gen disappears if we no longer assume that morphemes are pieces of representation.

In the present framework, morphemes are constraints like any other constraints, and themselves perform the work of eliminating those candidate representations that do not conform to their demands. In Optimality Theory, Gen must be brought in as a way of failing to generate those candidates in the first place. This is a two-step process. The fact that some conceivable representation is not a legal instantiation of a set of morphemes could be traced to either of two failures: it was not in the output of the constraint hierarchy (because it was not the most harmonic input candidate) or it was not in the output of Gen (for some reason yet to be determined).

The following diagram might serve to illustrate the difference between the two approaches. *U* stands for the universe of all possible representations, *MC* for the class of those representations that are consistent with the morphemes in question. *PS* is the candidate that is the actual phonological representation:



In the present framework, *U* is cut down to *MC* by the constraints that *are* the morphemes in question (and vacuously by all other morphemes of the language).²³ In Optimality Theory, *MC* must somehow be generated constructively from the pieces of representation that serve as morphemes. *U* plays no role in Optimality Theory.²⁴ If the *MC_i* that would normally be produced by Gen does not contain enough candidates (i.e., does not contain the actual empirically observed form), ad hoc language particular

²³Of course, morphemic and non-morphemic constraints have been separated in the diagram for comparative purposes only. There is no sense in which the morphemic constraints apply before the non-morphemic constraints, and the set *MC* plays no role in the framework.

²⁴This is not entirely accurate. For McCarthy and Prince, Gen does in fact produce *U*, the set of all possible representations, as output when faced with an empty input marked as being reduplicative. (Or, at least, removing the non-reduplicative parts of the generated candidates would leave behind the set *U*.)

devices, such as the rule proposed by McCarthy (in press), are needed in order to expand MC_i . These problems are not encountered by an approach that takes U as its starting point.

Despite these problems, it would have been a point in favour of the Gen architecture be if it managed to segregate universal from language particular information, the former living in the constraint hierarchy and the latter living in the lexicon from which the underlying representations are drawn (as well as in the language-particular ranking of constraints). Even within the analyses offered by Prince and Smolensky (1993) and McCarthy and Prince (1993), this type of segregation will not be possible. Several constraints that play a role in these analyses apply only to a handful of morphemes in a particular language, or even to a single morpheme. A typical morpheme-sensitive constraint can be found in McCarthy and Prince (1993: 106):

- (7.25) AFX-TO-FT (Ulwa)
Base of 'possessive' is foot

This type of constraint severely compromises the ideal that the hierarchy should consist of constraints supplied by Universal Grammar and that the only language-particular aspect should be the ranking. The very concept of universality is made vacuous if we must believe that there could be “universal” constraints like “Base of ‘third person dual feminine/zoic indirect object in a subjunctive subordinate clause’ is foot.”²⁵

Prince and Smolensky (1993: 101) seem aware of the problem, though not of its full seriousness. Discussing a constraint proposed for Lardil:

- (7.26) FREE-V
Word-final vowels must not be parsed (in the nominative)

they write:

Although FREE-V takes the bull by the horns, it would not perhaps be put forth as the canonical example of a universal markedness principle... Any theory must allow latitude for incursions of the idiosyncratic into the

²⁵The problem is especially acute with reduplication. There are many languages where a single abstract morpheme, like ‘plural’, is instantiated by means of two or more different reduplicative patterns, each obeying a different set of prosodic constraints. I do not see how members of the constraint hierarchy could be made sensitive to the difference between the patterns. Reference to the abstract morpheme, as is apparently being attempted in (7.25), would not work, since the abstract morpheme is the same for all the patterns. Reference to the allomorphic underlying representations fed into Gen would not work either since, by the hypotheses of McCarthy and Prince, every pattern would have exactly the same underlying representation—an empty morph whose emptiness is what triggers reduplication in the first place. What would be needed is some kind of abstract diacritic feature, and, in an ideal world, a theory of the possible diacritics and an explicit account of the ways in which they may legally affect the application of the constraint hierarchy.

grammar. What is important for our program is that such incursions are best expressible as *constraints*; that they are (slightly) modified versions of the universal conditions on phonological form out of which core grammar is constructed; and that they interact with other constraints in the manner prescribed by the general theory.

Here the authors express pride in the homogeneity of universal and language-particular information, momentarily forgetting that one of the cornerstones of the theory was supposed to be the heterogeneity of the universal and the particular and their segregation into qualitatively different parts of the architecture and qualitatively different types of formal substance. The position of McCarthy and Prince (1993b) that ALIGN is a universal constraint schema rather than a universal constraint is another way of denying this original claim, and in doing so removing one of the last possible points in favour of the Gen architecture. The problem is not the homogeneity of universal and language-particular information, but the initial assumption of heterogeneity.

I believe that the most sensible solution to the problems of the Gen architecture is the one proposed in this dissertation for hard-constraint-based systems, namely that morphemes literally are constraints. The prosodic constraint on Ulwa possessives is part of the Ulwa possessive morpheme. The morpheme-particular prosodic constraint can interact with universal constraints in interesting ways, but this does not make it different from the rest of the information that makes up the possessive morpheme (the information that McCarthy and Prince want to reify into an underlying representation that would serve as Gen's input), since this information is also framed in constraints that can potentially interact with universal constraints in interesting ways. I agree fully with Prince and Smolensky that incursions of the idiosyncratic into the grammar are best expressible as constraints and would simply argue that the problems of the Gen architecture result entirely from not taking this observation seriously enough.

7.6 “Declarative” phonology

There has been an increasing amount of work lately in approaches to phonology that are based on hard, inviolable constraints. For lack of a better term, we can call these approaches “declarative phonology”. In this section, I shall look at only two works, the dissertations of Bird (1990) and Scobbie (1991), since these are two of the first sustained and detailed attempts to rework phonology within a declarative framework. Other examples can be found in the papers contributed to the volume edited by Bird (1992), in Coleman (1992), Bird and Ellison (1992), and increasingly throughout the computational linguistics literature.

Declarative phonology approaches vary greatly in detail. While they all aim to characterize legal phonological representations by means of constraints or descriptions, authors differ widely in what they propose to be the properties of the phonological representations that the constraints are applied to. The differences in the proposals reflect

the differences found in the wider field of phonology: the nature of segmental content, the inventory of features, hierarchies of dominance or association, syllabic structure, prosodic structure, and so forth. Some choices for representational properties clearly make the job of a constraint-based framework easier, but the fact that thoroughgoing constraint-based frameworks can still work independent of many of these choices is a promising sign. Another difference, the choice of how to represent the property of temporal precedence, will be more fully discussed in the sections on Bird (1990) and Scobbie (1991).

Despite the differences, declarative frameworks share many similarities beyond a common commitment to non-derivational approaches to morphology. One theme shared by much of the work in the area is a renewed sensitivity to and interest in the role of phonetics. Most constraint-based approaches take seriously the idea that a phonological representation is not an exhaustive embodiment of all the sound-related properties of an utterance. Phonology is partial. A much greater role is assigned to the phonetics and to the phonology-phonetics interface in fully interpreting a partial phonological structure. Much of the work that has often been assumed without argument to be the results of phonological rules is instead assigned to principles of phonetics or the interface. (For example, intrusive stops as in the English /prɪns/ [prints], often attributed to a phonological rule, are argued by Bird (1990) to be the result of phonetic fill-in. Cf. Browman and Goldstein’s discussions of overlapping gestures.)

Another trait shared by much of the work in declarative phonology is its close relationship with unification-based approaches to syntax, such as GPSG, HPSG, LFG, and so on.²⁶ Often this relationship is explicitly foregrounded, as with Bird and Ellison’s attempts to devise a theory of phonology within HPSG to deal with the PHONOLOGY value of HPSG signs. Even where the declared intent is not to contribute to some existing unification-based approach to grammar, it is common for the ideas and formalisms of these approaches to be borrowed in a highly visible manner. Scobbie (1991), for example, makes extensive use of the attribute-value matrices common in unification-based syntax.

Because of the intimate relationship between declarative approaches to phonology and unification-based syntax, before discussing Bird (1990) and Scobbie (1991) I shall briefly sketch some of the relevant shared ideas: the formalism of attribute-value structures, the unification operation, and the relation of unification to (the conjunction of) constraints.

7.6.1 Attribute-value structures and unification

It is probably easier to understand what unification does if we start with some examples than by plunging straight into a definition. Unification operates on a type of representation known as **feature structures** or **attribute-value structures** (AVSs). Since

²⁶Recent versions of most “unification-based” theories of syntax, most notably the HPSG of Pollard and Sag (1993), are also better described by the term “constraint-based”.

the term “feature structure” would probably be confusing in a subdiscipline where “feature” already has several meanings, I shall follow workers such as Johnson (1988) in referring to these objects as attribute-value structures.

To get an idea of what an attribute-value structure (AVS) is like, let’s look at the following example:

(7.27)

NAME:	Elizabeth	
NUMBER:	1	
BORN:	[YEAR: 1533 MONTH: DAY: PLACE: Greenwich]
DIED:	[YEAR: 1603 CAUSE:
FATHER:	[1]	[NAME: Henry NUMBER: 8]
MOTHER:	[NAME: [FIRST: Anne LAST: Boleyn] DIED: [YEAR: 1536 PLACE: London CAUSE: [1]]]

About the basic structure of the AVS, we notice first that there are several constants or atoms: *Boleyn*, *8*, *Greenwich*, *1603*, *NAME*, *CAUSE*. We also have **attribute-value pairs**, with the name of an attribute to the left of a colon and the value for that attribute to the right: *NAME* is an **attribute**, *Elizabeth* is its **value**. There are complex structures: sets of attribute value pairs, enclosed in square brackets (these being sets, the order of the pairs is irrelevant). The value in an attribute-value pair may be either an atom or a complex structure.²⁷ Simply, an AVS is one of these complex structures—it is a set of attribute-value pairs, which may recursively contain other AVSs.

A couple of general points about AVSs are brought out in (7.27). First, it is possible for an attribute in an attribute-value pair to be missing its corresponding value, e.g., *MONTH* in the AVS that is the value for *BORN*. (Using the terminology introduced in chapter 3, we can also say the value of the *BORN*|*MONTH* **path** is missing). Technically, in most unification-based frameworks, the “missing” value of this attribute is in fact \perp , the universal object consistent with any object.


Second, two parts of an AVS may share the same structure. This phenomenon is usually referred to as **structure-sharing** or as **re-entrancy**. The usual way to diagram this is with a boxed numeral representing an index for the sub-structure that

²⁷Johnson’s (1987) definition allows the attribute of an attribute-value pair to be a complex structure as well.

is shared.²⁸ In (7.27), the sub-AVS that is the value of FATHER has a boxed 1 as an index. There is another boxed 1 as the value for the CAUSE of Anne Boleyn's death, which means that exactly the same sub-AVS occurs in both these places. If, by unifying (7.27) with another AVS, we learn more information about Henry, we automatically learn more information about the cause of Anne Boleyn's death. The extra information in the result would be accessible through both the FATHER path and the MOTHER|DIED|CAUSE path.

Before proceeding, we should show the relationship between the attribute-value matrices we have been looking at and the sorts of diagrams used for phonological structures that are used in the rest of this dissertation. We can do this by looking at an alternative notation for attribute-value structures, **directed acyclic graphs**, or **dags**. For example, the matrix in (7.28) corresponds to the dag in (7.29).

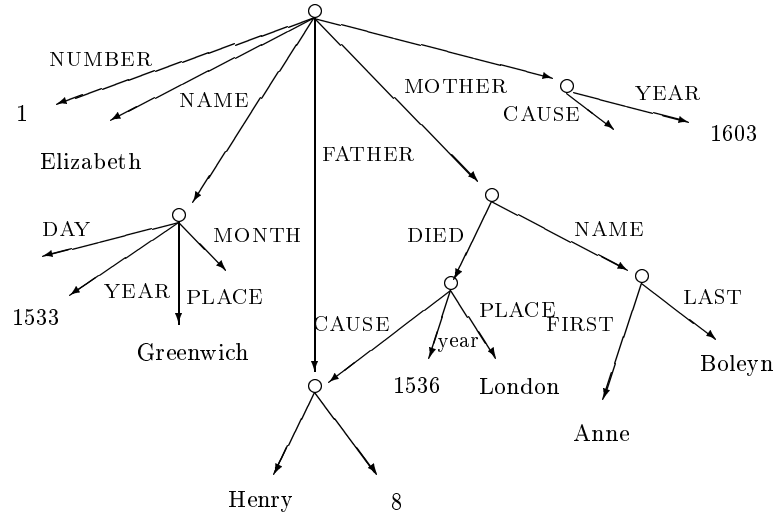
$$(7.28) \left[\begin{array}{l} \text{NAME:} \quad \text{Elizabeth} \\ \text{NUMBER:} \quad 1 \\ \text{BORN:} \quad \left[\begin{array}{l} \text{YEAR:} \quad 1533 \\ \text{PLACE:} \quad \text{Greenwich} \end{array} \right] \end{array} \right]$$



Every attribute-value matrix has a dag equivalent. Each object in a matrix (atomic or complex) corresponds to a node in a dag. An attribute corresponds to a labelled arc. Reentrancy or structure-sharing is represented in dag notation by having the same node as the tail of two different arcs. The entire Elizabethan AVS given in matrix notation in (7.27) can be equivalently drawn using the dag in (7.30):

²⁸Here, I will use square brackets in place of a box.

(7.30)



The similarities between a dag and the PSs we have been dealing with should be apparent. The equivalence between these two notations was exploited during the exposition of the segmental theory in chapter 2, where most segments were diagrammed using both an SPE-style feature matrix (an attribute-value matrix) and an autosegmental-style feature hierarchy (a dag).

An extremely important relation between AVSs is that of **subsumption**, or inversely **extension**. This can be seen as essentially a subset-superset relation. An AVS A is an **extension** of an AVS B if A has all the attribute-value pairs in B , and possibly more besides. If A is an extension of B , we also say that B **subsumes** A .

$$(7.31) \quad a \quad \left[\begin{array}{l} \text{NAME:} \quad \text{Elizabeth} \\ \text{NUMBER:} \quad 1 \\ \text{BORN:} \quad \left[\text{YEAR:} \quad 1533 \right] \\ \text{FATHER:} \quad \left[\begin{array}{l} \text{NAME:} \quad \text{Henry} \\ \text{NUMBER:} \quad 8 \end{array} \right] \end{array} \right]$$

$$b \quad \left[\begin{array}{l} \text{NAME:} \quad \text{Elizabeth} \\ \text{BORN:} \quad \left[\text{YEAR:} \quad 1533 \right] \end{array} \right]$$

Here, a is an extension of b , and b subsumes a . This is clearly a reflexive relation: an AVS always subsumes itself. (If we want to rule out this possibility in a certain case, we can talk about proper subsumption or proper extensions.)

A small complication is introduced when we consider AVSs that contain smaller AVSs. The example in (7.31) should provoke a premonition of the complication. While

the attribute `NAME` had an atom, `Elizabeth`, as its value, the attribute `BORN` is not so lucky. Its value is the complex sub-AVS `[YEAR:1533]`. So far, there are no problems, since it is exactly the same sub-AVS in both *a* and *b*. But what if it hadn't been? What if the sub-AVS in *a* had been more complicated, perhaps including information about the place of birth as well, as *c* does:

(7.32) *c*

NAME:	Elizabeth					
NUMBER:	1					
BORN:	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="padding-right: 10px;">YEAR:</td> <td>1533</td> </tr> <tr> <td>PLACE:</td> <td>Greenwich</td> </tr> </table>	YEAR:	1533	PLACE:	Greenwich	
YEAR:	1533					
PLACE:	Greenwich					
FATHER:	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="padding-right: 10px;">NAME:</td> <td>Henry</td> </tr> <tr> <td>NUMBER:</td> <td>8</td> </tr> </table>	NAME:	Henry	NUMBER:	8	
NAME:	Henry					
NUMBER:	8					

Now does *b* subsume *c*? Intuitively, we should want it to. *c* certainly contains all the information *b* does, and then some.

To cover cases like this, we should extend our definition of extension and subsumption to be recursive. We no longer require exactly the same attribute-value pairs in both AVSs. Instead *A* subsumes *B* if for every attribute-value pair in *A*, *B* has a pair with the same attribute, where (recursively) the value of the pair in *A* subsumes the value in *B*. (Recall that two AVSs subsume each other if they are equal.)

Using this definition, *b* in (7.31) does subsume *c* in (7.32). The values of the `BORN` attribute are different in *b* and *c*, but *b*'s value, `[YEAR: 1533]` subsumes *c*'s value, `[YEAR: 1533, PLACE: Greenwich]`, so *b* as a whole can still subsume *c* as a whole.

The notion of subsumption (or extension) allows us to define the central operation of **unification**. Two AVSs may unify if there is an AVS that is an extension of both of them (or, alternatively put, an AVS that both of them subsume). The result of the unification operation is the smallest common extension of the two inputs. The practical effect of this is that two AVS may unify if they have not got conflicting values for any path of attributes.

In the following rather perverse examples, (7.33a) can unify with (7.33b) without any conflicts, to result in (7.33c), their smallest common extension.

(7.33)

- a. $\left[\begin{array}{l} \text{NAME: Henry} \\ \text{NUMBER: 8} \\ \text{BORN: [YEAR: 1491]} \\ \text{RELIGION: Anglican} \end{array} \right]$
- b. $\left[\begin{array}{l} \text{NUMBER: 8} \\ \text{RELIGION: Anglican} \\ \text{BORN: [PLACE: Richmond Park]} \\ \text{ABDICATED: [YEAR: 1936]} \end{array} \right]$
- c. $\left[\begin{array}{l} \text{NAME: Henry} \\ \text{NUMBER: 8} \\ \text{BORN: [YEAR: 1491} \\ \quad \text{[PLACE: Richmond Park]} \\ \text{ABDICATED: [YEAR: 1936]} \\ \text{RELIGION: Anglican} \end{array} \right]$

(7.33c) is the smallest AVS that both (7.33a) and (7.33b) subsume.

(7.34a), on the other hand, cannot unify with (7.34b), because they disagree on the value of the NUMBER attribute:

(7.34)

- a. $\left[\begin{array}{l} \text{NAME: James} \\ \text{NUMBER: 6} \\ \text{MOTHER: Mary Stuart} \\ \text{COUNTRY: Scotland} \end{array} \right]$
- b. $\left[\begin{array}{l} \text{NAME: James} \\ \text{NUMBER: 1} \\ \text{MOTHER: Mary Stuart} \\ \text{DIED: [YEAR: 1625]} \end{array} \right]$

In other words, there is no AVS that both (7.34a) and (7.34b) subsume.

As an over-simplified example of how AVSs and unification have been applied to linguistic problems, we can consider the requirement in most unification-based approaches to syntax that the sets of “agreement” features of a subject and a predicate should be able to unify. The English lexical items *we*, *she*, and *goes* might have the following agreement AVSs as part of their representations:

(7.35)

- a. *we* $\left[\begin{array}{l} \text{PERSON: } 1 \\ \text{NUMBER: } \text{plural} \end{array} \right]$
- b. *she* $\left[\begin{array}{l} \text{PERSON: } 3 \\ \text{NUMBER: } \text{singular} \\ \text{GENDER: } \text{feminine} \end{array} \right]$
- c. *goes* $\left[\begin{array}{l} \text{TENSE: } \text{present} \\ \text{PERSON: } 3 \\ \text{NUMBER: } \text{singular} \end{array} \right]$

She goes is a legal sentence, because the agreement AVSs of (7.35a) and (7.35b) can unify to:

$$(7.36) \quad \left[\begin{array}{l} \text{PERSON: } 3 \\ \text{NUMBER: } \text{singular} \\ \text{GENDER: } \text{feminine} \\ \text{TENSE: } \text{present} \end{array} \right]$$

**We goes* is not legal, because there is no way to unify (7.35a) and (7.35c) without a contradiction in the values of the PERSON and NUMBER attributes.

Unification is commutative, that is

$$(7.37) \quad A \sqcup B = B \sqcup A$$

Under ideal circumstances, it is also associative,²⁹ that is

$$(7.38) \quad A \sqcup (B \sqcup C) = (A \sqcup B) \sqcup C$$

This means that the information (in the form of AVSs) that we get from various parts of the linguistic system—lexical entries, constraints, node admissibility conditions, and so on—can be unified together in any order whatsoever and the result will be the same.

So far we have been assuming that there is only one possible result of a unification $A \sqcup B = C$. This is obtained by identifying the root of A's dag with the root of B's dag, giving the root of C's dag (or, equivalently, the outer brackets of A are lined up with the outer brackets of B). But this is not the only conceivable way unification might be defined.

Given:

²⁹As a result of wanting to extend the use of unification beyond the characterization given so far, we shall see some less than ideal circumstances below, in our discussion of unifications with ambiguous results. The consequence will be that (the extended version of) unification cannot be associative if it is taken to be a function mapping from a pair of AVSs to a third AVS, though other formalizations will be able to preserve associativity.

$$(7.39) \quad A \left[\begin{array}{l} \text{colour: red} \\ \text{size: big} \\ \text{sits-on: } A' [\text{shape box}] \end{array} \right] \quad B \left[\begin{array}{l} \text{colour: red} \\ \text{shape: box} \end{array} \right]$$

the only unification our definition allows identifies the roots of A and B, adding B's shape feature to the outer level of A:

$$(7.40) \quad C \left[\begin{array}{l} \text{colour: red} \\ \text{size: big} \\ \text{shape: box} \\ \text{sits-on: } A' [\text{shape box}] \end{array} \right]$$

But it is intuitively conceivable that the merger might have take place at the lower level, identifying the roots of B and the sub-AVS A':

$$(7.41) \quad C' \left[\begin{array}{l} \text{colour: red} \\ \text{size: big} \\ \text{sits-on: } \left[\begin{array}{l} \text{colour: red} \\ \text{shape box} \end{array} \right] \end{array} \right]$$

In a unification-based phonology, there are many situations where it would be desirable to have this kind of unification at a level other than the highest. As a simple example, though one quite similar to some proposals in the area, including my own, consider a model where we represent phonological sequence as a kind of list structure:

$$(7.42) \quad \left[\begin{array}{l} \text{C:} \\ \text{V:} \\ \text{next-syll: } \left[\begin{array}{l} \text{C:} \\ \text{V:} \\ \text{next-syll: } \left[\begin{array}{l} \text{C:} \\ \text{V:} \\ \dots \end{array} \right] \end{array} \right] \end{array} \right]$$

Now, if morphemes are seen as AVSs (that is, as real representations) we would want to be able to give a morpheme the ability to “float” up and down the levels of this list structure looking for the right place to unify. Otherwise we would be faced with the unworkable situation where the phonological representations of all the morphemes in a word were required to stack up at the highest possible level of the list structure. In short, simple concatenation would be impossible.

It turns out that actually implementing the kind of freedom we have been considering within a unification framework is a non-trivial matter. Particularly, unification would no longer be a function,³⁰ it would no longer be guaranteed to give a unique

³⁰or at least a function from representations to representations. It could still be treated as a function from sets of representations to sets of representations.

AVS as a result. The unification of A and B in (7.39), as demonstrated, could have two different AVSs as its result. We might decide to arbitrarily choose just one of these possibilities as the “correct” result of the operation, but then unification would no longer be guaranteed to obey the property of associativity — in the unification $A \sqcup B \sqcup C$, there is no reason to suppose that the “correct” result that is chosen under the $A \sqcup (B \sqcup C)$ order of application would be the same one chosen under the $(A \sqcup B) \sqcup C$ order.

These problems are the result of trying to treat morphemes directly as AVSs, that is, as pieces of representation, and therefore using unification as the operation necessary for combining morphemes. For this, and for other reasons to be discussed in the next subsection, it is better to see morphemes not as AVSs, but as constraints on AVSs.

Constraints on attribute-value structures

As syntacticians worked with the attribute-value formalism in the 1980s, many began wanting to use more complicated types of structures involving disjunction and negation. This was an especially pressing issue because most subscribed to what I referred to in section 5.1 as the “one-word, one-terminal-node” assumption. For example, a bare English verb could be entered in the lexicon several times with each of its person and number combinations (1sg, 2sg, 1pl, 2pl, 3pl), but a more concise way would be to list it once and simply require it not to be 3sg:

$$(7.43) \left[\begin{array}{l} \text{pred: swim} \\ \text{subj: } \left[\begin{array}{l} \text{agr: } \sim \left[\begin{array}{l} \text{num: sg} \\ \text{pers: 3rd} \end{array} \right] \end{array} \right] \\ \text{tense: pres} \end{array} \right]$$

Similarly, the German article *die* could be more concisely represented using a single AVS with disjunction:

$$(7.44) \left[\begin{array}{l} \text{cat} \quad \text{determiner} \\ \text{agr} \quad \left\langle \left[\begin{array}{l} \text{number: sg} \\ \text{gender: fem} \\ \text{number: pl} \end{array} \right] \right\rangle \\ \text{case} \quad \left\langle \begin{array}{l} \text{nom} \\ \text{acc} \end{array} \right\rangle \end{array} \right]$$

These kinds of structures posed problems for the assumption that AVSs were linguistic objects or representations. It is difficult to see an object as consisting of two mutually incompatible alternatives at the same time or as containing the negation of some other object. The proper way to deal with disjunction and negation in AVSs was the focus of much of the formal theoretical research in unification in the 1980s. Kasper and Rounds (19xxx) designed a logic for expressing constraints on AVSs, in

which negation and disjunction were expressed in the constraints and not in the AVSs themselves. Johnson (1988) showed that there was no need to design a new logical language for these constraints — and hence prove from scratch all the mathematical properties of the language, such as decidability — they could instead be expressed in first-order logic, whose mathematical properties are already well-understood.

Since the formal and empirical advantages of separating representations and descriptions has been a recurring theme throughout much of this dissertation, in this appendix I shall only point out the relationship between a description-based approach to AVSs and the earlier conception of unification as an operation on representations.

Consider the following example, taken from Johnson (1991), involving the unification of the lexical items *salmon* and *swims* in the sentences *The salmon swims*. In a representation-based unification framework, the *salmon* and *swims* would have the AVSs in (7.45) and (7.46) in their lexical entries. The noun *salmon* has an agreement value specified only for person [*pers:3rd*], but not for number, as it is compatible with either [*num:sg*] or [*num:pl*]. *swims* is compatible with only [*num:sg*], so this will be the value of the entire sentence.

$$(7.45) \quad \textit{salmon} \quad e' \left[\begin{array}{l} \text{pred: } \textit{salmon} \\ \text{agr: } f' \left[\text{pers: } 3\text{rd} \right] \end{array} \right]$$

$$(7.46) \quad \textit{swims} \quad g'' \left[\begin{array}{l} \text{pred: } \textit{swim} \\ \text{subj: } e'' \left[\text{agr: } f'' \left[\begin{array}{l} \text{num: } \textit{sg} \\ \text{pers: } 3\text{rd} \end{array} \right] \right] \\ \text{tense: } \textit{pres} \end{array} \right]$$

In a description-based unification framework, though (7.45) and (7.46) might be used for notational convenience, the lexical entries of the words consist of constraints written in the description language, such as:

(7.47) Lexical constraint/description for *salmon*:

$$e' \xrightarrow{\text{pred}} \textit{salmon} \wedge e' \xrightarrow{\text{agr}} f' \wedge f' \xrightarrow{\text{pers}} 3\text{rd}$$

(7.48) Lexical constraint/description for *swims*:

$$g'' \xrightarrow{\text{pred}} \textit{swim} \wedge g'' \xrightarrow{\text{tense}} \textit{pres} \wedge g'' \xrightarrow{\text{subj}} e'' \wedge e'' \xrightarrow{\text{agr}} f'' \wedge f'' \xrightarrow{\text{num}} \textit{sg} \wedge f'' \xrightarrow{\text{pers}} 3\text{rd}$$

To get the entire sentence in the representation-based framework, the operation of unification would apply to the matrices in (7.45) and (7.46), identifying (7.45) with the *subj* value of (7.46), to yield the matrix in (7.49):

$$(7.49) \quad \textit{The salmon swims} \quad \left[\begin{array}{l} \text{pred:} \quad \textit{swim} \\ \text{subj:} \quad \begin{array}{l} e' \\ e'' \end{array} \left[\begin{array}{l} \text{agr:} \quad \begin{array}{l} f' \\ f'' \end{array} \left[\begin{array}{l} \text{num:} \quad \textit{sg} \\ \text{pers:} \quad \textit{3rd} \end{array} \right] \\ \text{pred:} \quad \textit{salmon} \end{array} \right] \\ \text{tense:} \quad \textit{pres} \end{array} \right]$$

(7.49) is simply the matrix that satisfies the descriptions of both (7.47) and (7.48) simultaneously. In other words, (7.49) satisfies the *conjunction* of (7.47) and (7.48), with the stipulation that the object denoted by e' is the same as the object denoted by e'' :

$$(7.50) \quad e' = e'' \wedge \\ e' \xrightarrow{\text{pred}} \textit{salmon} \wedge e' \xrightarrow{\text{agr}} f' \wedge f' \xrightarrow{\text{pers}} \textit{3rd} \wedge \\ g'' \xrightarrow{\text{pred}} \textit{swim} \wedge g'' \xrightarrow{\text{tense}} \textit{pres} \wedge g'' \xrightarrow{\text{subj}} e'' \wedge \\ e'' \xrightarrow{\text{agr}} f'' \wedge f'' \xrightarrow{\text{num}} \textit{sg} \wedge f'' \xrightarrow{\text{pers}} \textit{3rd}$$

Put briefly, the result of unifying of two matrices is the result of conjoining the descriptions of the two matrices.

7.6.2 Bird (1990)

The main goal of Bird (1990) is to give the graphical illustrations used every day by phonologists the kind of formal foundation that would enable phonology to be integrated into the constraint-based approaches to linguistics that have been developing in syntax and semantics. Integrating the attribute-value logic of Johnson with the temporal logic of van Benthem (1983), Bird argues for a model-theoretic view of phonology: grammars are made up of descriptions framed in a logical language, linguistic representations are their models, the domain objects that satisfy the descriptions and which the descriptions denote. He is the first researcher I am aware of to offer a complete logical description language for phonological representations and illustrate its application.

Bird posits a rather unique kind of phonological representation. His proposals for segmental representation are based largely on the work of Browman and Goldstein on articulatory gestures. For suprasegmental structure he uses a moraic representation along the lines of Hyman (1985), where onset consonants belong to the first mora of the syllable.

In chapter 1, Bird (1990) introduces the basic concepts of a model-theoretic or constraint-based approach to phonology, contrasting them with the assumptions that have underlied most of generative phonology. He argues for the possibility of having a monostratal phonology and presents defences against the usual generative arguments against monostratality (e.g., feature changing rules in Pasiego and Chumash).³¹ He

³¹While Bird manages to find principled constraint-based analyses of Pasiego and Chumash, I have argued in chapters 3 and 4 that the data are not entirely as advertised.

discusses the sign-based aspects of his proposals and the ways in which it could be integrated into a more general sign-based theory of linguistics (e.g., HPSG).

In chapter 2, Bird presents a formal first-order language \mathcal{L} for the description of phonological representations. Many aspects of this language have been used in the present dissertation, for example, the notion of sorts. The core of this chapter is a discussion of the hierarchical organization and the temporal organization of phonological structures, the relationship between the two, and their proper formalization.

The relation δ represents immediate dominance: $x\delta y$ is true if the nodes denoted by x and y are connected by an association line and x is higher in the feature hierarchy or prosodic hierarchy than y is. Immediate dominance is subject to the constraints imposed by the following axioms.³²

(7.51) Immediate dominance is irreflexive

$$\forall x \sim (x\delta x)$$

(7.52) Immediate dominance is asymmetric


$$\forall x, y (x\delta y) \rightarrow \sim (y\delta x)$$

(7.53) Immediate dominance is intransitive

$$\forall x, y, z (x\delta y \wedge y\delta z) \rightarrow \sim (x\delta z)$$

These are the kinds of defining characteristics of association lines that are usually assumed, but seldom made explicit, in work in generative phonology. Using these mechanisms, other commonly desired (or assumed) constraints can also be formulated. For example, a model that assumes direct association of an onset root node to a syllable node, as in (7.54), would want a constraint like (7.55) to require all root nodes immediately dominated by a syllable node to be consonants. *syl*, *root*, and *consonantal* are sorts on nodes:³³

(7.54)



³²For readability and consistency with the notation used in the rest of this dissertation, slight cosmetic changes have been made to Bird's statement of axioms (e.g., insertion of parentheses and commas).

³³Though (7.55) is based closely on Bird's illustrative constraint (2-14) (1990: 22), it should be noted that Bird does not ultimately adopt this version of moraic theory.

$$(7.55) \quad \forall x, y \delta y \wedge \text{syl}(x) \wedge \text{root}(y) \rightarrow \text{consonantal}(y)$$

Just as important as hierarchical structure is the representation of time. For this, Bird uses an interval-based temporal logic based on van Benthem (1983), with the relations \circ (overlap) and \prec (precedence), which are subject to the following axioms.

$$(7.56) \quad \text{a. Overlap is reflexive} \\ \forall x \ x \circ x$$

$$\text{b. Overlap is symmetric} \\ \forall x, y \ x \circ y \rightarrow y \circ x$$

$$\text{c. Precedence is asymmetric} \\ \forall x, y \ x \prec y \rightarrow \sim y \prec x$$

$$\text{d. Precedence is disjoint from overlap} \\ \forall x, y \ x \prec y \rightarrow \sim x \circ y$$

$$\text{e. Precedence is transitive (through overlap)} \\ \forall w, x, y, z \ w \prec x \wedge x \circ y \wedge y \prec z \rightarrow w \prec z$$

$$\text{f. Time is linear} \\ \forall x, y \ x \prec y \vee x \circ y \vee y \prec x$$

From these axioms, a version of the No Crossing Constraint can be derived:

$$(7.57) \quad \text{No Crossing Constraint} \\ \sim \exists w, x, y, z \ w \prec x \wedge y \prec z \wedge w \circ z \wedge x \circ y$$

This particular theorem deals only with conflicting precedence relations between nodes whose intervals overlap. It says nothing directly about nodes and association lines. This can be remedied with the addition of the following constraint, which intuitively would be needed by an autosegmental framework anyway, saying that the temporal interval of a node overlaps that of any node it dominates (δ^* , or dominance, is the transitive closure of δ , immediate dominance):

$$(7.58) \quad \text{Locality Constraint} \\ \forall x, y \ x \delta^* y \rightarrow x \circ y$$

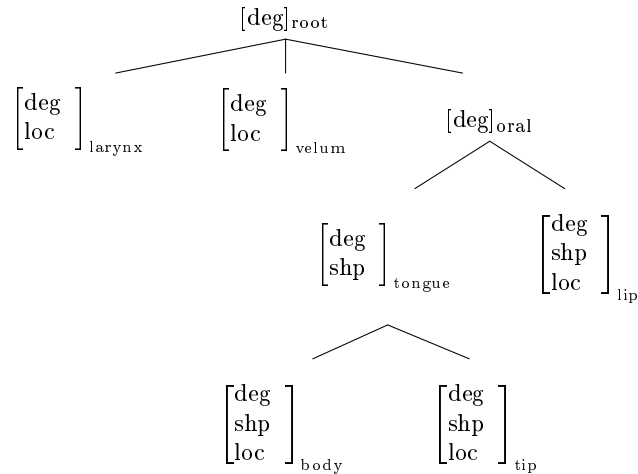
This much machinery, it would seem, would be shared by any constraint-based approach.³⁴ Chapter 3 of Bird (1990) applies this machinery to formalize a particular

³⁴Except for the one presented in this dissertation. Ironically, one of the most common uses of the dominance relation in the present framework (as applied to the nuclear spine) is to achieve the effect of Bird's \prec relation. See the comparison below.

model of phonological representation. After examining some of the feature hierarchies on the market (e.g., Clements 1985, Sagey 1986), Bird proposes a more articulatorily based model of segmental structure drawing heavily from the work of Browman and Goldstein (1989, 1990) in *Articulatory Phonology* (already discussed briefly in section 2.3 and 3.6). Browman and Goldstein analyse segments into articulatory constriction gestures, specified for the constriction articulator, the constriction location (e.g., palatal, velar, labial, pharyngeal), and the constriction degree (wide, mid, narrow, critical, closure).³⁵ As well, the tongue tip and the lips may have a shape specification (for rounding and laterals).

To this catalogue of possible gesture types, Bird (1990: 53) adds the idea of combining them into a hierarchical structure based on the active articulator. The proposed hierarchy is:

(7.59)



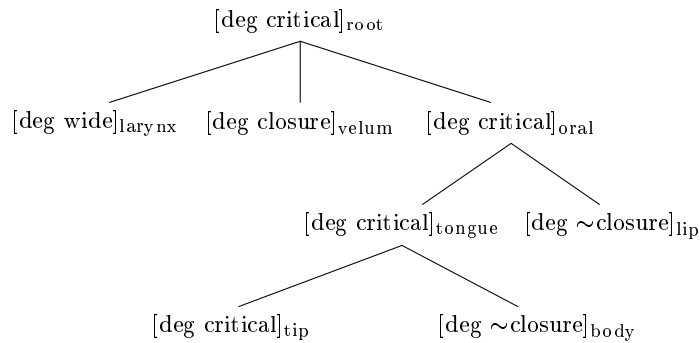
Constriction degrees are something like the manner features of Clements, but are marked on every node of (7.59). For any single segment, one of these nodes' degrees is especially important and provides the "manner" feature for the whole structure. We can intuitively grasp which node it will be by considering the "tube geometry" of Browman and Goldstein (1989), where the vocal tract is considered as a collection of tubes hooked up in series and in parallel, with a number of valves that can cut off the

³⁵Collapsing these constriction degrees into three, interpreted differently in consonants and vowels, is so far as I can tell a novel claim of this dissertation. Browman and Goldstein's degrees correspond to mine as follows: closure = [d:0], critical = $C[d:1]$, narrow = $V[d:1]$, mid = $V[d:2]$. Browman and Goldstein's wide gesture of the tongue body or tip would correspond to a gesture with pharyngeal site articulated by the tongue body or root, i.e., [a:Dor, s:Pha] or [a:Rad, s:Pha].

airflow to certain constriction degrees. We want a rough measure of the overall airflow out of the vocal tract, in terms of the constriction degrees. Clearly, the airflow through the oral cavity is determined by the valve (tongue body, tip, lips) with the greatest constriction (a complete closure for the tongue tip makes a complete closure for the entire oral tube). The total supralaryngeal airflow will be the maximum of the oral airflow and the nasal (velum) airflow (i.e., the lesser constriction). The output of the whole tube will be the lesser of the supralaryngeal airflow and the glottal airflow (i.e., the greater constriction).

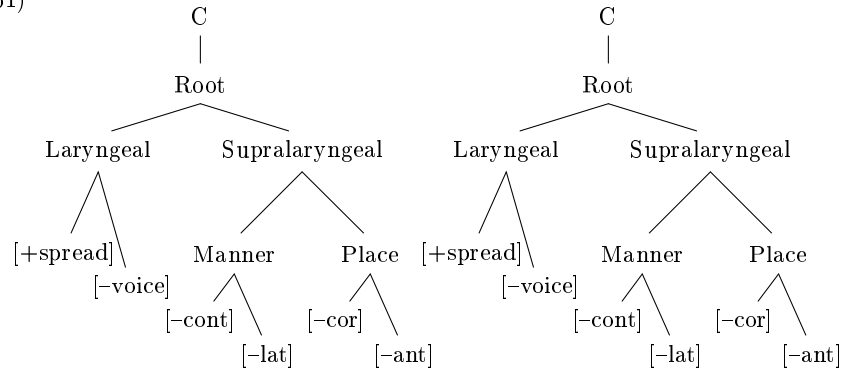
Bird captures these relationships by percolating the constriction degree feature through the tree in (7.59) according to the maximum and minimum rules just outlined. For the segment [s], the critical degree feature of the tongue tip constriction percolates up through the entire tree, as marked by the heavy lines, till it becomes the “manner” feature for the tree as a whole.

(7.60)



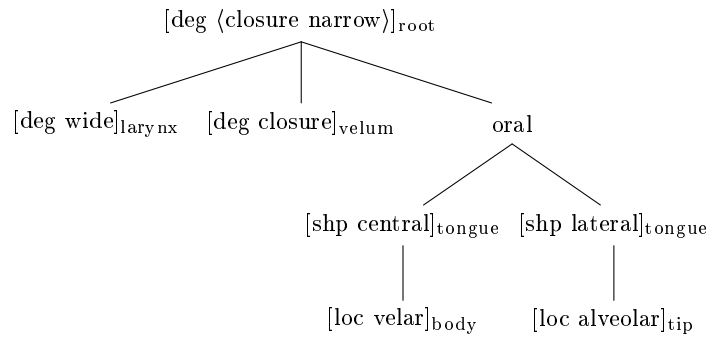
Bird also allows a node to dominate two “heads” (i.e., nodes whose degree features are percolated up) of the same type, resulting in a **sequence** of degree specifications. Where a Clements-style feature hierarchy would represent the onset cluster *kl* along the lines of:

(7.61)



Bird proposes a single multiply-headed root node:

(7.62)



The two tongue nodes are ordered by the \prec relation. The root node is specified for a sequence of degrees, <closure narrow>. These specifications are percolated down the tree, the left element of the sequence down the left branch and the right element of the sequence down the right branch of the multiply-headed oral node. While a multiple root node representation like (7.61) would have to resort to OCP stipulations or dissimilation rules in order to ban clusters like *tl* or *bw*, Bird argues that his model automatically predicts the absence of these clusters as the result of trying to percolate incompatible feature specifications through the tree.

This proposal for dealing with complex onsets captures the insight that onsets, no matter how complex, usually behave as single “segments” or, put another way, that branching onsets are more like light diphthongs than like heavy diphthongs. This is the same insight that lies behind Steriade’s treatment of consonantal closure and release that was adopted in chapter 2.

Bird further extends this type of hierarchical structure to the syllable level, adopting a Hyman-style moraic theory (where onsets attach to the first mora of the syllable).

Representing phonological sequence

Phonological sequentiality are handled directly in Bird (1990) with the temporal relations of \prec (precedence) and \circ (overlap). The relations hold between actual nodes of a phonological representation. In the present framework, temporal organization is in a sense irrelevant for phonology. Temporal relations like precedence and overlap are vital parts of Phonetic Event Structures (see section 3.6) and play a role in the phonology-phonetics interface. However, they are not taken to hold between pieces of a *phonological* representation, but between the PES intervals that are phonetic interpretations of those phonological pieces. Instead, the central relation used here is immediate dominance, δ for Bird and in the present framework the *arc* predicate and its $\overset{S}{\rightarrow}$ abbreviation. What is interpreted as precedence in the phonetics is in the phonology just a particular kind of dominance, represented by abstract government arcs that are in themselves atemporal.

I believe the difference is one of the most important conceptual differences between Bird’s and Scobbie’s approach and my own. It might be said that I believe sequential relations to be phonologically “active.” For me, a sequential relation is simply a government arc, and as a government arc it can behave in phonologically interesting ways. Phonologically interesting properties can be predicated of it. It can create a local domain between the nodes it joins, causing spreading. It can participate in natural classes with other types of government arcs. For example, in Rotuman metathesis, a morpheme could underspecify government arcs simply as being members of the class {nuclear-licence, release, secondary}, resulting in alternations among sequential syllables, light diphthongs, and unlauted monophthongs. This kind of morphological alternation between a “temporal relation” and an “association line” is hard to imagine in a model that treats them as fundamentally different kinds of phonological things.

7.6.3 Scobbie (1991)

Oversimplifying matters somewhat, we might say that where Bird (1990) showed the possibility of doing phonology in a declarative way, Scobbie (1991) tried to show its desirability (for phonologists). Where Bird concentrates on more or less “static” phenomena, such as phonotactic restrictions, Scobbie deals with the more “dynamic” phenomena (things that seem to involve changes through time, such as spreading rules) that have traditionally been the focus of attention for phonologists. Being interested in a direct comparison between the proceduralist claims of autosegmental phonology and the declarativism of constraint-based approaches, Scobbie works with diagrams that are as similar as possible to those used in autosegmental phonology and addresses many of the same problems that have been recurring in the autosegmental literature

for the past several years.

One of the recurring themes of this section is the different mechanisms for dealing with temporal ordering. Besides the different assumptions concerning other aspects of representations, Scobbie and Bird also have different views on ordering. Where Bird allowed the possibility of temporal relations between members of every type of node (indeed his temporal axiom (7.56f) *requires* that every pair of nodes whatsoever stand in some temporal relation), Scobbie argues forcefully that only the root node tier should be participate in linear ordering relations. Since every non-root node is connected by a path to a root node, the linear position of every node can be determined derivatively.

Scobbie's analogy of a bead curtain may be helpful:


The value of PHON can be likened to a bead curtain, consisting of strings of beads hanging vertically from a single rail. Each attribute is like a differently coloured bead and its value is the entire string below it. Since the hierarchical position of attributes is fixed, rows of a single colour are the analogue of a sequence of instances of the same attribute. And just as we can describe the relative order of red beads in the red row, so we can refer to the sequence of PLACE attributes, say. But it is crucial to note that this is merely shorthand. If we think of the bead curtain we can see that to refer to the absolute locations of two red beads we must in fact refer to the locations of the string of beads (the paths) leading down to them. ... We know that if the curtain is sent swinging the beads do not change location, nor even if the strings are tied in knots — the true order of beads does not alter unless we cut the strings and detach them from the track, both of which are nonmonotonic operations.

In other words, a PS has no inherent representation of linear order between two non-root nodes, say [high] features. The only way to determine their order is through the mediation of the root node tier and the paths that connect the [high]s to the root node tier. There is no [high] tier that exists independently of the root node tier.

Scobbie handles the mechanics of representing sequentiality somewhat differently than Bird. Instead of positing primitive relations like \prec between root nodes, Scobbie takes the value of the PHON attribute to be an unordered set³⁶ of ordered pairs $\langle i, S \rangle$, where i is an index and S is a root node. Precedence is defined by imposing a weak ordering on the indices. Because many constraints are sensitive to adjacency, that is, to immediate precedence relations between indices, constraints need to be able refer to these indices and their temporal relations.

One result of this way of doing things is an automatic ban on one-to-many associations:

³⁶See, e.g., Johnson (1991) on how sets can be represented in the attribute-value formalism.



Since any node in an attribute-value matrix can have only one value for a given attribute (in my terms, every attribute/arc is of the sort *unique*), (a) is an incoherent structure. Nothing is wrong with (b), however—it is simply another example of re-entrancy or structure sharing, represented in the set of (index, root node) pairs with a double occurrence of the same token:

$$(7.64) \{ \langle i, [\text{MELODY}|\text{PLACE} \quad [1]] \rangle, \langle j, [\text{MELODY}|\text{PLACE} \quad [1]] \rangle \} \quad i \prec j$$

AVP can represent geminates, it cannot represent the contour segments that autosegmental phonologists have often used for affricates.

Syllabic structure is also represented on (or “under”) this sequence of root nodes. A syllable has the general structure represented by the AVS:³⁷

$$(7.65) \left[\text{SYLL} \quad \begin{bmatrix} \text{ONSET} & [1] \\ \text{NUCLEUS} & [2] \end{bmatrix} \right]$$

[1] will be the melody of the onset, [2] the melody of the nucleus. This representation of the entire syllable will be a feature of each root node in the syllable, through structure sharing. Each root node has a MELODY attribute and a SYLL attribute. So the melody of a nuclear root node is accessible by both the MELODY path and the SYLL|NUCLEUS path. The nuclear root node would also have access to the onset’s melody by the SYLL|ONSET path. A prosodic licensing condition can be expressed by requiring every melody to be dominated by the SYLLABLE attribute.

For example, AVP would give the English word *Andy* the following syllabic structure, where the $\langle a, b, c, d \rangle$ notation represents the root nodes of four index-root node pairs with successive indices:

³⁷In addition to attributes for onset and nucleus, Scobbie also proposes an attribute he calls ϵ -SYLL, whose value is essentially the second mora of a heavy syllable. This second mora also has the nucleus and onset attributes relevant for the first: “This assignment effectively divides the syllable into two parts; an obligatory CV mora and an optional mora *mora*2. This second mora is a reversed and weakened echo of the first; hence “ ϵ -SYLL’. This assignment incorporates four syllabic functions: onset, nucleus, off-glide and coda consonant.” (1991: 31). I shall not address the similarities and differences between this concept and my own proposal for strict CV structures. In what follows, I shall simply abbreviate paths involving ϵ -SYLL as CODA.

(7.66) AVP syllabification of *Andy*

$$\left\langle \begin{array}{l} \left[\begin{array}{ll} \text{MELODY} & [1] \text{'a'} \\ \text{SYLL} & [2] \left[\begin{array}{ll} \text{NUCLEUS} & [1] \\ \text{CODA} & [3] \end{array} \right] \end{array} \right], \left[\begin{array}{ll} \text{MELODY} & [8] \left[\begin{array}{l} \text{NASAL} + \\ \text{PLACE} & [4] \text{'coronal'} \end{array} \right] \\ \text{SYLL} & [2] \end{array} \right], \\ \left[\begin{array}{ll} \text{MELODY} & [5] \left[\text{PLACE} & [4] \right] \\ \text{SYLL} & [6] \end{array} \right], \left[\begin{array}{ll} \text{MELODY} & [7] \\ \text{SYLL} & [6] \left[\begin{array}{ll} \text{ONSET} & \{[5]\} \\ \text{NUCLEUS} & [7] \text{'i'} \end{array} \right] \end{array} \right] \end{array} \right\rangle$$

Note that the melody 'a' of the first segment is present as the value of both the MELODY path and the SYLL|NUCLEUS path, a fact represented by the multiple occurrence of the tag [1]. Since *a* and *n* are in the same syllable, they both share the same AVS for the value of their syllable attribute, represented by the tag [2].³⁸ Similarly, *d* and *i* share the same syllable value, [6]. We can note in passing that the *n* and the *d* share their place of articulation, as represented by the double occurrence of the tag [4].

The Sharing Constraint

Structure-sharing can be a very powerful formal device, and allowing it to operate unconstrained may give us results that we do not want. Scobbie tries to constrain the power of structure sharing by limiting the possibility only to nodes that are adjacent. Formally, he states this requirement with the Sharing Constraint (1991: 64):

(7.67) Sharing Constraint

If a structure $\mathcal{M} =_s []$ is dominated by two paths of type \mathcal{P} with indices i and j , where $i \prec^* j$, then for every index n where $i \prec^* n \prec^* j$ there is a path $\langle n, \mathcal{P} \rangle$ dominating \mathcal{M} .

In other words, if any two root nodes have the same value \mathcal{M} for some path, then every root node between them must also have the value \mathcal{M} for that path.

Most of Scobbie's dissertation is essentially an examination of the consequences of the Sharing Constraint. Chapter 3 compares the theoretical and empirical differences between the Sharing Principle and autosegmental phonology's No Crossing Constraint. Chapter 4 looks at the integrity of shared structure and compares the predictions of the Sharing Constraint with the proposals of Schein and Steriade (1986), Hayes (1986), and Itô (1986, 1989) on geminate inalterability. Chapter 6 looks at the arguments presented in the autosegmental literature in favour of long-distance dependencies, structures that would necessarily violate the Sharing Constraint.

Several different proposals have been advanced to explain the resistance of geminates to rules that apply to non-geminates, none of them entirely satisfactory. Scobbie instead questions the status of these inalterability effects as an interesting problem


³⁸So the melody 'a' is also accessible by the SYLL|NUCLEUS path of the second segment *n* as well.

that exists apart from a commitment to procedural phonology. Hayes (1986) characterizes inalterability effects by saying, “Long segments often resist the application of rules that *a priori* would be expected to apply to them.” Scobbie shows the lengths that phonologists have sometimes gone to in order to write a rule in such a way that they can be surprised when it fails to apply. He proposes that inalterability effects are best handled by default rules and a few very simple phonotactic constraints that are needed in the language anyway. For example, Latin *l* was velarized in coda position, except when that coda consonant was the first member of a geminate. Autosegmental phonology has tried to capture this distribution with a rewrite rule like:

$$(7.68) \quad /l/ \rightarrow \text{Dorsal} \left[\begin{array}{c} \text{Coda} \\ | \\ \hline \end{array} \right]$$

and then try by various theoretical machinery to block the application of this rule just in case the */l/* is associated to more than one syllabic position. Instead, Scobbie argues that the distribution should be dealt with by a phonotactic constraint against velarized *ls* in onset position, a constraint which Latin needs anyway, and a default rule stating that coda */l/s* are velarized, with an “all else being equal” clause that comes for free from the nature of default rules. Geminates are the one case when all else is not equal. The default rule will allow us to infer velarization only if it is consistent with other already-known facts. But a geminate */l/* shares its place features between a coda and an onset, and we already know (from the phonotactic constraint) that a velarized *l* cannot be attached to an onset position, so the default rule cannot apply. In this way, inalterability effects can be obtained from the interaction of independently needed phonotactic constraints and default rules, with no need for an extra stipulation added to the mechanics of rule application to handle just those cases of multiple attachment.

One of Scobbie’s reasons for not wanting sharing between non-adjacent positions is that he wants to interpret feature nodes as representing articulatory gestures. In general, it is not desirable to admit non-convex phonetic intervals, that is, intervals that are split into two discontinuous segments.³⁹ This is what would be required to interpret a multiply linked structure like (7.69), where the doubly attached place node would correspond to the discontinuous phonetic gesture *g*, interrupted by the gesture of the intervening consonant:



³⁹See Bird and Klein (1990) for a discussion of the properties of **convex** and **concave** temporal intervals.

Scobbie prevents phonetic situations like this from arising by preventing the phonological situations that would give rise to them. Instead of forcing non-adjacent nodes to share values for certain paths, the most a constraint could do would be to require the same sort of value on both paths — type identity rather than token identity. This type of analysis is forced, even for situations like vowel harmony. Only elements adjacent on the root node tier (the only tier there is) can share structure, e.g., geminate consonants, long vowels.

There is some phonetic evidence supporting this conception of convex (contiguous) phonetic gestures. For example, it is usual to analyse tone “spreading” in Japanese by means of a multiply associated tone feature, a situation not permitted by Scobbie’s Sharing Constraint. But Pierrehumbert and Beckman (1988) show that the analysis that best captures the phonetic facts is one where the high tone is attached to only one vowel and the F0s of the other vowels are determined by interpolation (cf. section 3.5).

Unfortunately, Scobbie does not go into detail on how the type-sharing that would have to be involved in phenomena like vowel harmony would work.

Scobbie’s Sharing Constraint is similar in many respects to the proposal on locality presented in this dissertation. The framework developed here is similar to AVP in having essentially only one tier on which linear ordering is defined, though for me it is the “moraic” tier rather than the root node tier. A major difference is that the present framework allows more ways than simple moraic precedence by which two nodes may be “adjacent”. Specifically, nodes may be adjacent through any local-domain-creating government arc, whether this be a nuclear-licence (moraic precedence), an onset-licence, a coda-licence, a release-licence, or adjacency along one of the metrical lines (a “subset” of the moraic tier). Two nodes may also form a local domain if they inherit locality properties from their parents, who in turn form a local domain (eventually reaching a pair of ancestors who are joined by a local-domain-creating government arc).

Actually, the comparison of the two proposals is more subtle. The only structure that two nodes in a local domain share outright are the atomic gestural specifications like *Pal*, *Dor*, *R*, or *1*. There is no problem in the framework for quite distant parts of a PS to share these specifications if they are in a local domain, though this kind of apparently non-local structure sharing would violate Scobbie’s Sharing Constraint. On the other hand, Scobbie’s intention to disallow non-convex phonetic gestures (as in (7.69)) is carried over into the present framework. Under the hypothesis of recursive locality developed in chapter 4, the sort of node that is phonetically interpreted as an articulatory gesture, a *constriction* node, is not something that is shared outright by two different non-adjacency parents (or for that matter by two adjacent parents — another difference from Scobbie’s proposal). All that is shared is the “type” of node, a sharing that is implemented by the two nodes themselves forming a local domain for gestural features.

I think much of this difference between Scobbie's approach and my own is in our conception of what phonological objects are phonetically interpreted as gestures. In addition to accepting fairly uncritically the standard catalogue of binary features used in generative phonology, Scobbie adds the assumption that these phonological features correspond to phonetic gestures. In the framework developed here, this is not the case. For example, the analogue of the feature [+high], [d:1], is emphatically *not* a gesture. It is a property of a gesture. Phonetically, it is meaningless without some indication of the other properties of the gesture, like the primary articulator and the site. Two **constriction** nodes in a PS can share these properties without thereby becoming interpreted as a single non-convex phonetic gesture.

I believe that the hypothesis of recursively-defined locality (with a required basis of adjacency under government) can capture all the positive aspects of Scobbie's Sharing Constraint (ruling out unconstrained structure sharing between any two arbitrary nodes of a PS, keeping the gesture intervals produced by phonetic interpretation convex), while at the same time providing a natural and constrained way to deal with those clear cases where two nodes not adjacent on a root node tier share significant properties of gestures, e.g., vowel harmonies that spread [d:1] or [s:Pal], reduplication.

References

- Anderson, John, and Colin Ewen (1987); *Principles of Dependency Phonology*; Cambridge: Cambridge University Press.
- Anderson, Stephen R. (1985); *Phonology in the Twentieth Century: Theories of Rules and Theories of Representations*; Chicago: University of Chicago Press.
- Anderson, Stephen R. (1992); *A-Morphous Morphology*; Cambridge: Cambridge University Press.
- Archangeli, Diana (1984); *Underspecification in Yawelmani Phonology and Morphology*; PhD dissertation, MIT.
- Archangeli, Diana (1985); "Yokuts harmony: evidence for coplanar representation in nonlinear phonology"; *Linguistic Inquiry* 16, 335–372.
- Archangeli, Diana (1988); "Aspects of underspecification"; *Phonology* 5, 182–207.
- Archangeli, Diana (1991); "Syllabification and prosodic templates in Yawelmani"; *Natural Language and Linguistic Theory* 9, 231–283.
- Archangeli, Diana, and Doug Pulleyblank (1989); "Yoruba vowel harmony"; *Linguistic Inquiry* 20, 173–219.
- Archangeli, Diana, and Doug Pulleyblank (1992); *Grounded Phonology*; ms, University of Arizona and University of British Columbia.
- Azra, Jean-Luc (1992); "Variation phonétique et ambiconstituance: la coupe syllabique en français"; ms, Université de Paris VIII.
- Bagemihl, Bruce (1991); "Syllable structure in Bella Coola"; *Linguistic Inquiry* 22, 589–646.
- Baker, Mark (1988); *Incorporation: A Theory of Grammatical Function Changing*; Chicago: University of Chicago Press.
- Banner Inouye, Susan (1989); "Flap as a contour segment"; *UCLA Working Papers in Phonetics* 72, 40–82.
- Behzad, Mehdi, Gary Chartrand, and Linda Lesniak-Foster (1979); *Graphs and Digraphs*; Boston: Prindle, Weber & Schmidt.
- Belvin, Robert S. (1990a); "Person markers and patterns of case assignment in Nisgha"; *MIT Working Papers in Linguistics* 12, 14–30

- Belvin, Robert S. (1990b); "Ergativity and accusativity in Nisgha syntax"; pp. 3–18 in David Costa, editor, *Proceedings of the Sixteenth Annual Meeting of the Berkeley Linguistics Society: Special Session on General Topics in American Indian Linguistics*; Berkeley: Berkeley Linguistics Society.
- Berge, Claude (1983); *Graphes*; 3^e édition; Paris: Gauthier-Villars.
- Besnard, Philippe (1989); *An Introduction to Default Logic*; Berlin: Springer-Verlag.
- Bird, Steven (1990); *Constraint-Based Phonology*; PhD dissertation, University of Edinburgh.
- Bird, Steven (1991); "Feature structures and indices"; *Phonology* 8, 137–144.
- Bird, Steven, ed. (1992); *Declarative Perspectives on Phonology*; Edinburgh: University of Edinburgh Centre for Cognitive Science.
- Bird, Stephen, and T. Mark Ellison (1992); "One level phonology: autosegmental representations and rules as finite-state automata"; Research Paper 51; Edinburgh: University of Edinburgh, Centre for Cognitive Science.
- Bird, Steven, and Ewan Klein (1990); "Phonological events"; *Journal of Linguistics* 26, 33–56.
- Bloomfield, Leonard (1933); *Language*; New York: Holt, Rinehart, Winston.
- Booij, Geert (1984); "Neutral vowels and the autosegmental analysis of Hungarian vowel harmony"; *Linguistics* 22, 629–641.
- Bosch, Anna (1991); *Phonotactics at the Level of the Phonological Word*; PhD dissertation, University of Chicago.
- Boyce, Suzanne E., Rena A. Krackow, and Fredericka Bell-Berti (1991); "Phonological underspecification and speech motor organisation"; *Phonology* 8, 219–236.
- Brentari, Diane (1990); *Theoretical foundations of American Sign Language phonology*; PhD dissertation, University of Chicago.
- Bresnan, Joan, editor (1982); *The Mental Representation of Grammatical Relations*; Cambridge, MA: MIT Press.
- Brewka, Gerhard (1991); *Nonmonotonic Reasoning: Logical Foundations of Commonsense*; Cambridge Tracts in Theoretical Computer Science 12; Cambridge: Cambridge University Press.
- Bromberger, Sylvain, and Morris Halle (1989); "Why phonology is different"; *Linguistic Inquiry* 20, 51–70.
- Browman, Catherine P., and Louis Goldstein (1989); "Articulatory gestures as phonological units"; *Phonology* 6, 201–251.

- Browman, Catherine P., and Louis Goldstein (1990); "Tiers in articulatory phonology, with some implications for casual speech"; pp. 341–376 in John Kingston and Mary Beckman, editors, *Papers in Laboratory Phonology 1*; Cambridge: Cambridge University Press.
- Bures, Tony (1989); *The Structure of the Syllable: Principles and Parameters in Syllabification*; MA thesis, University of Toronto.
- Bynon, Theodora (1977); *Historical Linguistics*; Cambridge: Cambridge University Press.
- Carrier, Jill (1979); *The Interaction of Morphological and Phonological Rules in Tagalog*; PhD dissertation, MIT.
- Carrier-Duncan, Jill (1984); "Some problems with prosodic accounts of reduplication"; pp. 260–286 in Mark Aronoff and Richard T. Oehrle, editors, *Language Sound Structure: Studies in Phonology Presented to Morris Halle by his Teacher and Students*; Cambridge, MA: MIT Press.
- Carstairs, Andrew (1988); "Some implications of phonologically conditioned suppletion"; *Yearbook of Morphology* 1, 67–94.
- Charette, Monik (1984); "The appendix in parametric phonology"; *Studies in African Linguistics* Supplement 9, 49–53.
- Charette, Monik (1988); *Constraints on Governing Relations in Phonology*; PhD dissertation, McGill University.
- Charette, Monik (1989); "The minimality condition in phonology"; *Journal of Linguistics* 25, 159–187.
- Charette, Monik (1990); "Licence to govern"; *Phonology* 7, 233–253.
- Chomsky, Noam (1965); *Aspects of the Theory of Syntax*; Cambridge, MA: MIT Press.
- Chomsky, Noam (1970); "Remarks on nominalization"; pp. 184–221 in R.A. Jacobs and P.S. Rosenbaum, editors, *Readings in English Transformational Grammar*; Waltham, MA: Ginn.
- Chomsky, Noam (1981); *Lectures on Government and Binding*; Dordrecht: Foris.
- Chomsky, Noam (1986); *Barriers*; Cambridge, MA: MIT Press.
- Chomsky, Noam (1992); "A minimalist program for linguistic theory"; *MIT Occasional Papers in Linguistics* 1.
- Chomsky, Noam, and Morris Halle (1968); *The Sound Pattern of English*; New York: Harper and Row.
- Churchward, C. Maxwell (1940); *Rotuman Grammar and Dictionary*; Sydney: Australasia Medical Publishing Company.

- Churma, D. (1984); "On explaining morpheme structure"; *Ohio State University Working Papers in Linguistics* 29, 12-29.
- Clements, G.N. (1985); "The geometry of phonological features"; *Phonology Yearbook* 2, 225-252.
- Clements, G.N. (1989); "On the representation of vowel height"; ms, Cornell University.
- Clements, G.N. (1990); "Place of articulation in consonants and vowels: a unified theory"; to appear in B. Laks and A. Rialland, editors, *L'architecture et la géométrie des représentations phonologiques*; Paris: Editions du CNRS.
- Clements, G.N., and S.J. Keyser (1983); *CV Phonology*; Cambridge, MA: MIT Press.
- Coleman, John (1990); "Charm theory defines strange vowel sets"; *Journal of Linguistics* 26, 165-174.
- Coleman, John (1992); "The phonetic interpretation of headed phonological structures containing overlapping constituents"; *Phonology* 9, 1-44.
- Crowley, Terry (1982); *The Paamese Language of Vanuatu*; Pacific Linguistics, Series B, no. 87; Canberra: Pacific Linguistics.
- Dell, François (1973); *Les règles et les sons: Introduction à la phonologie générative*; Paris: Hermann.
- Dell, François (1984); "L'accentuation dans les phrases en français"; in F. Dell, D. Hirst, and J.-R. Vergnaud, editors, *Forme sonore du langage*; Paris: Hermann.
- Dell, François, and Mohamed Elmedlaoui (1985); "Syllabic consonants and syllabification in Imdlawn Tashlhiyt Berber"; ms, CNRS Paris and Université Mohamed I, Oujda.
- Demirdache, Hamida (1988); "Transparent vowels"; pp. 39-76 in Harry van der Hulst and Norval Smith, editors, *Features, Segmental Structure and Harmony Processes*, vol. 2; Dordrecht: Foris.
- den Dikken, Marcel, and Harry van der Hulst (1988); "Segmental hierarchic architecture"; pp. 1-78 in Harry van der Hulst and Norval Smith, editors, *Features, Segmental Structure and Harmony Processes*, vol. 1; Dordrecht: Foris.
- Di Sciullo, Anna-Maria, and Edwin Williams (1987); *On the Definition of Word*; Cambridge, MA: MIT Press.
- Dixon, R.M.W. (1972); *The Dyirbal Language of North Queensland*; Cambridge: Cambridge University Press.
- Emeneau, Murray B. (1984); *Toda Grammar and Texts*; Philadelphia: American Philosophical Society.
- Everett, Daniel, and K. Everett (1984); "On the relevance of syllable onsets to stress placement"; *Linguistic Inquiry* 15, 705-711.

- Fudge, Erik (1969); "Syllables"; *Journal of Linguistics* 5, 253–286.
- Gazdar, Gerald, Ewan Klein, Geoffrey Pullum, and Ivan Sag (1985); *Generalized Phrase Structure Grammar*; Cambridge, MA: Harvard University Press.
- Gibb, Lorna (1992); *Domains in Phonology: With Evidence from Icelandic, Finnish and Kikuyu*; PhD dissertation, University of Edinburgh.
- Goad, Heather (1991); "Dependency and complementarity in vowel geometry: on [high], [low] and [atr]"; *The Linguistic Review* 8, 185–208.
- Goad, Heather (1993); *On the Configuration of Height Features*; PhD dissertation, University of Southern California.
- Goldsmith, John (1976); *Autosegmental Phonology*; PhD dissertation, MIT.
- Goldsmith, John (1985); "Vowel harmony in Khalkha Mongolian, Yaka, Finnish and Hungarian"; *Phonology Yearbook* 2, 253–275.
- Goldsmith, John (1990); *Autosegmental and Metrical Phonology*; Oxford: Blackwell.
- Gorecka, Alicja (1989); *The Phonology of Articulation*; PhD dissertation, MIT.
- Guerssel, M., and Jean Lowenstamm (in preparation); "Verbal stem derivation in Arabic"; ms, Université du Québec à Montréal.
- Haiman, John (1980); *Hua: A Papuan Language of the Eastern Highlands of New Guinea*; Amsterdam: John Benjamins.
- Hale, Kenneth (1973); "Deep-surface canonical disparities in relation to analogy and change: an Australian example"; pp. 401–458 in T.S. Sebeok, editor, *Current trends in linguistics*, volume 11, *Diachronic, areal, and typological linguistics*; The Hague: Mouton.
- Hall, Beatrice L., R.M.R. Hall, Martin D. Pam, Amy Myers, Stephen A. Antell, and Godfrey K. Cheron (1974); "African vowel harmony systems from the vantage point of Kalenjin"; *Afrika und Übersee* 57, 4, 241–267.
- Halle, Morris, and Jean-Roger Vergnaud (1981); "Harmony processes"; pp. 1–22 in W. Klein and W. Levelt, editors, *Crossing the Boundaries in Linguistics*; Dordrecht: Reidel.
- Halle, Morris, and Jean-Roger Vergnaud (1987); *An Essay on Stress*; Cambridge, MA: MIT Press.
- Hammond, Michael (1984); *Constraining Metrical Theory*; PhD dissertation, UCLA.
- Hammond, Michael (1988); "Lexical opacity in Kalenjin vowel harmony"; ms, University of Wisconsin-Milwaukee and University of Arizona.
- Harrell, Richard S. (1962); *A Short Reference Grammar of Moroccan Arabic*; Washington: Georgetown University Press.

- Harrell, Richard S., editor (1966); *A Dictionary of Moroccan Arabic: Moroccan-English*; Washington: Georgetown University Press.
- Harrington, John P. (1974); "Sibilants in Ventureño"; *International Journal of American Linguistics* 40, 1-9.
- Harris, John (1990); "Segmental complexity and phonological government"; *Phonology* 7, 255-300.
- Harris, Zellig (1942); "The phonemes of Moroccan Arabic"; *Journal of the American Oriental Society* 62, 309-318.
- Hayes, Bruce (1980); *A Metrical Theory of Stress Rules*; PhD dissertation, MIT.
- Hayes, Bruce (1985); "Iambic and trochaic rhythm in stress rules"; *Proceedings of the Eleventh Annual Meeting of the Berkeley Linguistics Society*, 274-289.
- Hayes, Bruce (1986); "Inalterability in CV phonology"; *Language* 62, 321-351.
- Hayes, Bruce (1989); "Compensatory lengthening in moraic phonology"; *Linguistic Inquiry* 20, 253-306.
- Hayes, Bruce (1991); *Metrical Stress Theory: Principles and Case Studies*; draft ms, UCLA.
- Heath, Jeffrey (1987); *Ablaut and Ambiguity: Phonology of a Moroccan Arabic Dialect*; Albany: State University of New York Press.
- Hewitt, Mark S. (to appear); "Binarity and ternarity in Alutiiq"; to appear in *Proceedings of the Arizona Phonology Conference 4*.
- Hoeksema, Jack (1985); *Categorial Morphology*; New York: Garland.
- Hoeksema, Jack, and Richard D. Janda (1988); "Implications of process-morphology for categorial grammar"; pp. 199-247 in Richard T. Oehrle, Emmon Bach, and Deirdre Wheeler, editors, *Categorial Grammars and Natural Language Structures*; Dordrecht: Reidel.
- Hooper, Joan Bybee (1976); *An Introduction to Natural Generative Phonology*; New York: Academic Press.
- Hudson, Grover (1980); "Automatic alternations in non-transformational phonology"; *Language* 56, 94-125.
- Hyman, Larry (1985); *A Theory of Phonological Weight*; Dordrecht: Foris.
- Inkelas, Sharon (1989); *Prosodic constituency in the lexicon*; PhD dissertation, Stanford University.
- Inkelas, Sharon, and Draga Zec, editors (1990); *The Phonology-Syntax Connection*; Chicago: University of Chicago Press.
- Itô, Junko (1986); *Syllable Theory in Prosodic Phonology*; PhD dissertation, University of Massachusetts, Amherst.

- Itô, Junko (1989); "A prosodic theory of epenthesis"; *Natural Language and Linguistic Theory* 7, 217–259.
- Jelinek, Eloise (1984); "Empty categories, case, and configurationality"; *Natural Language and Linguistic Theory* 2, 39–76.
- Jelinek, Eloise (1986); "The ergativity hypothesis and the argument hierarchy in Nisgha"; *Papers of the 21st International Conference on Salish and Neighboring Languages* Seattle: University of Washington.
- Johnson, David E., and Paul Postal (1980); *Arc Pair Grammar*; Princeton: Princeton University Press.
- Johnson, Mark (1988); *Attribute Value Logic and the Theory of Grammar*; CSLI Lecture Notes, Number 16; Chicago: University of Chicago Press.
- Johnson, Mark (1990); "Features, frames, and quantifier-free formulae"; pp. 94–107 in Patrick Saint-Dizier and Stan Szpakowicz, editors, *Logic and logic grammars for language processing*; New York: Ellis Horwood.
- Johnson, Mark (1991); "Features and formulae"; *Computational Linguistics* 17, 131–151.
- Kager, Rene (1992); "Shapes of the generalized trochee"; paper presented to the West Coast Conference on Formal Linguistics, Los Angeles.
- Kaplan, Ronald, and Joan Bresnan (1982); "Lexical-functional grammar: a formal system for grammatical representation"; pp. 173–281 in Joan Bresnan, editor, *The Mental Representation of Grammatical Relations*; Cambridge, MA: MIT Press.
- Kari, James (1976); *Navajo Verb Prefix Phonology*; PhD dissertation, University of New Mexico.
- Kari, John (1989); "Affix positions and zones in the Athapaskan verb complex: Ahtna and Navajo"; *International Journal of American Linguistics* 55, 424–454.
- Kay, Martin (1982); "Parsing in Functional Unification Grammar"; pp. 251–278 in D.R. Dowty, L. Karttunen, and A. Zwicky, editors, *Natural Language Parsing*; Cambridge: Cambridge University Press.
- Kaye, Jonathan (1990a); "Government in phonology: the case of Moroccan Arabic"; *The Linguistic Review* 6, 131–159.
- Kaye, Jonathan (1990b); "Coda licensing"; *Phonology* 7, 301–330.
- Kaye, Jonathan (1990c); "The strange vowel sets of Charm theory: the question from top to bottom"; *Journal of Linguistics* 26, 175–181.
- Kaye, Jonathan, Malika Echchadli, and Souad El Ayachi (1986); "Les formes verbales de l'arabe marocain"; *La phonologie des langues sémitiques; Revue québécoise de linguistique* 16, 61–99.

- Kaye, Jonathan, and John Harris (19xx); "A tale of two cities"; .
- Kaye, Jonathan, and Jean Lowenstamm (1984); "De la syllabicit  "; in Fran  ois Dell, Daniel Hirst, and Jean-Roger Vergnaud, editors, *Forme sonore du langage: structure des repr  sentations en phonologie*; Paris: Hermann.
- Kaye, Jonathan, and Jean Lowenstamm (1985); "Compensatory lengthening in Tiberian Hebrew"; pp. 97–132 in Leo Wetzels and Engin Sezer, editors, *Studies in Compensatory Lengthening*; Dordrecht: Foris.
- Kaye, Jonathan, Jean Lowenstamm, and Jean-Roger Vergnaud (1985); "The internal structure of phonological segments: a theory of charm and government"; *Phonology Yearbook* 2, 305–328.
- Kaye, Jonathan, Jean Lowenstamm, and Jean-Roger Vergnaud (1990); "Constituent structure and government in phonology"; *Phonology* 7, 193–231.
- Keating, Patricia (1988a); "The window model of coarticulation: articulatory evidence"; *UCLA Working Papers in Linguistics* 69, 3–29.
- Keating, Patricia (1988b); "Underspecification in phonetics"; *Phonology* 5, 275–292.
- Kenstowicz, Michael, and Charles Kisseberth (1979); *Generative Phonology*; New York: Academic
- Kisseberth, Charles (1970); "On the functional unity of phonological rules"; *Linguistic Inquiry* 1, 291–306.
- Koskenniemi, Kimmo (1983); phdTwo-level Morphology: A General Computational Model for Word-form Recognition and Production, University of Helsinki.
- Koutsoudas, Andreas, Gerald Sanders, and Craig Noll (1974); "The application of phonological rules"; *Language* 50, 1–28.
- Kowalski, Robert (1979); "Algorithm = logic + control"; *Communications of the ACM* 22, 424–436.
- LaCharit  , Darlene
(1993); "On the Need for Negative Constraints and Repair: Consonant Mutation in Setswana"; *Canadian Journal of Linguistics* 38, 257–278.
- Ladefoged, Peter, and Ian Maddieson (1986); *Some of the Sounds of the World's Languages*; UCLA Working Papers in Phonetics 64, November 1986.
- Lahiri, Aditi, and Vincent Evers (1991); "Palatalization and coronality"; in Carole Paradis and Jean-Fran  ois Prunet, editors, *Phonology and Phonetics 2: The Special Status of Coronals*; San Diego: Academic Press.
- Lakoff, George (1988); "A suggestion for a linguistics with connectionist foundations"; pp. 301–314 in D.S. Touretzky, G.E. Hinton, and T.J. Sejnowski, editors, *Proceedings of the 1988 Connectionist Models Summer School*; San Mateo, CA: Morgan Kaufmann.

- Lapointe, Steven (1980); *A Theory of Grammatical Agreement*; PhD dissertation, University of Massachusetts, Amherst.
- Leben, William (1973); *Suprasegmental Phonology*; PhD dissertation, MIT.
- Levin, Juliette (1985); *A Metrical Theory of Syllabicity*; PhD dissertation, MIT.
- Levin, Juliette (1988); "Deriving ternary feet"; *Texas Linguistic Forum* 29, 97–113.
- Lieberman, Mark (1975); *The Intonational System of English*; PhD dissertation, MIT.
- Lieberman, Mark, and Alan S. Prince (1977); "On stress and linguistic rhythm"; *Linguistic Inquiry* 8, 249–336.
- Lieber, Rochelle (1987); *An Integrated Theory of Autosegmental Processes*; Albany: State University of New York Press.
- Lieber, Rochelle (1989); "On percolation"; *Yearbook of Morphology* 2, 95–138.
- Lieber, Rochelle (1992); *Deconstructing Morphology: Word Formation in Syntactic Theory*; Chicago: University of Chicago Press.
- Lightner, Theodore M. (1963); "A note on the formulation of phonological rules"; *MIT Quarterly Progress Report of the Research Laboratory of Electronics* 66, 187–189.
- Lombardi, Linda (1991); "Laryngeal features"; ms.
- Maddieson, Ian (1988); "Linguo-labials"; *UCLA Working Papers in Phonetics* 68, 21–45.
- Marantz, Alex (1982); "Re reduplication"; *Linguistic Inquiry* 13, 435–482.
- Marcus, Mitchell (1980); *A Theory of Syntactic Recognition for Natural Language*; Cambridge, MA: MIT Press.
- McCarthy, John J. (1979); *Formal Problems in Semitic Phonology and Morphology*; PhD dissertation, MIT.
- McCarthy, John J. (1984); "Theoretical consequences of Montañes vowel harmony"; *Linguistic Inquiry* 15, 291–318.
- McCarthy, John J. (1986); "OCP effects: gemination and antigemination"; *Linguistic Inquiry* 17, 207–265.
- McCarthy, John J. (1988); "Feature geometry and dependency: a review"; *Phonetica* 43, 84–108.
- McCarthy, John J. (1989a); "Linear order in phonological representation"; *Linguistic Inquiry* 20, 71–99.
- McCarthy, John J. (1989b); "On gutterals"; ms, Brandeis University.
- McCarthy, John J. (1993); "A Case of Surface Constraint Violation"; *Canadian Journal of Linguistics* 38, 169–195.

- McCarthy, John J., and Alan Prince (1986); "Prosodic Morphology"; ms, University of Massachusetts, Amherst, and Brandeis University.
- McCarthy, John, and Alan Prince (1990a); "Prosodic Morphology and templatic morphology"; pp. 1–54 in Mushira Eid and John McCarthy, editors, *Perspectives on Arabic Linguistics II: Papers from the second annual symposium on Arabic linguistics*; Amsterdam: John Benjamins.
- McCarthy, John J., and Alan S. Prince (1990b); "Foot and word in Prosodic Morphology: the Arabic broken plural"; *Natural Language and Linguistic Theory* 8, 209–283.
- McCarthy, John J., and Alan S. Prince (1993); *Prosodic Morphology I: Constraint Interaction and Satisfaction*; ms, University of Massachusetts, Amherst, and Rutgers University.
- McCawley, James D. (1968); *The Phonological Component of a Grammar of Japanese*; The Hague: Mouton.
- McDonough, Joyce M. (1990); *Topics in the Phonology and Morphology of Navajo Verbs*; PhD dissertation, University of Massachusetts, Amherst.
- Mester, R. Armin (1986); *Studies in Tier Structure*; PhD dissertation, University of Massachusetts, Amherst.
- Mester, R. Armin, and Junko Itô (1989); "Feature predictability and underspecification: palatal prosody in Japanese mimetics"; *Language* 65, 258–293.
- Miller, Philip H. (1991); *Clitics and Constituents in Phrase Structure Grammar*; PhD dissertation, University of Utrecht.
- Mohanan, K.P. (1986); *The Theory of Lexical Phonology*; Dordrecht: Reidel.
- Mohanan, Tara (1989); "Syllable structure in Malayalam"; *Linguistic Inquiry* 20, 589–625.
- Nespor, Marina, and Irene Vogel (1986); *Prosodic Phonology*; Dordrecht: Foris.
- Ola, Olanike (1992); *Yoruba Vowel Harmony*; MA thesis, SOAS.
- Ouhalla, Jamal (1991); *Functional Categories and Parametric Variation*; London: Routledge.
- Padgett, Jaye (1991); *Structure in Feature Geometry*; PhD dissertation, University of Massachusetts, Amherst.
- Paradis, Carole (1988a); "On constraints and repair strategies"; *The Linguistic Review* 6, 71–97.
- Paradis, Carole (1988b); "Towards a Theory of Constraint Violations"; *McGill Working Papers in Linguistics* 5, 1–43.
- Paradis, Carole (1990); "Focus in Gere configurational constraints"; *Current Approaches to African Linguistics* 7, 53–62.

- Paradis, Carole (1993); "Ill-Formedness in the Dictionary: A Source of Constraint Violation"; *Canadian Journal of Linguistics* 38, 215–234.
- Paradis, Carole, and Fatimazohra El Fenne (to appear); "French verbal inflection revisited: constraints, repairs and floating consonants"; to appear in J. Durand, M.-A. Hintze, and A. Battye, eds., *New Trends in French Phonology and Phonetics*, *Lingua*.
- Paradis, Carole, and Jean-François Prunet (1993); "On the validity of Morpheme Structure Constraints"; *Canadian Journal of Linguistics* 38, 235–256.
- Penny, Ralph J. (1969a); *El habla pasiega: ensayo de dialecológica montañesa*; London: Tamesis Books.
- Penny, Ralph J. (1969b); "Vowel-harmony in the speech of the Montes de Pas (Santander)"; *Orbis* 18, 148–166.
- Petros Banskira, Degif (1992); "Vowel licensing and empty onsets"; *Proceedings of the 1992 Annual Conference of the Canadian Linguistic Association; Toronto Working Papers in Linguistics*, 1–14.
- Pierrehumbert, Janet (1993); "Dissimilarity in the Arabic verbal roots"; *NELS* 23.
- Pierrehumbert, Janet, and Mary Beckman (1989); *Japanese Tone Structure*; Albany: State University of New York Press.
- Piggott, Glynne (1988); "The parameters of nasalization"; ms, McGill University.
- Piggott, Glynne (1989); "Variability in feature dependency: the case of nasality"; ms, McGill University.
- Piggott, Glynne (1991); "Empty onsets: evidence for the skeleton in prosodic Phonology"; ms, McGill University.
- Pollard, Carl, and Ivan Sag (1987); *Information-Based Syntax and Semantics, Vol. I*; CSLI Lecture Notes, Number 13; Chicago: University of Chicago Press.
- Poser, William J. (1982); "Phonological representations and action-at-a-distance"; pp. 121–158 in Harry van der Hulst and Norval Smith, editors, *The Structure of Phonological Representations, Part II*; Dordrecht: Foris.
- Prince, Alan S. (1983); "Relating to the grid"; *Linguistic Inquiry* 14, 19–100.
- Prince, Alan S., and Paul Smolensky (1993); *Optimality Theory: Constraint Interaction in Generative Grammar*; Technical Report 2, Rutgers Center for Cognitive Science.
- Pulleyblank, Douglas (1983); *Tone in Lexical Phonology*; PhD dissertation, MIT. [published 1986, Dordrecht: Reidel]
- Pulleyblank, Edwin G. (1989); "The role of coronal in articulator based features"; *Chicago Linguistics Society* 25, 379–394.
- Reiter, Raymond (1980); "A logic for default reasoning"; *Artificial Intelligence* 13, 81–132

- Ringen, Catherine O. (1977); *Vowel Harmony: Theoretical Implications*; PhD dissertation, Indiana University.
- Ringen, Catherine O. (1988); "Transparency in Hungarian vowel harmony"; *Phonology* 5, 327–342.
- Russell, Kevin (1993); "The internal structure of feet: generalized iambs in Cayuvava"; ms, University of Southern California; paper presented to the annual meeting of the Linguistic Society of America, Los Angeles.
- Sadanandan, Suchitra (1990); "Malayalam syllabification: a Government Phonology approach"; ms, University of Southern California.
- Sadock, Jerrold M. (1991); *Autolexical Syntax: A Theory of Parallel Grammatical Representations*; Chicago: University of Chicago Press.
- Sagey, Elizabeth (1986); *The Representation of Features and Relations in Non-linear Phonology*; PhD dissertation, MIT.
- Sagey, Elizabeth (1988a); "On the ill-formedness of crossing association lines"; *Linguistic Inquiry* 19, 109–118.
- Sagey, Elizabeth (1988b); "Degree of closure in complex segments"; pp. 169–208 in Harry van der Hulst and Norval Smith, editors, *Features, Segmental Structure and Harmony Processes*, vol. 1; Dordrecht: Foris.
- San Duanmu (1990); *A Formal Study of Syllable, Tone, Stress and Domain in Chinese languages*; PhD dissertation, MIT.
- Schane, Sanford (1984); "The fundamentals of particle phonology"; *Phonology Yearbook* 1, 129–155.
- Schein, Barry, and Donca Steriade (1986); "On geminates"; *Linguistic Inquiry* 17, 691–744.
- Schindwein, Deborah "A segmental moraic theory of representation: a reply to Hayes"; ms, University of Southern California.
- Scobbie, James (1991); *Attribute Value Phonology*; PhD dissertation, University of Edinburgh.
- Scobbie, James (1993); "Constraint Violation and Conflict from the Perspective of Declarative Phonology"; *Canadian Journal of Linguistics* 38, 155–167.
- Scott, Graham (1978); *The Fore Language of Papua New Guinea*; Canberra: Pacific Linguistics.
- Selkirk, Elisabeth (1982); *The Syntax of Words*; Cambridge, MA: MIT Press.
- Selkirk, Elisabeth (1984); *Phonology and Syntax: The Relation Between Sound and Structure*; Cambridge, MA: MIT Press.

- Selkirk, Elisabeth (1986); "On derived domains in sentence phonology"; *Phonology Yearbook* 3, 371–405.
- Selkirk, Elisabeth (1988); "Dependency, place and the notion 'tier'"; ms, University of Massachusetts, Amherst.
- Selkirk, Elisabeth (1991); "Major place in the vowel space"; ms, University of Massachusetts, Amherst.
- Shaw, Patricia A. (1985); "The role of prosodic structure in reduplication"; paper presented to the annual meeting of the Canadian Linguistic Association, Montreal, and ms, University of British Columbia.
- Shaw, Patricia A. (1991); "Consonant harmony systems: the special status of coronal harmony"; pp. 125–157 in *Phonetics and Phonology, volume 2: The Special Status of Coronals*; New York: Academic.
- Shieber, Stuart M. (1986); *An Introduction to Unification-Based Approaches to Grammar*; CSLI Lecture Notes, Number 4; Chicago: University of Chicago Press.
- Smolensky, Paul (1986); "Information Processing in Dynamical Systems: Foundations of Harmony Theory"; pp. 194–281 in David E. Rumelhart, James L. McClelland, and the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, volume 1: Foundations*; Cambridge, MA: MIT Press..
- Speas, Margaret (1990); *Phrase Structure in Natural Language*; Dordrecht: Reidel.
- Spencer, Andrew (1991); *Morphological Theory*; Oxford: Blackwell.
- Spring, Cari (1990a); "The phonological components of reduplication"; ms, University of Arizona.
- Spring, Cari (1990b); "Unordered morphology: the problem of Axininca reduplication"; *Proceedings of the Sixteenth Annual Meeting of the Berkeley Linguistics Society: Special Session on General Topics in American Indian Linguistics*, 137–157.
- Stabler, Edward P., Jr. (1990); "Parsing as logical constraint satisfaction"; pp. 72–93 in Patrick Saint-Dizier and Stanislaw Szpakowicz, editors, *Logic and Logic Grammars for Language Processing*; New York: Ellis Horwood.
- Stabler, Edward P., Jr. (1992); "Implementing Government Binding theories"; pp. 243–275 in Robert Levine, editor, *Formal Grammar: Theory and Implementation*; Oxford: Oxford University Press.
- Stabler, Edward P., Jr. (1993); *The Logical Approach to Syntax*; Cambridge, MA: MIT Press.
- Stanley, R. (1967); "Redundancy rules in phonology"; *Language* 43, 393–436.
- Stemberger, Joseph Paul (1981); "Morphological haplology"; *Language* 57, 791–817.

- Steriade, Donca (1982); *Greek Prosodies and the Nature of Syllabification*; PhD dissertation, MIT.
- Steriade, Donca (1987); "Redundant values"; *Papers from the Parasession on Autosegmental and Metrical Phonology, Chicago Linguistic Society 23*, 339–362.
- Steriade, Donca (1988); "Reduplication and syllable transfer in Sanskrit and elsewhere"; *Phonology 5*, 73–155.
- Steriade, Donca (1993a); "Complex Onsets as Single Segments: The Mazateco Pattern"; ms, UCLA.
- Steriade, Donca (1993b); "Segments, contours, and clusters"; ms, UCLA; reduced version to appear in Proceedings of the 15th International Congress of Linguists.
- Steriade, Donca (1993c); "Preoralization, postoralization and markedness"; paper presented to the Berkeley Linguistics Society, and ms, UCLA.
- Steriade, Donca (1993d); "Closure. Release, and Nasal Contours"; pp 401–470 in Huffman and Krakow, eds., *Phonetics and Phonology 5: Nasals, Nasalization, and the Velum*; New York: Academic Press.
- Stevens, Kenneth (1989); "On the quantal nature of speech"; *Journal of Phonetics 17*, 3–45.
- Tarpen, Marie-Lucie (1987); *A Grammar of the Nisgha Language*; PhD dissertation, University of Victoria.
- Touretzky, David S., and Deirdre W. Wheeler (1989); "A connectionist implementation of Cognitive Phonology"; paper presented at the Berkeley Workshop on Cognitive Phonology, May 1989.
- Tucker, A.N. (1964); "Kalenjin phonetics"; pp. 445–470 in David Abercrombie, D.B. Fry, P.A.D. MacCarthy, N.C. Scott, and J.L.M. Trim, editors, *In Honour of Daniel Jones: Papers Contributed on the Occasion of his Eightieth Birthday, 12 September 1961*; London: Longmans, Green and Co.
- Vago, Robert M. (1980); *The Sound Pattern of Hungarian*; Washington: Georgetown University Press.
- Vago, Robert M. (1988); "Underspecification in the height harmony system of Pasiego"; *Phonology 5*, 343–362.
- van Benthem, Johan (1983); *The Logic of Time*; Dordrecht: Reidel.
- van Benthem, Johan (1988); *Manual of Intensional Logic*; second edition; CSLI Lecture Notes No. 1; Chicago: University of Chicago Press.
- Vance, Timothy J. (1987); *An Introduction to Japanese Phonology*; Albany: State University of New York Press.

- Waksler, Rochelle (1990); *A Formal Account of Glide/Vowel Alternations in Prosodic Theory*; PhD dissertation, Harvard University.
- Walker, Douglas C. (1984); *The Pronunciation of Canadian French*; Ottawa: University of Ottawa Press.
- Wheeler, Deirdre (1981); *Aspects of a categorial theory of phonology*; PhD dissertation, University of Massachusetts, Amherst.
- Wilbur, Ronnie *The Phonology of Reduplication*; PhD dissertation, University of Illinois.
- Williams, Edwin (1976); "Underlying tone in Margi and Igbo"; *Linguistic Inquiry* 7, 463–484.
- Wiltshire, Caroline R. (1992); *Syllabification and rule application in Harmonic Phonology*; PhD dissertation, University of Chicago.
- Woods, W.A. (1970); "Transition Network Grammars for natural language analysis"; *Communications of the Association for Computing Machines* 3, 591–606.
- Young, Robert W., and William Morgan, Sr. (1987); *The Navajo Language: A Grammar and Colloquial Dictionary; revised edition*; Albuquerque: University of New Mexico Press.