

6

Further Issues

6.1 INTRODUCTION

In this final chapter, we shall discuss briefly several remaining areas of current interest. First, we shall consider the relation of feature analysis to autosegmental structure, and look at recent proposals for feature trees in phonology. Second, we shall briefly consider the nature of some vowel systems and the treatment of vowel harmony within an autosegmental framework. Third, we shall examine the Obligatory Contour Principle and the Morpheme Tier Hypothesis, two proposals governing the organization of feature specifications on autosegmental tiers. Finally, we shall consider some proposals concerning the nature of phonological rule application and their relation to phonotactic conditions.

6.2 FEATURES

6.2.1 *Introduction*

What are features? The term has been used to cover a number of related notions, and closely related terms (such as *components*) have been used in similar ways. The term *feature* covers two notions which are logically distinct, but which have considerable overlap in actual practice.

First, 'feature' refers to a notion that organizes a *classificatory* scheme; it provides a way of establishing what the 'natural classes' in phonological statements will be. This use has been most clearly described in, e.g., Halle (1962). The idea is by now a familiar one: statements regarding distribution and phonological processes typically involve not a single segment, but larger and smaller groups of segments, which we call *natural classes*. The sense in which these classes are 'natural' is that they recur across the generalizations of a single language, and across the various human languages of the world. Features are used to define and, to some degree, predict what possible natural classes are. Each feature F

defines two sets of segments, those that are +F and those that are -F, and the claim of classical generative phonology has been that the natural classes of human languages are defined by the various intersections of these sets.¹ For example, the feature [voice] defines the set of segments that are [+voice] and the set of segments that are [-voice], and the feature [sonorant] similarly defines two sets, those that are [+sonorant] and those that are [-sonorant]. The intersection of two or more of these, such as the set of segments that are [-sonorant, -voice], is a natural class, and such sets can be found undergoing phonological rules as a group. This conception of features is often called *classificatory*.

Second, features may be viewed as a way of specifying the several and simultaneous characteristics that comprise what is, from the point of view of the flow of time, a single articulatory or acoustic event. This latter conception of features (or 'components') is most clearly developed in the post-Bloomfieldian phonology² of the 1940s and 1950s, especially in work of Harris (1944), Hockett (1947, 1955) and Bloch (1948), and it has recently re-emerged in discussions of autosegmental phonology (Sagey 1986, Halle 1988, *inter alia*), as well as in Anderson (1974) in a slightly different context. We will refer to this as the *componential* notion of the feature.

By and large, the post-Bloomfieldian tradition was steeped in positivism, the view that scientific discourse and practice has as its goal the classification and organization of particular experiences (rather than, for example, methodically speculating about the internal workings of a physical or biological device which we will never actually be able to peer into). Scientific knowledge, on one version of this empiricist view, is pyramidal in shape, and rests on a broad base of observations and experiences accessible to all interested scientists; in this case, those observations and experiences are particular linguistic utterances in space and time. The true foundations of an empiricist phonology, then, will be the principles that govern how descriptions of linguistic utterances are consistently classified and described, and one view of features fits naturally into this perspective.

The process of classifying and encoding particular linguistic utterances was traditionally viewed as a problem of making 'vertical cuts', and sometimes 'horizontal cuts' as well – slicing the continuously divisible flow of time into chunks, which could be viewed as more or less homogeneous throughout. As a number of linguists observed, if we record the linguistic act in question from the point of view of the articulatory apparatus, then it makes considerable sense – both common sense and linguistic sense – to focus separately on individual aspects of the speech event, to make horizontal cuts for each such aspect. Thus the sequence of actions in time for each subact might be viewed as forming a

stream in time, and this stream would be sliced vertically into segments. If it should turn out that there was a reasonably simple relationship between the slicing, or segmentation, for each articulator whereby the cuts would more or less line up appropriately in time, like the playing of a well trained orchestra, then we could think of the chunks of the smaller-scale actions as the features, or components, of the complete unit, the total action of the articulatory apparatus. In (1), I have reproduced Hockett's (1955: 134) tiered account of a Nootka utterance, in a passage immediately followed by the quotation that appears in the Introduction to the present volume.

(1) q' a· q' a n a λ' a ʔ i· ɸ

	q'	a·	q'	a	n	a	λ'	a	ʔ	i·	ɸ
Position	bd		bd		ap		lt				
Spirant/glottal	gl		gl				gl		gl		
Spirantal release	—		—				+				+
Nasal					+						
Pharyngeal constriction									+	(+)	+
Height		lo		lo		lo		lo		hi	
Round/front										fr	
Shortness		—									—

(bd = back dorsal; ap = apical; lt = lateral)

Hockett observed, as well, that the horizontal slicing induced in (1) is not the only one possible; it is imposed not by the data, but by our analytical process and choice. The features that arise in this fashion roughly correspond to the features familiar in most of classical generative phonology, though the point-of-articulation feature in (1) is not binary. However, another account is possible, in which each row corresponds to a separate articulator. Such an arrangement was avoided (for better or worse) in the representation in (1) because in Nootka, as Hockett observes, oral articulators function almost exclusively one at a time. An alternative tiered account includes a tier for each of the following: the lower lip; the apex of the tongue; the blade of the tongue; the tongue as a whole; the front part of the dorsum; the back part of the dorsum; the velic; the pharynx; and the glottis. Hockett then arrives at the account given in (2) of the same Nootka word as in (1), where *k* stands for 'closure', *kl* for what we might call 'lateral closure', *c* for 'constriction', *o* for 'open', and *v* for 'voicing'.

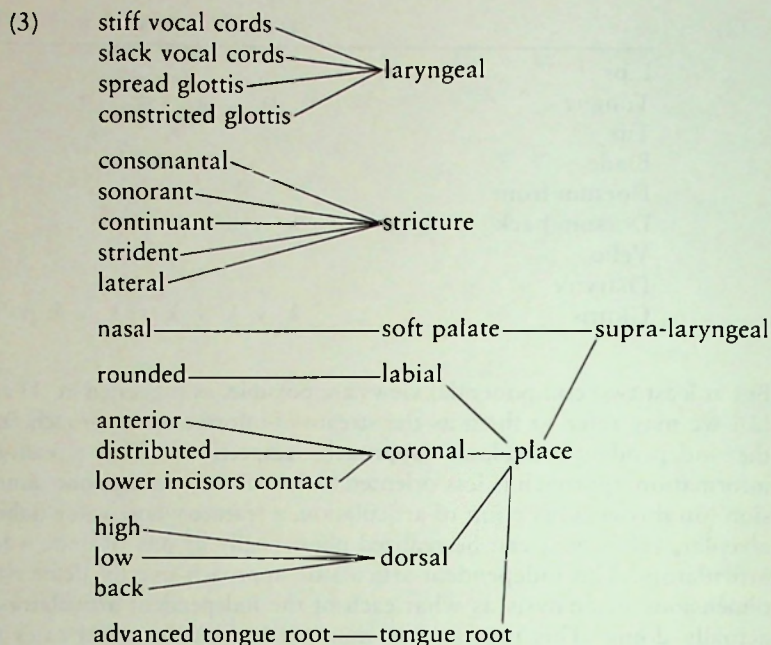
This componential notion of the feature takes its inspiration from an exploration of articulatory phonetics, and not from acoustic phonetics, which provides no such natural way of slicing spectrograms horizontally.

(2)	q'	a	q'	a	n	a	λ'	a	ʔ	i	h
Lips											
Tongue		a		a		a		a		i	
Tip					k						
Blade								kl			
Dorsum/front											
Dorsum/back		k		k							
Velic						o					
Pharynx									c		c
Glottis		k	v	k	v	k	v	k	v	k	v

But at least two componential views are possible, as suggested in (1) and (2): we may refer to them as the stream-of-information approach, and the independent-articulator approach, respectively. The stream-of-information approach is less oriented to articulatory reality; one dimension (on this view) is point of articulation, a feature whose value (labial, alveolar, velar, etc.) can be realized phonetically by any of quite a few articulators. The independent-articulator approach overtly defines the dimensions of analysis as what each of the independent articulators is actually doing. This may seem a minor point, but in actual cases the difference between the two as regards the choice of phonological features can be significant. A recent proposal (Halle 1988) synthesizing current work steers a middle ground between the two, offering a system (3) of eighteen features, organized into groupings whose significance we will return to in the next section.³

Certainly there is a close kindredness of spirit between the componential view of segmentation and autosegmental phonology. From the point of view of the phonologists who have developed and applied it, though, autosegmental phonology does not derive its multi-tiered structure from a decision as to how best to translate a fine-grained description of an articulatory event into one consisting of the discrete units called segments, autosegments, or components.⁴ No doubt the structure of the articulators, and the neurological network that governs their behavior, serves as the starting point in the development of the phonological system (it is no coincidence that most features correspond to independent articulators); but the anatomic system proposes, while the phonology disposes.

The statement that the segment is a 'simultaneous actualization of a set of attributes' (Halle 1964) emphasizes the componential view of the segment, but it should be clear that not all characteristics that are features in the *classificatory* sense are *ipso facto* features in the componential sense, and the ways in which the two can differ are several, of



which we shall mention only a few. First, a classificatory feature may have a purely temporal realization; for example, one may establish a feature [\pm long] (though we have seen empirical reasons not to); such a feature is not componential. Second, along a similar line, the stress of a vowel may be analyzed as a feature (as it was in the *SPE* tradition); but, as was emphasized in early metrical work, there is little or no simple phonetic *manifestation* of such a classificatory feature. Once again, such a feature manifests itself, in part, in a more global set of temporal properties. Third, a consistent phonological difference may have no specific phonetic basis that can be impressed into servitude. The classic example is the feature that separates the final segment of *wife* and *knife* from *cliff* and *dove*; if the last two are voiceless and voiced, respectively (i.e. /f/ and /v/), then what is the final segment of *wife*, a segment that undergoes a rule voicing it when followed by a voiced segment within the same word (cf. the plural forms *wi[vz]*, but *cli[fs]*)? Some feature must distinguish it from both *f* and *v*. In a purely classificatory scheme, a new feature could be set up, and much debate ensued as to just how far removed from the componential sense classificatory features should be allowed to be.⁵ Virtually all linguists agree that little divergence between the two should be permitted, but disagreement still arises as to how little is too much.⁶

In the discussion in the literature to date, relatively little work has gone into the differences in feature systems that might arise from the division of phonology into lexical and post-lexical branches; the two components naturally lean toward classificatory and componential views of features, respectively. As we noted in the preceding chapter, much work on lexical phonology has proceeded on the assumption that all features are gradually changed from privative to equipollent⁷ during the lexical phonology, and that a number of (equipollent) features are unavailable for each given language during its lexical phonology. Relatively little concern has been expressed in recent work as to whether the differences in the sets and the organization of features may vary in the two components of the grammar.⁸

6.2.2 *Two models of feature organization*

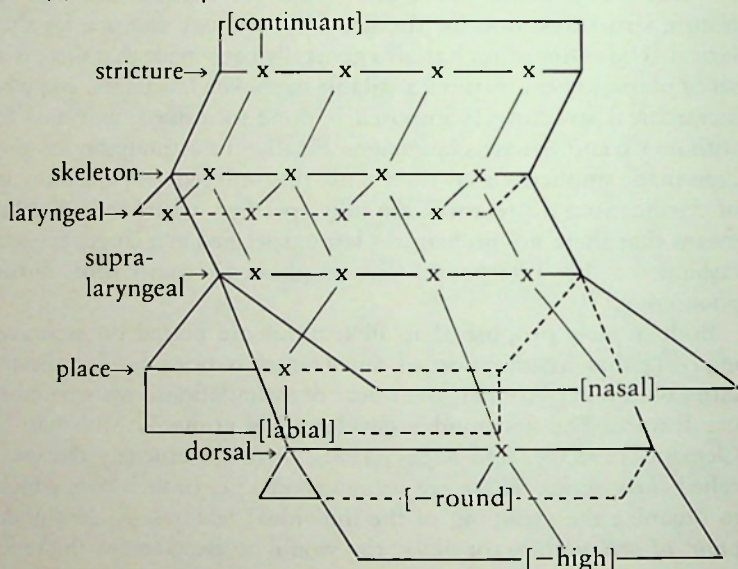
In this section, we will sketch two models that are currently being considered for the organization of features.⁹ Most of the processes that have appeared to be relevant to the subject have been post-lexical, so one may interpret these discussions as being aimed primarily at providing a model of the post-lexical phonology. Discussions in the literature for the most part have not addressed the question as to whether a particular proposal is intended to be interpreted as holding in the lexical or the post-lexical phonology, presumably on the unspoken assumption that, all other things being equal, one would not want to posit two distinct feature structures, one for the lexical phonology and one for the post-lexical. The assumption has also generally been made that there is a fixed set of phonological features available to spoken languages, and whatever hierarchical structure is imposed is done to a fixed, universal format, with no variation across languages. Finally, the assumption has generally been made implicitly that (in a sense that still requires a certain amount of clarification) all features are fully specified, where 'full specification' means that there are no features left unspecified in a language-particular fashion to be interpreted by a phonetic component outside of phonology.¹⁰

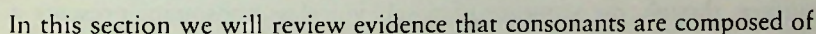
Both models propose that all features are placed on separate tiers, implying that assimilation of any feature is possible. But the theories differ with regard to their treatment of assimilations involving more than one feature. The first model, developed by primarily Mohanan (1983), Clements (1985b), and Sagey (1986), involves crucially the use of so-called *class nodes*. These are segments on a tier of their own which serve to organize the grouping of the individual features. A class node on a *point-of-articulation* (or *place*) tier would be associated with the feature-autosegments that determine point of articulation, and it itself might in

turn be associated with a higher class node on yet another tier (in (3), this is the 'supra-laryngeal node'), one that gathered together all of the autosegmentalized features of the segment to its left in (3).) All of the segment's specifications would be associated to a *root* node, and that segment would, in turn, be associated to the skeletal position; for purposes of simplicity, we have conflated the skeletal and root tier in (4)-(5), but we return to this on pp. 292-4). When the point of articulation of one segment associates leftward, for example, onto a preceding segment, it is the class-node autosegment that reassociates, on this view. All this is exemplified in the sketch of the word *spin* in (4), assuming a set of features as in (3). In (4a), I indicate only the distinctive feature specifications; in a framework that did not use underspecification, full specification of each feature would be required, and the representation would be quite difficult to place on a piece of paper. In (4b), I give an example of what one vertical 'slice' of that representation would look like, for the segment *s*, the first segment in *spin*. Each of the four segments in (4a) would be equally spelled out for each feature, on such a view.

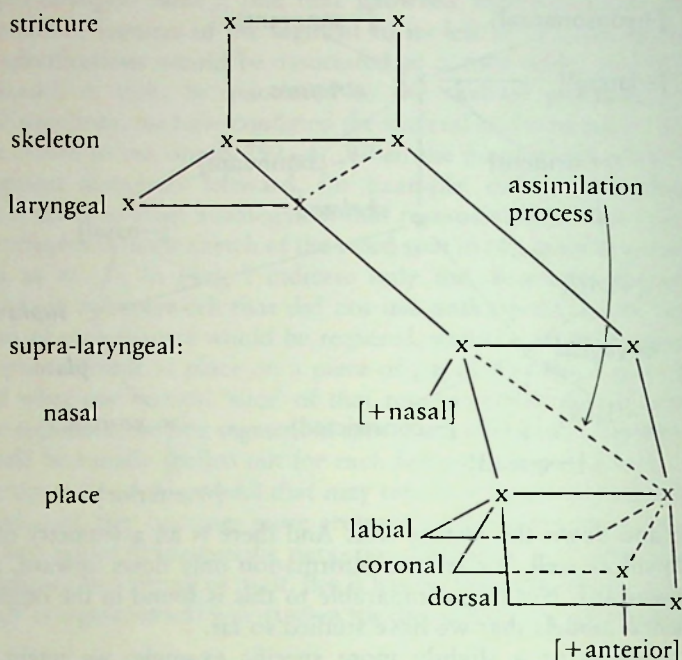
One characteristic of this model that may render it less transparent is that the notion of 'tier' becomes more abstract, and it plays more than one role. Class nodes define specific tiers; the class node for 'place', for example, appears on a tier all to itself. But it has no features specific to it: it serves only as a geometrical way-station for the passage of information

(4) (a) the word *spin*





(5) Assimilation of a nasal's point of articulation

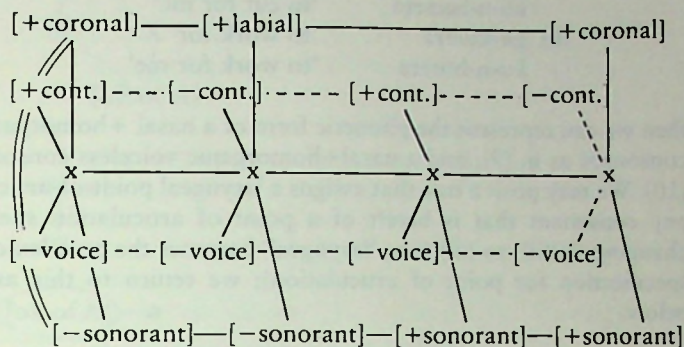


separate subsections corresponding to material on separate autosegmental tiers, of which the most striking is point of articulation, a notion that we have encountered quite a few times in the discussions of Spanish, Catalan, Selayarese, English, and several other languages. Most of the phenomena that we observe in this section concern rules of assimilation, and it is natural to hypothesize that all rules of assimilation must be treated autosegmentally – that is, that all rules of assimilation must be analyzed as the spreading of an autosegment over a larger domain.¹¹

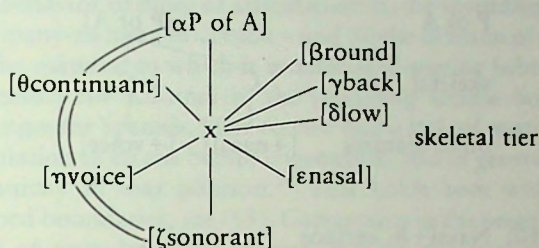
As already noted (chapter 1), KiRundi is a Bantu language, and as in all Bantu languages, a nasal must be homorganic with a following voiced consonant, as illustrated in (7a). Underlyingly, there is a contrast among the nasals between three of the four points of articulation, as can be seen in the nasals in onset position in (7b); there is no underlying velar nasal. In this respect it is much like the situation in Catalan that we considered in chapter 5.

The contrast among the nasals is lost, however, before a voiced consonant; there we find homorganic sequences as in (7a) only. On the surface, though, the sequence of nasal+voiceless consonant presents a different pattern. Here, we find all four surface nasals (*m*, *n*, *ɲ*, *ŋ*)

(6) (a) Spiral-notebook model



(b) Side view



contrasting, but we do not find the nasal followed by a voiceless stop; where we might expect that, we find instead an *h*. It is quite obvious what is happening here, doubly so in the light of alternations as in (8). Clearly, in traditional segmental terms this would be described in two steps. A nasal is always homorganic to a following consonant, and a voiceless consonant following a nasal loses its oral point of articulation, leaving only its aspirate character behind, what we transcribe as an *h* in syllable onset position.

If we establish a separate autosegmental tier for point of articulation,

(7) Kirundi

(a) Nasals homorganic to following voiced consonant

ku-m-bona 'to see me'

i-n-dobo 'bucket'

umu-ganga 'doctor'

(b) ku-nanira 'to be hard for'

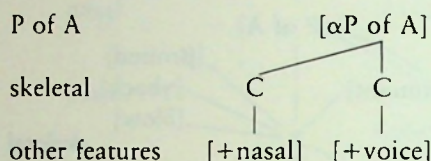
ku-mera 'to grow'

ku-ñwa 'to drink'

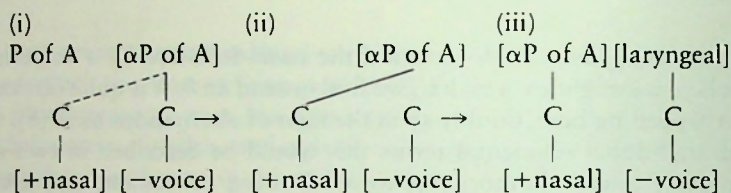
- (8) (a) ku-temera 'to cut for someone'
 ku-n-hemera 'to cut for me'
 (b) gu-korera 'to work for'
 ku-ŋ-horera 'to work for me'

then we can represent the phonetic form of a nasal + homorganic voiced consonant as in (9), and a nasal + homorganic voiceless consonant as in (10). We may posit a rule that assigns a laryngeal point-of-articulation to any consonant that is bereft of a point of articulation specification, changing (10(ii)) to (iii) (i.e., 'laryngeal' becomes the post-lexical default specification for point of articulation); we return to this assumption below.

(9) Nasal + voiced consonant: surface

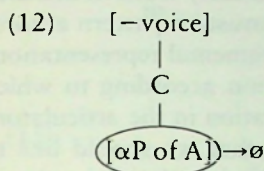
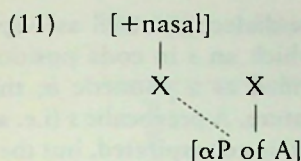


(10) Nasal + *h*: surface



It follows, then, that if the surface sequence of nasal + *h* derives from a deeper representation in which the nasal is followed by a voiceless consonant with a real point of articulation, then we must have two rules at work here. One spreads, or assimilates, the point-of-articulation autosegment from the obstruent leftward to the nasal; the second dissociates the point-of-articulation autosegment from the voiceless consonant, so that it becomes, so to speak, an *h*.

How do we formulate the rule that assimilates the point of articulation of the nasal to that of the following consonant? This question is easy to answer, in light of our autosegmental formalism: it is as in (11) (a rule that we have already seen; cf. chapter 5, rule (8), for Catalan). The rule that deletes the point of articulation of the voiceless consonant is given in (12).



We have seen so far, then, two of the phonological characteristics of autosegments in the behavior of point of articulation: (i) the spreading of the autosegment – a many-to-one association – and (ii) the deletion of the autosegment, with the material to which it is associated staying behind.

The analysis suggested for KiRundi in the preceding section holds equally with few changes for Spanish. Here too we find a lack of contrast in the point of articulation of nasals before consonants, and in general a pattern of homorganicity in that position.¹² This holds both within words and across word boundaries; see (13). Consonants in the onset do not lose their point of articulation in Spanish like the voiceless consonants in KiRundi, but Spanish does show one further development. In many Caribbean dialects, nasals show no contrast in point of articulation when in coda position, but manifest this by having a consistent velar articulation instead of being homorganic to a following consonant. In still others, as Harris (1969: 15–16) observes, a contour nasal is created. Harris notes that ‘many Cubans pronounce *enfermo* as [e_n^hmfermo], where the first nasal, presumably a systematic phoneme *n*, is realized with no alveolar contact at all, but rather with a labio-dental articulation superimposed on a dorso-velar articulation.’

- (13) Indefinite article *un* preceding consonants in Spanish (after Harris 1969)

u[m] beso	‘a kiss’
u[m̠] foco	‘a focus’
u[n] dedo	‘a finger’
u[ŋ] gato	‘a cat’

Spanish does not line up with KiRundi, however, in affecting a voiceless consonant in the onset of its syllable, whether a nasal precedes

or not. However, many Western Hemisphere dialects, as well as a few continental ones, do possess a process by which an *s* in coda position loses its oral gesture, and is typically identified as a phonetic *h*; this process is referred to as *aspiration* in the literature. A prevocalic *s* (i.e. an onset *s*) is not aspirated; thus the *s* in *eso* 'that' is not aspirated, but the *s* in *frasco* 'bottle' is aspirated.

The KiRundi and Spanish examples, along with early work in this area by Thráinsson (1978) on Icelandic, point to a consistent pattern according to which consonants weaken to *h*, and autosegmental representations have attempted to integrate the traditional notion according to which this weakening is the loss of any gesture-specification in the articulatory tract above the larynx. There is a sense, then, that we should like to capture in which the laryngeal *h* is truly unspecified for its oral point of articulation, though in a somewhat different sense from the way the term 'unspecified' has been used up to now. In the lexical phonology, all underspecified consonants that we have seen have been realized as alveolar consonants; now, in the post-lexical phonology, we find that, when a segment loses its point-of-articulation specification, it is realized as a glottal segment. Why this difference? One simple answer would be that there is a rule *P* that specifies obstruents as alveolar if unspecified for point of articulation. Any rule deleting a point-of-articulation specification *after* rule *P* has its chance to apply will create an *h*; any before, an alveolar. If rule *P* is a post-lexical rule, this will have the desired results. This conception is certainly too simplistic to serve as an ultimate solution to the observed difference in the two types of 'unspecification', but it serves as an adequate way to summarize our observations.

6.2.4 Toba Batak

Hayes (1986b) offers an account of several assimilation and weakening processes in Toba Batak that is rich in consequences for the theory of placement of features on separate tiers, and for underspecification theory, as well as for the theory of licensing and the Conjunctivity Condition (referred to as the 'Linking Constraint' in Hayes (1986a)).

Syllables in Toba Batak are maximally of the form CVC, with a restricted range of consonants appearing in the coda. Of the consonants in the phonemic inventory (see (14)), only the three nasals, *l*, *r*, and the four non-low voiceless obstruents appear in the coda before various optional sandhi rules apply. The onset may be host to any of the fourteen consonants; cf. (14). If we attempt to analyze these segments into their component features, we may arrive at the system in (15), should we follow Hayes, who takes all features to be equipollent at this post-lexical level.

(14)	Bilabial	Alveolar	Palato- alveolar	Velar	Glottal
	p	t		k	
	b	d		g	
			j		
		s			(h)
	m	n		ŋ	
		r			
		l			

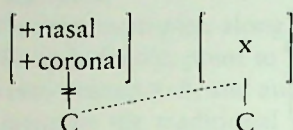
(15)		p	t	k	b	d	g	j	s	h	m	n	ŋ	r	l	ʔ
	Anterior	+	+	-	+	+	-	-	+	-	+	+	-	+	+	-
	Coronal	-	+	-	-	+	-	+	+	-	-	+	-	+	+	-
	Continuant	-	-	-	-	-	-	-	+	+	-	-	-	-	-	-
	Strident	-	-	-	-	-	-	-	+	-	-	-	-	-	-	-
	Sonorant	-	-	-	-	-	-	-	-	-	+	+	+	+	+	-
	Nasal	-	-	-	-	-	-	-	-	-	+	+	+	+	-	-
	Voice	-	-	-	+	+	+	+	-	-	+	+	+	+	+	-
	High	-	-	+	-	-	+	+	+	-	-	-	+	-	-	-
	Glottis:															
	Spread	-	-	-	-	-	-	-	-	+	-	-	-	-	-	-
	Constricted	-	-	-	-	-	-	-	-	-	-	-	-	-	-	+

Hayes's chart (16) specifies the possible changes that may occur to sequences of consonants appearing across syllable boundary, whether within the same word or across word boundary. To determine the surface form of a C_1C_2 sequence, we find the first consonant (the one in the coda position) by reading down the side, and the second by reading across.¹³ There are four major processes at work here, and Hayes

(16)	C_2															
	p	t	k	h	s	b	d	j	g	m	n	ŋ	r	l		
C_1	p	ʔp	ʔt	ʔk	pp	ʔs	ʔb	ʔd	ʔj	ʔg	ʔm	ʔn	ʔŋ	ʔr	ʔl	
	t	ʔp	ʔt	ʔk	tt	ʔs	ʔb	ʔd	ʔj	ʔg	ʔm	ʔn	ʔŋ	ʔr	ʔl	
	k	ʔp	ʔt	ʔk	kk	ʔs	ʔb	ʔd	ʔj	ʔg	ʔm	ʔn	ʔŋ	ʔr	ʔl	
	s	sp	st	sk	ss	ss	sb	sd	sj	sg	sm	sn	sŋ	sr	sl	
	m	pp	pt	pk	pp	ps	bb	md	mj	mg	mm	mn	mŋ	mr	ml	
	n	pp	tt	kk	kk	ss	bb	dd	jj	gg	mm	nn	ŋŋ	rr	ll	
	ŋ	kp	kt	kk	kk	ks	ŋb	ŋd	ŋj	ŋg	ŋm	ŋn	ŋŋ	ŋr	ŋl	
	r	rp	rt	rk	rh	rs	rb	rd	rj	rg	rm	rn	rŋ	rr	ll	
	l	lp	lt	lk	lh	ls	lb	ld	lj	lg	lm	ln	lŋ	lr	ll	

formulates them as shown in (17)–(20). *n*-assimilation, written with only two tiers (a skeletal tier, and a tier for all other features), instructs us that a segment following an *n* will be associated to the *n*'s skeletal position. In rules (18)–(20), Hayes assumes that the features are split onto three tiers:

(17) *n*-assimilation

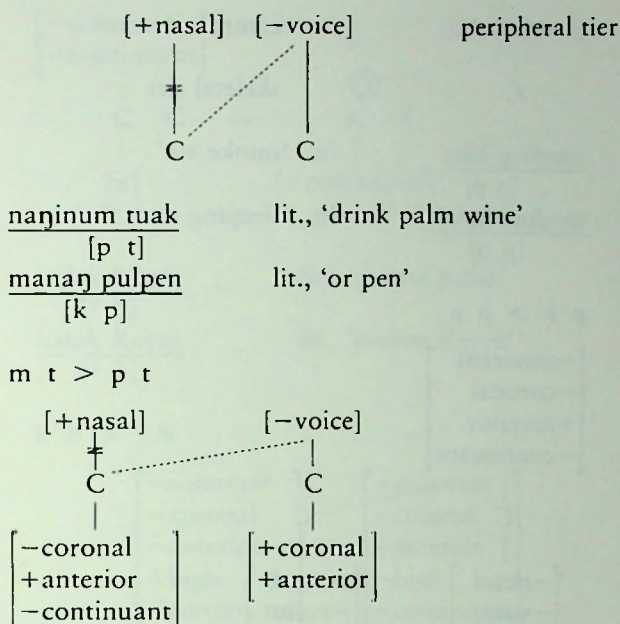


<u>maŋan baɔa an</u>	lit., 'eat man that'
[b b]	
<u>baɔa an pɛdɛk</u>	lit., 'man that short'
[p p]	
<u>lɛan lali</u>	lit., 'give hen-harrier'
[l l]	

a skeletal tier, a *central* tier for the point-of-articulation features and manner features (continuant, sonorant, etc.), and a *peripheral* tier, for the nasal and laryngeal features. Denasalization (18) shifts a nasal to its corresponding voiceless stop when a voiceless consonant follows. *h*-assimilation (19) spreads a voiceless obstruent from a coda position to a following onset position when the onset is filled only by an *h*. In the second example in (19), the word-final *m* is first denasalized by (18) before (19) applies. Finally, a rule of *glottal formation* (20) applies, which converts all obstruent stops in coda position to glottal stops.

Hayes observes that an autosegmental account of the assimilation processes in (17), (18), and (19) predicts that structures that have undergone such assimilations will no longer be eligible for (20), glottal weakening, on the basis, he suggests, of the Conjunctivity Condition. This prediction is correct. *tt* created from *th*, for example, does not then undergo glottal formation (and thereby produce ²*t*). If we analyzed the data from a purely segmental and linear point of view, this result would require some special ordering statement; viewed as autosegmental restructuring, the result is expected.¹⁴ Similarly, if denasalization involves the spreading of [–voice] on the peripheral tier, then a linked structure is created, and glottal formation will not apply to any *kp* structure, for example, if it has been derived from an underlying *ŋp* structure, though of course the rule does apply to a non-derived *kp* structure. Finally, *n*-assimilation will create geminates such as *tt* derived from a deeper *nt*; these geminates are not subject to glottal formation, and no surface ²*t* results from the structure, again as expected.¹⁵

(18) Denasalization

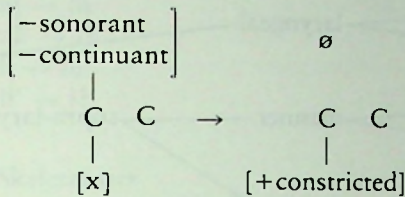


6.2.5 Feature hierarchies and class tiers

The conclusion that the point of articulation of a consonant acts as a unit in many languages – in particular, as an autosegment on a separate tier – raises as many questions as it answers. For example, (i) do we interpret 'point of articulation' now as a single feature, on a single autosegmental tier? Or as more than one feature, but still on a single tier? (ii) If we interpret it as several features, each on their own tiers, how do we express the fact that these feature specifications act so consistently as a single unit, a bundle assimilating as a group? (iii) If we take that analytic route, we will then also ask: do other subsets of features have such a property?

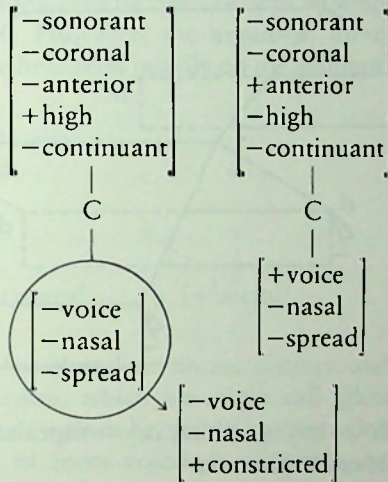
In this section we will look at the feature-hierarchy approach of the sort sketched in (3), or the earlier proposal of Clements (1985b) in (21), an approach that establishes a fixed set of features and an organization of them on separate tiers. This organization, in turn, defines which sets of features may assimilate together as a single group, with point of articulation being the prototypical example of a class node. These are well presented in diagram (22), from Clements (1985b). This feature-hierarchy model allows for rules that assimilate individual features, or a set of features found under specific class nodes, or all features of a given

(20) Glottal formation



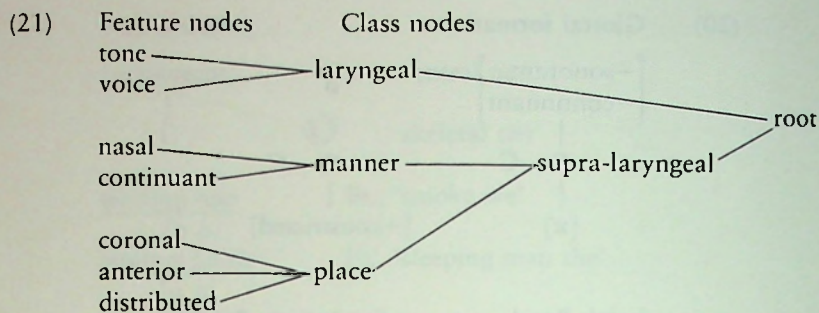
<u>halak Batak</u>	lit., 'person Batak'
[^h b]	
<u>halak Korea</u>	lit., 'person Korea'
[^h k]	

k b > ^h b

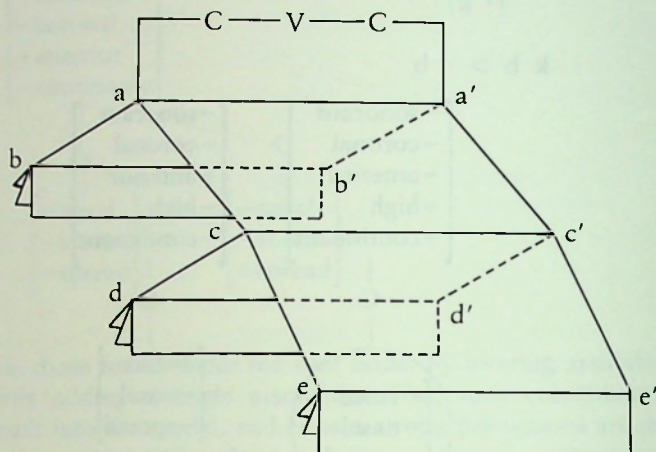


segments. Some way of indicating this dependence is necessary, and a feature representation as in (21) does not quite do that. More generally, any serious argument for setting up two features, F and G, under a class node should include an argument that both F and G act as individual features, which is to say, that each feature can participate in a natural assimilation process on its own. This is not always possible to do.

Other questions about the independence of the features proposed under a given class node can arise in other ways. Clements (1985b) puts forward an argument for the structure offered in (21) illustrating the autosegmental coherence of the laryngeal features responsible for aspiration and glottalization. He provides clear evidence of the stability (and



(22)



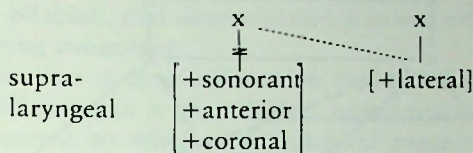
aa' = root tier, bb' = laryngeal tier, cc' = supralaryngeal tier,
dd' = manner tier, ee' = place tier

thus the autosegmental status) of the laryngeal features of [spread glottis] and [constricted glottis], which characterize voiceless sonorants and *h*, and glottalized sonorants, respectively, in Klamath. Segmentally viewed, the alternations given in (23) are found. This set of alternations reduces to two closely related processes involving the reassociation of the oral gesture that comprises the three versions of *l* (the plain, the glottalized *l'* and the voiceless *l̥*). Any sequence of anterior, coronal sonorant, when followed by an *l̥*, will undergo the rule in (24).

Clements further assumes, for purposes of simplicity, that the *l* which is neither glottalized nor voiceless is not associated to any segment on the

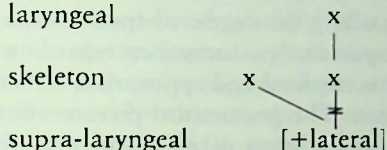
- (23)
- | |
|-----------------------|
| $nl \rightarrow ll$ |
| $nl_0 \rightarrow lh$ |
| $nl' \rightarrow l?$ |
| $ll_0 \rightarrow lh$ |
| $ll' \rightarrow l?$ |

- (24) Skeletal tier



laryngeal tier, in accord with underspecification theory.¹⁷ Under this assumption, the *l* geminated by rule (24) will be simplified by rule (25) in just the right cases. However, the argument for the existence of the laryngeal class node here rests heavily on the assumption that the features

- (25) Degemination
laryngeal

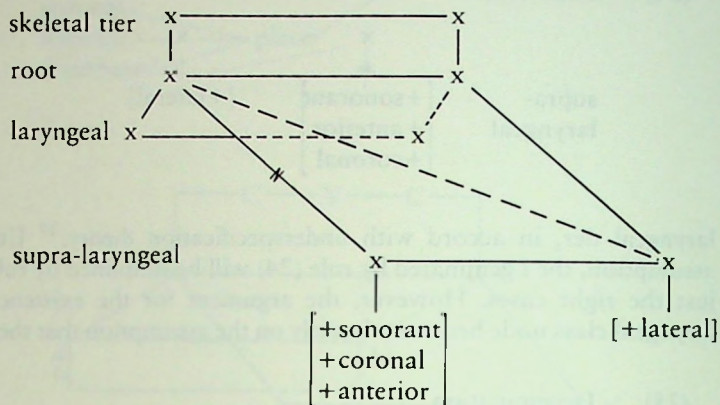


for aspiration and for glottalization are distinct, and that there is not a single laryngeal feature, which we might call [glottalic width], which takes on the feature [-glottalic width] to form glottalized consonants, [+glottalic width] to form voiceless sonorants, and no marking for 'normal' sonorants. If we have such a binary feature, which may take on a third unmarked value, then the argument for a laryngeal class node is considerably weakened. Similarly, the argument for a laryngeal class node (in Klamath or elsewhere) would be strengthened if it were shown that the feature(s) governing glottalization and voicing formed a true natural class in an autosegmental rule along with tonal features, but such arguments have not been forthcoming.

In the example from Klamath above, I simplified matters by not explicitly taking into account the presence of the root-node tier, referring in the representations here to associations directly to the skeletal tier. In this hierarchical account, however, it is crucial that the skeletal positions associate directly to a root-node tier, not to other, lower, class nodes,

because we would not be able to account for the very simple and natural cases of total assimilation with the addition of a single association line, an important goal to the project defined in Clements (1985b). The assimilation of *nl* to *ll*, for example, would more properly be represented in this framework as in (26).

(26) *nl* to *ll* with root tier



One final observation regarding the degree of specification in feature hierarchies. Much depends upon one's assumptions regarding the degree of featural specification that is required and appropriate at various levels of phonological representation. The geometrical picture offered in (26) assumes more or less full specification of all features at the level of representation we are concerned with, and much work has implicitly or explicitly assumed that all features are represented and specified at a phonetic level. Many of the traditional arguments in classical generative phonology for highly specified phonetic representations go by the boards within the more modern context of autosegmental phonology. Consider the following typical example.

Many languages display the pattern that we have already observed for English, Catalan, and Spanish: there are both bilabial and labio-dental consonants, but there is no contrast as such between these two points of articulation, since the continuants are labio-dental and the stops are bilabials. However, in Spanish and Catalan, a nasal that assimilates to a following consonant in point of articulation will agree with that consonant down to this non-contrastive difference, and a nasal preceding a bilabial will be bilabial, just as a nasal in front of a labio-dental will be labio-dental. If we write this assimilation process as a segmental rule, then there will have to be a feature available in the grammar of Spanish or Catalan whose value can be 'transmitted' to the preceding nasal. If,

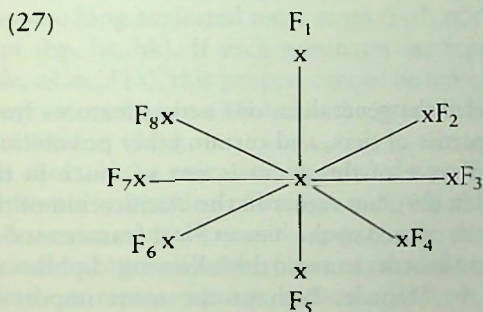
however, the process of nasal assimilation is one that spreads the consonant's point-of-articulation autosegment leftward onto the nasal, then it does *not* follow that there must be a feature to distinguish the labio-dental position from the bilabial. Nonsegmental rules of phonetics may clarify the nature of the gestures available in the language, and define the articulators available for stop and fricative production – this is determined on the basis of the contrastive points of articulation, and the manner of articulation. But once the labial position is determined (labio-dental vs. bilabial), that determination is spread over both the nasal and the following consonant.

There is thus a close connection between the degree of over- or underspecification in a phonological representation, on the one hand, and the extent to which a hierarchical organization of features is motivated on the other. The more features there are, and the more specified they are, the more we stand in need of an explicit organization of them in our phonological representation.

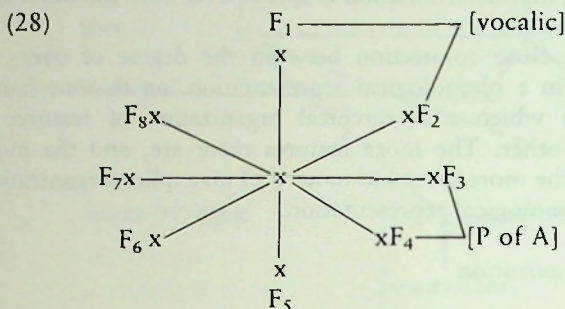
6.2.6 Feature organization

As noted above, an alternative geometrical model for features is that given in (6), called variously the 'spiral-notebook model', the 'rolldex model', and the 'bottlebrush model'.¹⁸ On this view, there is something much closer to the traditional segment as a 'bundle of features', rather than as a hierarchically organized structure of features. Instead of viewing features-with-values as the elements of an unordered set, as in the classical generative model, we take features to be autosegments, each on their own separate tier, forming charts with the skeletal tier in each case. Viewed end-on, we arrive at a picture as in (27), in a system with eight features, with the skeletal tier in the center, and all other tiers facing it.

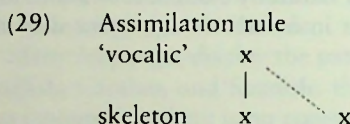
As Hayes (1988) notes, features that naturally assimilate as a unit may be identified as forming a constituent in a fashion consistent with the



geometry of (27) by establishing a constituent structure on the tiers, rather than by establishing additional class-node tiers to do the structuring, as in the hierarchical model discussed in the previous section. For example, if features F_1 and F_2 in (27) are the features [low] and [round] which define the vowel system of a language, and if these features typically assimilate together in the language, we would naturally like to develop a way of expressing this connection between the features, and a way of expressing the rule of assimilation; cf. (28).



One way to express such a rule would be as follows. We recognize that there is some property – a feature, in effect, though not an autosegment – which is necessary to indicate the class of tiers to which the nucleus of a syllable associates. This property is shared by the tiers that define vowel quality; let us call this *vocalic*; this term therefore specifies the set of tiers with autosegments that freely associate to syllable nuclei (F_1 and F_2 , in the example at hand). We may then write an assimilation rule as in (29), where the tier(s) contributing the autosegment(s) that spread is defined indirectly; (29) thus abbreviates two spreading rules, one for the tier F_1 and the other for tier F_2 .

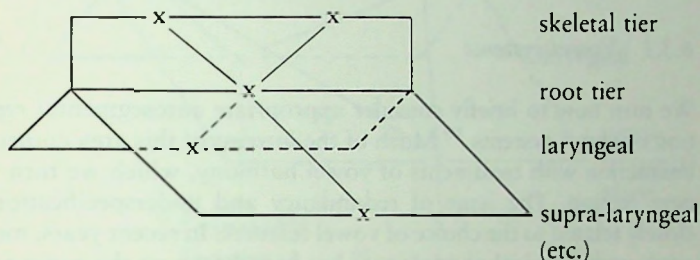


In effect, we have transferred the generalizations across features from the geometry of tiers to properties of tiers, and certain other possibilities can be envisaged. The significance of this point is not so much in the particular proposal sketched in (29), but rather in the clarification of the kind of alternative that may be offered to the hierarchical feature model.

An alternative may be sought in order to avoid the following 'diphthongization paradox', observed by Steriade. Perhaps the most important

single claim that distinguishes the feature-hierarchy model and the rollodex model is the constraint offered by the feature-hierarchy model that all rules of assimilation can be expressed by the addition of a single association line. If some common assimilations require more than one association line, then there is no particular need to organize features into class nodes; the class nodes serve the purpose of making wholesale association of separate features possible at a low geometrical cost (i.e. by adding only one association line). This constraint leads directly to the postulation of a root tier, distinct from the skeletal tier; for otherwise it would not be possible to express total assimilation with the addition of a single association line and still maintain that each feature is on a separate tier. (Total assimilation is the case where a segment completely assimilates, in every feature, to a neighboring segment, as with *n*-assimilation in Toba Batak in (17).) Such a total assimilation is illustrated in a hierarchical scheme as in (30). Thus, total assimilations will always create a structure in which one root-tier position is associated to two skeletal-tier positions.

(30) Geminates



However, there are a number of cases in which structures as in (30) occur (either underlyingly or as the result of a rule) in which *one* of the halves of structure undergoes a change, the creation of a diphthong from a long vowel being a typical example. Other cases include one discussed by Clements (1985b), originally analyzed by Thráinsson (1978), whereby Icelandic long aspirated tense stops (pp^h , tt^h , kk^h) become pre-aspirated stops (hp , ht , hk). If such geminates are represented with a single root node, as in (31a), this process cannot be represented; only if the geminate is represented as in (31b) can the process be naturally represented, as illustrated in (31c).

In short, if the unity and identity of a geminate is to be represented as the double association of a single node – in this case, the root node – then we have no natural way to specify any change that could take place internal to one half or the other of that long segment. This suggests that, indeed, a representation such as (31b) is correct for Icelandic (as

Clements does propose), from which we may safely infer that not all geminates are formed as in (30) with a single root node. But now the root node is doing no work for us – nothing the skeletal tier itself could not do. This, in turn, weakens to a considerable degree the *prima facie* attractiveness of the feature-hierarchy-cum-root-tier model.

In sum, current work is actively pursuing a number of alternative approaches to the issue of feature geometry. Of all the issues that heavily influence the ultimate decision in this matter, without a doubt the most important remains the question of the degree of specification appropriate for post-lexical phonology. Intimately tied up with this is the question of the universality of features, and the extent to which features may be only binary. To the extent that features may be multivalent, taking on several values (as, for example, Hockett's feature of 'position' (i.e. point of articulation) in (1)), several arguments for hierarchical structure become significantly less compelling. Much work remains to be done in this area.

6.3 VOWEL SYSTEMS AND VOWEL HARMONY

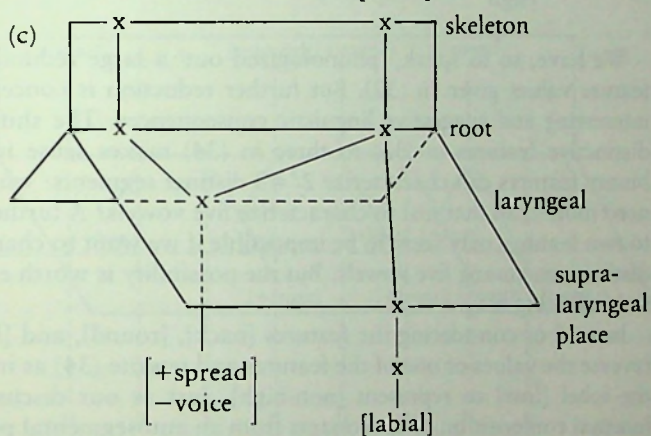
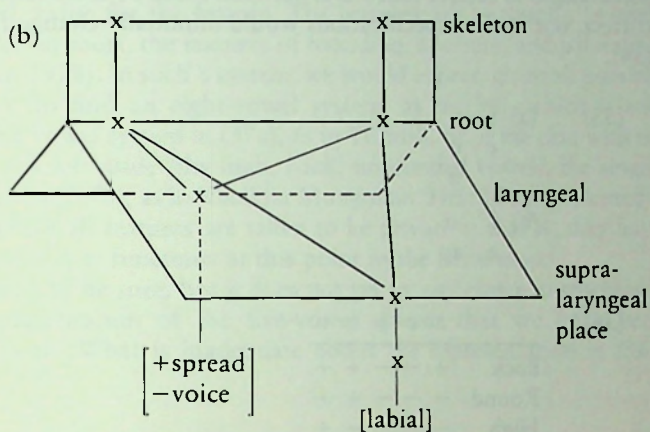
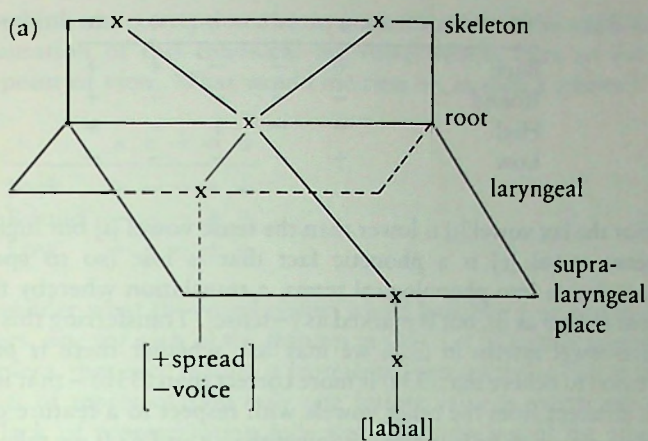
6.3.1 *Vowel systems*

We turn now to briefly consider appropriate autosegmental representation of vowel systems.¹⁹ Much of the interest of this area comes from its interaction with treatments of vowel harmony, which we turn to in the next section. The issue of redundancy and underspecification is also closely related to the choice of vowel features. In recent years, most of the work within lexical phonology²⁰ has been based on the assumption that the core features of vowel space are those given in (32), where the feature specification of the canonical five vowel system is presented.

There is a thoroughgoing redundancy in such a system that permeates all representations and rules: a vowel cannot be specified as both [+low] and [+high], and, as we have observed many times, redundancy in features is an aspect of the representation that underspecification theory aims to eliminate. It has often been observed, in perhaps too offhand a way, that the restriction against [+high, +low] segments, while incorporated into phonology, has its origins in a simple phonetic fact: the tongue cannot be both high and low at the same time, just as any physical object cannot be in two places at the same time.

A more appropriate response to the matter might just as well be to rethink the matter of these features, for the dimensions that we use to analyze vowel space phonologically are not simply present in the data, passively open to inspection: to the contrary, the traditional observation

(31)



(32)	a	e	i	o	u
Back	+	-	-	+	+
Round	-	-	-	+	+
High	-	-	+	-	+
Low	+	-	-	-	-

that the lax vowel [ɪ] is lower than the tense vowel [i] but higher than the tense vowel [e] is a phonetic fact that is lost (so to speak) in the translation into phonological terms, a translation whereby the lax [ɪ] is just as *high* as [i], but is marked as [-tense]. Transferring this issue to the five-vowel system in (32), we may ask whether there is *phonological* reason to believe that (33a) is more correct than (33b) – that is, whether *a* is different from the other vowels with respect to a feature of height. Is there reason to believe that *a* is lower than *e* and *o*? If we take (33b) to be correct, our feature specifications would minimally change from (32) to (34).

- (33) (a) i u
 e o
 a
- (b) i u
 e ao

(34)	a	e	i	o	u
Back	+	-	-	+	+
Round	-	-	-	+	+
High	-	-	+	-	+

We have, so to speak, ‘phonologized out’ a large redundancy in the feature values given in (32). But further reduction is conceivable, with interesting and suggestive linguistic consequences. The shift from four distinctive features in (32) to three in (34) makes sense in that three binary features can characterize $2^3=8$ distinct segments: why would we need more than that just to characterize five vowels? A further reduction to two features may seem to be impossible if we want to characterize the distinctions among five vowels. But the possibility is worth exploring, in the following way.

Instead of considering the features [back], [round], and [high], let us reverse the values of one of the features and rewrite (34) as in (35), using the label [low] to represent [non-high]. Just as our discussion of the internal composition of consonants from an autosegmental point of view

helped us rethink our conception of consonantal features, let us approach the representation of this canonical five-vowel system from an autosegmental point of view. What would the tiers be, in such a scheme?

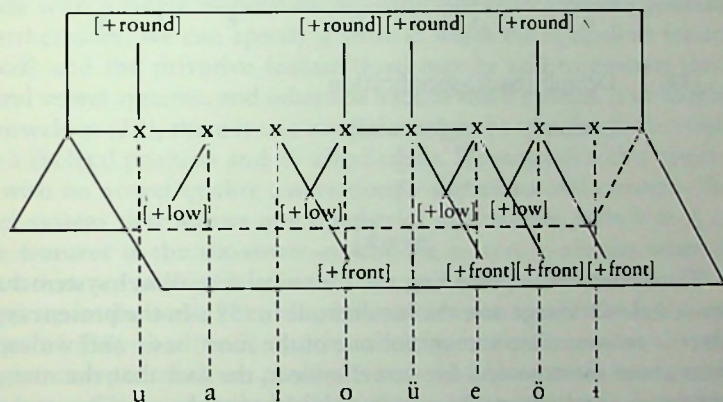
(35)

	a	e	i	o	u
Back	+	-	-	+	+
Round	-	-	-	+	+
Low	+	+	-	+	-

Let us consider what the vowel system would look like if we had three separate tiers, one for each of the features in (35). Let us further assume, for the moment, that each feature is represented privatively, so that at the present level of representation only one feature value is overtly represented. A lack of representation indicates the equivalent of the other, non-marked, value for the feature. The features are [+round], [+low], and [-back] (in short, the features of rounding, aperture, and palatality; see Donegan 1978). In such a system, we would expect, given all possible associations, to find an eight-vowel system, as in (36), which would represent the vowel system in (37a), as in Turkish, or, if the case with no associations is left aside (the high, back, unrounded vowel), the seven-vowel system in (37b), as in Khalkha Mongolian. Here I have presented a situation where all features are taken to be privative; that is, they have only one value that functions at this point in the derivation.

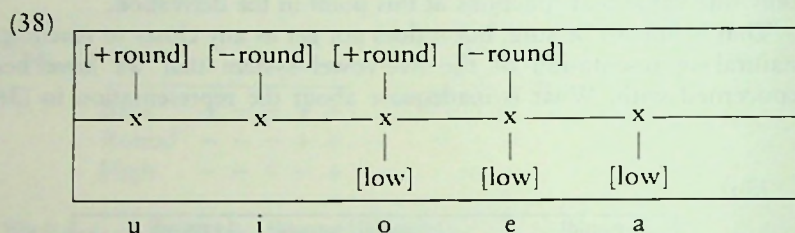
That is fine, to be sure, but it does not get us any closer to reaching a natural representation of the five-vowel system that we have been concerned with. What is inadequate about the representation in (36),

(36)

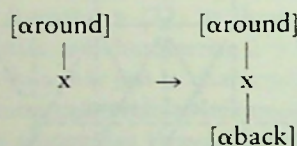


- (37) (a) i ü ɨ u
 e ö a o
 (b) i ü u
 e ö a o

where the features [round] and [front] are on separate tiers, is precisely the freedom the representation gives to those two features. [round] and [front] do *not* combine freely in the canonical five-vowel system: rather, the one is essentially predictable from the other except in the case of the vowel *a*. We propose, then, the representation in (38), which expresses directly an intuition that was only covertly expressed in the traditional chart in (32), with four 'distinctive' features for vowels. The notion expressed in (38) is that *a* does not minimally contrast with *o* with respect to height, nor with *e* with respect to fronting. (38) expresses the idea that the vowel *a* steps out of the system of front/back and round/non-round contrasts that the vowels {i, e, o, u} participate in. This is expressed autosegmentally by using an equipollent feature [round], with front/backness being nondistinctive and fully predictable from rounding, but not requiring that all vocalic positions be associated with one value or the other. Front/backness will be predicted by the general post-lexical rule (39): [back] is not a distinctive feature of the system all.²¹ In short, *a* is neither round nor non-round, and thus neither front nor back.



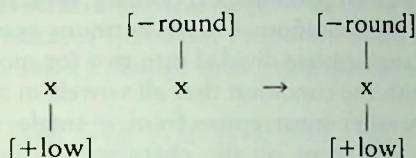
- (39) Default [back]-specification



This system for representing the canonical five-vowel system has an immediate advantage over the familiar one in (32). In the present system, there is an immediate account for one of the most basic and widespread facts about the canonical five-vowel system, the fact that the merger of the vowels *i* and *a* forms the vowel *e*, while that of the vowels *u* and *a* is *o*.

This recurring pattern, seen, for example, in the Kirundi example in chapter 5, is perhaps the most common vowel merger pattern found in languages, but our traditional feature account, as in (32), offers no explanation for it. On the present account, it is a matter of merger of skeletal positions, as in (40).

(40) *a+i to e*



What is surprising about this system, if anything is, is the way autosegmental phonology allows a natural niche for something like a three-way contrast when a binary feature is involved. Within the tonal realm, this very natural distinction can often be seen, as in the case of Sukuma, discussed in chapter 1, where vowels could be associated with a High autosegment, a Low autosegment, or no autosegment. We can see that the same situation arises in the case of other underlyingly equipollent features, such as [round] in the canonical five-vowel system.²²

On the present treatment, then – which does not rule out the traditional account in (32) *per se*, though it suggests an alternative possibility – there is a formal naturalness to the process, often observed, of neutralization in unstressed position to the extent that, from a larger five-, seven-, or ten- vowel system, only the three cardinal vowels {i, u, a} may appear in unstressed position. On the account in (32), these vowels do not form a natural class;²³ on the reanalysis in (38), they are the vowels with a single association of vowel quality to a skeletal position.

Furthermore, we can specify a sense in which the equipollent feature [round] and the privative feature [low] may be said to generate three natural vowel systems, and other less natural vowel systems. If we look at the vowels in (38), there is one more that might be considered: the vowel with a skeletal position and no associations. We may call such a vowel – one with no vowel quality associations – the *schwa* of the system. The vowel system that allows any number of associations, from 0 to 2, of these features is the six-vowel system, {u, i, o, e, a, ə}; this range of association we will refer to as (0, 2), and such a six-vowel system is *complete*, in the sense that all combinations are found. The five-vowel system of (38) is a (1, 2) system, allowing either one or two vocalic associations per skeletal position; and the three-vowel (sub-)system consisting of {i, u, a} is the (1, 1) system, with no more and no less than

one association per vowel position on the skeleton. The notion of *completeness* for a vowel system is an important one, one which we take to be a strong desideratum of an analysis of a vowel system.²⁴

6.3.2 Vowel harmony

Vowel harmony is a term used to describe a restriction on the set of vowels possible within a given phonological domain, typically the word. We may offer the following definition: a vowel harmony system is one in which the vowels of a language are divided into two (or more) (possibly overlapping) subsets, with the condition that all vowels in a given word (or *domain*, more generally) must come from a single such subset. Such a definition does not focus on the character of the restriction, though, and in most cases of vowel harmony the restriction is relatively transparent or natural from a phonological point of view. In such cases, we find that all the vowels in the domain share a particular phonological feature that is distinctive for vowels, such as [back], [tense], or [round]. More to the point, vowel harmony systems are best understood in general as cases where vowel features act strikingly autosegmentally, spreading over a domain that is greater than a single segment. Put slightly differently, a vowel harmony system is what arises when a vocalic feature starts to lose its strict one-to-one association with the skeletal tier, and begins to behave more like tone.

A well-known example of vowel harmony is found in Turkish, where the examples in (41) (from Clements and Sezer 1982, from which I draw heavily here) illustrate the pattern of agreement of vowels in a word. Based on the behavior of vowels in the suffixes, we might arrive at the following statement, the traditional one: all vowels in the word agree with respect to backness, and a high vowel, such as in the genitive suffix, will be round if it follows a round vowel.

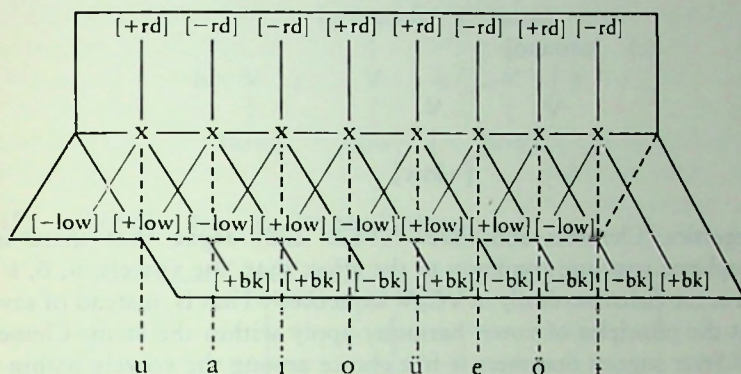
Vowel harmony in Turkish consists, we see, of two distinct spreading

(41)	Turkish				
	Gloss	Nom. sg.	Gen. sg.	Nom. pl.	Gen. pl.
	rope	ip	ip-in	ip-ler	ip-ler-in
	girl	kız	kız-ın	kız-lar	kız-lar-in
	face	yüz	yüz-ün	yüz-ler	yüz-ler-in
	stamp	pul	pul-un	pul-lar	pul-lar-in
	hand	el	el-in	el-ler	el-ler-in
	stalk	sap	sap-ın	sap-lar	sap-lar-in
	village	köy	köy-ün	köy-ler	köy-ler-in
	end	son	son-un	son-lar	son-lar-in

processes of vocalic features: one involving the feature [back], and the other involving the feature [round]. [round] spreads under more restricted conditions, in that only the high vowels, and not the low vowels, act as 'receptors' to such a spreading feature, and the feature will not 'jump over' non-low vowels to spread the feature [round], as we see from a form like *sonlarin*.

Should this system be represented as in (36), with three privative vowel features? Clements and Sezer suggest instead representing it with three equipollent features ([±back], [±round], and [±high]), as in (42). For consistency's sake, I change their [αhigh] to [-αlow].

(42)



The process of back/front harmony would work as in (43a) (*el-ler*), where the feature [back] spreads to a suffix whose vowel is itself specified only for the feature [low]. Rounding harmony spreads the feature [round] rightward, but only onto an adjacent [-low] segment, as in (43b) (*pul-ın*). Rounding harmony could be written in a familiar notation as in (43c) (though Clements and Sezer indicate it somewhat differently).

However, while such generalizations hold for suffixal vowels in general, Clements and Sezer argue that within the lexical stem these generalizations no longer hold true for the modern language. The stem itself is not governed by these harmony principles, we may say, though the affixal material is. But it is not the case that *any* vowel can appear in any position in the stem; the generalization, they suggest, is that in polysyllabic stems, any of the vowels from the set {a, e, i, o, u} may co-occur; in fact, these are the only vowels that can occur underlyingly in the suffixes as well, once we abstract away from the harmony processes.

This suggests that the naturalness of the canonical five-vowel system that we discussed in the previous section is indeed one that arises from its phonological character, not from (or not simply from) its phonetic

- (43) (a) $[-\text{back}]$
 | |Vr
 | |
 |+low| |+low|
- (b) $[+\text{round}]$
 | |Vn
 | |
 |−low| |−low|
- (c) $[\alpha\text{round}]$
 | V
 | |
 |−low|

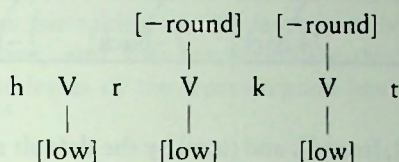
properties. Clements and Sezer (1982: 228) argue that there is a morpheme structure condition to the effect that 'the vowels / \ddot{u} , \ddot{o} , \ddot{i} / do not occur disharmonically in VC_0V sequences'. That is, instead of saying that the principles of vowel harmony apply within the stem, Clements and Sezer suggest that there is free choice among the vowels within the stem, except that the three vowels of the system that do *not* belong to the canonical five-vowel system may not freely appear; they appear only if they are 'harmonic', i.e. if they could be derived by a vowel harmony rule from a simpler form.

Clements and Sezer's account of Turkish suggests the following reanalysis, focusing on the one hand on the close connection between the presence of the canonical five-vowel system and on the other on the redundancy of the features $[\pm\text{round}]$ and $[\pm\text{back}]$, as suggested in rule (39) (default $[\text{back}]$ specification) in the canonical five-vowel system.²⁵ I suggest that words with vowels chosen entirely from the system ($\{i, e, a, o, u\}$) do not contain a specification for the feature $[\text{back}]$ underlyingly, but rather are represented as in (38). Since the feature $[\text{back}]$ is not present at this level, one cannot speak of vowels violating or respecting back/front harmony; all combinations of vowels are permitted within a stem, as Clements and Sezer illustrate in (44).

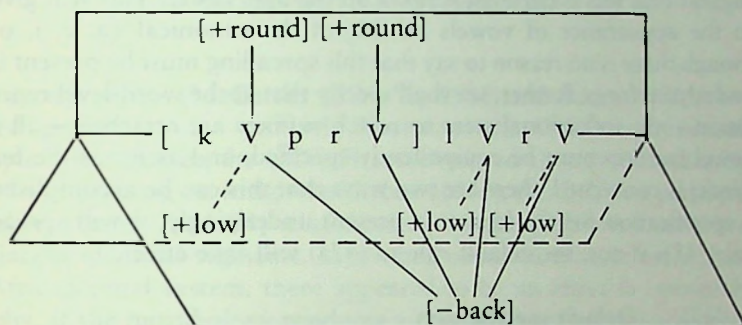
The three vowels of Turkish that may appear in a stem that are not of this system are $\{\ddot{o}, \ddot{u}, \ddot{i}\}$. In words containing these vowels, the feature $[\text{back}]$ must also be present in the lexical representation of the stem, as in (45). However, as Clements and Sezer inform us, such words do not permit violation of backness harmony. Put another way, when the

(44) Clements and Sezer's canonical nonharmonic stems

(a)	a, i:	va:li	'governor'	izmarit	'sea-bream'
(b)	a, e:	hareket	'movement'	hesap	'bank account'
(c)	o, i:	orkinos	'tunny fish'	sifon	'toilet flush'
(d)	o, e:	rozet	'collar pin'	metot	'method'
(e)	u, i:	zigurat	'ziggurat'	muzip	'mischievous'
(f)	u, e:	su:ret	'copy'	mebus	'member of parliament'

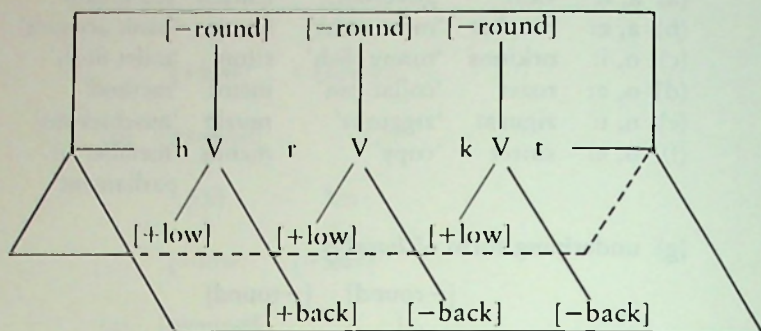
(g) underlying form of *hareket*

(45) köprü – ler – e 'to the bridges'



feature [back] is present, it *must* spread across a word, as indicated in (45). In short, when the feature is not necessary to represent the vowels of the stem, as in (44), it is redundant, and is not used. The default specification for rounding, (39), is put into effect, and all [−round] vowels are marked as [−back], while all others are marked as [+back] (including the low vowel *a*), giving us the representation in (46), where we have the appearance (and, indeed, the reality) of backness violation. This arises, however, out of the fact that no backness specification was present underlyingly.

Let us summarize so far. Apparent violations of harmony within a stem that contain only the vowels {*a*, *e*, *i*, *o*, *u*} are marked for the features [round] and [low] underlyingly, but become fully specified for

(46) Derived form of *bareket*

the features [back], [round], and [low] by the default rule, summarized in (47). Stems that contain a specification for the feature [back] will have that feature value spread over the entire word, as in (45). This specification may be present only once in a stem, and simplicity considerations suggest that this is permissible *only* on the first vowel. This will give rise to the appearance of vowels outside of the canonical {a, e, i, o, u}, though there is no reason to say that this spreading must be present in the underlying form. Rather, we shall specify that at the word-level representation – the inflectional stem to which suffixes are attached – all three vowel features must be equipollently specified, and as far as the feature [back] is concerned, there are two ways that this can be accomplished. If a specification for the feature is present underlyingly, it will spread, by rule (48); if not, the default rule in (47a) will take effect.

(47) Default specification

- (a) back: [-round] → [-back]
(otherwise, [+back])
- (b) vowel unmarked for low → [-low]
- (c) vowel unmarked for round → [-round]

(48) Spread [back] to the right



One more thing needs to be said about rounding harmony (43c). There is one respect in which the canonical nonharmonic stems of (44) differ from the more familiar stems as in (45). In the former, the feature [±round] could appear contrastively on any vowel, while in the latter –

in stems where the feature [\pm back] is marked – the feature [\pm round] can occur only on a vowel that is also marked [+low]. Elsewhere, among the high vowels, rounding is nondistinctive, and is specified by the ‘roundness harmony’ rule (43c).

In sum, the Turkish vowel harmony system illustrates one way in which vowel features act autosegmentally, producing the effect that is known traditionally as ‘vowel harmony’ by spreading. In addition, underspecification theories lead to natural accounts of apparent violations of harmony, as when a feature is filled in by a default rule, and is thus not subject to any autosegmental spreading that would create a harmonic span. Finally, we see from within the vocalic features some of the kinds of reasons for taking certain features to be equipollent and others to be privative, and also for permitting this parameter to be different at different levels of the representation (underlying and word level, in this case).²⁶

6.4 THE OBLIGATORY CONTOUR PRINCIPLE AND THE MORPHEME TIER HYPOTHESIS

Three recent papers (McCarthy 1986, Yip 1988, and Odden 1988) have focused attention on a principle known as the Obligatory Contour Principle – hereinafter, the OCP. First formulated as such in Goldsmith (1979) and drawing on insights of Leben (1973), the OCP is a principle (or rather, a family of closely related principles) that prohibits consecutive or adjacent identical segments. Leben had observed that, in more than one African tonal system, there appeared to be an effect in operation whereby, if the morphology produces a concatenation of two adjacent, identical tones, the two fuse into a single tone before the tones are ‘mapped onto’ their corresponding vowels. In Tiv, for example, following earlier work by Arnott (1964), Leben suggests that the pattern for the imperative verbal form is ‘BHL’, meaning the ‘base’, or underlying, tone of the radical, followed by H and L. If the radical is underlyingly High, then this sequence is ‘HHL’, and Leben suggests that, before this sequence is mapped onto a string of syllables, such an HHL sequence will be simplified to HL, as in (49).

In Goldsmith (1979), where autosegmental phonology was first proposed and explored, the possibility was raised that this could be a general property of autosegmental systems. The issue in its earliest form focused on two matters. First, there are the potential ambiguities (or uncertainties) of representation in autosegmental phonology *vis-à-vis* segmental

(49)	Syllables:	ti	re	{basic tone=H}
	Imperative	B	H	L
	Fill in B	H	H	L
	OCP	H	L	
	Mapping	ti	re	
		H	L	

phonology: given a bisyllabic word with two high-toned vowels, how can we determine whether the proper representation is as in (50a) or (50b)? Second, if we focus on languages in which tones and vowels are mapped onto one another in a straightforward one-to-one fashion as discussed in chapter 1, can we make the strong claim that no such language will have trisyllabic words with a HHL pattern, as in (51)?

(50)	(a)	CV	CV	(b)	CV	CV
					↘	↙
		H	H		H	

(51)	CV	CV	CV
	H	H	L

Since there are languages with words that *seem* like they have a HHL tone pattern – English, for example, in the neutral pronunciation of a word such as *linguistics*, or any word with the same stress pattern – the question was posed as to whether we could immediately draw the conclusion that the language possessed some kind of accent system from the mere presence of a word possessed of a HHL pattern (or, equally, LLH, or any other violation of the OCP). An accent system would allow for the association of a distinguished ('accented') tone with the primary-accented syllable, wherever it might occur in the string of syllables, as in (52).

(52)	CV	CV*	CV		CV	CV*	CV
		*		becomes		*	
		H	L			H	L

In Goldsmith (1979), it is argued that the OCP is *not* operative actively in the phonology, in the sense that adjacent identical segments are not automatically and universally reduced to one, on the basis of a small number of languages in which there was no independent evidence of an

accent system, and in which there were found to be adjacent same-toned vowels which were apparently *not* associated with the same tone; see (53).

(53) Etung (Edmundson and Bendor-Samuel 1966)

e se be 'sand'

| | |
H H L

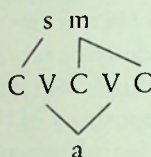
More importantly, perhaps, it was argued there that there are good and general reasons to expect that the OCP will appear to operate as a design strategy in the phonology of natural languages – not because the principle is overtly present, but because it follows from basic considerations of how phonology is learned. The segments in our underlying phonological representations do not, after all, come from nowhere, and we must bear in mind that they actually come from the phonetic representations presented to us during the acquisition stage. Details aside, it is obvious that underlying representations of morphemes (at any given stage of language acquisition) will match surface forms, except insofar as underlying forms may leave out redundant information, and insofar as they must differ from surface forms to account for perceived allomorphic variation. In short, underlying representations match surface forms as much as possible; but this is not a principle either inside a grammar or, for that matter, in a repository of Universal Grammar: it is a general property of how a system learns, when its inner representations are set up to correspond to outer form, as a child must do when acquiring a language.

What does this have to do with the OCP? When a language acquirer approaches the phonetic signal, he or she may break it down into various channels of information (as indicated in (1) and (2) in this chapter). In the case of (50a,b), the tonal part of the phonetic signal is a period of high pitch, and nothing else – which, when represented with phonological segments, is represented as (50b). In the absence of any reason to analyze the form differently, then, the underlying form will be like the phonetic form, and we will see the effects of the OCP on the underlying representation of morphemes – not as an absolute, inviolable principle, but rather as a strong tendency. This interpretation has as its consequence that any clear cases where two successive, identical tones from *separate* morphemes merge into a single segment during the derivation must be cases of language-particular rules of merger. We will refer to this interpretation of the OCP as the 'naturalness' interpretation, alluding to Postal's (1968) 'naturalness principle' in phonology.²⁷

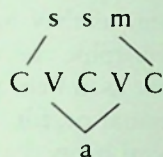
Our awareness of the significance of the OCP has been heightened by

work by McCarthy (especially 1979a, 1981, 1982). McCarthy argued that, given his autosegmental account of Arabic morphology (which we studied in chapter 1), the joint assumption of the OCP and left-to-right association of consonants to skeletal positions has as its consequence that we may find stems such as *samam* in Arabic, as in (54a), but never forms like *sasam*; there would be no way to derive such forms, except from underlying *ssm* consonantal roots, which in turn would be ruled out by the OCP. (Forms like *samam*, of course, would derive from *sm*, not *smm*, which McCarthy argues is correct.) The 'naturalness' interpretation of the OCP does not, strictly speaking, rule out the possibility of an underlying root such as *ssm*, but McCarthy observed that phonological theory may indeed prefer a stronger version of the OCP, one that will unconditionally rule out such a possibility; cf. (55). This suggestion is based on the generalization made by Greenberg (1960) that in Arabic adjacent consonants in the root may not be homorganic. As McCarthy observes, much (though by no means *all*) of Greenberg's generalization will be predicted by the stronger interpretation of the OCP as a strict constraint on underlying representations of individual morphemes.

(54) (a) *samam*



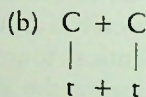
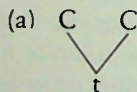
(b) **sasam*



(55) OCP-1: There can be no adjacent identical segments on the melody tier in underlying representation of morphemes.

It should be clear that one consequence of the generalization in (55) is that all tautomorphic (i.e. morpheme-internal) geminate consonants must be 'true geminates', as mentioned in chapter 2 (56a). Accidentally identical consonants across morpheme boundary will *not* be true geminates; they will be only apparent geminates, at least, given what we have said so far; more remains to be said.

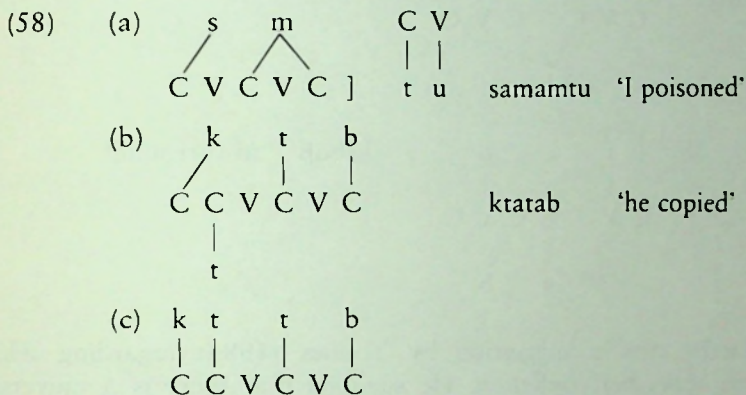
(56)



McCarthy (1986) proposes a version of the OCP in (57), one that is somewhat different from the OCP-1 in (55) above. A number of serious questions arise in connection with the deceptively simple word 'pro-

- (57) Ocp-2: At the melodic level (i.e. on non-skeletal tiers), adjacent identical elements are prohibited.

hibited' (to which we return in connection with our discussion of Yip 1988), such as what the consequences are taken to be of a rule 'attempting' to apply and thereby creating a violation. But McCarthy notes that there is a close connection between his use of the OCP and his analysis of Arabic, in which vowel and consonants are represented on separate autosegmental tiers. He is at pains to show that, while there are clear reasons to interpret multiple copies of a single *root* consonant in Arabic as multiple associations of a single consonant on a separate tier (58a), this effect disappears when we face two consonants that are phonologically identical but which come from separate morphemes, as in (58b). McCarthy suggests that there is clear evidence that the two *ms* of (58a) are simply two realizations or associations of the same autosegment *s*, while the two *ts* of (58b) reflect two distinct autosegments *t*. In this way, the morphological origin of a segment may have consequences as far as the autosegmental geometry is concerned. If all consonants were on the same tier underlyingly (if the underlying representation were as in (58c), then at the deepest level, only OCP-1 (55) would hold – not OCP-2.



An alternative view, and the one that McCarthy prefers, is to adopt version OCP-2, and to do as I have indicated in (58a,b): to place separate morphemes onto separate melodic tiers at the deepest level at which the morphemes are joined together, the underlying representation of the phonological derivation.

This proposal is known in the literature as the *Morpheme Tier Hypothesis*, although, like the OCP, it is a family of closely related variant hypotheses, all of which aim at assigning separate tier status to separate morphemes.²⁸ One immediate consequence for McCarthy's

The general principle is this, then, McCarthy suggests: before tier conflation, phonological segments are organized morphologically, in accordance with the Morpheme Tier Hypothesis; after tier conflation, the only autosegmental effects are the 'local' ones that we have considered up to now, involving geminate consonants and long vowels.

It is not too hard to see that, although the OCP-2 and the Morpheme Tier Hypothesis are logically independent, there is nonetheless a close connection between them. If the segments of separate morphemes are on separate tiers, then only tautomorphemic segments stand as a test to the OCP, and adjacent identical consonants from separate morphemes simply will not be relevant to determining the truth or falsity of the OCP (e.g. as in *cool-ly*). While the strongest evidence proposed for the OCP may be with respect to underlying forms, this may simply be because the morpheme-tier structure makes other effects of the OCP-2 less visible.

McCarthy suggests that the OCP-2 should be understood not only as a condition on possible underlying forms (as in OCP-1), but also as part of the algorithm involved in rule application – in the sense that, if a rule's structural description is met but its output would contain a violation of the OCP, the rule will fail to apply, and the derivation will continue as before.

This suggestion makes the prediction, for example, that rules of vowel deletion will automatically be constrained so as not to apply to vowels flanked on either side by identical consonants. Tonkawa provides an example of a language with a vowel-deletion rule subject to such a constraint, as illustrated in (61). A version of the rule is given in (61a), and its effects are illustrated in (61b). (See Kenstowicz and Kisseberth 1979 for further discussion.) In (61c), we see that, when the consonants on either side of the consonant that is subject to deletion are identical, the vowel does not delete.

On McCarthy's account, then, the OCP is not violated in underlying morphemic representations. It certainly is not violated in the deep phonological representations, because separate morphemes are kept on separate tiers; and at tier conflation, identical segments that are brought together are fused by convention, i.e. by an active, dynamic version of the OCP – a rule (or rather, a convention) that fuses any two adjacent identical segments precisely at the point of tier conflation (but cf. n. 30).

Odden (1988) undertakes a reanalysis of McCarthy's range of observations, and points out a number of problematic features, several of which I shall mention here.

1 There are simple empirical counter-examples to OCP-1 in its strict form. Odden offers Chukchi (citing Krause 1980) and Hua (Haiman 1980) as languages with rules of schwa-insertion that break up clusters of consonants, even when these consonants are identical. Presumably, if

(61) (a) $V \rightarrow \emptyset / VC - CV$

(b) /picena/ 'cut'

<i>he Xes it</i>	<i>progressive</i>	<i>he Xes them</i>	<i>progressive</i>
picno ^ʔ	picnano ^ʔ	wepceno ^ʔ	wepcenano ^ʔ

*he Xes me*kepceno^ʔ kepcenano^ʔ

/notoxo/ 'hoe'

notxo ^ʔ	notxono ^ʔ	wentoxo ^ʔ	wentoxono ^ʔ
--------------------	----------------------	----------------------	------------------------

kentoxo^ʔ kentoxono^ʔ

(c) /hewawa/ 'die'

hewawo^ʔ 'he is dead'/ham²ama/ 'be burning'ham²amo^ʔ 'he is burning'

adjacent consonants are identical and the OCP-1 holds, then integrity would not allow this schwa-insertion. For example, in Chukchee there is a rule inserting schwa between two word-final consonants when the penultimate consonant is not glottalized: cf. (62). The final example in (62) illustrates the behavior of an underlying stem /ekke/, where the sequence of *kk* does not display the expected behavior of a geminate.

(62)	Abs. sg.	Abs. plural	Gloss
	miməl	miml-ət	water
	wiwər	wiwri-t	board
	ekək	ekke-t	son

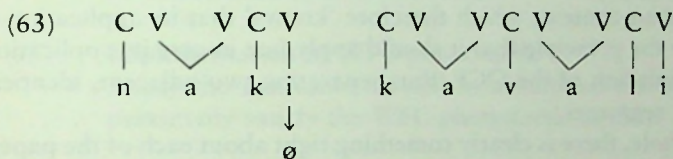
2 The original motivation for taking the OCP as an absolute restriction on underlying forms rather than a matter of simplicity and tendency was the Greenberg generalization cited above regarding the strict prohibition against consecutive homorganic consonants in the Arabic root. But the Greenberg generalization is larger than the OCP; it rules out sequences of homorganic consonants even when they are not identical. Thus the OCP, if it is to do the necessary work, must be informed (so to speak) to focus on just one sub-tier, the point-of-articulation sub-tier. But that is not motivated by OCP-1, and implementing such a proposal takes us far beyond the original spirit of the principle.

3 No examples have ever been found where consonant deletion is blocked when that would create sequences of identical vowels. And

counter-examples are attested; Odden cites the case of Estonian, and the consonant-weakening processes in several Eastern Bantu languages that have produced sequences of identical vowels as their output. One may sense a reflection of the same problem here as in the previous problem; many OCP effects do revolve around point-of-articulation specifications, which of course vowels do not possess. In any event, once again, to the extent that the OCP is proposed as a general theoretical property of the geometry, the asymmetry of vowels and consonants in this respect is disturbing.

4 Odden reports several cases where vowel deletion applies regardless of whether it creates geminates. For example, he reports that in Hindi (citing Bhatia and Kenstowicz 1972) there is a schwa syncope rule that applies regardless of whether the flanking consonants are identical (e.g. *kaanən+i* > *kaann+i* 'garden') or not (*daanəw+i* > *daanwi* 'demon'). He suggests that this rule is 'phonological' enough to distinguish between stem vowels, where it applies, and vowels to the right of the stem, and he concludes that any attempt to characterize as 'merely phonetic' those rules that fall outside the constraining effect of the OCP is not justified at this point.

5 There are phrase-level (post-lexical) rules of vowel deletion which apply only when their output *creates* geminate consonants. The fact that the rules may apply at phrase level, and between words, confirms the notion that the two consonantal autosegments on either side of the deleting vowel are distinct, and are not the multiple association of a single consonantal autosegment. Thus, for example, in Koya (Taylor 1969), there is a phrase-level rule that deletes word-final vowels when the consonants on either side are identical. Thus, *na:ki ka:va:li* 'to me it is necessary' surfaces as *na:kka:va:li*, and *a:ru ru:pa:yku* as *a:rru:pa:yku* '6 rupees'. This is illustrated in (63).



Odden offers a number of additional counter-arguments, leaving little doubt that in its strongest form McCarthy's proposal is not tenable. Odden's conclusion is that, to the extent that there are OCP effects active in phonology, these are language-particular, and rule-particular, effects. There are language-particular rules that achieve the fusion of two adjacent, identical segments, rules that have been informally dubbed 'OCP effects' in the literature; but they are not different in *kind* from other rules.³⁰

Yip (1988) also approaches McCarthy's suggestions, but from a perspective considerably different from Odden's. While Yip's position is fundamentally more sympathetic to McCarthy's argument, it leads to a proposal that is itself far more radical than McCarthy's, a proposal that ultimately leads us to the heart of the final section of this chapter and the discussion of harmonic rule application.

Yip suggests that the OCP, as in (57) (OCP-2), is a well-formedness condition on representations, and that there is a class of rules that is triggered to apply to a given representation just in case it violates the OCP. She suggests that these rules are of a special type; they are rules with no structural description, applying when and only when they are needed to repair violations of the OCP. More tentatively, Yip also suggests that there is a late point at which an active merger of any two identical adjacent autosegments, following, perhaps, McCarthy's account (see note 30). However, Yip is at pains, as we shall see, to show that there can be representations in the phonology of a language *after* tier conflation (which is to say, in the 'normal' part of the phonology, where vowels and consonants are properly intercalated) where the OCP is violated; her point is that there may be various strategies at hand which undo OCP violations that are present – epenthesis, metathesis, deletion, and so forth. On Yip's account, merger of adjacent identical segments might just as well be yet another language-particular strategy for avoiding OCP violations, although she does not choose that particular approach, leaving fusion as the automatic and final solution to the OCP's demands.

Thus, the following example, which is offered by Odden (1988) as a counter-example to McCarthy's position, may be interpreted as a clear case working as Yip would have it. In Lenakel (Lynch 1978), a schwa is inserted between identical consonants, as when underlying *i-ak-ken* 'I eat' becomes *yagəgen*. On Yip's view, this would be the result of a rule written with no context, which therefore 'knows' that its application is governed by the principle that it should apply just in case its application resolves a violation of the OCP, thus separating two adjacent, identical consonantal segments.

On the whole, there is clearly something right about each of the papers cited in this section. I believe that McCarthy is correct in drawing our attention to the importance of the OCP as a principle expressing well-formedness of representations at several levels. Odden (1988) correctly informs us that the extremely strong claims offered in McCarthy (1986) cannot be accepted at face value, but Yip (1988) suggests a radically different perspective from McCarthy's, and from most familiar generative accounts. It is very much in line with various ideas regarding rule application that have come up several times in the course of this book.

6.5 HARMONIC RULE APPLICATION AND AUTOMATIC SPREADING

The final topic that we shall consider is one of the most far-reaching, and we cannot do it full justice in just a few pages. On a number of occasions during the course of the preceding five chapters, we have alluded to the notion that certain processes must be understood as applying just in case they encounter a violation of a well-formedness condition (i.e. a phonotactic) which will be removed by the application of the rule. The rule, in short, constitutes a particular 'repair strategy' as far as that phonotactic is concerned, and it acts only in that capacity.

This notion has a special place in the development of autosegmental theory, where in early years it was associated with the question as to whether there is 'automatic spreading' in tone languages. Let us review how this question, and its treatment, arose in the development of the theory.

In the earliest work on autosegmental phonology (Goldsmith 1979), a principle known simply as the 'Well-formedness Condition (WFC)', was suggested, and much later work generally assumed the validity of the condition. The WFC consisted of the statement in (64), and an algorithm (65) that utilizes that condition. I have added to (64) the phrase 'phonotactic', for reasons which we shall see below.

(64) **Well-formedness Condition (-Phonotactic)**

1. All vowels are associated with at least one tone.
2. All tones are associated with at least one vowel.
3. Association lines do not cross.

(65) **Implementation of WFC-phonotactic**

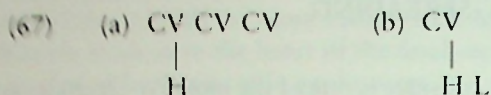
Apply the operation in (66) in a minimal fashion so as to maximally satisfy the WFC-phonotactic in (64).

(66) **Repair operation**

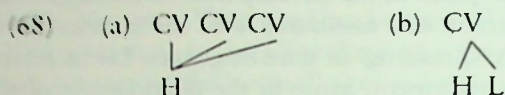
V
⋮
T

The WFC itself, in (64), describes a state of affairs that may or may not be met in a given representation; the representations in (67) illustrate cases violating each of the first two clauses. According to this theory,

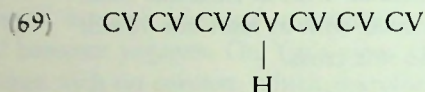
violations of the WFC trigger the implementation of the algorithm (taken in these works to be universal) given in (65).



Taken altogether, this meant that the grammar would add the minimum number of association lines (in ways that would not violate the WFC itself) but always in such a way as to maximally satisfy the WFC. This would have the effect of changing the forms in (67) into the corresponding forms in (68); the rule in (66), we might conveniently say, acted like a particular kind of repair strategy for the WFC-phonotactic.



Later work³¹ emphatically showed that languages could have surface forms that were in some cases massively in violation of the WFC. That observation was not in itself too surprising, for the original formulation had clearly left open the possibility of such cases (to allow, most importantly, for floating tones underlyingly and floating tones on the surface, in the latter case to act as downstep triggers). What was surprising was that cases like Sukuma (see chapter 1) could exist, where, as in (69), High tones could show no tendency at all to spread.



This observation was widely taken to show an inadequacy in the WFC itself. However, the assumptions of early autosegmental phonology that led to the automatic spreading in question here were three in number, as we have seen: the WFC-phonotactic (64), the 'instruction' or 'rule' in (66), and the universal algorithm (65) instructing how to apply the 'rule' in a minimal fashion to maximally satisfy the WFC. If any one of these failed to be universal, then the spreading effect would no longer be universal.

Now, the notion of rule that was adopted by the early works in autosegmental phonology was in all important respects that of traditional generative phonology. A rule was, in this light, a language-particular statement, and it would relate two adjacent representations in a deriva-

tion just in case the deeper of the two representations satisfied the particular structural description of the rule. In fact, the static picture of derivations in which rules relate adjacent stages in the derivation is sometimes less effective as a metaphor than the more common active metaphor according to which rules actually come along and modify representations, since, when a representation satisfies the structural description of a rule, it must 'undergo' the rule, which is to say there must be another stage in the derivation corresponding to the 'output' of the rule.

With this much borne in mind, it should be clear that the tripartite nature of the Well-formedness Condition with its implementation algorithm simply did not fit into the picture of phonological derivations of classical generative phonology. If accepted, it had to be viewed as something overlain upon the true phonological rules, a universal mechanism that stood outside the set of phonological rules that constitute the phonological grammar of the language. More than for any other reason, this was because phonological rules in the classical generative picture were not conceived of as applying or not applying in a fashion dependent on whether or not their output achieved a specifiable output structure. But that was precisely what governed the implementation of the association line additions demanded by the Well-formedness Condition.

Contemporaneous with the proposal of autosegmental phonology, Sommerstein (1974) suggested that a wide range of generative phonological rules (though how wide he was silent about) could best be analyzed into two parts: a set of changes that operated upon a representation – we may refer to this as the conditional rule; and a set of surface phonotactic conditions linked to one or more (conditional) rules in the following fashion.³² A conditional rule will apply if and only if its input violates one of its phonotactic conditions and its output satisfies that condition. His arguments on the point are quite straightforward, and address traditional segmental problems of Latin phonology. Sommerstein observes, for example, that a rule of final coronal obstruent deletion can be written in a complex fashion, if we choose to do so; but positive statements on possible word-final clusters are simpler to state, and allow us the following possibility: we can express our rule of final coronal deletion with no 'environment' in the rule, other than to say that it applies word-finally, as long as we specify that the rule is one that applies if and only if its input violates a phonotactic condition and its output satisfies the condition.

A more intricate example given by Sommerstein concerns the process of fricative deletion in Latin, which applies if any of five independently motivated phonotactics are violated. By indicating that the rule applies if

and only if it repairs a violation of such a phonotactic, a single, simple rule can be formulated, even though in some cases its effect is to resolve a violation of a voicing-agreement constraint, in others a violation of an obstruent-resonant cluster constraint, in still others a violation of a constraint against obstruent-glide clusters, and so on. In short, to write separate rules where each specifies the particular *way* in which a phonotactic can be violated – and to call that, then, the ‘structural description’ of the rule, as if it were that *particular* sequence that caused the rule to apply, rather than the representation’s failure to satisfy the phonotactic – is to miss a string of important generalizations.

Although the connection was not remarked upon at the time, Sommerstein’s conception of language-particular rule application and the procedure for implementing repairs of the WFC-phonotactic were fundamentally the same. And Sommerstein’s work has by no means gone unnoticed. Singh (1987), for example, explicitly argues in favor of adopting a strong version of Sommerstein’s view, emphasizing once again the importance of phonotactics, and the insights gained in trading off rule complexity against phonotactic specifications.³³

This suggests the following reconstruction of the organization of phonology.³⁴ A phonological level will be defined as a set of phonotactics placed on representations. The *word-level* (W-level) in a particular language, for example, will consist of a set of phonotactics, or well-formedness conditions, that apply to phonological representations in that language. A general theory of word-level phonotactics will constrain the technical language in which such phonotactics can be specified, and the work discussed in this book suggests the following hypothesis: language-particular word-level phonotactics consist entirely of syllable structure-conditions, including autosegmental licensing specifications and autosegmental restrictions on the minimum/maximum number of associations. Other word-level phonotactics are universal. We return to some cases of this sort below.

Along with a set of (universal and language-particular) phonotactics for the W-level, each language will contain a set of rules that operate as repair strategies, applying just in case their output eliminates the violation of a phonotactic in their input. There is no guarantee that all violations will, in fact, be resolved by the time all the rules have done their work; in fact, it seems quite clear that it will *never* be the case that all such W-level phonotactics are perfectly resolved. Rather, the W-level phonology attempts to achieve a maximal satisfaction of its constraints, subject to the resources it has for fixing problematic violations.³⁵

We may understand the word-level, then, as a series of representations $\{W_1, \dots, W_n\}$, where the last one satisfies the W-level phonotactics as well as the language can manage, and the first is supplied by the

morphology in a way that we shall return to momentarily. If we think of well-formedness – or its opposite, ill-formedness – as a matter of degree, then the path that the representation takes as it moves, so to speak, from W_1 to W_n may be conveniently thought of as a downhill path towards a ‘local minimum’ of ill-formedness, where the rules of the language define what an allowable path is. The W-level representation of a given form is then the entire sequence of representations $\{W_1, \dots, W_n\}$, and we may refer to the ‘repair strategy’ rules that apply internally to that level as ‘(W, W)’ rules, in the sense that their input and output are both parts of the W-level representation. Schematically, this may be represented as in (70).

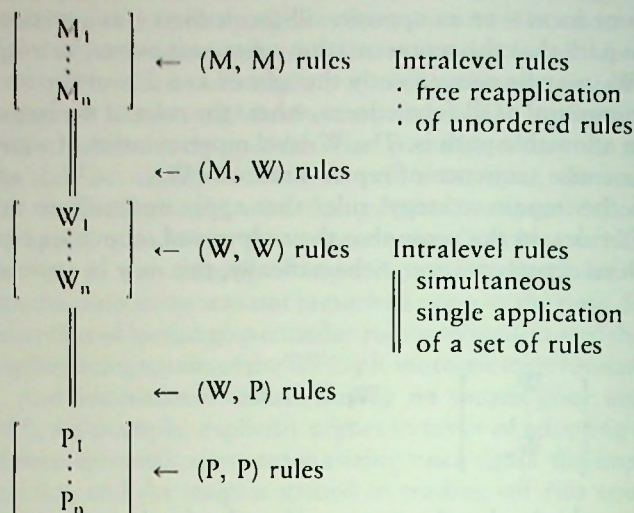
$$(70) \quad \left[\begin{array}{c} W_1 \\ \vdots \\ W_n \end{array} \right] \leftarrow (W, W) \text{ rules}$$

We hypothesize that there are two more levels relevant to the phonology: one essentially morphological in character (therefore, an *M-level*), and one of systematic phonetics (a *P-level*). As with the W-level, these other levels consist of a sequence of representations aimed at achieving maximal well-formedness in accordance with level-specific tactics. We may furthermore take there to be one further set of rules aligning the levels: one set of (M, W) rules aligning the M-level with the W-level, and one aligning the W-level with the P-level. We then arrive at the diagram in (71), which we shall refer to as a *harmonic phonology*.³⁶ P_n serves as the representation of systematic phonetics, and as the interface with the phonetic component. M_1 is the representation that interfaces with the morphosyntax.

Current work suggests that, within a level, rules apply in the manner generally referred to as ‘free reapplication’, subject, unsurprisingly, to the Elsewhere Condition, in the sense that, when a language has two competing repair strategies for a phonotactic violation within a given level, it chooses the one that is more specific for the task at hand. Inter-level rules (M, W) and (W, P) operate in non-interactive ways, i.e. simultaneously. Typical examples of various processes are sketched in (72).

I will conclude by reviewing several significant advantages to this conception of rule application within phonology. The issue that is involved is a broad, difficult, and important one, and, while it goes beyond the bounds of the present book, I will spell out some of the important differences that have come to light in distinguishing between traditional and harmonic modes of rule application.

(71) Harmonic phonology



(72)	Type of rule	Example
	(a) (M, M)	Melody spreading before tier conflation
	(b) (M, W)	Tier conflation
	(c) (W, W)	Syllabification; epenthesis
	(d) (W, P)	Default feature specification
	(e) (P, P)	Flap formation in English

I will sketch eight areas where this approach shows a solid, coherent advantage over other approaches. These are intended as illustrative, not exhaustive, cases, as indicated above. If the suggestions considered here are correct, then the general principle of harmonic application governs all essentially phonological rule application.

(1) We often arrive at a considerable simplification of individual rules, as noted in part by Sommerstein, Singh, and Paradis, among others, when we do the following three things: (i) remove the structural description from the rule itself; (ii) invert it, specifying not what is *disallowed*, but rather what tactics must positively be met; and (iii) note that the positive conditions determining whether a rule will apply involve reference to the output of the rule, not the input – though, of course, an element in the input of the rule may be deleted in order that the output satisfy a condition.

Many languages have rules of epenthesis and of cluster simplification (i.e. consonant deletion) whose target structure is the well-formed

syllable of the sort we have discussed.³⁷ Wiltshire (1988), for example, discusses the syllable structure of IṛuLa, a Dravidian language, on the basis of materials in Diffloth (1968). She analyzes the phonotactics as deriving from a W-level coda capable of licensing only the feature [nasal] and vowel quality features – apparently only the feature [\pm round]. As in Selayarese and several other languages we considered in chapter 3, geminate consonants are permitted intervocalically, as are nasal-stop clusters. Vowels may be contrastively long or short regardless of whether the coda is associated with a consonant or not. There is considerable modification of the phonological form between the underlying representation – our M_1 – and the surface form, but virtually all of the complexity derives from various strategies pursued by IṛuLa to achieve well-formed syllable structure, as determined by the coda licensing condition. An (M, W) rule applies between coronal-final verb roots and the past-tense suffix *t*, creating a geminate, as in (73).³⁸

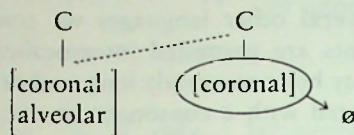
As the reader will notice, other processes come into play when there are consonants that cannot be licensed. The strategy of deletion is used for word-final sonorants, but only for them. Elsewhere, as we see in (73d), epenthesis of U, a short, centralized vowel, applies in order to create a licensing environment – the syllable node – for that consonant. Of course, where the gemination-formation process of (73b) applies, the epenthesis rule does not need to apply to create a well-formed position for the first consonant; no epenthesis occurs after the stem-final consonant that spreads rightward above – the property of ‘geminate integrity’ that we have discussed. We return to this point below. Wiltshire suggests a third rule, which may be operative in the (W, W) component, a rule deleting the first of three consonants, as *kol-nd-en* > *koṇḍe*, again aiming at W-level well-formedness.

(2) Yip’s interpretation of the OCP as a motivator for a certain class of phonological rules is automatically derived. We differ from her account only in that we take all intra-level rules (i.e. (M, M) rules, (W, W) rules, and (P, P) rules) to have the character that they apply just in case they improve the well-formedness of their input, and we take the OCP to be only one of several such tactics that may hold of levels. We furthermore interpret Odden’s impressive scholarship as establishing that the OCP is a tactic that must be specified in a language-particular way for each of the three levels of the phonological grammar.

(3) As suggested in note 16 of chapter 2, the principle of rule application in harmonic phonology – that rules apply only if their application improves the well-formedness of a representation along a certain ‘dimension’ – when combined with the theory of autosegmental licensing

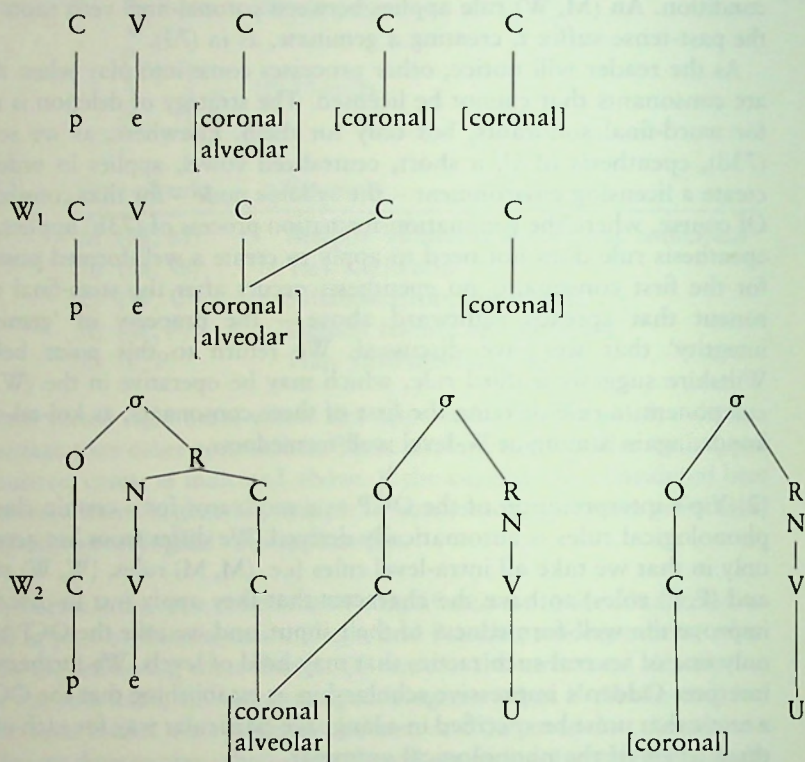
- (73) (a) Underlying form: $pe\tau + t + t$
 Surface: $[pe\tau \tau UdU]$ 'give birth + past + 3rd person'
 τ is an alveolar stop; voicing is non-distinctive word-medially.

(b) Rule



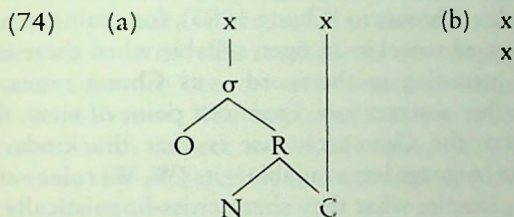
(c) Rule: insert U

(d) Schematic derivation



proposed in chapter 3 serves to account automatically for the most compelling examples of geminate integrity and inalterability. With regard to inalterability, the clearest examples are all of the general character that a coda-weakening process fails to apply to geminates. Klingenberg's Law in Hausa, readers will recall, is a typical example of this sort, according to which obstruents in coda position become sonorants. This shift is entirely conditioned by licensing considerations. Hausa does not license point of articulation in its coda, a W-level phonotactic; however, the coda may associate with a point of articulation autosegment just in case that autosegment is also associated with an onset position, which licenses it. Thus, it follows that geminate obstruents do not violate the W-level phonotactic, and Klingenberg's Law will not apply – will not be even be *tempted* to apply, so to speak. Precisely parallel considerations arose in the case of Toba Batak earlier in this chapter; see note 15.

(4) We may capture significant 'soft' cross-linguistic universals which formerly eluded formal capture. One of the most striking of these arose several times in chapter 3, in connection with the natural relationship between heavy syllables and prosodic prominence. Heavy syllables are syllables with a coda that licenses association with a second Row 0 grid mark, as in (74a); a prosodically prominent, or stressed, syllable, is one with a Row 1 grid mark, as in (74b). The two are distinct, but nonetheless there is a clear connection between them. We may express this as a universal W-level³⁹ phonotactic, which specifies that syllable-grid associations are well-formed in those cases where the syllable is heavy if and only if the syllable is stressed.⁴⁰ This leads to four relations of inequality with regard to weight and prominence, as shown in (75).



Each of these relative statements of well-formedness can serve as triggers for simple rules of grid or syllable adjustment. Case (a) is, of course, just the principle that governs the rule of quantity-sensitivity (QS), discussed in chapter 4. Case (b) arose several times in our discussions in chapter 3, in connection with Selayarese, with the Scandinavian languages, and with Zoque. In each case, a rule added a mora to a

- (75) (a) $\begin{smallmatrix} x \\ x \ x \end{smallmatrix}$ is better-formed than $\begin{smallmatrix} o \\ x \ x \end{smallmatrix}$

A stressed heavy syllable is better-formed than an unstressed heavy syllable.

- (b) $\begin{smallmatrix} x \\ x \ x \end{smallmatrix}$ is better-formed than $\begin{smallmatrix} x \\ x \end{smallmatrix}$

A stressed heavy syllable is better-formed than a stressed light syllable.

- (c) $\begin{smallmatrix} o \\ x \end{smallmatrix}$ is better-formed than $\begin{smallmatrix} o \\ x \ x \end{smallmatrix}$

An unstressed light syllable is better-formed than an unstressed heavy syllable.

- (d) $\begin{smallmatrix} o \\ x \end{smallmatrix}$ is better-formed than $\begin{smallmatrix} x \\ x \end{smallmatrix}$

An unstressed light syllable is better-formed than a stressed light syllable.

stressed syllable, just in case the syllable *needed* that mora in order to be heavy – in other words, the language would lengthen a vowel in a stressed open syllable. But not all languages do such things, and that is an aspect that the treatment provided by harmony phonology deals with especially well. A language such as the three mentioned above may take an especially simple route to make its W-level structures better-formed. The rule will be simply: add a coda position – and it will apply only in the right cases, those where it improves the well-formedness of certain syllables with respect to (75b). Other languages may contain rules that are somewhat more complex. Chamorro (Chung 1983), for example, has a rule that lengthens a stressed vowel in an open syllable when there is a stressed, closed syllable preceding in the word – as Chung notes, a harmony principle of a rather abstract sort. From our point of view, the important conclusion from the Chamorro case is that the kinds of phonological resources the language has available – its (W, W) rules – are typically, but not always, simple; what they share cross-linguistically is their common direction of improvement, as specified by principles like those given in (75).

Case (75c) represents the motivation for all languages that shorten vowels in unstressed position. This is a common process, though often not recognized for what it is. A particularly interesting example of this is given by Selkirk (1986) for the Bantu language Chimwiini. Finally, case (75d) represents all cases where light syllables are destressed, a not

uncommon process at word-periphery, where it does not wreak havoc with the permissible foot structure of a language.

What is important to see in all these cases is that to specify the precise environment for each rule is sometimes formally difficult and always unnecessary insofar as it simply recapitulates the universal tactic. Rule (61) of chapter 3, for example, in effect adds a mora to a stressed syllable in case the syllable is monomoraic, but geometrical representations are notoriously poor ways of representing what something is *not*. As we have noted on several occasions, our autosegmental and metrical forms of representations and of rules are not well suited for expressing a *lack* of associations. From the point of view of harmonic phonology and its conception of rule application, this is as it should be, because those rules whose function is to add association lines or metrical structure when it is not yet there are always rules aiming at satisfying a 'completeness' or a 'saturation' of a representation, or some other related kind of phonotactic, such as that in (75).

A similar perspective is offered in Goldsmith (1987c, 1990), with respect to the interaction of tone assignment and metrical structure.

(5) This naturally brings us back to a consideration of the original 'Well-formedness Condition' of autosegmental phonology, in (64). We may now re-ask the question: first, is there a WFC in autosegmental phonology? and second, is spreading of the sort that it induces universal? The answer is that the Well-formedness Condition in (64) is just one of many W-level or P-level phonotactics that can be stated in terms of the (minimum, maximum) notation discussed in chapter 1. In particular, (64) says that, on the skeleton-tone chart, the skeleton is specified for a (1, x) value, and the tone tier is likewise specified for a (1, x) value (where 'x' means not specified). Is this universal? At this point, the answer is uncertain. It may be that in all cases where less than the minimum association is provided there simply is no rule available in the language that would allow the representation to become well-formed in this respect. In short, it may well be universal. The implementation procedure (65), we suggest, *is* universal with respect to (M, M) rules, (W, W) rules, and (P, P) rules. However, rule (66) itself is not universal: it is a language-particular rule.

(6) On a related point, in our discussion of Kiparsky's analysis of Catalan, we noted that a proper phonological account of point-of-articulation assimilation for nasals needed to be specified as a rule that applied only to nasals that were not already specified for a point of articulation, as sketched in (15) and (16) of chapter 5. The present notion of harmonic application provides just that notion. The phonotactic in (16) there required that all consonants be specified for a point of

articulation. The autosegmental spreading rule of assimilation would accomplish that end, if relevant; otherwise, a context-free default specification would assign a point of articulation.

(7) The naturalness of compensatory lengthening when an empty coda position is produced can be captured in the same way as the other cross-linguistic 'soft' constraints that we have discussed, from the point of view of harmonic application. Readers will recall that the generalization we wish to capture is as follows. Syllable and coda structure is established on the basis of segmental material that may later undergo deletion. If such a deletion process leaves a coda position unassociated with any melodic material (consonantal or vocalic features), then there is a strong (but soft) universal tendency for an element, on either the left (a vowel) or the right (a consonant), to reassociate to that coda position. The phonotactic may be as simple as this: that a licenser must license at least one melodic (vocalic, consonantal) autosegment at the W-level and one at the P-level.

(8) Finally, the distinction used here – between M-level, the level at which segmentally represented morphemes are represented, and W-level – motivates those uses of the Morpheme Tier Hypothesis that can be empirically motivated. The M-level representation is essentially devoid of phonological motivation; its representations may violate every conceivable phonotactic, every conceivable phonologically oriented constraint of the language. Its sole function is as a repository of the minimal information necessary to capture the sound characteristics of the morpheme. It is a structure that incorporates the morphemes that provide the realization of the morphosyntactic information. The W-level, on the other hand, is the level at which such phonological information is restructured in order to maximally satisfy the language-particular organization principles which we call syllable and autosegmental phonotactics, of which licensing is an important, though not a unique, member. The W-level representation thus expresses the form the language squeezes its morphemes into in order to satisfy the alternating rhythm of consonants and vowels, of properly licensed coda and syllable material, of tonal association, and so on. The phonological rules of the language are its ways of manipulating the phonological substance present at the deeper M-level, and they express the options open to the language with regard to how much the language can 'deform' the underlying representation in order to maximally satisfy the multitude of competing demands of well-formedness at the W-level.

Thus, it seems reasonable that the morphological procedures responsible for constructing an M-level representation may produce a 'pseudo-phonologized' representation in which morphemes are placed on separ-

ate tiers. The process of forming a W-level representation, then, requires what McCarthy calls tier conflation, which is one particular technique for restructuring an M-level representation into one that satisfies the universal and language-particular demands of W-level.

6.6 CONCLUSION

A theory of phonology is built of three parts: it is a theory of the nature of phonological representations; it is an inventory of levels of representation, and a characterization of each level; and it is a theory of phonological rules, the statements that relate representations on each level.

This book is aimed primarily at the first part, the nature of phonological representations. We have explored the nature of autosegmental representations, metrical grids, and syllable structure. We have offered autosegmental licensing as a characteristic that determines the essential properties of syllable structure.

At the same time, we have had to develop a certain number of ideas concerning levels of representation, and we have emphasized the importance of W-level structure, that structure over which licensing conditions serve as the primary phonotactic, or well-formedness condition. We have explored lexical phonology as one explicit account of several levels within generative thought, and have tried to separate some of the more useful from the less useful ideas in that area.

With respect to the notion of rules, throughout most of this book we have retained the traditional generative conception, according to which rules come with a structural description and apply if that description is met. As indicated briefly in the last two chapters, and especially in the preceding section, I believe that this notion stands in need of serious revision, although, as we have seen, ongoing research in phonological theory has been able to enunciate a powerful conception of phonological representations, independent of any further changes in the theory of rules. Now, however, with this new theory in hand, we may proceed to a novel and even more compelling picture of the nature of phonology, in which rules interact with phonotactic conditions on a small number of levels to develop representations at each level satisfying the conditions stated there. This picture has much in common with current work in a number of other areas of linguistic theory.

In phonology, the model we arrive at is one that looks much more like a model of chemistry than the models of classical generative phonology, in which the phonological grammar resembled nothing more than a

computer program. In the model that is emerging currently, representations have a complex geometric structure, but relatively few degrees of freedom in the changes they may undergo. Rules define possible changes in the structure of the phonological material, and in each and every case, the changes are motivated by an attempt to achieve a greater satisfaction of well-formedness conditions. This bears a striking similarity to the notion that chemical systems tend toward a lower energy level, consistent with the physical properties that they have. The application of this kind of model has been urged elsewhere in cognitive studies by Smolensky (1986), for example, and the convergence of work in phonology with that in other areas of cognitive science offers great hope for continued advances of the sort that we have seen in phonology in the last fifteen years.

- 54 This formulation of the restriction is tree-oriented in its statement. A less constituent-oriented formulation would be that the right-to-left perfect grid (quantity-insensitive) application that constitutes stress retraction applies only to stretches of unstressed syllables, which is the way Perfect Grid always works; in addition, there is no forward clash override, again the unmarked case. Kiparsky (1982a) offers one example of a case where stress retraction does not work this way: the case of *solidify*, where he suggests that the 'unfooted' *i* of the suffix *-ify* is enough to trigger stress retraction. It is equally reasonable to suppose that this is a case of close juncture, precisely as Kiparsky proposes for such 'irregular' forms as *democratize*, where stress retraction does indeed appear to have overridden the stress pattern of the base.
- 55 The same point is made on independent grounds in Fabb (1985).
- 56 See the typological remarks in Booij and Rubach (1987).
- 57 The skeptic who was trying to establish a case against treating trisyllabic shortening as a phonological process might proceed suffix by suffix, looking a bit more closely to see whether alternations (like the one cited by Kiparsky 1982a, *omen/ominous*) holds up across the range of words in the English lexicon. A glance through a backwards-alphabetized dictionary, such as *Walker's Rhyming Dictionary*, suggests that there is no large class of words formed with an *-ous* suffix attached to an independently existing word base. Some exist, to be sure, such as *humorous*, related to *humor*, or *scandalous*, or *perilous*. Others, like *ominous*, the example mentioned by Kiparsky, are rather distantly related to the word that looks like it might serve (or might once have served) as a base, since calling a sky ominous is not to call it an omen. Still others look like they are formed with an *-ous* suffix but have no plausible base; the pattern *felicitous/felicity* is not matched by *jealous/jelly*.

A glance, then, at the eight or nine hundred *-ous* suffixes in English turns up only one clear case where the adjectival form has a short vowel, but the related base has a long vowel: *libido/libidinous*, but the short vowel [ɪ] found in the derived form is not the expected vowel (we would expect ɛ) – nor, of course, is the *in* expected. A good number exist that violate the shortening prediction, such as *cretinous*, *mountainous*, *libellous*, *poisonous*, *sonorous* (for some speakers, such as myself), *scrupulous*, *numerous*, *gratuitous*, *odorous*, or *cumulous*. It may not be accidental that most of these examples involve long round vowels, which may simply not fall under the generalization expressed by trisyllabic shortening; but be that as it may, the *-ous* forms cannot be said to provide positive support for the existence of trisyllabic shortening as a rule in English. What the examples do suggest, however, is that the direction of change that is induced in the stem by the juxtaposition of affixes is in the direction of what would be found in nonderived forms.

Chapter 6 Further Issues

- 1 Or rather, the intersection of these intersecting sets with the set of segments in the language at hand. On a related point, one interpretation of a theory of

privative feature values is that each feature *F* defines only *one* set of segments; natural classes would then consist of the intersections of the various sets thus defined.

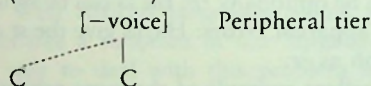
- 2 I use the term 'post-Bloomfieldian' to describe the set of views on phonological theory outlined in various ways in the now classic reader Joos (1957). (See Hymes and Fought 1981 for an excellent discussion of the term and the trend.)
- 3 Several of these features deserve some mention, because they are not very familiar. Voiceless obstruents are [+stiff vocal cords]; voiced obstruents are [+slack vocal cords]. Aspirated obstruents are [+spread glottis], as is *h*; glottalized obstruents are [+constricted glottis], as is ʔ.
- 4 At least one statement appears in the literature that adopts such a view (Sagey 1988), though that note misconstrues the basis of a theory of phonology, in my opinion. Sagey discusses a model of autosegmental phonetics – i.e. a model for the description of articulatory events in time – rather than a theory of phonology. In addition, she attempts to show that properties of an autosegmental model of phonology (or, as I suggest, phonetics) may derive from 'extralinguistic knowledge' (p. 109). Again, this seems to me to be mistaken in principle, not in detail. If we take the term 'knowledge' in a strictly cognitive and reflective sense, then such knowledge is irrelevant to the structure of phonological representations; if we take it in some other sense (though what sense that might be is difficult to imagine), a sense that would extend to the phonetic events that take place in time spans measuring no more than 10–100 milliseconds, then the axiomatization of our common sense notion of time (p. 110) is certainly false – in a wide range of areas, subjective events at the micro-level do not organize themselves in a fashion that respects our common-sense view of time; deriving phonetic principles from an *a priori* axiomatization of time in such a case does not show that the principles derive from some external knowledge in that case (even leaving aside, as I have said, the problematic notion of 'knowledge' that is involved).
- 5 See, for example, Kiparsky (1968).
- 6 Another case in which a classificatory feature has seemed appropriate – though it is not matched, it would seem, by a phonetic manifestation in any direct way – are the features of juncture, such as the featural difference between a '+' boundary and a '#' boundary, in the SPE analysis.
We not infrequently find segments that are identical (for our practical purposes) in two or more different languages, but whose phonological behavior is distinct in an unexpected way. Both *b* and *d* have sonorant-like properties in several West African languages, while *v* and *w* are also segments that may act like a sonorant in one language, an obstruent in another. A common way to deal with this problem is by changing the specification of this segment for the feature [sonorant], but that is just a way of saying that a phonological use of a feature may diverge from a fixed and constant phonetic realization.
- 7 Readers will recall that a feature is used as a *privative* feature if only one value of that feature is permitted in a representation, and it is used as an

equipollent feature if two values (+ and –) are permitted in the representation.

- 8 This question is raised, though not answered, in the interesting discussion in Hockett (1961: esp. 41). Fudge (1967) considers an interesting, but currently unpopular, view. On the general subject of the difference between 'phonetic features' and 'phonological features', see Vennemann and Ladefoged (1973), and the apposite remarks of Hayes (1986b: 477).
- 9 A number of useful papers will appear in van der Hulst and Smith (to appear), of which I have seen only Dikken and van der Hulst (1988).
- 10 This remark may deserve some further elaboration. Some aspects of the sound signal go unrepresented in the phonological and phonetic representations. Aspects that are universal and difficult to represent at these levels are prime candidates for characteristics to be left out of such representations, such as the effect of vowel height on fundamental frequency or on duration. In this way, then, phonetic representations unabashedly underrepresent the speech signal, but that is not problematic. The question becomes thornier with respect to characteristics of a speech signal that may be language-particular, and yet which we do not need to represent explicitly in a phonological representation at any level, as far as we can see. An example of this sort might be voicing of vowels in English. To my knowledge, there is no evidence, or reason to believe, that vowels in English are marked for voicing, though cross-linguistically this feature may well be contrastive for vowels. In short, underspecification theories of the sort we considered in ch. 5 drive out a good deal of the featural specification in underlying representations; our question now is to determine precisely what 'overspecification' theory (so to speak) requires that such feature specification should be put back in, and at what point. As should be apparent, I believe that considerable caution is in order with respect to a strong 'full specification' or 'overspecification' position, as of the sort mentioned in the text above. I have been influenced here by unpublished work by Osamu Fujimura on these issues from a phonetician's point of view; cf. Keating (1988), which appeared shortly before this book went to press.
- 11 Cf. Goldsmith (1981), Halle and Vergnaud (1980).
- 12 Complications arise in the palatal member of the series, because the palatal consonants to which the nasal assimilates are laminal, rather than apical, and the assimilating *n* remains apical; see Harris (1969: 9–13). For a phonological account of this general area, see Carreira (1988). Harris (1984) discusses the general problem of nasal assimilation in Spanish from an autosegmental point of view in much greater detail than I do here.
- 13 There are three optional alternative forms given by Hayes: the *ph*, *th*, *kh* sequences can be optionally ^ɰ*h*; the *ts* can be *ss*, and the *ms* can be *ss*. I have changed two apparent errors: Hayes give the *st* combination as *sp*, and the *kl* combination as ^ɰ*r*.
- 14 This is an excellent example of the more general proposition that geometrized autosegmental and metrical analyses tend to require far less extrinsic ordering, all other things being equal, than purely segmental analyses.
- 15 While Hayes has undoubtedly presented an elegant and insightful account of Toba Batak, certain questions do remain regarding the degree of 'overspeci-

fication' of the representation he employs. That is, the kinds of generalization that we observed in chs 3 and 5 regarding weakenings of consonants in coda positions arise here in Toba Batak, and such processes can be described only in the context of an underspecification theory. Assimilation processes, such as those by which coda consonants assimilate to onset positions for point of articulation, are motivated and guided by licensing restrictions that block a coda position from licensing a point of articulation. The formation in Toba Batak of a glottal stop in coda positions that are not otherwise geminated (on Hayes's account, protected therefore by the Conjunctivity Condition) is highly reminiscent of the effects that we observed in ch. 3, if we assume that in Toba Batak there is a phrase-level syllable representation at which certain licensing conditions are imposed. Such conditions would have to permit in coda position a glottal stop (an obstruent unspecified for point of articulation), but to rule out a voiceless stop specified for point of articulation, a condition very similar to what we saw in a number of languages in ch. 3. However, that cannot be *quite* right for the case at hand, because an independent point of articulation *is* permitted in coda position just so long as the consonant is nasal (i.e. *m*, *n*, and *ɲ* appear contrastively in the coda). Licensing does not provide an account of why point of articulation may not be licensed except in the presence of a nasal autosegment in the coda. However, there is an alternative possibility worth considering. We assume that, despite the fact that these rules apply post-lexically, they apply to representations that satisfy underspecification criteria; as we have just observed, this is a necessary condition for using a licensing approach. Nothing prevents us, however, from assuming as well that the value of the feature [nasal] that is operative in Toba Batak is [–nasal]. On this account, a segment unspecified for this feature is nasal, and oral obstruents must be explicitly marked as [–nasal]. Three factors suggest that this is indeed correct. First, this interprets an *n* as the totally unspecified consonant, and the rule of *n*-assimilation (17) becomes formally more natural, as it is interpreted as the assimilation of the total unspecified consonant to its right-hand consonantal neighbor. Second, the rule of denasalization (18) clearly demonstrates that the feature [–nasal] is present and can spread autosegmentally; its doubly-linked character in the output representations of (18) is what serves to block the application of (20), glottal formation. Denasalization (18) will be simplified further by eliminating the change whereby a [+nasal] autosegment is deleted, since there will be no such autosegment present. The modified form of the rule will be as in (i). Third,

(i) Denasalization (reformulated)



the rule of glottal formation (20) now is revealed as a coda restriction blocking the simultaneous licensing of a [–nasal] and point of articulation on a coda consonant. If we assume the following feature specification in (ii), then we may identify the coda in Toba Batak as licensing a maximum of *one*

of the following distinctive features at the phrase level in question (not including major class features, as before). Segments specified with no more than one feature are: *m, n, ŋ, s, l, r, ʔ*. The other segments, the oral stops, are combinations of point of articulation and [–nasal]. (We have little information about the liquids *l* and *r* on which to base their featural analysis.)

(ii)	p	t	k	b	d	g	s	h	m	n	ŋ	r	l	ʔ
P of A	lab.	cor.	vel.	lab.	cor.	vel.	alv.		lab.	cor.	vel.			
Nasal	–	–	–	–	–	–	–							–
Rhotic												+		
Lateral													+	
Voice				+	+	+								

- 16 See Cook (1987) and Poser (1982), for two examples.
- 17 It is, unfortunately, not clear whether this process is taken to apply in the lexical phonology or in the post-lexical phonology, a point bearing on the suggestion in n.15. If the creation of geminate *ls* creates a form that is not present in underived forms, then the well-known version of lexical phonology discussed in ch.5 would predict that the rule is post-lexical, since structure preservation would prevent a rule from creating a type of structure within the lexical phonology that was not found in underived forms. If this is the case, then this represents a case of underspecification in post-lexical phonology as well.
- 18 This section is heavily influenced by Hayes (1988), who attributes the discussion on the 'diphthongization paradox' to Steriade (forthcoming), which I have not seen. I have modified Hayes's notation a bit, substituting a geometric model for an algebraic formalism.
- 19 This section is based on Goldsmith (1985b, 1987a); for more details, see these references.
- 20 Most notably, the papers by Kiparsky (1982a, 1982b, 1985), Pulleyblank (1986a, 1986b). On the other hand, several studies have considered more substantive revisions of vowel representations, including Goldsmith (1985b), Rennison (1985), a number of papers by van der Hulst and Smith (cf. 1985), and papers by Schane (1984), and Kaye, Lowenstamm, and Vergnaud (1985), among others.
- 21 Given the symmetry of [round] and [front], we could in principle choose to call the feature [front] and make rounding predictable post-lexically, in the simple five-vowel system.
- 22 There has been considerable discussion in the literature as to how one might allow underspecification of one sort or another without allowing anything that smacked of a three-way formal distinction for features. Put another way, the discussion has addressed the question, if binary features cannot be used in a ternary fashion, what are the consequences for formal phonology? As the discussion in this chapter suggests, this seems to me to be premature; binary features, used in an equipollent fashion, *can* give rise to distinctions that are in effect three-way.

- 23 One can establish ordered default rules in such a way as to make these vowels arise out of the fewest underlyingly marked feature specifications. That is hardly the point; one can *make* any set of vowels be the least marked. The task is to establish a representational system in which the correct result has a natural basis.
- 24 It is similar to the familiar notion of *symmetry*, to be sure, but formalized in a slightly different fashion.
- 25 This analysis is very similar in spirit to the accounts given in Goldsmith (1985b), where I argue that both Hungarian and Finnish should be viewed synchronically as having a canonical five-vowel system, with one equipollent feature, [round], and one privative feature, [low]. In Hungarian and Finnish, the privative feature [back] may be present as well in the lexical entry. See also Goldsmith (1987a), and Ringen (1988b).
- 26 See Clements and Sezer (1982) for further discussion of Turkish, and the important matter of its interaction with consonantal specification.
- 27 This proposal was first made in Goldsmith (1979, ch. 4). It is aptly discussed by Singler (1980) and Odden (1986) (both excellent studies of this principle), but is not correctly represented in McCarthy (1986: 253–4), at least as I read it; McCarthy cites only Odden's (1986) paper (not yet published).
- 28 Pulleyblank (1988) offers several appealing arguments for the Morpheme Tier Hypothesis.
- 29 McCarthy actually suggests that (60b) will maintain the prefixal material *li* on a separate tier from the stem material *shob*, for reasons that need not concern us here concerning the ordering of processes.
- 30 McCarthy's (1986) position regarding fusion is guarded, it appears. As Odden notes, he suggests that the function of the OCP 'is not that sporadically assumed in the tonal literature ... [that of] a process that fuses adjacent identical tones into a single one' (208); 'I reject the fusion interpretation of the OCP' (222). Yip (1988) interprets McCarthy as including OCP fusion effects as part of the tier conflation process, though McCarthy actually offers this at the end of his paper as a notion 'in the realm of speculation' (257); through the substantive part of the paper, he is clearly at pains to avoid any such suggestion.
- 31 See Liberman (1979), Halle and Vergnaud (1982), Haraguchi (1977), and most forcefully Pulleyblank (1986a).
- 32 Sommerstein (1977) offers a broader discussion of the issue, though in less detail, and suggests (73) that his conception is 'to some degree under the influence of' stratificational grammar, as articulated, for example, in Lamb (1966).
- 33 Paradis (1988) more recently, following up on Singh (1987), has extended and developed some of these ideas. I have also profited from Bosch (1988) and Wiltshire (1988), who explore these issues with respect to Scottish and Irish respectively.
- 34 There are more than a few parallels to central considerations of stratificational phonology, it may be noted; see Sommerstein (1977). Of course, even lexical phonology is considerably more stratificational than classical generative phonology.

- 35 This type of notion of a 'soft' – a violatable – well-formedness condition is extremely important to the approach being suggested here, and in outlook is at odds with the classical generative approach. Nonetheless, it has clear antecedents in the literature that we have mentioned. For example, this is precisely the claim of the 'Well-formedness Condition' of Goldsmith (1979) discussed above; it is noted in Liberman and Prince (1977: 311), who distinguish between situations that produce 'pressure for change', and language-particular specifications of when and how permission is granted to change a representation. Yip (1988) also observes this point, though she takes it to be the case that one 'repair strategy' (for the OCP, in the case at hand) will always be available, though no evidence is presented for this.
- 36 One especially obvious aspect that is overlooked in this representation is the characterization of cyclic morphology, as discussed in the last section of ch. 5. For purposes of clarity, I will leave the diagram as it is, recognizing that additional complexity is required. The term 'harmonic' alludes to work by Smolensky (1986), to which we will briefly return below.
- 37 The pervasiveness of this process and its linkage to well-formed syllabification was the basis of a large part of Kisseberth's influential notions concerning 'conspiracies' (Kisseberth 1970).
- 38 I simplify Wiltshire's presentation in (71b); she argues for a coplanar representation of the various coronal points of articulation, along the lines suggested in Archangeli (1985).
- 39 There is some evidence that this should hold of the P-level in some languages.
- 40 Clearly there is more to be said about how such well-formedness conditions should be properly stated, but this question takes us well beyond the bounds of this chapter.