A tier geometry for vowel systems[*]

PHILLIP BACKLEY

1 Introduction

Over the past two decades the literature has born witness to a sizeable array of different approaches to the organisation of phonological information. Yet amid this theoretical diversity, the ultimate aims of the discipline have remained largely unaltered. Perhaps the most important of these goals can be summarized by the following three points. First, the necessity for observational adequacy requires little in the way of independent justification, since a model which lacks the capacity to capture a full range of attested lexical contrasts must be deemed to fail even at a descriptive level. Second, a degree of naturalness should somehow be encoded in a representation. In the case of vowel systems, which will be the primary focus of this paper, naturalness is reflected in the relative ease or simplicity with which an unmarked melodic expression or a widely encountered vowel inventory can be expressed within the formalism. And third, the construction of a representation system should allow a characterization of attested phonological processes to be accessed with the same facility that is displayed in the description of static melodic properties. Within the context of vowel systems, the latter point must encompass the full range of observed harmony phenomena, together with related issues such as transparency and opacity.

In recent years phonology has seen a notable shift in approach with regard to the way representational models have attempted to fulfil these criteria. The powerful rule-based analyses that followed the traditions established by Chomsky and Halle 1968 (henceforth SPE) achieved a descriptive adequacy, but clearly at the expense of any real explanatory worth. It has since been acknowledged that the inherent problems of this approach stem from the burden of description having been carried by the rule component, rather than by the structures to which these rules were applied. In marked contrast, most current approaches have developed a shift in emphasis towards a richer, yet more highly constrained form of both melodic and prosodic structure, to which a vastly impoverished set of dynamic operations is permitted to apply. The benefits of such a move are most clearly manifest in the increased explanatory power that these models possess. So, given a particular input structure, it is the structure

[*]My thanks go to John Harris and to my postgraduate colleagues for their insightful comments on an earlier draft of this paper.

itself which constrains the nature of permitted phonological phenomena that may potentially be observed. Typically this is accomplished through a set of very general, cross-linguistic principles which impose restrictions on the well-formedness of both structures and operations. Thus, these models succeed in offering some degree of explanatory power, due to the necessity of a process finding a local cause in an input representation. Conversely, the impossibility of non-attested phenomena follows naturally from the absence of a suitable trigger.

One of the most restrictive and highly constrained approaches to phonological representation owes its theoretical origins, at least in part, to the *Theory of Charm and Government* first presented in Kaye, Lowenstamm and Vergnaud (KLV, henceforth) 1985. This model has subsequently found support in various works including Kaye (1989), Charette (1991), and KLV (1990). The core assumptions of KLV (1985) have also fuelled the development of a related, but clearly distinct phonological model as set out in Harris (1992) and Harris and Lindsey (in press), and which is summarized and evaluated in Harris (1994a). The latter sees the notion of licensing as a central property which plays a pivotal role in the amalgamation of all melodic and prosodic components into a single, coherent structure. The theoretical standpoint assumed throughout this paper to a large extent mirrors that presented in Harris (1994a), where the following tenets are argued for. Melodic primes are monovalent, and are fully interpretable, either alone or in combination with one another. And the vowel space is marked out by the three resonance elements [A], [I], and [U], which are most commonly interpreted as the basic corner vowels **a**, **i**, and **u** respectively.

However, there are several issues where the present discussion fails to find agreement with Harris (1994a), the first of which concerns the asymmetric relation of headship existing between melodic units within an expression. Harris exploits this property in the description of both static melodic contrasts and also long-distance harmony phenomena. Yet I shall follow Takahashi in suggesting that head/dependent relations at the melodic level do present problems of incompatibility with the widely established phonological principle of Structure Preservation (first proposed by Selkirk (1982) in connection with syllable structure), and that, consequently, it may be desirable to investigate an alternative means of expressing the facts which commonly refer to headship distinctions in their description. I shall propose that a closer examination of the way in which individual melodic tiers are arranged in relation to one another may offer an insight into the representation of melodic contrasts and harmonic phenomena in vowels. I take as a starting point the tier geometry proposed in Rennison (1987, 1990), and consider the ways in which such a configuration could be adapted to the theoretical environment assumed here.

The second point at which my proposed model departs from the representational structure assumed by Harris concerns the mechanism which allows melodic material to be interpreted across a domain larger than that for which it has been lexically specified. The traditional autosegmental notion of spreading will prove somewhat awkward for this task, given the way I intend to portray the presence or absence of melody in a structure. Specifically, I adopt the notion presented in Takahashi (in preparation) of a full set of elements being available to all prosodic positions, their interpretation being dependent on lexical activation. In this way, a range of assimilatory phenomena can be captured by activating melodic units at a relatively 'high' level in the prosodic hierarchy.

My discussion will be structured as follows. I begin in §2 by arguing that melodic representation need not necessarily make reference to properties of headship and, in this vein, I examine the melodic tier configurations appropriate to a number of commonly attested vowel inventories. I shall propose that the variation which may exist between these configurations is restricted according to the scope of only two parametrically controlled axes of variation. The first controls the extent of tier division, and is based on the work of Rennison cited above. The second determines whether or not a complement may be licensed by a tier, and owes its existence to the interesting notion of *persistent complement*, developed in Takahashi (in preparation) from the idea of *persistent non-head* employed in Harris (1994a). Then the suitability of the proposed tier geometry will be demonstrated in §3, where it is employed in a re-analysis of ATR harmony in Akan. This will be followed in §4 by a short discussion of the mechanism by which melodic units are multiply-linked to prosodic structure. The appropriateness of spreading as a means of achieving many-to-one associations will be questioned, and the case for melodic alignment will be proposed as an alternative.

2 The representation of vocalic melody

2.1 Introduction

As already indicated, I shall assume the three monovalent elements [A], [I], and [U] as the only melodic primes necessary for the description of contrasts in (oral) vowels. Harris and Lindsey (in press) present a summary of arguments to support the selection of these as an exhaustive set of resonance elements, although their arguments are not entirely new to the literature. Indeed, the same tricorn vowel space had also been employed much earlier within a variety of different frameworks. These include Particle Phonology (Schane 1984), Dependency Phonology (Anderson and Ewen 1987), and Government Phonology (KLV 1985, 1990). For as long as these atoms have been utilised in the composition of melodic expressions it has been recognized that the three elements alone are inadequate for capturing the entire set of attested contrasts. Consequently, an additional property has been required in order to expand the generating capacity of the model to a level which can match the vowel inventories it attempts to describe. In Anderson and Jones (1974) an asymmetric headship device (in their terms, a *subjunctive dependency relation*) was encoded into melodic representation in order to carry out this function, and this was subsequently inherited by KLV (1985) and developed by the majority of those who have followed the KLV tradition.

2.2 The role of headship

According to the notion of headship, a single element within a complex expression is identified as the head of that expression, while the remaining primes assume dependent status. The head is deemed to make a greater contribution to the melody than do its dependents; indeed, the prominence of the head is often apparent in the interpretation of the expression, as shown in the comparison of $(\underline{U},A)$ with $(U,\underline{A})$, where heads are underlined. The respective interpretations as **o** and **ɒ** are explained by the saliency of labiality in the former and openness in the latter, as reflected in the headship distinction.

The identification of an element's status as either a head or a dependent has proved crucial in the description of a number of harmonic phenomena. In these cases, the participation of an active element [ε] is restricted to those instances where it is lexically present as the head of an expression, or alternatively as a dependent element, according to the exact specification of the harmony process at work. An example of this is found in the analysis of Chichewa height harmony as proposed in Harris (1994b). Here, headship is employed in order to capture the generalisation that the element [A] spreads from the mid vowels **e** $(\underline{I},A)$ and **o** $(\underline{U},A)$ but not from the low vowel **a**, where it assumes head status $(\underline{A},@)$[1]. The harmonic pattern is thus described by referring to the spreading of dependent [A].

The analysis of Tigré vowel harmony found in Lowenstamm and Prunet (1988) also relies on the notion of headship, but exploits it in a rather different way. Instead of utilising head status or dependent status to identify the active unit in a harmonic

[1]It should be noted that a 'dummy' element [@], devoid of phonologically significant properties, is conventionally used to indicate the absence of any 'true' melodic content in a given head or dependent position. In the representation of the low vowel $(\underline{A},@)$, therefore, there are no dependent elements.

process, it is the head status of a particular element which becomes the property of an entire domain, as the following examples demonstrate.

(1)	(a)	sälsälät			'bracelet'
	(b)	sälsälät + u	→	sälsälätu	'his bracelet'
	(c)	sälsälät + a	→	salsalata	'her bracelet'

Briefly, the alternation between **a** and **ä** depends on the presence of a low vowel later in the word (vowel harmony applies leftwards in this system). The addition of the low vowel suffix *-a* in (1c) has the effect of lowering the central vowels of the noun, resulting in the interpretation *salsalata*. In contrast, the central quality **ä** is retained in the absence of a following **a**, as in (1a) and (1b). Lowenstamm and Prunet account for this in the following way. The low vowel is taken to be the interpretation of the element [A] as a head, and is therefore represented as ($\underline{A}$,@). The central vowel **ä**, on the other hand, consists of the same sole element [A], but in a less salient guise. Hence, it assumes the role of a dependent element and the expression is represented as (A,$\underline{@}$)[2]. The harmonic pattern shown in (1c) is then characterized by appeal to the notion of 'head harmony', whereby all the vocalic expressions within a domain must contain the same element as a head. In *sälsälät* the vowels are all [@]-headed[3]. However, when a suffix containing an [A]-headed expression is added, as in (1c), a head-switching operation is observed in the harmonising vowels, causing lexical (A,$\underline{@}$) to be interpreted as ($\underline{A}$,@).

So, headship has been invoked in the description of harmonic phenomena in two different ways. In Chichewa (Harris 1994b) it was used to distinguish two distinct representational objects, headed [A] and dependent [A], only one of which was shown to be active in the harmony process. On the other hand, in Tigré (Lowenstamm and Prunet 1988) all the vowels in a given environment are required to share the same element as a head, which is achieved via a head flipping operation that allows the lexically given expression (A,$\underline{@}$) to be interpreted as ($\underline{A}$,@). I wish to claim that the asymmetric headship relation is problematic with respect to both of these applications, each for its own independent reasons. Below I shall return to the problems inherent

---

[2]In Tigré there is a lexical contrast between the two central vowels **ä** and **ɨ**. The latter can be described in terms of the absence of all resonance elements, resulting in the representation [@]. For details regarding the interpretation of melodically empty nuclear expressions, the reader is referred to Charette (1991).

[3]An expression which employs the 'empty' or 'null' element (see footnote 1) as its head may be appropriately described as 'headless', given the absence of any lexically specified melodic content to act as a head.

in Harris' utilisation of the head/dependent distinction, while here I shall focus on the consequences of 'headship harmony' as proposed by Lowenstamm and Prunet.

In evaluating the head-switching mechanism proposed for the analysis of Tigré we are forced to consider a much wider issue which concerns the restrictions that may be placed on the possible interpretation of a given lexical form. With respect to those frameworks couched within the SPE tradition, this question became a central area of debate, and was formalised under the notion of abstractness. In contrast, within the context of a more tightly constrained model of representation such as that assumed here, where the scope for structural alteration is drastically limited, the issue of how a lexical form can be manipulated in its interpretation provokes little argument. In the KLV tradition it is assumed that only two structural operations are permitted: linking (the insertion of an association line) and delinking (the deletion of an association line). The former rules out the possibility of adding to an expression any melodic material which does not have a local source, while the latter allows the non-interpretation of lexically given material under certain conditions. The question we must now address is whether a mechanism which, as in Tigré, converts a non-head [A] into its headed counterpart [$\underline{A}$], can also be considered a permissible phonological operation, even though it does not involve either of the procedures just mentioned. Charette and Göksel (1994) consider the issue of the vowel harmony found in various Turkic languages, and offer an analysis which involves the switching of elements between head and dependent status. They claim that this is possible under the strict control of language-specific licensing constraints; however, it is the precise nature of these constraints, rather than their theoretical implications, which are the focus of their discussion.

Harris (1994b) is clear that head-switching cannot be successfully accommodated into his restrictive model, citing the principle of Structure Preservation (henceforth SP) as the potential barrier. The notion of SP was first invoked by Selkirk (1982) in order to restrict the output of resyllabification rules, requiring the derived structures to conform to the syllable template of the language in question. However, it has subsequently been acknowledged as a valuable generalisation which may be applied to other levels of structure too. Harris (1994a) proposes to extend the scope of SP to melody, with the effect that *lexically established dependency relations remain stable under spreading* (p.535). So, in the case of Tigré, where the lexically specified expression (A,$\underline{@}$) is interpreted as ($\underline{A}$,@) under the influence of a dominant [A]-headed vowel, it is reasonable to claim that this swapping of heads within an expression amounts to an illegal operation and, specifically, a violation of SP. A similar criticism can be levelled against the widely accepted analysis of ATR harmony within the same framework. The representational difference between an ATR expression and its non-ATR counterpart is taken to be one of headship, such that the former is headed by a 'true' resonance element while the latter is empty headed. So **e** ($\underline{I}$,A) contrasts with **ɛ** (I,A,$\underline{@}$). But in alternations between ATR and non-ATR under harmony conditions, the same head-switching operation must be invoked, since the phenomenon cannot be explained in terms of either spreading or delinking. Takahashi (in preparation) makes the same observation in relation to ATR harmony in Maasai, and he too treats this as a violation of SP.

It seems, then, that although headship is taken to play an important contrastive role in the representation of vowels, it is unable to participate in phonological processes in the way that other melodic properties do. Yet in some ways this amounts to an unsatisfactory conclusion, and it would be desirable to find some explanation as to why head/dependent relations should be necessary in static, but not dynamic descriptions. One solution may be to investigate the feasibility of melodic representation without headship, and I shall explore this possibility in the following section. In fact, further support for this move is highlighted by the example of ATR harmony just considered. The notion of headship was taken up in KLV (1985) as a property found in melodically complex expressions which served to capture the unequal contribution made by the different elements in that expression. Its function was to identify a single element, the head of the expression, as having relative prominence within the resulting complex melodic unit. In other words, headship was used to describe the way in which elements were fused together. But by exploiting headship in the distinction between ATR and non-ATR expressions, the original conception of headship as *an asymmetric relationship between two (or more) units* can no longer be maintained, since melodic units consisting of only a single element must also be accounted for. The contrast between **u** and **ʊ**, for instance, is illustrated in their respective representations as headed [$\underline{U}$] and non-headed [U]. Yet our instinctive conception of headship refers not to a property of a single element, but to an asymmetric *relation* between elements co-occurring within an expression. And in view of this established usage, it is difficult to view headship as something of an absolute value which is identifiable by examination of a single element in isolation. I propose that the notion of headship ought not to be preserved as a phonologically significant property if, on the one hand, it is employed in its original capacity of identifying an asymmetrical relation between elements and, on the other hand, it is used with reference to a single element to distinguish, for example, [$\underline{U}$] and [U] as distinct representational objects. I shall discuss an alternative way of representing the ATR distinction in §2.3.

2.3 Tier geometry for vowels

In this section I suggest some proposals which will characterize, in a restrictive manner, the vowel inventories we should wish to generate on the basis of observed contrasts. In light of the above comments, this will be attempted without recourse to any notion of headship. The discussion owes much to the insights of Rennison (1987, 1990), who focuses his attention not on the elements themselves, but rather, on the melodic tiers on which those elements reside. In fact, my starting point will be Rennison's claim that it is the geometric arrangement of these tiers, under parametric control, which both determines and restricts the systems of vocalic contrast that may be generated. Before proceeding any further, however, I should mention one particular assumption made below which departs from the standard position adopted in frameworks employing monovalent primes. Systems of phonological contrast which are privative in nature can allow access to only one term of any melodic opposition. That is, a unit can be either present in, or absent from a representation, but specific reference to its absence (as in [-ε]) is deemed impossible[4]. So, in the representation of the vowel ɔ, for example, the elements [A] and [U] are specified, while [I] is omitted altogether, since the latter makes no contribution to the vowel in question. In this way, the absence of palatality in ɔ means that a process affecting this vowel cannot refer to palatality present in the expression, thus complying fully with the attested facts.

Within a privative representation scheme, then, a melodic prime may be interpreted if it is present in the structure, while its absence makes interpretation impossible. However, in the analyses which follow it will become apparent that the most suitable way of expressing the fact that an element is to be (potentially) interpreted may not be derived simply from its presence in a structure. The inappropriateness of the presence/absence form of elemental representation is highlighted in those instances where an element is seen to 'spread' across a domain larger than that of a single timing unit (i.e. under the effects of assimilatory phenomena). This issue will be pursued further in §4, where I shall question the entire notion of elemental spreading. An alternative method of representing the same melodic information is put forward by Takahashi (in preparation), who proposes that all skeletal positions contain a full set of elements, and that those which are lexically specified as interpretable at any given point are marked as ACTIVATED, while the remainder lie dormant on their respective tiers. To illustrate, the representation of the Akan word *sika* 'money' is given here (data from Akan will be discussed in §3.2 below):

(2) `I/U tier:----[I]----[ ]-----`

[4]Arguments in support of monovalency are set out in Harris and Lindsey (in press); further discussion is to be found in den Dikken and van der Hulst (1988).

```
               |      |
            s  V   k  V
               |      |
    A tier:----[ ]----[A]-----
```

In keeping with this revised means of expressing melodic structure, the permitted operations earlier referred to as linking and delinking are replaced by the two mechanisms ACTIVATE and DEACTIVATE respectively. It is the form of melodic representation shown in (2) which I shall adopt throughout this paper, the benefits of this move becoming clearer when I examine the nature of vowel harmony systems below.

Returning to the issue of tier geometry, I wish to maintain the basic tenet of current phonological thinking which requires the burden of explanation to be largely the responsibility of the structural representation itself, rather than of any independent set of conditions or rules. Accordingly, I propose that the range of vowel systems to be found across the world's languages should be generated, to the exclusion of non-existing inventories, by the parametrically controlled arrangement of elements and their tiers. Ideally this should be achieved with as little recourse to parametric variation as possible, given the degree of arbitrariness and 'expense' that each parameter introduces into the analysis. Furthermore, we would wish to find some indication of markedness inherently present in the proposed structures. In characterizing those systems which are rarer, or less easily acquired, or in some way divergent from the symmetrical norm, additional mechanisms should be required beyond those necessary for capturing the more widely attested vowel systems. As already mentioned, I shall begin by following up the notion of tier division as presented in Rennison (1987, 1990).

One of the defining characteristics of the resonance elements [I], [U], and [A] refers to the way in which their interpretation closely mirrors many of the simplest vowel inventories to be found. These typically consist of the three 'corner' vowels {i,u,a}, and constitute the vocalic system of numerous languages such as Warlpiri and Classical Arabic. As pointed out by Harris (1994b), this simple system also defines the initial stage in first language acquisition. In these three-vowel systems the elements are unable to fuse, and we may capture this fact, in accordance with the conventions of autosegmental phonology, by assuming that all three primes reside on the same melodic tier, as shown in (3).

(3)

```
     a    i    u

     V    V    V
     |    |    |
```

```
----[A]---[I]---[U]----      (resonance tier)
```

In order to generate a system any more complex than that just defined, the resonance tier must be allowed to divide, thus permitting one or more elements to occupy their own independent tier(s). Following tier division, the primes belonging to separate tiers may fuse to create complex melodic units. Empirical evidence seems to weigh heavily in favour of the [A] element being the first to break away from the single resonance tier, leaving the remaining primes [I] and [U] to share what I shall refer to as a 'colour' tier. The resulting configuration is given in (4).

(4)

```
                    i     u     e     o     a

                    V     V     V     V     V
                    |     |     |     |     |
[I/U] tier:--[I]--[U]--[I]--[U]--[ ]---
                    |     |     |     |     |
  [A] tier:--[ ]--[ ]--[A]--[A]--[A]---
```

This constitutes the very common five-vowel system which typically contrasts two heights, front and back, with a single low vowel[5]. With some degree of phonetic variation, this inventory is found in countless languages such as Spanish, Japanese, Georgian, and Hawaiian. It should be noted that no reference need be made here to headship properties, since the tier arrangement in (4) generates all and only those contrasts which make up the symmetrical five-vowel system being described[6].

At present there appears to be little phonological motivation for choosing [A] as the element which most readily breaks away to occupy its own tier. There is evidence to suggest that [I] and [U] do share some acoustic characteristics which are absent from [A]; in fact, this distinction was encoded by Jakobson in terms of the compact~diffuse opposition. And Rennison (1987) notes this point by grouping together [I] and [U] as 'colour' elements, while [A] is distinguished for its 'sonority' properties. But although this observation highlights the nature of the element split, it remains intuitively lacking in terms of *explanatory* value. In fact, without any convincing phonological explanation as to why it should be the [A] element which resides on an

---

[5]In all of the vowel systems under discussion we should also allow for the possibility of an additional vowel (of varying quality) which may be interpreted in the absence of any lexically specified resonance elements. The interpretation of empty nuclei is examined in Charette (1991).

[6]In spite of this, the convention within Element Theory remains one of marking certain elements as heads, even though their function is redundant in this context.

independent tier in (4), then our model predicts that both [I] and [U] should be equally likely to detach in the same manner within other vowel systems. Yet the briefest survey of possible inventories will demonstrate that this is clearly not the case. Nevertheless, there is some indication that other tier configurations are indeed exploited in five-vowel systems, albeit in a very small number of languages. For example, the vowel inventory of the Altaic language Chuvash, as given in Lass (1984), comprises the set {i,ü,u,ɛ,a}[7]. In this case the elements [A] and [U] are apparently unable to fuse, given the absence of **o**, ɔ, or ɒ. Yet the system contains a front, rounded vowel, requiring [I] and [U] to reside on separate tiers. In view of these facts, the following configuration seems appropriate.

(5)

```
                 i    ü    u    ɛ    a    ɯ

                 V    V    V    V    V    V
                 |    |    |    |    |    |
[A/U] tier:-[ ]--[U]--[U]--[A]--[A]--[ ]--
                 |    |    |    |    |    |
  [I] tier:-[I]--[I]--[ ]--[I]--[ ]--[ ]--
```

The arrangement of tiers in (5) generates the system of contrasts found in Chuvash by exploiting all permissible combinations of elements. Given the possibility of this configuration, we should also predict that the [U]-tier may break away to leave [A] and [I] as tier sharers. This would generate a small (i.e. maximally six) vowel system which lacks non-high front vowels (since [I] and [A] cannot fuse) but which does contain a front rounded vowel (as [I] and [U] reside on separate tiers and are thus free to combine). Mandarin appears to qualify as such a system, with its inventory consisting of {i,u,a,ü,ɤ}. However, it should be borne in mind that both Mandarin and Chuvash represent very unusual (perhaps unique) vowel systems, and that, consequently, the tier configurations required to generate these inventories must be considered highly marked. The fact remains that statistically the most favoured five and six vowel inventories involve a separate [A]-tier, and thusfar the motivation for this appears somewhat elusive.

As already indicated, a particular feature of the melodic structure to be assumed throughout this paper supposes a full set of elements to be present under each skeletal position. And I shall further assume that these are arranged on their respective tiers

---

[7]Chuvash also contains the high, back, unrounded vowel ɯ, which I take to be the interpretation of an empty nucleus in this system (see footnote 5). Further investigation is needed in order to verify this.

according to the language-specific parameter settings which control the organisation of those tiers. In the following discussion there will emerge a structural condition which seems to require that a melodic tier must be licensed before any of the elements it contains can be interpreted[8]. I shall assume that the licensing of melodic tiers takes place in the way described by the notion of Licensing Inheritance, as introduced in Harris (1992). Here Harris proposes that a licensed unit inherits its licensing potential from its licensor, and that the stock of licensing power becomes increasingly depleted each time it confers potential on to individual units and is carried further down the licensing path away from the domain head (the ultimate licensor). Let us assume that, in order to be active in a structure, all phonological units must be licensed — a requirement set out in the Phonological Licensing Principle (Kaye 1990). In accordance with this tenet, I shall posit that a necessary condition for the activation (and thus, interpretation) of a melodic element is the licensing of the tier on which that element resides. Referring to the configuration in (4), the colour tier is directly a-licensed[9] by the prosodic structure, and this sanctions the activation of a melodic unit on that tier. Now that the colour tier is licensed (though not necessarily active), it may then act as a licensor for the [A]-tier. So, in order for the element [A] to be interpreted, there must be sufficient licensing potential bestowed upon the [I/U]-tier to license it, and furthermore, to allow it to act as a licensor for its licensee tier.

It appears, then, that the notions of element activation and tier licensing are inextricably interwoven. Indeed, they are both indispensable in the attainment of a common goal — the interpretation of a melodic prime. So do we really need to maintain a distinction between the two? Or are they merely different instantiations of the same structural principle? I shall suggest that it is only tiers, and not elements, which are licensed; this means that it is principally the tiers, rather than the elements, which are under the control of the licensing potential bestowed upon them by their licensor (even though both contribute to the depletion of licensing potential). Once a tier is licensed, then it may act as a licensor for another tier (assuming enough potential is available), and it may also become active itself (and interpret melodic material). The latter I shall assume to be largely a matter of lexical stipulation; yet the successful interpretation of an element which is lexically activated remains at the

---

[8]In fact, this is predicted by the Phonological Licensing Principle (Kaye 1990), which states that all units in a representation must be licensed, with the exception of the ultimate head of the domain. In my discussion it will become evident that melodic tiers should indeed be viewed as phonological units on a par with elements or prosodic constituents, rather than as a mere notational device.

[9]The term a(utosegmental)-licensing is employed in Harris (1994a), and refers to the sanctioning of melodic units by a skeletal position. He makes a distinction between this and another licensing relation, p(rosodic)-licensing, which exists between units in the prosodic hierarchy.

mercy of its tier and the licensed/unlicensed status of that tier. In short, for a melodic unit to be interpreted, there are both lexical conditions (i.e. activate [ε]) and phonological conditions (i.e. a licensed tier) which need to be satisfied. The failure of only one of these must result in a failure to interpret the relevant melodic material.

Let us suppose that licensing potential is consumed whenever a melodic tier becomes licensed, regardless of whether or not it contains any active elements. Furthermore, suppose that this depletion is independent of the consumption of licensing potential which results from the interpretation of an element. Up to this point, my argument for the licensing of tiers and elements is unable to compete with a far simpler account of the consumption of licensing potential which involves straightforward element counting. By following this less complex alternative, however, we encounter difficulties when we begin to consider the individual behaviour of the elements in very general terms. In particular, I refer to the apparent asymmetry observed in the static distribution of the resonance elements, as highlighted by the Theory of Charm proposed in KLV (1985). According to the stipulations of Charm Theory, the colour elements inherently possess a 'zero' charm value (represented as I° and U°) which, amongst other things, renders them relatively weak licensors. In contrast, the element [A] is assigned a 'positive' charm value (A^{+}). In what was essentially no more than a description of observed cross-linguistic patterns, KLV proposed that the head (in their terms, the *governing*) position within a branching nucleus could only be occupied by either a complex expression or a positively charmed simplex expression. Thus **ei** and **au** make possible diphthongs, whereas **iu**, for example, cannot form a branching nucleus. These are illustrated in (6).

```
(6)     N           N           N           N
        |\          |\          |\          |\
        | \         | \         | \         | \
        x  x        x  x        x  x        x  x
        |  |        |  |        |  |        |  |
        I°I°                | U°     I°U°             U°U°
        |           |                          |
        A+          A+                         A+

        [ei]        [au]        *[iu]10     *[uo]
```

In this way KLV captured the generalisation that either a low vowel, as in **au**, or a mid vowel, as in **ei**, could license (via a governing relation within the nuclear constituent) a lone I° or U° in the nuclear complement position.

Evidently, the low vowel **a** displays behaviour which has more in common with complex expressions than simplex ones. And this finds further support in the distribution of elements in the licensed position of the examples in (6), where simplex **i** and **u** are permitted but **e**, **o**, and **a** are banned. For some time it has been recognized that the arbitrary values assigned to elements by Charm Theory could offer little towards an explanation for the asymmetry observed between the three primes. In contrast, the model of tier division proposed here can provide a clearer insight into these distributional patterns by exploiting the requirement that a tier must be licensed in order to act as a licensor for the tier immediately below it on the licensing path. Referring back to the configuration in (4) we see that the mid vowels require two licensed tiers, since each tier contains melodic material which is interpreted. But crucially, the low vowel also demands that both tiers be licensed — the skeletal position a-licenses the colour tier which, although inactive, is then able to license the [A]-tier to become active. Hence, complexity may be calculated not only in terms of the number of elements interpreted, but also the number of tiers licensed: mid and low vowels require two licensed tiers, while the high vowels (in a five vowel system, at least) need only one. The advantages of this approach can be observed in the way that the amount of a-licensing potential inherited by a position is reflected in the capacity of that position to support either a simplex or complex expression. In the nuclear head slot it is complex (i.e. two licensed tiers) expressions which abound, while in nuclear

[10] There are, in fact, isolated cases of [iu] constituting a well-formed branching nucleus, as in both Welsh and Welsh English. In this, and several other examples of infrequently attested diphthongs, it seems that the well-formedness constraints in KLV (1985) need to be relaxed. Harris (1990b) confronts the problematic nature of such cases and highlights the need to refer to segmental complexity.

complement positions, which are located further down the licensing path where proportionately less licensing potential is available, only the simplex (single tier) expressions **i** and **u** are permitted.

To summarize, there are grounds for assuming that the least marked vowel inventories are those comprising the three corner vowels; the facts of language acquisition would seem to support this position. The next least marked systems involve a split within the resonance tier, most often resulting in the breaking away of the [A] element to occupy its own melodic tier. This then generates the canonical five vowel system. The mechanism which creates both the three and five vowel systems also allows for the possibility of an additional (usually central) vowel represented by a melodically empty position. According to the findings of the Stanford Phonology Archiving Project, which collected data on the segment inventories of over 200 languages and which provides the source for a typology of vowel systems presented in Crothers (1978), this mechanism generates well over half of the vocalic systems surveyed. It has also been noted that independent motivation for the tier geometry illustrated in (4) above may be found in the asymmetric behaviour exhibited by the resonance elements, such that [A] appears to pattern with mid vowels rather than with other single-element expressions in long nuclei.

In the absence of headship relations, the two-tier model of representation already illustrated can maximally generate a six-vowel system, as already illustrated in Chuvash. In order to take us beyond this number of contrasts we must initiate a second tier division which, in almost all cases, will involve a split in the [I/U]-tier. The resulting configuration then allows for each resonance element to occupy its own separate tier. The potential fusion of [I] and [U] so created predicts the presence of contrastive rounding in front vowels. Statistical data suggest that only a small minority of known languages actually exploit this possibility of employing three licensed tiers, thus making the arrangement in (7) below a relatively marked structure. Pursuing further the question of the naturalness, we also find purely linguistic evidence which similarly indicates a fairly high degree of markedness: this stems from the facts of language acquisition and refers to the comparatively late stage at which the front rounded vowels **ü** and **ö** are acquired in comparison with those generated by the two-tier configuration. A typical inventory generated by a second tier split is found in languages such as Turkish. This is illustrated in (7).

(7)

```
           i    u    e    o    ü    ö    a    ɨ

           V    V    V    V    V    V    V    V
           |    |    |    |    |    |    |    |
      I:-[I]--[ ]--[I]--[ ]--[I]--[I]--[ ]--[ ]--
           |    |    |    |    |    |    |    |
      A:-[ ]--[ ]--[A]--[A]--[ ]--[A]--[A]--[ ]--
           |    |    |    |    |    |    |    |
      U:-[ ]--[U]--[ ]--[U]--[U]--[U]--[ ]--[ ]--
```

As in the examples of two-tier systems given above, the geometric relations between the tiers in (7) should show some degree of variation according to the parametric settings of the language in question. In any case, each permutation should at least characterize a *possible* vowel system, given the undesirability of having to stipulate impossible configurations on the basis of non-occurrence. In the case of Turkish it is the effects of vowel harmony which give us an insight into the most appropriate geometric arrangement of melodic tiers. Briefly, [I] harmonizes in an uninterrupted fashion across a domain, while labial harmony is restricted by the fact that [U] cannot be harmonically activated either in or across an expression where [A] is already lexically active. As a result, suffix vowels display the following alternations.

(8)	(a)	*kɨz-lar*	daughter + pl.	'daughters'
	(b)	*diš-ler*	tooth + pl.	'teeth'
	(c)	*ev-ler*	house pl.	'houses'
	(d)	*kul-lar*	slave + pl.	'slaves'
	(e)	*kɨz-ɨm*	daughter + 1st poss.	'my daughter'
	(f)	*diš-im*	tooth + 1st poss.	'my tooth'
	(g)	*kol-um*	arm + 1st poss.	'my arm'
	(h)	*göz-üm*	eye + 1st poss.	'my eye'

In (8d) labial harmony is ruled out due to the presence of [A] in the suffix — hence, the ill-formed **kul-lor.* Let me posit that the tier arrangement for Turkish is that shown in (7). Then the blocking effect of the [A] element falls out quite naturally from the representation if we further assume that suffix vowels occupy prosodically weak positions which are unable to a-license as much melodic material as the vowels of the root. In order to interpret the low vowel of the plural suffix *-lar* it is necessary for two tiers to be autosegmentally licensed: the [I]-tier, although inactive, must be licensed so as to sanction its licensee [A]-tier. The latter also consumes an additional amount of licensing potential in the activation of the interpreted element [A]. If the vowel of *-lar* were to harmonize (to give **kul-lor*) then this would involve the licensing of a third tier and another active element, for which there appears to be insufficient licensing potential in this context.

As a very crude characterization of the consumption of a-licensing potential, I shall suggest that a similar amount is used up in the licensing of a tier as is consumed in the activation of an element. On this basis the low vowel of the suffix *-lar* consumes three 'units' of licensing potential, and this would increase to five if labial harmony were allowed to take place[11]. The latter amount seems too great to be supported by the stock of potential inherited by the suffix vowel, and therefore [U] harmony does not succeed. On the other hand, the suffix vowel of *kol-um* (lexically *kol-ɨm*) in (8g) does allow [U] to be activated, even though three licensed tiers are required. In this case only four units of a-licensing potential are consumed in the interpretation of the vowel in *-um*, which appears to be within the capacity of this position. I readily acknowledge that a more thorough investigation of the same facts would, no doubt, call for a rather more subtle method of calculating the consumption of licensing potential. In the meantime, however, this demonstration at least serves to highlight the necessity of recognizing some 'cut-off' point, beyond which a given degree of melodic complexity cannot be supported by the amount of a-licensing potential available.

This illustration of Turkish and its harmonic phenomena has at least provided an indication of the scope with which a three-tiered representation may be employed in the description of some larger vocalic systems. In addition, it has been demonstrated that some explanatory power may lie in the notion of tier licensing as a pre-requisite for element activation, to the extent that the former, as well as the latter, contributes to the depletion of licensing potential inherited by a skeletal position. Clearly, the possibilities created by this three-tier geometry require extensive investigation, and I believe that this will prove to be a fruitful research area both in terms of the generation of some less common systems and also the characterisation of harmonic phenomena. However, I regret that limitations of space force me to postpone this undertaking until a later time when the appropriate attention can be given to this line of enquiry.

Thusfar my discussion has focused on tier division, and has illustrated some of the most commonly found inventories consisting of three, five, six, and eight vowels. Yet this is only one of the parameters along which vowel systems may vary. And

[11] These 'units' are not to be understood as absolute values — I am convinced that reference to any such precise measurements must prove detrimental to the restrictiveness of the model. Instead, what this example attempts to show is the *more than* or *less than* relation which is at issue here. The formulation of such a relation has been achieved without recourse to any 'counting' methods within Optimality Theory (see discussion of *Harmonic Ordering* in Prince and Smolensky (1993), §3).

evidently we must find another direction in which to formalise the kinds of systemic diversity not yet encountered. This will be taken up in the following section, where I shall build on the notion of tier complement proposed by Takahashi (in preparation) in his discussion of the representation of ATR distinctions.

2.4 Tier complements

In §2.2 I motivated a representation of melodic structure which makes no reference to any asymmetric relation of headship between elements within a single expression. However, immediately this proves problematic for the standard assumption within Element Theory that the contrast between ATR and non-ATR is captured solely via headship distinctions[12]. It is now widely held that, for example, ATR **u** and non-ATR **ʊ** are distinguished phonologically by the headedness of the former and the non-headedness of the latter: hence, **u** [$\underline{\text{U}}$] contrasts with **ʊ** [U]. Yet it is precisely this deployment of headship within the scope of a single melodic tier which proves difficult to reconcile with the original role of headship as an *asymmetric* association existing between *different* elements in a complex expression. Further problems with the established representation of ATR are encountered when we turn to harmonic phenomena involving this property. Common to many West African systems, for instance, is the division of the vowel inventory into two sets — one ATR and the other non-ATR — where one category occurs to the exclusion of the other in a given domain. Within an ATR span, those vowels which are lexically headless (i.e. non-ATR) are compelled to align themselves and are accordingly interpreted as headed expressions. This operation is most usually achieved by switching head status from a 'dummy' dependent element [$\underline{@}$,ɛ] to a resonance element [$\underline{ɛ}$] which is lexically present (see §2.2 for a discussion of Tigré). It has already been noted, however, that this move ought to be questioned on two counts: first, it may constitute a violation of the principle of Structure Preservation; and second, it forces us to expand the repertoire of permitted operations within autosegmental phonology beyond those involving activation and de-activation (or, more conventionally, spreading and delinking).

In view of these comments I shall introduce the notion of tier complement, put forward by Takahashi (in preparation) as an alternative means of representing the property ATR, and attempt to integrate this into the tier-geometric model established above. If integration is to be achieved successfully, then a number of criteria should ideally be fulfilled. For instance, the lexically given structure must remain intact, in

[12]See Charette and Kaye (1993) for arguments in support of this position.

accordance with SP, and the mechanism by which ATR harmony is characterised must comply with the activate/de-activate operations already mentioned. Also, the model should capture some of the phonological differences observed between ATR and non-ATR vowel pairs. These differences manifest themselves as distinct distributional patterns in languages such as English and German. In English long vowels, the head position of the branching nucleus cannot be occupied by a high non-ATR vowel; hence, the long lax vowels **ɪ:** and **ʊ:** are not observed. In contrast, the long low vowel **a:**, which is also non-ATR, does create a well-formed branching nucleus. In other words, a lax [I] or [U] cannot license a nuclear complement position[13], whereas [A] acts as a sufficiently strong licensor to be able to sanction a licensee. Interestingly, a similar set of facts was addressed above in the discussion of the asymmetrical behaviour observed between the three resonance elements. And in §2.2 I concluded that the patterning of [A] with complex (rather than simplex) expressions could be accounted for by assuming that the low vowel was also a complex expression, by virtue of the *two* tiers which needed to be licensed in order for [A] to be interpreted. Now let me apply the same analysis to the ATR distinction. The low vowel **a** and the high ATR vowels **i** and **u** are seen to behave in a uniform way, in that they are all sufficiently strong licensors to assume the head position of a branching nucleus; in contrast, the non-ATR high vowels are banned from this position. I shall propose that the former group are all complex expressions, while lax **ɪ** and **ʊ** are simplex. As for the three complex expressions, it is clear that only one element is active in each case. Therefore I posit that their complexity is derived from the licensing of more than a single tier in their representation.

The necessity of at least two licensed tiers in the structure of an ATR vowel reflects Takahashi's proposal that systems which employ an ATR contrast should have at their disposal a mechanism whereby a complement may be activated as a dependent of a colour element. He claims that the role of a complement is to enhance the characteristic acoustic patterns of an element — thus, in ATR expressions the relevant complement is active, while non-ATR is marked by a dormant complement. This notion may be translated into the formalism of the current model by portraying the apparent complexity of ATR expressions using a tier complement, as in the configuration in (9). The structure I propose does depart slightly from the model presented by Takahashi, in that here the complement merely constitutes an enhancement of the relevant tier, rather than a new tier in its own right. This will be more fully described below.

---

[13]The sole exception to this generalisation is the set of 'centering' diphthongs comprising [ɪə], [ɛə], [ʊə], and [ɔə]. I suggest that the uniform patterning shown by all of these diphthongs perhaps reveals more about the nature of [ə] than the nature of lax vowels.

(9)

```
                complement
                  /
------[I/U]-tier-------
                |
-------[A]-tier--------
```

In systems which employ an ATR contrast the colour tier complement will be latently present under each prosodic position, as is the case with all other units in melodic representation, and activated either via lexical stipulation or as a result of harmonic influence. In the seven vowel system of Italian, for instance, an ATR distinction is found in mid vowels, where **e** and **o** contrast with **ɛ** and **ɔ** respectively. This may be represented as in (10).

(10)

```
  comp:      [I]      [ ]       [U]       [ ]
             /        /         /         /
[I/U]:  --[I]--  --[I]--  --[U]--  --[U]--
            |        |        |        |
  [A]:  --[A]--  --[A]--  --[A]--  --[A]--

            e        ɛ        o        ɔ
```

I propose that both the colour complement and the [A]-tier are directly licensed by the [I/U]-tier. Although difficult to depict diagrammatically, the relationship between the colour tier and its complement should be viewed as one which introduces a third dimension to the structure. Accordingly, an active complement has the function of enhancing the saliency of a colour element by affording it 'depth', and not by inserting an additional plane into the melodic representation. It should be noted, however, that the relationship between the colour tier and its complement is not identical to that existing between the colour tier and the [A]-tier. In the former association there is no new elemental material added to the structure when the complement is activated; instead, the same plane is merely expanded in another direction. This being the case, we expect a close affiliation between the two units[14], which is manifest in the requirement that the colour tier be active (not only licensed) in order to license an active complement. In contrast, the relationship between the colour tier and the [A]-tier is predicted to be somewhat looser, given that they contain different elements

[14]I shall henceforth refer to a tier complement as if it were a separate tier; my motivation for this stems from the fact that a complement, just like the kind of licensed tier already seen, needs to be licensed and consumes licensing potential in the process. Strictly, complements do not reside on their own tiers, but expand a 'true' melodic tier already present in the structure.

which display a certain amount of autonomy and independence from each other. This more distant association is reflected in the claim, already stated above, that the colour tier need only be licensed, but not necessarily active, in order to act as a licensor to the [A]-tier.

Returning now to the apparent complexity of ATR high vowels, this can be represented straightforwardly as the activation of the colour tier and its complement. In this way, the interpretation of a single element can be treated as a complex expression, to match its phonological behaviour. As outlined in §2.2, this is applicable to low vowels as well as to ATR ones. In (11) the high ATR vowels require two active tiers (see footnote 14) in order to be interpreted, while the low vowel needs two licensed tiers (this is independent of whether or not a colour tier complement is required in the system). (11) illustrates the maximum number of contrasts generated by a system which licenses both a colour tier complement and the colour/aperture division given in (4).

(11)

```
  [I]  [ ]    [U]  [ ]    [I]  [ ]    [U]  [ ]    [ ]  [ ]
  /    /      /    /      /    /      /    /      /    /
[I]-[I]     [U]-[U]     [I]-[I]     [U]-[U]     [ ]-[ ]
 |   |       |   |       |   |       |   |       |   |
[ ]-[ ]     [ ]-[ ]     [A]-[A]     [A]-[A]     [A]-[ ]

 i   ɪ       u   ʊ       e   ɛ       o   ɔ       a   ɨ
```

It seems relatively rare for a vowel system which licenses a colour tier complement to exploit the full contrastive range in (11), although the inventory of the Ugandan language Lugbara does confirm the existence of such a system. More usual is a language which selects an ATR contrast either in high vowels or in mid vowels. Two West African systems provide an example of each in (12).

(12)	Kpelle:	i		u	Ewe:	i		u
		ɪ		ʊ		e		o
		ɛ		ɔ		ɛ		ɔ
			a				a	

In Kpelle the colour tier licenses an active complement only when the [A]-tier is not licensed — that is, only in high vowels. It may be suggested that this system owes its overall shape to a condition which maximally allows two tiers to be a-licensed within a single expression. In other words, a nuclear position can supply sufficient licensing potential to license no more than two tiers. This permits all of the attested vowels and simultaneously rules out ATR **e** and **o**. Of course, such a condition contributes little to an *explanation* of why the Kpelle system has adopted the shape it has, but at least the relatively unmarked status of this system can be captured by the fact that we need refer to only a single generalisation in its characterization. In a system such as that of Ewe, on the other hand, the only ATR distinction is found in mid vowels — in fact, this constitutes the more common of the two types. To exclude the high lax vowels from the inventory we could surmise that only complex vowels are in some sense 'stable' enough to be interpreted, thus ruling out ɪ and ʊ, which require only a single tier to be licensed. Naturally, the same caveat regarding the lack of explanatory value is equally applicable here.

A question that arises as a consequence of the above discussion is whether it is only a colour tier which can potentially license a complement, or whether the [A]-tier may also be able to do so. I shall make the assumption that a significant amount of licensing potential is needed for a melodic tier to a-license its own complement. This predicts that an aperture tier is less likely to license a complement than is a colour tier, since the former is located further down the licensing path than the latter and is thus expected to have less a-licensing potential at its disposal. At the same time, there appears to be no motivation for entirely ruling out the possibility of an [A]-tier complement. Indeed, a structure which exploits this configuration is required in those (relatively few) systems that employ a three-way contrast among central vowels, where **a**, ə and ɨ are all lexically contrastive. The Austronesian language Cham (spoken in Vietnam) provides an illustration of this tier arrangement in which there is a colour/aperture split, but the [A]-tier, rather than the [I/U]-tier, licenses a complement.

```
(13)    i     e     ɛ     a     ɨ     ə     ɔ     o     u

      [I]--[I]--[I]--[ ]--[ ]--[ ]--[U]--[U]--[U]
       |    |    |    |    |    |    |    |    |
      [ ]--[A]--[A]--[A]--[ ]--[A]--[A]--[A]--[ ]
        \    \    \    \    \    \    \    \    \
        [ ]  [ ]  [A]  [A]  [ ]  [ ]  [A]  [ ]  [ ]
```

Such a system is not widespread, given that an impoverished [A]-tier is predicted not to have sufficient a-licensing potential at hand to license a complement easily. On this basis, I suggest that a language which licenses both a colour tier complement and an [A]-tier complement should be considered very rare indeed, since the licensing of four tiers simultaneously must constitute a burden too heavy to be borne by most, if not all, vowel systems. In fact, I have yet to discover a system that requires such a tier configuration.

At the outset I suggested that the evaluation of a vocalic representation system should be partially based on the degree of success with which it could depict attested vowel harmony phenomena. §3 will take up this topic and offer an analysis of harmony in Akan to demonstrate the appropriateness of tier complements with respect to harmonic agreement. My argument will follow Takahashi (in preparation) in characterising ATR harmony as the ACTIVATION of a complement throughout a harmonic domain; this analysis will show certain advantages over a more widely established account which employs a headedness distinction to capture the same ATR contrast. Briefly, by invoking the notion of tier complement we avoid the problematic head-switching operations inherent in the process of head alignment, by simply activating or de-activating a tier which is already lexically present in the structure. The postulation of an enhanced (i.e. three-dimensional) tier also proves to be more generally appropriate than headship, in that it is equally suited to expressions containing a single element as it is to melodically complex units. This is, of course, not the case with head/dependent relations, which are founded on an asymmetry between different primes interpreted simultaneously. Finally, a lack of equivalence between the concepts of headedness and tier complements can be seen in the characterisation of, for example, the canonical five-vowel systems such as Spanish, discussed in §2.2 above. Here, the redundant status of licensed complements is reflected in their absence from the tier geometry of these languages — indeed, the simplicity of the melodic structure of these systems mirrors their relatively unmarked status cross-linguistically. On the other hand, if the role of headship is maintained as a valid one, then every element must be specified as either a head or a dependent, even in those systems where this information is phonologically redundant.

I have now presented the two structural characteristics which, I claim, are sufficient for capturing the variation which is observed across vowel systems: in §2.2 I discussed some of the possibilities generated by the geometric arrangement of melodic tiers, and in this section the notion of tier complement was introduced. The latter will now be implemented in the analysis of Akan ATR harmony.

3 Akan vowel harmony

3.1 Introduction

Formerly designated *Twi*, the Akan language of Ghana is typical of many West African systems in that it displays a form of vowel harmony which recognizes two corresponding sets of vowels that may be distinguished by the presence or absence of ATR. Several comprehensive analyses of the facts relating to Akan harmony can already be found in the literature, and it is not my intention here to attempt a restatement of these facts with any comparable degree of thoroughness[15]. Instead, my main concern will be to illustrate the appropriateness of the tier geometry established in §2.2 above by examining the way in which it accommodates the harmonic patterns observed in the examples to be introduced. Below I will outline the conditions under which harmony takes place, and refer briefly to two existing analyses of the data, those of Stewart (1967) and Clements (1981). My re-analysis will be based on a geometrical arrangement of vocalic primes which requires the postulation of a colour tier complement. Unless stated otherwise, the examples discussed below belong to the Asante dialect of Akan, the system which has generally received attention in the literature.

3.2 Vocalic patterns in Akan

The vocalic system of Akan consists of ten surface vowels, which may be arranged into two corresponding sets, as in (14), according to their appearance in, or exclusion from ATR harmony spans. That is, the vowels belong to a basic five vowel system, which has an additional ATR contrast.

(14)	*set 1:*	i	u	*set 2:*	ɪ	ʊ

[15]For detailed discussions on the matter, the reader is referred to Stewart (1967, 1983) and Clements (1981, 1984); it is from these works that the data below has largely been taken.

e	o	ɛ	ɔ
ɜ		a	

Of the ten phonetically distinct vowels, nine are potentially contrastive[16], while the vowel transcribed as **ɜ** only ever behaves as a surface variant of **a** in predictable contexts — there is never a lexical contrast between low vowels. I shall follow Stewart (1967) in assuming ATRness to be the property of a morpheme, lexically marked as such. This position then allows us to posit only five underlying vowels, each of which has two reflexes — ATR and non-ATR — depending on the harmonic domain.

In general, the harmony system of Akan dictates that all sequences of non-low vowels within a verb root or noun morpheme agree in ATRness; that is, all are selected from only one of the harmonic categories shown in (14). In addition, the ATR value of any verbal suffix is determined by the rightmost vowel of the root, while prefixes adopt the ATR quality of the leftmost vowel. Harmony may thus be said to be bi-directional, and in many cases results in the concatenated domain displaying complete agreement in terms of ATR.

(15)	*group 1*		*group 2*	
	o-kusi-e	'rat'	ɔ-kɔdɪ-ɛ	'eagle'
	o-fiti-i	'he pierced it'	ɔ-cɪrɛ-ɪ	'he showed it'
	o-be-je-i	'he came & removed it'	ɔ-bɛ-jɛ-ɪ	'he came and did it'
	mi-tie	'I listen'	hoɟo	'to loosen'
	wu-be-numʔ	'you will suck it'	wʊ-bɛ-nʊmʔ	'you will drink it'

The items of the first group in (15) contain vowels drawn exclusively from the ATR vowel set; since the verbal or nominal roots are lexically assigned to the harmonic category labelled *set one* in (14) above, it follows that all affix vowels adopt the set one reflex of their underlying vowels in these cases. On the other hand, the items in group two have no lexical ATR specification in their root morphemes, and this is reflected in the selection of affix vowels from set two.

The presence of low vowels in a representation, however, has the effect of interrupting harmony spans, with the result that an ATR and a non-ATR sequence may co-exist within a single concatenated domain, eg. *oobegwyar*ɪʔ 'he is going to

---

[16]Although nasalization adds another contrastive dimension to the vowel system of Akan, nasal vowels behave no differently from oral vowels with respect to harmony, and will therefore be disregarded in this study. Similarly, tonal markings will be omitted from the examples given.

wash'. As given in (14), there are two reflexes of the low vowel in Akan. But unlike other corresponding pairs of vowels (one taken from each set), these are always in complementary distribution. Whereas the shape of a non-low root vowel is determined lexically, both from its own melodic representation and any ATR marking on the morpheme, the choice between **a** and **ɜ** may be predicted from the environment according to the conditions set out below.

(16) **ɜ** occurs in the following contexts to the exclusion of **a**:

(i) preceding a high ATR vowel (verbs and nouns) or a mid ATR vowel (nouns only):

yɜfunu	'belly'
kaŋkɜbi	'millipede'
pɜtiri	'to slip'
kɜri	'to weigh'
wɜ-tu	'he has dug it'
mɪ-kɔ-kɜri	'I go and weigh it'
pɪrɜko	'pig'
wɜ-bisa	'he has asked'

(ii) preceding a root beginning with a palatal/labialized consonant, followed by **a**:

ɜ-ɟʷa	'assembly'
e-cʷɜɜ-cʷa	'scars'
mɜ-cʷa	'I have cut (it)'
wɜ-ɲ anɪ	'he has awakened'

(17) **a** appears to the exclusion of **ɜ** in all other contexts:

sika	'money'
kosua	'egg'
a-mɪna	'hole'
e-cʷa	'scar'
kasa	'to speak'
a-nɜɲcʷi-e	'cows'
baɲcɪ	'cassava'

yarɪ	'to be ill'
wa-tʊ	'he has thrown (it)'
wa-be-tu	'he has come and dug (it)'
a-tene	'it (news) has spread'

This outline of the vocalic patterning to be found in Akan will prove quite adequate for the purposes of the present illustration. In fact, as regards the Asante dialect, there are merely two disharmonic roots thusfar discovered which fail to conform to the above generalisations: *piɲc*ɛ 'to come close', and ɲ*ins*ɛɲ 'to be pregnant'. A number of attempts have been made to analyse these as morphologically complex words, which would then render them well-formed.

Stewart (1967) correctly notes that occurrences of the low vowel serve to mark the boundaries of an ATR domain: in other words, low vowels act as harmonic blockers. Furthermore, he adds the provision that the low vowel on the left of a domain forms part of the span (thus interpreted ɜ) while that to the right does not (and hence, is realized **a**). This point is illustrated by the form *w*ɜ*-bisa* 'he has asked'. The same facts may also be captured by the observation that, whenever the low vowel is the final vowel of the noun or verb stem, it is always interpreted as **a**, even within an ATR morpheme. Elsewhere the low vowel is interpreted as ɜ whenever predicted to do so within the contexts outlined in (16i) and (16ii). It may be noted that both low vowel reflexes act as opaque segments, halting the progression of harmony; this is demonstrated in *wa-k*ɜ*ri* (**w*ɜ*-k*ɜ*ri)* 'he has weighed it', where leftward spreading to the prefix is blocked, and in *o-bisa-*ɪ (**o-bisa-i*) 'he asked', where the low vowel prevents rightward spreading to the suffix.

Stewart (1967) offers what is essentially a phonemic analysis of Akan harmony, but incorporates an autosegmentalized ATR feature which may be lexically present in roots. He proposes a set of rules to determine the exact position of this prosodic feature within the root, while further rules are invoked to mark out the extent of the harmonic span. Although Stewart's account does provide for all of the necessary facts, it also serves to highlight the nature of the contemporary literature with regard to the understanding of non-segmental phenomena. In particular, his rules serve no more than a descriptive function. Their arbitrary and unconstrained formalism is representative of the theoretical environment in which they were conceived; and they are to be viewed largely as language-specific, showing little relevance to the characterisation of harmony systems on a cross-linguistic basis.

In contrast, Clements (1981) proposes a representation for Akan which dispenses with the linearity of Stewart's phonemic string in favour of a more fully autosegmentalized arrangement of features along the lines of Goldsmith (1976). And clearly, Clements' response does reflect the extent to which an awareness of non-segmental phenomena had developed since the time of Stewart's earlier analysis. For instance, he employs language-independent Association Conventions to map ATR values on to vowels, under the assumption that such mechanisms are already motivated elsewhere in the phonology (eg. in the assignment of tone). Yet Clements' account is not without its own shortcomings. In order to describe the numerous examples of harmonic spread being interrupted, he invokes the unqualified, *ad hoc* statement that low vowels are opaque to harmony. But given the significance of this fact as a defining characteristic of the Akan harmony system, not to mention the frequency with which this same condition applies in other languages with ATR harmony, it would seem desirable to derive opacity from the representation itself. Another problem that emerges from Clements' study — his use of a bivalent ATR feature — must be seen simply as a by-product of the feature representation system which he employs. Finally, when translated into a restrictive system of melodic representation which recognizes independently interpretable elements, the inappropriateness of an object that functions only to contribute ATRness quickly becomes apparent[17].

3.3 A geometrical analysis

In terms of the model developed in §2.3, a contrast in ATR between two vowels that are otherwise phonologically identical may be captured by referring to the licensing of a colour tier complement. That is, in a system where ATR is a contrastive property of vowels, the I/U tier will license a complement, and that complement will be active in those expressions which are either lexically designated as ATR or which assimilate to an ATR quality when they fall within the appropriate harmonic domain. Thus, the vowel harmony process in Akan may be straightforwardly referred to as *the activation of an [I/U]-tier complement*[18].

[17]For arguments against the postulation of an independent ATR element, see Harris and Lindsey (in press).

[18]This tier configuration for the representation of ATR was first proposed by Takahashi in his analysis of Maasai.

(18) (a) o - be - je - i (b) ɔ - bɛ - jɛ - ɪ

```
   [U]    [I]    [I]    [I]        [ ]    [ ]     [ ]    [ ]
   /      /      /      /          /      /       /      /
 -[U]--[I]--[I]--[I]-          -[U]--[I]--[I]--[I]-
   |     |     |     |           |     |     |     |
 -[A]--[A]--[A]--[ ]-          -[A]--[A]--[A]--[ ]-

   V - bV - jV -  V              V - bV - jV -  V
```

'he came and removed it' 'he came and did it'

The examples in (18) illustrate a minimal distinction which refers only to ATR for lexical contrast. That is, their melodic representations are identical, save the fact that the verb root in (18a) contains an active complement; this complement is then licensed to become active throughout the concatenated domain, yielding ATR vowels in the affixes too.

At this point it may be noted that, within the present model, the mechanism by which harmony proceeds has made something of a departure from the more traditional notion of spreading that has been widely adopted in other autosegmental and feature geometric frameworks. In (18a) there is clearly some form of agreement between the vowels within the domain. But this agreement cannot easily be captured by referring to any kind of spreading, since this has come to imply the presence of an independent ATR element or feature (rather than, as here, a particular configuration of tiers) which has been allowed to associate to the skeletal tier in a one-to-many fashion. In cases where we wish to characterize harmony in terms of the activation of a tier, the notion of *alignment* appears to fit the descriptive mould rather more comfortably. In this way the unified behaviour of a particular property (or properties, since here we are concerned with enhanced [I] and enhanced [U]) throughout a domain may be achieved by bringing into line the material which acts as a target for harmony, without recourse to any association operations.

The straightforward harmonic divisions represented by the examples in (15) may thus be formally captured as in (18), where an ATR domain is marked by an active tier complement. Now let us turn to the task of incorporating the many instances of mixed morphemes (i.e. those containing a mixture of low and non-low vowels) into the analysis, without losing sight of the fact that low vowel opacity is relatively common and apparently not coincidental within systems featuring ATR harmony. In fact, given the geometric configuration required for the contrastive system of Akan, shown in (9) and repeated here,

(19)

```
                complement
                    /
-------[I/U]-tier----------
              |
--------[A]-tier-----------
```

the opacity of low vowels can be shown to fall out directly from the representation.

We have established that the domain of ATR harmony is marked by an active complement. And in §2.3 I proposed that, while an active licensee tier could be licensed merely by the *presence* of its dominant tier (such is the relationship between the [A] tier and the colour tier in a basic five vowel system), a complement, on the other hand, required an *active* dominant tier in order to be active itself[19]. But in the case of low vowels it is only the [A] tier which is active; the inactive status of the colour tier automatically excludes the possibility of a complement being active, and hence, excludes these expressions from participating in ATR harmony. What is more, if the extent of a harmony span is dictated by the extent of an active complement, then it is quite predictable that the presence of a low vowel, which is unable to license an active complement, should halt the progress of harmony. This is illustrated in (20):

(20) o-bisa-ɪ (*o-bisa-i)

```
   [U]   [I]   [ ]   [ ]
   /     /     /     /
--[U]--[I]--[ ]--[I]--
   |     |     |     |
--[A]--[ ]--[A]--[ ]--

   V - bV  s V -  V
```

'he asked'

In (20) the root *bisa* is clearly an ATR morpheme: the high vowel belongs to the set one category, while the low vowel retains its **a** quality by virtue of its root-final position. Although the vowel of the prefix aligns successfully to the leftmost root vowel, the suffix is prevented from aligning (i.e. activating its complement), owing

[19]This condition need not be interpreted as a stipulation. Since the complement does not introduce any new melodic material, but simply enhances what is already present on the colour tier, it is to be expected that a stronger relationship should hold between these two tiers than exists between tiers housing different elements that show greater mutual independence. See §2.3.

to the absence of an active complement in the low vowel, the latter interrupting the harmonic span. The example in (21) illustrates another ATR root morpheme.

(21) wa-kɜri (*wɜ-kɜri)

```
      [ ]     [ ]      [I]
      /       /       /
---[ ]---[ ]---[I]---
    |       |       |
---[A]---[ ]---[ ]---

 w V - k V - r V
```

'he has weighed it'

In the case of (21), however, the low root vowel is not in final position and may therefore adopt its centralized reflex **ɜ**. As predicted, the same low vowel brings to an end the ATR harmony span, with the result that the low vowel of the prefix cannot belong to an ATR domain and is thus realized as **a**.

The example in (21) also serves to highlight the final issue to be addressed here. The low vowel has two phonetic variants, **a** and **ɜ**, the choice depending on whether or not it occupies a position within an ATR span. So how is each to be represented phonologically? In the case of **a**, which is taken from the non-ATR vowel set, we may simply posit a sole [A] element. As for the other variant, the object to be described appears to be an ATR version of the low vowel — something which the geometric arrangement in (19) is unable to generate. In fact, as mentioned above, there are several instances in the phonological literature of representational systems for vowels which similarly disregard the theoretical significance of such objects. To further weaken the status of **ɜ** as a low ATR vowel, it seems that the realizational possibilities across different dialects for this vowel transcribed as **ɜ** range from [æ] to [ɛ] and [e], not all of which are even phonetically ATR articulations. As an alternative analysis I propose that the object indicated by **ɜ** is nothing more than the interpretation of an empty nuclear position, thus emphasizing the inappropriateness of ATRness being associated to a low vowel. And the fact that it receives a range of different, yet phonologically insignificant phonetic interpretations according to dialect does no more than reinforce the claim that the expression is melodically unspecified.

In order to arrive at this conclusion we must assume that Akan interprets the conditions for vowel harmony in the strictest of ways. It has been shown that all vowels within an ATR domain must license an active complement. We must then assume that it is the low vowel's inability to align itself in this manner that causes its failure to be interpreted. That is, strict adherence to harmony conditions seems to dictate that all melodic expressions in an ATR span should involve an active complement; so in an expression where the latter is ruled out (as in (21)), its presence in the structure cannot be sanctioned. And if the lexically specified [A] cannot be interpreted, then it is only the remaining melodically empty nucleus which is audible. As for those instances of the low vowel which are realized as **a** even within an ATR span, we can make the assumption that there is some prosodic motivation for their resistance to decomposition, since their position is invariably the same - root-final. We may conjecture that Akan does not permit domain-final nuclei to be empty[20], with the result that, in this position, a sole melodic element must resist any pressure to deactivate. And the data above do suggest that this parametric requirement is indeed active in the phonology of Akan, although it would be necessary to note its extended application beyond phonetic interpretation to the phonological structure itself.

In spite of having offered only an outline of the way in which vowel harmony operates in Akan, I hope to have demonstrated, at least with respect to ATRness, the appropriateness of the vocalic geometry proposed above. It has been shown how the rather special status of low vowels in ATR harmony systems need not be treated as an entirely stipulative matter. And an alternative analysis was offered for the representation of the 'tenth' vowel which attempts to avoid the unsatisfactory juxtaposition of ATRness and openness within the same vocalic expression. As a final note I shall now attempt to evaluate the notion of alignment, which has been assumed throughout this section as the most appropriate medium by which the effects of harmonic phenomena are achieved.

4 Spreading versus alignment

Within the recent literature that focuses on the representation of melodic processes, perhaps the most effective contributions have been those which place the heaviest restrictions on both the size of the inventory of melodic primes and the number of ways in which phonological processes are permitted to operate. In the version of Element Theory assumed throughout this paper, the set of melodic atoms primarily used in the representation of nuclear expressions has been limited to three — the elements [A], [I], and [U] — a set which is now acknowledged as the most basic reference point in the description of vowel sounds within many different approaches to phonology. The grounds for its widespread acceptance are due in part to the restrictions it places on the generation of vocalic expressions, such that the

[20]Thanks are due to John Harris for making this suggestion to me.

undesirable over-generation of unattested expressions is largely avoided.[21] With respect to the variety of ways in which processes are able to take place, a survey of the current literature reveals that the most highly constrained models are those that propose only two mechanisms for the description of all possible phenomena: the loss of melodic material from an expression (i.e. decomposition), and conversely, the gain of material from a specified source elsewhere in the representation (i.e. composition). This is the position taken in, for example, Harris (1994a), and is the one adopted here. It is the latter operation, that of melodic composition, which is central to the current issue, since it is the operation of spreading that is most often invoked in the description of element or feature association to positions where these units are not lexically specified.

Throughout this discussion I have assumed that all melodic primes (while respecting language-specific tier configurations) are latently present at every position on the timing tier, and that in the event of an element being lexically activated, it can (potentially) be interpreted. Immediately, this appears problematic for the notion of spreading, which implies that an element [ε] is absent from a position x_2 prior to the spreading of [ε] to x_2 from another position x_1. Consider a typical assimilatory process such as palatalisation. One of the countless languages where this is found is Japanese, as the following example shows. In the noun *hito* [çito] 'person' the palatal properties of the first vowel are also to be observed in the preceding consonant, even though the latter is not lexically specified for palatality. This is customarily portrayed via a spreading process, as illustrated in (22).

(22) regressive spread of [I]

```
 x    x  t  o
 |    |
[h]   |
 |    |
  ⇦⇦[I]
```

hito [çito] 'person'

But since we are assuming that [I] is already latently present (and simply inactive) in word-initial position, any spreading process which serves to extend the association of the lexical [I] to another expression appears somewhat incongruous. A more

[21]In fact, most phonologists working within a tricorn system of vowel representation would claim that the primes alone are inadequate for capturing all of the contrasts to be observed. As a result, additional devices such as headship, complements, etc. have been found necessary in order to boost the capacity of these units to the required output level. See §2.2 for relevant comments.

precise description of the facts would need to make reference to the active status of [I] in the nuclear position, which in turn activates the latent [I] lying dormant in the word-initial onset. This alternative approach is exemplified in (23):

(23) (a) lexically given (b) alignment of active [I]-tier

```
   x    x   t   o           x    x   t   o
   |    |                   |    |
-[h]-[  ]------          -[h]-[  ]------
   |    |                   |    |
-[  ]-[I]------          -[I]⇦[I]------
```

Here, we may observe the active status of [I] in (23a) being inherited by the preceding position, with the result that a palatal fricative is interpreted word-initially. As already noted, we cannot employ the term 'spreading' with any real justification, since [I] is already present in all positions, making any new association between the element and the position unnecessary. And equally incongruous is the idea that an element should spread along its own tier, given that an element [ε] is taken to occupy an [ε]-tier to the exclusion of all other elements[22].

A rather different approach, in the form of alignment, has been adopted here with respect to the domain within which melody units are associated to timing units. (23b) demonstrates that the interpretation of the element [I] is dependent on whether or not the *tier* on which [I] resides is active at a particular point. So now we must address the question of how a tier is activated or de-activated. Let us suppose that activation is controlled largely via lexical stipulation, so it is the lexicon which determines that [h] is active at the beginning of the word *hito*. (This is essentially no different from claiming that the first position of the word *hito* has a melodic representation consisting of [h]). In the same way, the lexicon must be responsible for de-activating the [h]-tier before its intersection with the following timing unit. As for [I] in (23), on the other hand, activation enjoys a wider scope and this fact must somehow be encoded in the grammar. In sum, by choosing to follow this tack we are faced with the burden of stipulating the activation and de-activation points (in effect, the span of the active tier) in every case of melodic activity.

Clearly, this is less than satisfactory when we consider that, in the vast majority of cases, an element is active only within the scope of a single timing unit. Only in instances of assimilation does this pattern vary, and in the case of assimilatory

[22]In the event of tier sharing, each element residing on the same tier can be assumed to occupy its own 'span' along the tier, which is broken by the lexical presence of another resident element.

phenomena, the domain of activation is often uniform within one system. It seems desirable, therefore, to develop some means of avoiding the arbitrary lexical placement of 'de-activation' in favour of a mechanism which captures the fact that a domain of activation always follows one of only a very limited set of forms - namely, activation within a single timing slot, within the scope of a syllabic constituent, within a morpheme, or within a (morphologically complex) word. Stated in these terms, we can immediately identify a link between the span of an active tier and various prosodic constituents. In the majority of cases an element is activated only within one timing position (e.g. the [h] in *hito* above), after which the relevant tier is de-activated and the prime is no longer interpreted. I shall propose that, in these instances, activation is determined and controlled at the lowest prosodic level, that is, the level referred to as p^0. In this way, a melodic unit that is activated by a nucleus, for example, can only be interpreted within the string which is a-licensed by that constituent.

The same argument can be extended to those cases where an element is active across a span that is larger than that of a single segment - here I refer primarily to instances of assimilation. Returning to the example of palatalisation in (23), the domain within which this process operates in Japanese can be identified as the syllable, since it is observed only between a nucleus and the onset to its left. This may be expressed by lexically activating [I] at the syllable level in those words where this phenomenon occurs, yielding an active [I]-tier throughout the span dominated by the syllable node (i.e. the onset/nucleus pair). Finally, the same treatment of tier activation may also be deployed in the description of the kinds of harmony processes which I have focused upon above, in which an active tier becomes the property of a morpheme or a whole word, given the required environment. Again, the domain of activation can be matched to a particular prosodic constituent - this time the prosodic word - which, I propose, bears the lexical information necessary for the activation of a melodic prime which participates in harmonic phenomena. So in the case of Akan ATR harmony (see §3 above), where a root morpheme is lexically labelled as either ATR or non-ATR, we may represent this by saying that the lexical marking *activate tier complement* comes into effect at the word level, where all participating nuclei are adjacent. In contrast, whenever the remaining melodic tiers in Akan are activated, this is carried out by referring to the lowest level of the prosodic hierarchy, since they are active only within a single timing slot.

To summarize, if we pursue a principle of activation whereby all melodic units are (latently) present at every timing slot, the notion of spreading appears quite inappropriate as a means of expressing many-to-one associations. As an alternative, melodic tiers must be activated in order for their respective elements to be interpreted. A tier is activated lexically at a particular prosodic level, the latter determining the location and the span of the active tier within the melodic string. As a result, the term 'de-activation' carries little phonological significance, as this merely amounts to the status of a tier which lies outside the scope of a lexically activated unit. Thus there is no motivation for an independent property 'de-activate' to counteract the effects of a lexical marking *activate [ε]-tier*: de-activation occurs as a natural consequence of the way activation is expressed. At this point it should be noted that the vowel harmony process of Akan was analysed above in terms of alignment, rather than activation. However, it now seems apparent that alignment simply refers to a span of activation along a single tier. In view of this proposed activation mechanism, an active complement is to be lexically specified at the word level in Akan, and the effects of alignment need not be stated separately. Instead, the tier is automatically activated throughout the entire domain, and from this fact we derive the necessary melodic agreement without recourse to any further stipulation.

So vowel harmony within a given domain is achieved via the activation of the relevant tier, which is lexically marked at the level of the prosodic word. As a final note I shall make the suggestion that each target vowel in the harmony process relies on licensing relations between prosodic units in order to become melodically aligned in the desired way. If a nucleus is projected to the word level where it receives the lexical label *activate [ε]-tier*, that constituent must possess a means of transferring this property to the nuclei lower down the hierarchy, in order for the lexical information to be interpreted in the signal. I propose that this be carried out through inter-nuclear licensing relations, such that when a nucleus N_1 licenses another nucleus N_2, the latter inherits the lexical activation marking from its licensor. This is then repeated along the licensing path until the point at which the timing units within the activation domain a-license the relevant melodic tiers. In the case of local assimilation phenomena such as that illustrated in (23), lexical activation is stipulated at the level of the syllable, and therefore only the licensees (direct and indirect) of this constituent — namely, the onset/nucleus pair — can dominate positions that a-license the harmonising tier. It is, of course, an empirical question as to whether a harmonic span (i.e. an activated tier) is only permitted to affect a string which marks the entire licensing scope of a single prosodic constituent. This matter is clearly worthy of further investigation.

5 Conclusion

In this paper I have assumed a model of vocalic structure which makes no reference to the relation of headship that is standardly employed in element-based representation systems. Instead, I have attempted to demonstrate the possibility of generating the full range of attested vowel inventories by the manipulation of only two structural variables. The first relates to the way in which the melodic tiers housing the three resonance elements are geometrically arranged. Some primes may share a tier, while others may occupy their own separate tier. The specific structural relations between those tiers have been shown to reflect some of the observed facts relating to melodic complexity; in addition, they have helped to characterize the vowel harmony phenomena displayed in some languages. In this paper I have focused on systems that have undergone only one division within the resonance tier (i.e. two tiers licensed). However, on the basis of this discussion it is clear that an investigation into those systems which involve the licensing of three separate tiers would proceed in a similar fashion. The second area of structural variation which contributes to the diversity found within the range of vowel systems refers to the parametrically controlled licensing of tier complements. As illustrated by the account of harmony in Akan, this proves to be particularly appropriate to the representation of ATR distinctions. Given that these two variation paths may potentially co-exist in a single system, it seems realistic to anticipate that their combined mechanism may be sufficient to characterize even the most marked or the most complex of attested inventories. This is, of course, an empirical issue, and one which I intend to pursue in the near future.

References

Anderson, J.M. and C.J. Ewen (1987). *Principles of Dependency Phonology*. Cambridge: Cambridge University Press.

Charette, M. (1991). *Conditions on Phonological Government.* Cambridge: Cambridge University Press.

Charette, M. and A. Göksel (1994). Switching and vowel harmony in Turkic languages. In *SOAS Working Papers in Linguistics and Phonetics* 4, 31-52.

Charette, M. and J. Kaye (1993). The death of ATR. Ms. School of Oriental and African Studies, University of London.

Chomsky, N. and M. Halle (1968). *The Sound Pattern of English*. New York: Harper and Row.

Clements, G.N. (1981). Akan vowel harmony: a non-linear analysis. In *Harvard Studies in Phonology* 2, 108-177.

Clements, G.N. (1984). Vowel harmony in Akan: a consideration of Stewart's word structure conditions. In *Studies in African Linguistics* 15:3, 321-337.

Crothers, J. (1978). Typology and universals of vowel systems. In J.H. Greenberg (ed.), *Universals of Human Language* 2: *Phonology.* Stanford: Stanford University Press.

den Dikken, M. and H. van der Hulst (1988). Segmental hierarchitecture. In van der Hulst and Smith (1988), part 1, 1-78.

Goldsmith, J.A. (1976). *Autosegmental Phonology*. PhD dissertation, MIT. Published 1979, New York: Garland.
Goldsmith, J.A. (1985). Vowel harmony in Khalkha Mongolian, Yaka, Finnish and Hungarian. In *Phonology* 2, 251-274.
Harris, J. (1990a). Reduction Harmony. Paper presented at the 13th GLOW Phonology Workshop, London.
Harris, J. (1990b). Segmental complexity and phonological government. *Phonology* 7, 255-300.
Harris, J. (1992). Licensing Inheritance. In *UCL Working Papers in Linguistics* 4, 359-406.
Harris, J. (1994a). *English Sound Structure*. Oxford: Blackwell.
Harris, J. (1994b). Monovalency and opacity: Chichewa height harmony. In *UCL Working Papers in Linguistics* 6, 509-547.
Harris, J. and G. Lindsey (in press). The elements of phonological representation. To appear in J. Durand and F. Katamba (eds.), *Frontiers of Phonology: atoms, structures and derivations.* Harlow, Essex: Longman.
Kaye, J. (1989). *Phonology: a cognitive view*. Hillsdale, NJ: Lawrence Erlbaum.
Kaye, J. (1990). 'Coda' Licensing. In *Phonology* 7, 301-330.
Kaye, J., J. Lowenstamm and J-R. Vergnaud (1985). The internal structure of phonological elements: a theory of charm and government. In *Phonology Yearbook* 2, 305-328.
Kaye, J., J. Lowenstamm and J-R. Vergnaud (1990). Constituent structure and government in phonology. In *Phonology* 7, 193-232.
Lass, R. (1984). *Phonology: an introduction to basic concepts*. Cambridge: Cambridge University Press.
Lowenstamm, J. and J-F. Prunet (1988). Tigré vowel harmonies. Rapport Annuel du Groupe de Recherche sur la Linguistique Africaniste au CRSH 1987-88. Université du Québec á Montréal.
Mtenje, A.D. (1985). Arguments for an autosegmental analysis of Chichewa vowel harmony. In *Lingua* 66, 21-52.
Prince, A. and P. Smolensky (1993). Optimality theory: constraint interaction in generative grammar. *Technical Report #2 of the Rutgers Center for Cognitive Science*. Rutgers University.
Rennison, J. (1987). Vowel harmony and tridirectional vowel features. In *Folia Linguistica* 21, 337-354.
Rennison, J. (1990). On the elements of phonological representations: the evidence from vowel systems and vowel processes. In *Folia Linguistica* 24, 175-244.
Schane, S.S. (1984). The fundamentals of Particle Phonology. In *Phonology* 1, 129-156.
Selkirk, E.O. (1982). The syllable. In H. van der Hulst and N. Smith (eds.), *The Structure of phonological representations* 2 (1982). Dordrecht, Holland: Foris.
Stewart, J.M. (1967). Tongue root position in Akan vowel harmony. In *Phonetica* 16, 185-204.
Stewart, J.M. (1983). Akan vowel harmony: the word structure conditions and the floating vowels. In *Studies in African Linguistics* 14, 111-139.
Takahashi, T. (in preparation). The theory of phonological licensing and elements. PhD dissertation. University College, University of London.