SOAS Working Papers in Linguistics, Vol. 16 (2013)

The Structure of Tone

Sean Jensen
phonologus@gmail.com

1. Introduction

The theory of tone proposed here has come about from a heretical idea: that melodic units may repeat. Which is to say, amongst other things, that the structure in figure 1a and the structure in figure 1b are both well-formed, and are both phonologically distinct.

(1a)

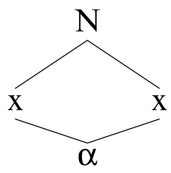

(1b)

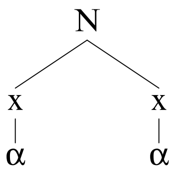

This distinction, we propose, is manifested as a 'tonal' distinction. In the case of a nucleus, this is the phenomenon of contrasting pitch.

The informal insight that the theory attempts to formalise is that 'high tone' is a lot of identical melody realised in a short period of time, and 'low tone' is less identical melody, realised in a longer period of time. The analogy is the commonplace observation of the relationship between periodicity and pitch: more frequent repetition is perceived as a higher pitch than slower repetition.

The phonological tools we have to express these notions are well known. Timing intervals are skeletal points, and melodic material consists of phonological expressions. Association is how we link the two together. The theory developed depends essentially on only these things, so any theory with which the notions of point, expression and association are compatible should be able to adopt at least the basic insight of the theory of tone.

Our theory is quite general, in that its structures and configurations are not restricted to a particular class of constituent, be it nucleus or onset. Of course, these constituents determine the ultimate interpretation of tonal structure: a tonal structure in a nucleus is naturally interpreted as pitch on a vowel, and in what follows, we will be looking at the nucleus in depth. The interpretation of tonal structures within the onset is not discussed here, but it is certainly a subject of great interest. That said, we do examine one place where vowel and consonant come into contact: trans-constituent government and 'codas'. This is to illustrate one of the ways in which phonological context and tonal structure may be expected to interact.

Our heresy is more or less incompatible with existing theories that rely on 'contour principles', at least at the skeletal level. Unfortunately this includes mainstream Government Phonology, in terms of which this theory is set out. What we offer, therefore, is a number of essential tweaks, chief of which is embracing the idea of multiple distinct instances of a phonological expression in the local skeletal structure. Another is more of a corollary to the the theory of tone: autosegmental elements, such as the H and L of mainstream Government Phonology, are not needed to represent 'tone'.

A major omission from this account is any discussion of so-called tonal harmony. What I present is a theory of the structure of tone, independently of any discussion of phonological processes. The success of the theory in its application to Chinese (we look in detail at Mandarin and Cantonese here) makes it worth our consideration, despite its having no need for autosegmental tonal elements. I will, however, offer some thoughts at the conclusion of this paper indicating where I believe an analysis of these phenomena lies.

Familiarity with the basic concepts, methods and praxis of standard Government Phonology is assumed (Kaye, Lowenstamm and Vergnaud, 1990; Kaye, 1990; Charette, 1990).

2. The Rules of the Game

In the following paragraphs we will explore the logical possibilites that association, combined with the idea of melodic repetition, gives us. We shall do this by asking ourselves how to solve a little series of games of the form: Given some number of skeletal points, and some number (possibly different) of phonological expressions, what patterns of association are possible?

The core idea of the theory of tone is repetition of identical melody. All the expressions that take part in a game are therefore all required to be identical instances of each other. In what follows instances of these identical expressions are symbolised by a lower case epsilon (ε).

We do not require that our skeletal points necessarily fall within a single constituent, say nucelus, or onset. We do require, however, that they at least be consecutive in the 'flow of time'. We also require that if a point is in the game, then it can't be left empty, and if an instance of an expression is in the game, it can't be left unassociated.

Skeletal points and association lines are constructed and represented in the familiar way. In particular, a skeletal point may have an exclusive one-to-one association with a single expression (figure 2);

(2)

A skeletal point may have exactly two expressions associated to it in a branching association (figure 3);

(3)

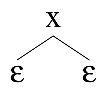

And a single expression may be shared between consecutive skeletal points (figure 4).

(4)

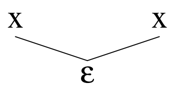

Note that it is perfectly within the rules for a point (boxed in figure 5) to be simultaneously a branching host and a sharing host.

(5)

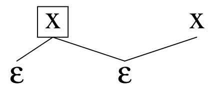

These are all and only the patterns of association that we allow.

In the following sub-sections we explore these Games in the abstract. In the section directly thereafter we examine in detail the possible interpretations of these structures.

2.1. The 1-Point Game

In the simplest cases we have a single skeletal point. When there is a single expression to associate, things are very straightforward. There is only one way to do it (figure 6).

(6)

When there are two expressions, the only possibility the rules allow is branching association (figure 7).

(7)

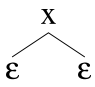

See how this gives us precisely ‘more identical melody in a short space of time’, which was our informal requirement for ‘high tone’. This then is the fundamental High Tone Structure. Anticipating slightly, we shall call the simple one-to-one structure the Neutral Tone Structure.

This exhausts the solutions to the 1-Point Game as there is no way to exhaustively associate any more melody.

2.2. The 2-Point Game

We now expand our game to include two skeletal points. With a single expression there is only one solution: the expression must be shared (figure 8).

(8)

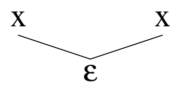

According to our working metaphor, this is exactly what we would expect ‘low tone’ to look like: it is a small amount of melodic material realised over a long extent. This is, then, the fundamental Low Tone Structure, about which we will have much to say in later sections.

Moving to two expressions, the rules of association permit three solutions. There is of course the simple one-to-one association of available expressions to available points (figure 9). This structure has more melodic material in it than the Low Tone Structure (over the same extent), so so we shall call this the Long Neutral Tone Structure (again anticipating ourselves a little).

(9)

There are two further solutions, both of which involve sharing one of the expressions. In figure 10 the non-sharing expression associates to the first point. This structure certainly spends most of its time as as a Low Tone Structure, but begins ambiguously as a cross between a High and a Neutral Tone Structure. We will discuss this structure, along with the following, in detail in the next section. We will call this structure the Falling Low Tone Structure.

(10)

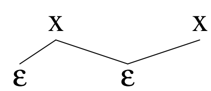

The last solution is the mirror image. The non-shared expression associates to the second point, and we have a structure which spends most of its time as a Low Tone Structure then rises at its end with the ambiguous High/Neutral Structure. We will call this structure the Rising Low Tone Structure (figure 11).

(11)

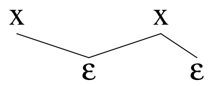

There are no more association possibilities with just two expressions, so we turn now to the game with three expressions. The rules of association permit exactly three solutions. Two solutions are unsurprising. The third is noteworthy. The first is a High Tone Structure followed by a Neutral Tone Structure (figure 12), which we shall call the Falling High Tone Structure.

(12)

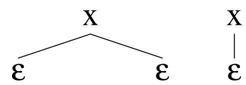

The second is a Neutral Tone Structure followed by a High Tone Structure (figure 13). We shall refer to this structure as the Rising High Tone Structure.

(13)

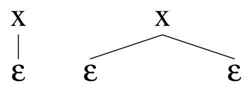

The third of the solutions, and the most noteworthy, is a variant of the Low Tone Structure (figure 14). This Low Tone Structure is flanked on either side by the ambiguous High/Neutral association we saw in the Low Rising and Low Falling Tone Structures. Again, we shall comment on this structure in the following section. We shall refer to this structure as the Dip Tone Structure.

(14)

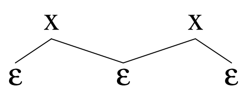

There are no more association possibilities with just three expressions. With four expressions there is a single solution: we can create two consecutive High Tone Structures (figure 15). Unsurprisingly, we shall refer to this as the Long High Tone Structure.

(15)

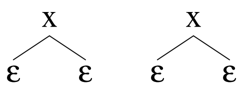

There are no further possible solutions involving two skeletal points. Five or more expressions have no exhaustive solution, as at least one expression will always be left unassociated.

This, then, completes the inventory of tonal structures created by the 1-Point and the 2-Point Games. One interesting observation is in order: it is possible to get a low dipping structure in the game with 2 points, but not a rising then falling structure, nor a dipping high-mid-high structure.

2.3. The 3-Point Game and Beyond

It is possible to continue adding points: the 3-Point Game, the 4-Point Game, and so on. As is traditional, this will be left as an exercise for the reader. What the diligent student will find is that no new tonal structures emerge. What will be found is simply various sequences of the tonal structures discovered above, and a generalisation of the Dip Tone Structure: a fall into a sequence of Low Tone Structures, the sequence ending with a rise.

We will not pursue these Games in this paper, instead confining our attention to the 1-Point and 2-Point Games, as most of the basic phonological contexts we are accustomed to dealing with at the skeletal level are two-point configurations: constituent government (branching onsets and rimes), and trans-constituent government (‘codas’). As will be mentioned at the conclusion of this paper, it is quite possible that these n-Point Games contain the solution to the problem of so-called tone harmony.

3. Pitch and Interpretation

Before trying to identify these structures in real-world languages, we need to say a few words about what expectations we should have about how they might be realised. It should be remembered that these are phonological structures, not phonetic or acoustic ones. We are certainly not claiming that a particular musical pitch can be definitively deduced from these structures. Nor would most avowedly phonological theories. What can we expect, then?

Let us consider the nucleus, and its interpretation as a vowel capable of bearing contrasting pitch. We will refer in the familiar way to these pitches using Chao tone numbers (Chao, 1968): 5 represents the highest pitch, 3 a middle pitch and 1 the lowest pitch, with 2 and 4 marking intermediate ‘notes’ on this rising scale of pitch. A long, sustained pitch is shown by doubling: 33 is a long mid pitch, and contour pitches are shown iconically: 35 is a mid pitch rising to a high pitch, and so on. This is a purely descriptive system which lets us talk easily about acoustic pitch, and has no other relation to the theory of tone.

Our basic expectation is naturally that High, Low and Neutral Tone Structures should map straightforwardly into pitches 5, 1, and 3 respectively, and this is a good rule of thumb.

However, and it is a big ‘however’, the theory of tone already has something to say about its own interpretive possibilities. The most important concerns the Low Tone Structure, which, realised within a branching nucelus, and with ε instantiated with the expression (A), takes on the form of figure 16. The canonical interpretation is a (long) low-pitched a-vowel, [a:^{11}]. I shall diagram all branching nuclei as in figure 16. This is a purely typographic conceit, as it makes the diagrams much easier to read, and can be safely understood as the familiar branching nucelus of Government Phonology.

(16)

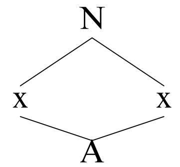

The Low Tone Structure is *necessarily* long: it can only exist when an expression is shared between two consecutive skeletal points. This tells us that a Low Tone Structure cannot exist at a single skeletal point, so if we encounter what we might describe as a low-pitched short vowel, or any low pitch of short duration such as a short, low point in a contour pitch, then **it is not the interpretation of a Low Tone Structure.**

Which raises the question, if a short, low pitch is not an interpretation of the Low Tone Structure, then what is it the interpretation of? The short answer is that it is one possible interpretation of a (short) Neutral Tone Structure that is not itself part of the Long Neutral Tone Structure. The long answer is as follows:

The Long Neutral Tone Structure has the same length as the Low Tone Structure, and the Long High Tone Structure. These three Tone Structures do not involve any contours either tonal or melodic, yet they are distinct: their tonal structure is all that distinguishes them. Therefore they define a *necessary* three-way tonal distinction, which must therefore be maintained in any (nuclear) interpretation strictly as three distinguishable pitches for long vowels: long mid [a:^{33}] (figure 17a); long low [a:^{11}] (figure 17b); and long high [a:^{55}] (figure 17c).

(17a)

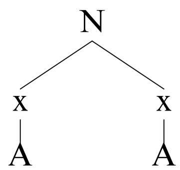

(17b)

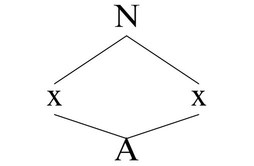

(17c)

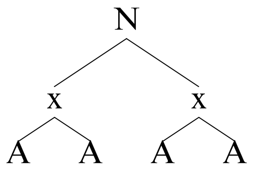

Now, outside of the Long Neutral Tone Structure, which we have just seen must always be interpreted by a distinguishable (long) mid-pitch, we find instances of the (single point) Neutral Tone Structure: trivially in the 1-Point canonical version (figure 18a); in the Rising High Tone Structure (figure 18b); and in the Falling High Tone Structure (figure 18c).

(18a)

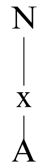

(18b)

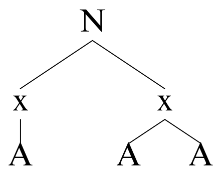

(18c)

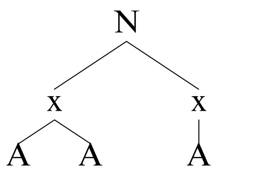

Since a Low Tone Structure cannot exist at a single point, the (single point) Neutral Tone Structure only ever need be distinguished from the (single point) High Tone Structure, so its interpretation only ever need be identifiably lower than the high pitch, and that could be either a (short) mid pitch like pitch 3, or a (short) low pitch like pitch 1. So, there is absolutely no *phonological* distinction possible between a short mid pitch and a short low pitch. They are both equally serviceable non-contrasting interpretations of the (short) Neutral Tone Structure. Reasonable interpretations for these structures are therefore anything from [a^{3}] through [a^{1}] (figure 18a); anything from [a:^{15}] through [a:^{35}] (figure 18b); and anything from [a:^{53}] through [a:^{51}] (figure 18c).

Similar reasoning applies to the interpretive possibilities of the Falling Low Tone Structure (figure 19a), the Rising Low Tone Structure (figure 19b), and the Dip Tone Structure (figure 19c).

(19a)

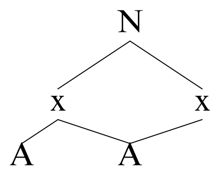

(19b)

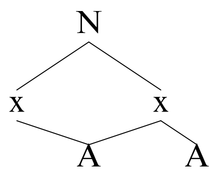

(19c)

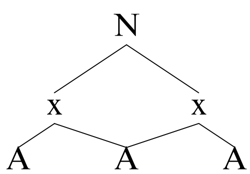

There is of course only one possible interpretation of the actual Low Tone Structure at the heart of these structures: it must be a long low pitch. However, the contour portion is forced to associate to a point which is independently host to part of a shared association, so there is only one possibility: associate as a branch. Therefore, there is only exactly one way in which the contour part can differ from the Low Tone Structure part. So the interpretation of the contour portion simply needs to be distingishable from the low pitch that interprets the Low Tone Structure, and that is all. Therefore a (short) high pitch (such as pitch 5) and a short mid pitch (such as pitch 3) or even a (short) middling low pitch (pitch 2) are all equally serviceable non-contrasting interpretations of the contours into and out of these various Low Tone Structures. There is no possible phonological difference between a (short) high pitch and a short (mid) pitch *when they are contours into or out of a long low pitch.*

Reasonable interpretations for these structures are therefore anything form [a:^{511}] through [a:^{211}] (figure 19a); anything from [a:^{115}] through [a:^{112}] (figure 19b); and anything from [a:^{5115}] through [a:^{2112}] through [a:^{5112}] through [a:^{2115}]. (figure 19c).

The only Tone Structure we haven't considered is the isolated High Tone Structure (figure 20), whose canonical interpretation is simply short high pitch, [a^{5}]. It is the polar opposite, in both pitch and length, of the Low Tone Structure.

(20)

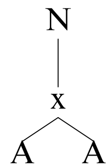

The reader is urged to keep these observations in mind as we proceed.

Before we leave this section, it is worth noting that a pretty result of the theory of tone is that a non-branching nucleus, being a 1-Point Game, can only ever contain a High Tone Structure or a Neutral Tone Structure, with, as we have just explained, expected interpretations [a^{5}] and something in the range (short) [a^{1}] through [a^{3}]. That is, the simplest tone-bearing system *necessarily* has only two tonal distinctions, a high pitched tone, and a non-high pitched tone. And the non-high pitched one cannot be an interpretation of the Low Tone Structure. Thus the theory of tone provides a proof of Kaye's (2000a) hypothesis that all two-tone systems must include a high tone (element H), and exclude a (true) low tone (element L).

4. Mandarin and Cantonese

We shall put the theory of tone to work in the study of the tone systems of Mandarin Chinese and Cantonese. There is not space to devote to a detailed critique of existing approaches, even in closely compatible frameworks. This is a particularly egregious omission given especially Yeng-seng Goh's pioneering work on the segmental phonology of Mandarin and Jonathan Kaye's work on tone and further work on Mandarin in Government Phonology (Goh, 1996; Kaye, 2000a,b,c).

However, the emphasis here is of necessity on conveying the ideas of the theory of tone. Assessing whether or not it is a more insightful replacement for Kaye's proposals must wait for a less 'working' version of this paper.

In addition to the tonal phenomena we are about to consider, both Mandarin and Cantonese possess secondary tonal phenomena in unstressed and so-called 'sandhi' contexts that we will not address here, keeping our account instead to citation forms.

All 'words', the general Pan-Chinese 字 *zi,* consist of an optional initial consonant, a single stressed vowel, or diphthong, and an optional final consonant. In Mandarin the final consonant can only be -n or -ng; in Cantonese, -m, -n, -ng or -p, -t, -k. The vowels, diphthongs and initial consonants vary, sometimes widely, between dialects.

Mandarin Chinese (Dow, 1980) is conventionally described as having four tones. In works in English they are called First Tone, Second Tone, Third Tone, and Fourth Tone. Their pitch characteristics are as follows: First Tone is high and level; Second Tone rises from mid to high; Third Tone falls quickly to low, stays there, and then quckly rises; Fourth Tone quickly falls from high to (fleetingly) low. All four tones can in general appear on all vowels and diphthongs, regardless of initial or final consonant.

In standard reference works in Chinese (Essentials, 1980; Lexicon, 1989) the Chao tone numbers are given as First Tone 11; Second Tone 35; Third Tone 214; and Fourth Tone 53. However, the Chao tone numbers take little account of length, and these symbols are usually accompanied by commentary to the effect that the low pitch of the Third Tone is predominant, and the Fourth Tone is a dramatic fall from high to low, crucially with no extended period of low pitch. In the notation of the previous section we would most probably describe the Third Tone as 2114 to emphasise this.

In contrast, Cantonese (Gao, 1980) is described as having variously eight, nine or even ten tones. The tones are traditionally separated into two classes: those that appear before the -p -t -k endings, called 入声 *ru-sheng* tones, and those that don't. Again, the Chinese language standard reference works, cited above, describe these with Chao numbers as follows. The pre-stop, or ru-sheng, tones are: short and high ru-sheng (pitch 5), long and mid ru-sheng (pitch 33), long and low ru-sheng (pitch 22) and short and low ru-sheng (pitch 2). The tones that appear elsewhere are: high level (55), high falling (53); low falling (21); high rising (35); low rising (23), mid level (33); and low level (22).

As in the case of the Mandarin Third Tone, it is important to realise that the low falling (pitch 21) and the low rising (pitch 23) tones are described as being tones which dwell in their low-pitch portions, and thus in our more length-sensitive notation would be more accurately notated as pitch 211 and pitch 223.

We shall assume, *contra* both Goh (1996) and Kaye (2000a,b,c), that what underlies all these in both Mandarin and Cantonese is a branching nucleus (figure 21).

(21)

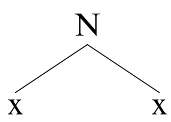

The two points of the branching nucleus provide the necessary arena for the 2-Point Game, which generates all the Tone Structures, as we have seen above.

For an illustrative expression, we will use (A), as [a]-vowels are found abundantly in both Mandarin and Cantonese. Therefore, the structures we will be looking at are exactly the structures we considered in the section on interpretation.

As we are dealing with a 'pure' branching nucleus, we will postpone discussion of the Cantonese pre-stop tones until a later section. For the remainder of this section the reader should understand that these ru-sheng tones are tacitly excluded from the discussion.

We shall proceed simply by considering each of the possible solutions to the 2-Point Game in turn, and together with their interpretive possibilities, see what sort of a fit there is with either Mandarin, Cantonese or both. As will be seen, a quite remarkable pattern emerges.

4.1. Games with 1 and 2 Expressions

So consider first the game when we have just one expression. There is one solution, the Low Tone Structure, whose interpretive possibilities are fairly conservative: a long low pitch, something like [a:^{11}] (figure 22).

(22)

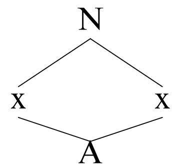

There is nothing in Mandarin that fits well with this, but it is certainly a perfect fit for the Cantonese low level tone (pitch 22, [a:^{22}]).

The game played with two expressions has three solutions, the Long Neutral Tone Structure (figure 23a), the Low Falling Tone Structure (figure 23b), and the Low Rising Tone Structure (figure 23c). The interpretations of these structures we would expect to be, respectively, [a:^{33}] (figure 23a); anything in the range [a:^{511}] through [a:^{211}] (figure 23b); and anything from [a:^{115}] through [a:^{112}] (figure 23c).

(23a)

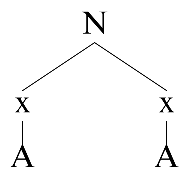

(23b)

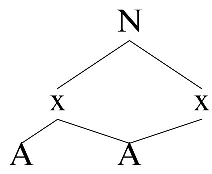

(23c)

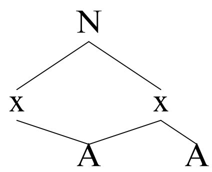

Again, Mandarin lacks any convincing realisations of these structures. In Cantonese, however, the opposite is true. They correspond quite obviously to the mid level tone (pitch 33, [a:^{33}], figure 23a); the low falling tone (pitch 211, [a:^{211}], figure 23b); and the low rising tone (pitch 223, [a:^{223}], figure 23c).

This exhausts the possible solutions to the one and two expression 2-Point Games. A poor showing for Mandarin, and a full house for Cantonese.

4.2. Overcrowded Games: 3 and 4 Expressions

So far every possible Tone Structure has found an instantiation in Cantonese, and not a single one appears in Mandarin. However, as soon as we move to the three and four expression 2-Point Games, Mandarin puts in a strong showing. This is particularly interesting because the only property that distinguishes the three and four expression games from the one and two expression games is whether or not the number of expressions exceeds the number of points. We could say that with three and four expressions, the 2-Point Game has become 'overcrowded', and that a distinguishing property of Mandarin, in contrast to Cantonese, is that all its branching nuclei are overcrowded.

So we continue with the 3-expression variation, and note its three possible solutions: the High Falling Tone Structure (figure 24a), the High Rising Tone Structure (figure 24b), and the Dip Tone Structure (figure 24c). The interpretations of these structures we would expect to be, respectively, anything form [a:^{53}] through [a:^{51}] (figure 24a); anything from [a:^{15}] through [a:^{35}] (figure 24b); and anything from [a:^{5115}] through [a:^{2112}] through [a:^{5112}] through [a:^{2115}]. (figure 24c).

(24a)

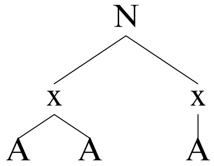

(24b)

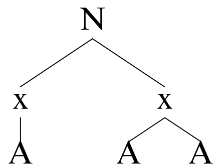

(24c)

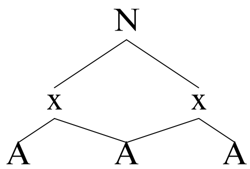

So here we have the dramatic appearance of Mandarin. Each of these is plainly identifiable in Mandarin as the Fourth Tone (pitch 51, [a:^{51}], figure 24a); the Second Tone (pitch 35, [a:^{35}], figure 24b); the the Third Tone (pitch 2114, [a:^{2114}], figure 24c).

Cantonese has two of these: the High Falling Tone Structure is the high falling tone (pitch 53, [a:^{53}], figure 24a), and the High Rising Tone Structure is the high rising tone (pitch 35, [a:^{35}], figure 24a). Letting the side down for Cantonese is the Dip Tone Structure, which it lacks. It is interesting to note that this one structure has the property that both its points are simultaneously a sharing host and a branching host. We will return to this observation at the end of this section.

And so we come finally to the four expression variant and its unique solution: the Long High Tone Structure (figure 25). Its interpretation we expect to be a conservative [a:^{55}].

(25)

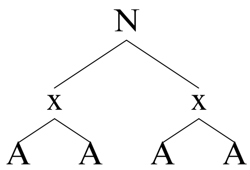

Both Mandarin and Cantonese have unequivocal instances of this Structure: Mandarin First Tone (pitch 55, [a:^{55}]), and the Cantonese high level tone (pitch 55, [a:^{55}]).

This completes all possible 2-Point Games. We have exhausted the possible Tonal Structures the theory provides, and, remarkably, we have exhausted all the Mandarin tones, and all the non-ru-sheng tones of Cantonese.

Let us briefly get our bearings. As far as all the Cantonese tones (not ru-sheng), are concerned, they instantiate all possible solutions to all possible 2-Point Games, apart from the Dip Tone Structure, which it lacks.

And the Mandarin tone system is in its entirety a perfect instantiation of all possible solutions to the Overcrowded 2-Point Game!

I would like to suggest that the differences in inventory show two axes that the phonology is sensitive to: one axis is crowdedness: in Mandarin, the nucleus must be crowded; in Cantonese there are no constraints on the crowdedness of the nucleus. The other axis is degree of symmetry with respect to the patterns of association present: Mandarin has no such symmetry requirements in the nucleus, whereas Cantonese requires that only one point can ever be simultaneously a branching host and a sharing host. We might say that Cantonese does not support doubly symmetrical, or 'over-symmmetrical', solutions.

We shall have more to say about symmetry when we come to consider the Cantonese ru-sheng (pre-stop) tones.

5. Interlude: Much Ado About Nothing

The theory of tone has at its core the idea that the possible repetition of phonological expressions manifests itself as tone. There is a very clear expectation, then, that empty skeletal positions are incapable of hosting any of the possible Tone Structures. *A fortiori,* we expect not to find any realised (that is, non-p-licensed) empty nuclei capable of expressing tonal contrasts.

Now, this is not in itself a problem, and indeed is a rather nice prediction of the theory of tone. However, it does set the theory of tone at loggerheads with Goh's (1996) and Kaye's (2000a,b,c) prevailing analyses of Mandarin in Government Phononology. In order to account for the very striking distribution of the I-element in Mandarin (and to a lesser extent, the U-element) it is claimed that an empty nucleus heads many Mandarin domains, and that in the phonology a process of I-spreading from the onsets containing the element I 'fills' these (lexically) empty positions. Naturally, in proposing the theory of tone I am therefore claiming that this analysis must be wrong, as all these purportedly empty nuclei express the full range of Mandarin tone. To be clear, the theory of tone, combined with the facts of tonal possibilities in Mandarin means that the head nucleus of the domain in Chinese *cannot ever be empty.*

It is not my intent to offer a full alternative analysis of the melodic structure of Mandarin here, as that would take us too far from our subject. However, I acknowledge that such an analysis is due, so briefly sketch here an outline which should at least show that reasonable alternatives exist.

A striking property of Goh's and Kaye's analyses is the extremely simple system of (lexical) nuclear melody, as against the comparatively rich (lexical) onset melody. This is perhaps a reflection of what were then current theoretical preoccupations with derivations and the directionality of various types of spreading. I suggest that there is in fact no theory-internal reason to expect an analysis which necessarily maximises sub-skeletal onset structure and minimises sub-skeletal nuclear structure. The theory of tone indicates that this bias towards a maximised onset in the analysis of Mandarin is in fact incorrect.

The essence of an alternative analysis is 'reversing the arrows' in Goh and Kaye's analyses: we want richly specified nuclear exressions spreading I and U outwards to sparsley specified onset positions. Or, as a process-agnostic alternative, onsets and nuclei are equally richly specified, but must possess certain 'melodic symmetries'. In both these alternative scenarios it is entirely possible to claim that the head nucleus of the domain in Mandarin can never be empty.

6. Diphthongs

Now, Mandarin and Cantonese have diphthongs, too. In particular, but not exclusively, both have [ai]-diphthongs. In terms of tone, [ai] behaves exactly the same way as the [a:] we have been looking at above. Given then that a branching nucleus with identical expression-instances is required for the tone structures, where does the (I) expression of the [ai]-diphthong go? Clearly there must be a following constituent, which, according to the rules of constituent licensing in Government Phonology, must be either a nucleus, or an onset-nucleus pair. Suppose then that the nucleus is the rising high tone of Cantonese, or the equivalent Second Tone of Mandarin, [a:^{35}], and (I) is in the second nucleus (figure 26).

(26)

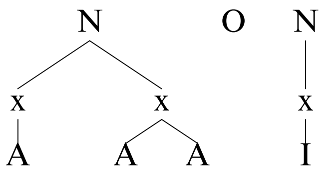

According to the theory of tone this (I) expression finds itself in an isolated Neutral Tone Structure, which we know should be realised as anything in the range [i^{1}] through [i^{3}]. Splitting the difference, we would expect this structure to be interpreted as [a:^{35}i^{2}]. Hence we would expect Mandarin and Cantonese to have no 35-pitched diphthongs: they would all have a rising and then falling contour. A tone with such a contour, we have seen, does not exist, in either Mandarin or Cantonese (although there are dialects of Chinese, such as Fujian Min, where it does). *A fortiori,* [a:^{35}i^{2}] does not exist in Mandarin or Cantonese. 35-pitched diphthongs abound. Thus we must reject the hypothesis that the (I) expression is in the second nucelus.

We therefore adopt the alternative hypothesis, that the (I) expression is in fact in the onset (figure 27). The second (empty skeletal position, domain final) nucleus is p-licensed by parameter, as is common for a large number of the World's languages.

(27)

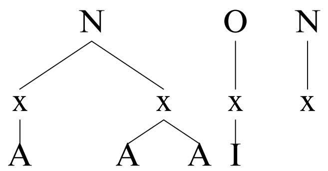

This correctly leaves the nucleus free to express all the tone structures of Mandarin and Cantonese, without the side-effects that occur if the (I) were in the second nucleus. This result neatly corroborates Kaye's (2000b,c) correction of Goh's (1996) assumption that the second element of the diphthong is to be found in the second nucleus.

The expected interpretation is (pedantically) [a:^{35}j], which is a completely reasonable description of the Mandarin and Cantonese [ai]-diphthong. (Which, to be consistently pedantic, we shouldn't now really call a diphthong.)

Identical arguments apply to the other finals of Mandarin, and the non- -p, -t, -k finals of Cantonese. We turn our attention to these latter in the remainder of this paper.

7. Games with Restrictions: Codas

Let us step back from particular examples and consider the 2-Point Game again in the abstract. Let us play a restricted version of this game, which should give us some insight into how tone interacts with other structural principles, such as coda-licensing.

The restriction will be a restriction on the patterns of association that a point allows. Up to now we have been considering the unrestricted case, where all points can be hosts to all three patterns of association.

Now, the most severe restriction of all is the complete opposite of this free-for-all: it is the one which forbids any and all association. A point subject to such a restriction within the context of a Game is a 'do-not-touch' point. Let us show this in our diagrams with a box. In this case the game reduces to the 1-Point Game, whose solutions are either the Neutral Tone Structure (figure 28)

(28)

or the High Tone Structure (figure 29).

(29)

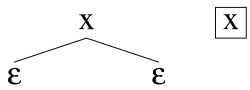

7.1. Strong Coda Licensing

Imagine now that that the licensing and governing relations that skeletal points find themselves in as a result of their position in the constituent structure can impose just such restrictions on the association possibilities locally available. This should not sound outlandish, as we are accustomed to thinking in terms of reduced possibilities in certain structural positions (Charette, 1990; Kaye, Lowenstamm and Vergnaud, 1990).

I should like to suggest that one such constraining configuration is the coda-licensing configuration (Kaye, 1990. Figure 30), and that the coda-licensed point (the 'coda' for short, shown boxed in the figure) has its possibilities for association curtailed. This will obviously constrain the way it plays the 2-Point Game with its governing nucleus (encircled in the figure). I am again abusing the usual Government Phonology diagramming conventions for branching constituents, again in the name of clarity. This can be safely read as the familiar branching rime of Government Phonology.

(30)

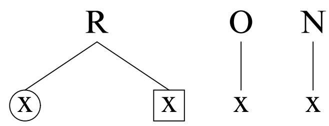

In one case, which I shall call Strong Coda Licensing, let this restriction be the do-not-touch restriction. A coda subject to Strong Coda Licensing therefore reduces the 2-Point Game with its governing nucleus to a 1-Point Game. Effectively it withdraws itself from the Game. Thus, before Strongly Licensed Codas, we will find only solutions to the 1-Point Game: the (short) Neutral Tone Structure (figure 31a); and the (short) High Tone Structure (figure (31b)

(31a)

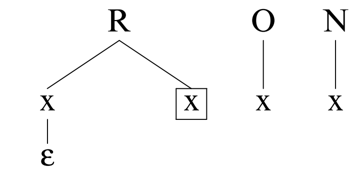

(31b)

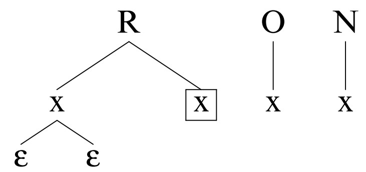

Now, a Coda Licence implies the presence of a stop consonant. In mainstream Government Phonology, the structure above is understood to be a short vowel followed by a geminate stop (Kaye, 1990); according to Jensen (1994), a short vowel followed by any type of stop; and, ignoring fundamental conceptual differences, according to Pöchtrager (2006), a short vowel followed by a fortis stop. So we have the rather nice result that before some kind of 'strong' stop consonant, we will only find a two-way (high level pitch versus some non-high level pitch) short-vowelled tone system.

And we find exactly this pre-stop subsystem in Cantonese: the two short ru-sheng (pre-stop) tones: short high ru-sheng (pitch 5) and short low ru-sheng (pitch 2). Perfect.

The two remaining ru-sheng tones are long mid ru-sheng (pitch 33) and long low ru-sheng (pitch 22), both long. These are both also followed by stop consonants, but they are, in contrast, invariably long tones, and at distinguishable mid and low pitch, never high pitch. Clearly there must be another type of coda-licensing which is not as restrictive as Strong Coda Licensing.

7.2. Weak Coda Licensing

There is no logical reason to suppose that do-not-touch is the only restriction possible. In fact, given that there are three types of association (one-to-one, branching, and sharing) we should certainly be investigating at least eight possible restrictions and the contexts that may or may not give rise to them.

We have also seen two new types of restriction, and these are restrictions on the possible solutions to a Game as a whole. We have seen that Mandarin is sensitive to crowdedness, but doesn't care about the symmetry of its solutions, whereas Cantonese doesn't care about crowdedness but is sensitive to over-symmetry in its solutions. Unfortunately space prohibits a thorough investigation of these two new axes, which seem symptomatic of deep principles at work.

Nonetheless, we do have the space to show one further case where the phenomena seem sensitive to these axes: the two remaining long ru-sheng tones of Cantonese.

We ended the previous subsection noting that there must be another form of coda-licensing, different from Strong Coda Licensing. Let this other form be called Weak Coda Licensing, which we shall define as a syndrome of two requirements on the entire rime constituent in which the coda resides. To illustrate the fact that the whole Game is constrained, not just a single point, we encircle the constrained constituent in diagrams (figure 32).

(32)

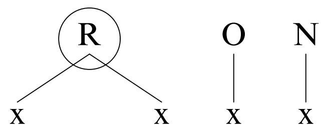

Weak Coda Licensing imposes two requirements on the solutions to the 2-Point Game between the nuclear head and the coda. Unsurprisingly, one is a restriction on the symmetry of the solutions, and one is a restriction on the crowding of the solutions. A Weakly Coda Licensed rime is a 2-Point Game which must be symmetrical, and must not be overcrowded.

The solutions to the Weakly Coda Licensed Game are the Long Neutral Tone Structure (figure 33a); and the Low Tone Structure (figure 33b). Note that no contour structures are possible because they break the symmetry requirement, and that a Long High Tone Structure is impossible because it is an over-crowded solution.

(33a)

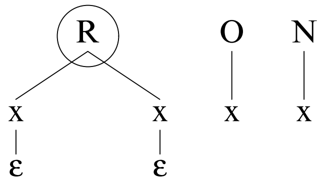

(33b)

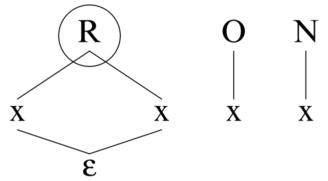

The expected interpretations of these structures are exactly the Cantonese long ru-sheng forms: long level mid pitch before a stop, and long low level pitch before a stop.

Note that we stated Strong Coda Licensing as a blanket 'no association' restriction on the coda point alone. However, it is almost certain that this is just a manifestation of crowding and symmetry requirements, as in the case of Weak Coda Licensing, and in the case of the tonal distributions in Mandarin and Cantonese. In fact we can see that Strong Coda Licensing satisfies the no overcrowding requirement of Weak Coda Licensing, so we can assume immediately that no-overcrowding is a common property of Coda Licensing in general. What distinguishes Strong Coda Licensing from Weak Coda Licensing is their respective symmetry requirements. Weak Coda Licensing requires symmetry; Strong Coda Licensing requires a pathological anti-symmetry. At present I do not have a satisfactory account of this pathological anti-symmetry, so I leave the Strong Coda Licensing condition stated as the descriptive do-not-touch requirement.

Now, as before, the Coda Licence implies the presence of a stop. Mainstream Government Phonology has not much to say about the structures above. In the system of Jensen (1994), they are again stops, but of an unidentified sort. The insight of Pöchtrager's (2006) theory of consonant voicing is that these are all lenis, or voiced, stops. Which is to say that Cantonese is more accurately described as has having *six* stop-consonant finals: Strong Coda Licensers, -p, -t, -k, and Weak Coda Licnesers, -b, -d, -g. The short-vowelled ru-sheng tones appear only before -p, -t and -k, and the long ru-sheng tones appear only before -b, -d, and -g. (Or, equivalently, short ru-sheng before -pp, -tt, -kk; and long ru-sheng before -p, -t, -k.)

This result in turn is further corroboration of Jensen's (1994) idea that stops are structurally complex, and in particular Pöchtrager's (2006) proposals regarding the representation of the strength of stops (the fortis/lenis distinction, in his terms). In fact we may draw a direct parallel between Pöchtrager's (2006) notion of m-command and our Strong and Weak Coda Licensing.

8. Summary

Given the general principles of the theory of tone outlined here, we have arrived at surprisingly concise descriptions of the apparently quite different tone systems of Mandarin and Cantonese. In both Mandarin and Cantonese, the head nucleus of the domain is always branching, and it may only contain identical instances of a single expression.

In Mandarin, the nucleus must additionally be overcrowded. This one extra condition is sufficient to describe the Mandarin tone system (at least in citation forms) in its entirety!

In Cantonese there is no constraint on crowding, but there is a constraint on over-symmetry. Additionally, the nucleus in Cantonese may be subject to both Strong and Weak Coda-Licensing. In both cases the nucleus must then not be overcrowded. In the Strong case nothing may associate to the coda point; in the Weak case, there must be symmetry. These complications are sufficient to describe the Cantonese tone system in its entirety.

All of these conditions are general properties of the theory of tone, and not a specific 'hack' for Cantonese or Mandarin. It is fully intended and expected that these conditions, and others, equally general, will be sufficient to describe the full of range of cross-linguistic tonal systems.

9. Going Forward

One important theoretical direction this work points to is a generalised theory of association, along the lines suggested in our discussion of Restricted Games and Coda Licensing. Such a general theory should certainly be aiming to subsume Strong and Weak Coda Licensing. Perhaps it may also be found ultimately to underpin the very idea of constituent structure itself.

The observation that the axes of symmetry and crowdedness seem to play a role in the phonology is certainly worthy of further scrutiny and testing. These notions are not part of the common vocabulary of Government Phonology, and it remains to be seen whether they are fundamental organisational principles of phonology, or whether they in fact follow from deeper principles that are already familiar.

The theory of tone raises questions too about the universality of branching constituents: I can imagine it being possible that all languages have branching constituents. The way the theory of association is restricted in a given language will manifest itself more or less as how much 'tonality' exists in the language. We might even speculate that there is a broad distinction between 'tonal languages' where constituents may only contain instances of a single expression, and 'diphthongal languages' where instances of more than one expression may exist within a constituent. And no doubt a spectrum between these two extremes.

Such is the long view. At a more immediate level, also to be explored is the interpretation of the Tonal Structures within the onset, and between the nucleus and both the preceding onset and the following onset. There are strong indications that nasality and low tone are kindred phenomena, suggesting perhaps that these latter two environments are precisely where to look for nasal phenomena. Nasal vowels of the French variety are possibly as structurally as complex as figure 34.

(34)

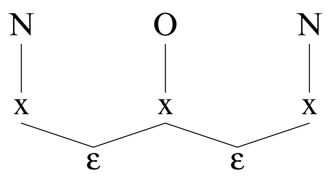

Tonal harmony has been completely unmentioned in this paper. Without any tonal autosegments, it seems like the theory of tone is doomed in the face of so-called harmonic phenomena. I would like to suggest that the correct way forward is to push the Games beyond just the 2-Point Game. Although no new tonal structures emerge, what we do find is that the solutions to the general n-Point Game contain linear strings of the basic tonal structures of the 1-Point and 2-Point Games, which interact at their edges in subtle ways. Combined with the constraining forces of constituent structure, we would expect quite complex pitch interactions between consonants and vowels and sequences of vowels. I believe, therefore, that a convincing account is within reach.

Acknowledgements
The Gods of Phonology are indeed strange. This work came about suddenly and completely unexpectedly while I was working on an unrelated project, and this paper and the accompanying analysis were hastily put together in the space of a week. My thanks, therefore to Connor Youngberg and the editors of the Working Papers for accepting it at such late notice and in truly Working condition. Many thanks also to Jonathan Kaye and Markus Pöchtrager for their ongoing constructive engagement and their patience in the face of a barrage of heat-of-the-moment e-mails from me at the beginning of the writing process.

References

Chao, Yuen Ren, 1968. *A Grammar of Spoken Chinese.* Berkley: University of California Press.

Charette, Monik, 1990. *Conditions on Phonological Government.* Cambridge: Cambridge University Press.

Dow, Francis D.M., 1984. *An Introduction to the Pronunciation of Chinese.* Edinburgh: Department of Chinese, University of Edinburgh.

Essentials, 1980. 汉语方言概要 (2ed), 1980 [1959]. Beijing: 文字改革出版社 [The Essentials of the Chinese Dialects].

Gao, Hua-nian 高華年, 1980. 廣州方言研究 Hong Kong: 商務印書館 [Studies in Cantonese].

Goh, Yeng-seng, 1996. *The Segmental Phonology of Beijing Mandarin.* Ph.D. School of Oriental & African Studies, University of London.

Jensen, Sean, 1994. Is ʔ an Element? Towards a Non-segmental Phonology. *SOAS Working Papers in Linguistics and Phonetics,* 4, pp.71-78.

Kaye, Jonathan, 1990. 'Coda' Licensing. *Phonology,* 7.2, pp.301-330.

Kaye, Jonathan, 2000a. *A Short Theory About Tones.* [manuscript] Guangzhou: Guangdong University of Foreign Studies. [pdf] Available through: `<http://www.unice.fr/dsl/tobweb/gppapers.htm>` [Accessed 26 May 2013].

Kaye, Jonathan, 2000b. *A Fresh Look at Putonghua Onset-Rime Pairs* [manuscript] Guangzhou: Guangdong University of Foreign Studies. [pdf] Available through: `<http://www.unice.fr/dsl/tobweb/gppapers.htm>` [Accessed 26 May 2013].

Kaye, Jonathan, 2000c. 普通话的声母与韵母 [manuscript] Guangzhou: Guangdong University of Foreign Studies. Chinese version of Kaye (2000b). In press 2001, but subsequent publication history uncertain (Kaye, p.c.).

Kaye, Jonathan; Lowenstamm, Jean; and Vergnaud, Jean-Roger, 1990. Constituent Structure and Governement in Phonology. *Phonology,* 7.2, pp.193-231.

Lexicon, 1989. 汉语方音字汇 (2ed), 1989 [1985]. Beijing: 文字改革出版社 [Lexicon of Chinese Dialect Pronunciation].

Pöchtrager, Markus, 2006. *The Structure of Length,* Ph.D. Universität Wien.