Computationally, tone is different*

Adam Jardine
University of Delaware

This paper establishes that unbounded circumambient processes, phonological processes for which crucial information in the environment may appear unboundededly far away on both sides of a target, are common in tonal phonology, but rare in segmental phonology. It then argues that this typological asymmetry is best characterised by positing that tone is more computationally complex than segmental phonology. The evidence for the asymmetry is based around attestations of unbounded tonal plateauing, but it is also shown how the 'sour-grapes' harmony pathology is unbounded circumambient. The paper argues that such processes are not weakly deterministic, which contrasts with previous typological work on segmental phonology. Positing that weak determinism bounds segmental phonology but not tonal phonology thus captures the typological asymmetry. It is also discussed why this explanation is superior to any offered by Optimality Theory.

1 Introduction

This paper establishes and characterises a typological difference between segmental and tonal phonology: unbounded circumambient processes are frequently attested in tonal phonology, but rare in segmental phonology. I define an UNBOUNDED CIRCUMAMBIENT PROCESS as one in which triggers or blockers appear on both sides of a target, and there is no bound, on either side, on the distance between these triggers or blockers and the target. This paper argues that these processes are more complex than those that are commonly attested in segmental phonology, and so the asymmetry can be characterised in a unified way by positing that tone is more computationally complex than segmental phonology.

A common unbounded circumambient process in tonal phonology is UNBOUNDED TONAL PLATEAUING (Kisseberth & Odden 2003, Hyman 2011; henceforth UTP), in which any number of tone-bearing units between two underlying high tones also become high. A simple example from Luganda (Hyman 2011: 231) is given in (1), where the plateau is underlined.

(1) bikópo byaa-walúsiimbi → bikópó byáá-wálúsiimbi
    cups   of-Walusimbi
    'the cups of Walusimbi'

UTP is an unbounded circumambient process because the triggering H tones can be any distance away from the affected tone-bearing units. Hyman (2011) observes that UTP is commonly attested in tonal phonology, but similar plateauing effects are, with one exception, unattested in segmental phonology.

The first contribution of this paper is to document this asymmetry in detail, and, by comparing UTP to the sour-grapes vowel-harmony pathology (Baković 2000, Wilson 2003, McCarthy 2010, Heinz & Lai 2013), to show that it is part of a typological asymmetry that applies to unbounded circumambient processes in general. The second contribution is to show that UTP and the sour-grapes pattern, by virtue of being unbounded circumambient, are *formally* similar. This is because for these processes, each target must 'look ahead' in either direction to see crucial information in the environment. It is argued that this places unbounded circumambient processes outside the WEAKLY DETERMINISTIC class of maps, a complexity class in formal language theory defined in terms of finite-state transducers. The third contribution is to offer an account of the unbounded circumambient asymmetry in terms of such classes of maps. Previous work has found segmental processes to be at most weakly deterministic (Chandlee *et al*. 2012, Heinz & Lai 2013, Chandlee 2014, Payne 2014). The unbounded circumambient asymmetry can thus be captured in terms of a complexity bound on segmental phonology: segmental phonology is restricted to weakly deterministic maps, but tonal phonology is not.

The structure of this paper is as follows. §2 establishes the unbounded circumambient asymmetry between tonal and segmental phonology. §3 introduces the formal language theory notions of complexity, and reviews previous work applying them to phonology. §4 shows that unbounded circumambient processes are not subsequential, and §5 argues that they are not weakly deterministic. This leads to the computational characterisation proposed in this paper: segmental processes are at most weakly deterministic, but tone is not restricted in this way. It is also discussed how two potential exceptions to this characterisation, Sanskrit *n*-retroflexion and Yaka vowel harmony, fit into this proposal. §6 discusses how Optimality Theory does not offer a unified way of characterising the typological asymmetry. §7 concludes, and mathematical definitions and a

proof that UTP is neither left-nor right-subsequential are given in an online appendix.

2 The unbounded circumambient asymmetry

This section defines unbounded circumambient processes, and shows how they are common in tonal phonology, but extremely rare in segmental phonology. A circumambient process is one whose application is dependent on the existence of triggers or blockers on both sides of a target; an unbounded circumambient process is one in which there is no bound, on either side, on the distance between these triggers/blockers and the target. These terms are discussed in detail below in §2.1.

The bulk of the evidence for the asymmetry comes from unbounded tonal plateauing. §2.2 surveys attestations of UTP in the tonal literature, as a thorough documentation of Hyman's claim. §2.3 reviews two known segmental unbounded circumambient processes: mid-vowel harmony in Yaka and *n*-retroflexion in Sanskrit. §2.4 and §2.5 summarise related generalisations in other typological work which support the conclusion that unbounded circumambient processes are rare in segmental phonology, and discuss why Sanskrit and Yaka are exceptional for these generalisations as well. Thus the asymmetry is not confined to UTP but rather to the *class* of unbounded circumambient processes. This is highlighted in §2.5 by way of the sour-grapes vowel-harmony pathology, which in Copperbelt Bemba H-spreading (Bickmore & Kula 2013, Kula & Bickmore 2015) has an attested correlate in tone.

2.1 Unbounded circumambient processes

A precise definition of an unbounded circumambient process is given in (2). Crucially, this property is atheoretical, and thus agnostic to specific theories of representation and processes.

(2) An unbounded circumambient process is a process:
    a. whose application is dependent on information (i.e. the presence of a trigger or blocker) on both sides of the target;
    b. in which, on both sides, there is no bound on how far this information may be from the target.

To illustrate the concept, a rewrite rule representation of this type of process is given in (3). In (3), X and Y are non-empty, they surround the target and there is no bound on the distance between them.

(3) $A \rightarrow B \mid X(U)\_\_(V)Y$
    (X and Y are non-empty; U and V may be of any length)

An example is the possible rule in (4).

(4) $[+\text{syll}] \rightarrow [+\text{back}] \;/ \begin{bmatrix} +\text{syll} \\ +\text{back} \end{bmatrix} (U) \underline{} (V) \begin{bmatrix} +\text{syll} \\ +\text{back} \end{bmatrix}$

(*U* and *V* are any strings of segments)

In (4), whether a vowel becomes [+back] depends on the presence of [+back] vowels on both sides of the target. However, these [+back] vowels may be separated from the target by strings *U* and *V*. Because *U* and *V* can be of any length, the process must be able to 'look back' for a [+back] segment over any distance in the left context (i.e. over *U*) and 'look ahead' over any distance for a [+back] vowel in the right context (i.e. over *V*).

Importantly, *X* or *Y* might contain 'blocking' information which prevents the process from applying. A possible such rule is given in (5).

(5) $[-\text{son}] \rightarrow [+\text{nas}] \;/\; [+\text{nas}]\,(U)\underline{}$
    (except in: [+nas] (*U*) __ (*V*) [+nas])

Here, the target will nasalise only when it is preceded by a [+nasal] segment and not followed by another [+nasal] segment. While (5) is given as a rule, such conditions are more intuitively expressed with Optimality Theory constraints (Prince & Smolensky 1993).

Furthermore, the definition of unbounded circumambient applies also to autosegmental representations. From an autosegmental standpoint, the 'target' referred to in (2) is any unit affected by the changing of association lines, and the distance from the 'trigger' is measured on the timing tier. These choices will become clear in §4.4.

Because 'unbounded' is critical to the definition in (2), we must set the criteria for an unbounded process. Intuitively, an unbounded process is one which operates over multiple units, like segments or tone-bearing units (TBUs), for which the correct generalisation does not refer to a bound on how many units over which it may operate. As linguists may differ as to what constitutes evidence for a process being 'unbounded', I use the criteria in (6).

(6) a. The source authors characterise the process as unbounded, and there is no evidence to the contrary.
  b. Examples exist of the process operating over multiple units.
  c. Examples exist of the process applying even when productive word- or phrase-formation processes extend its domain.

On its own, criterion (6a) does not constitute strong evidence for a process being unbounded, and thus no processes are included here which only meet (6a).
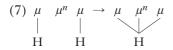
For some, (6b) is enough evidence that a process is unbounded, especially if there are examples of it operating over three or more units – as Kenstowicz (1994: 372) puts it, 'phonological rules do not count past

two'. In contrast, other researchers may consider only (6c) sufficient. However, the evidence presented here shows there is a typological asymmetry, regardless of which criterion one considers. By either (6b) or (6c), unbounded circumambient processes are far more common in tonal phonology than in segmental phonology.

2.2 Unbounded tonal plateauing

Having established the criteria for classifying a process as unbounded and circumambient, we can now discuss individual unbounded circumambient processes. This section surveys eight languages with some form of unbounded tonal plateauing, in which any number of L-toned or unspecified (Ø) TBUs (the TBU will be assumed to be the mora) surface as H if they are between two Hs, but as L otherwise.

Kisseberth & Odden (2003: 67) motivate UTP as a repair for a constraint against 'toneless moras between Hs'. Hyman & Katamba (2010) formalise UTP in Luganda (which they refer to as 'H-tone plateauing') as in (7).

$$(7) \quad \mu \quad \mu^n \quad \mu \quad \rightarrow \quad \mu \quad \mu^n \quad \mu$$
$$\qquad | \qquad\quad | \qquad\qquad \diagdown\!\!\diagup$$
$$\qquad H \qquad H \qquad\qquad H$$

UTP thus fits the definition in §2.1 of an unbounded circumambient process, because whether or not the process applies depends on two Hs that (a) are on both sides of the affected TBUs, and (b) can be arbitrarily far away from any one of the affected TBUs.[1]

Using data from a number of sources, the following subsections confirm Hyman's (2011) claim that UTP is a well-attested tonal process, and that by criterion (6b) it is unbounded in all cases and by (6c) in most.

2.2.1 *Luganda*.   UTP occurs in Luganda (Bantu, Uganda; Hyman *et al*. 1987, Hyman & Katamba 2010) both word-internally and in the phrasal phonology. Luganda TBUs can be either H or unspecified underlyingly. Lexically, a L tone is inserted after a H, causing a falling HL contour when the H is on a final syllable. In some cases, these intermediate L tones contrast with Ø, but not with regard to UTP (technically, they delete in the UTP environment), so they will not factor into the discussion. Luganda nouns are given in (8) (all Luganda forms in this section are taken from Hyman & Katamba 2010).[2] Underlying and surface H tones are indicated; underlying unspecified and surface L tones are unmarked.

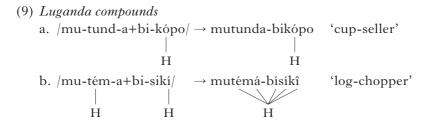[1] This can be contrasted with bounded plateauing, in which only one Ø TBU becomes H in between two Hs. This pattern is attested, for example, in Kihunde (Goldsmith 1990).

[2] The tones transcribed here are for what Hyman & Katamba call the 'intermediate' forms, before the superimposition of phrasal boundary tones. The boundary tones are inserted late in the derivation and play no role in the phonology.
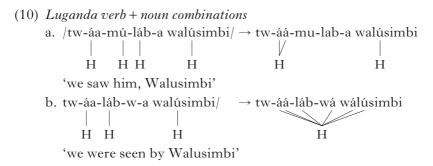
(8) *Luganda nouns*

/ki-kópo/   kikópo   'cup'
/ki-sikí/   kisikî   'log'
/ki-tabo/   kitabo   'book'
/mu-tund-a/ mutunda 'seller'   (/-tund-/ 'to sell')
/mu-tém-a/   mutéma 'chopper' (/-tém-/ 'to chop')

If a toneless noun and a noun with an underlying H tone form a compound, they are pronounced as in isolation, as in (9a). However, when both nouns have an underlying H, a H-tone plateau occurs between them, as in (9b). In compounds, the toneless plural /bi-/ is used for nouns from the /ki-/ noun class IV.

(9) *Luganda compounds*

  a. /mu-tund-a+bi-kópo/ → mutunda-bikópo   'cup-seller'
             | |
             H   H

  b. /mu-tém-a+bi-sikí/   → mutémá-bísíkî   'log-chopper'
       | | \|/
      H   H   H

In (9b) there is a plateau over three unspecified TBUs. This shows target toneless TBUs three TBUs away from their triggers – the mora represented by /a/ in /mu-tém-a/ is two TBUs from its right trigger H /í/ in /bi-sikí/, and the second /i/ in /bi-sikí/ is three TBUs from its left trigger /é/ in /mu-tém-a/. Thus Luganda satisfies the 'multiple unit' criterion for unboundedness in (6b).

UTP operates over syntactic phrases as well, satisfying (6c). Under certain morphosyntactic conditions, noun–verb sequences can also form a domain for UTP. The noun in (10a) is an adjunct, and thus outside of the domain of plateauing. In (10b), however, the noun and verb form a phonological phrase, and plateauing occurs across all TBUs between the surviving Hs.

(10) *Luganda verb + noun combinations*

  a. /tw-áa-mú-láb-a walúsimbi/ → tw-áá-mu-lab-a walúsimbi
         | | | | \|/ |
        H   H H   H   H   H
    'we saw him, Walusimbi'

  b. tw-áa-láb-w-a walúsimbi/   → tw-áá-láb-wá wálúsimbi
         | | | \|/
        H H   H   H
    'we were seen by Walusimbi'

A more extreme example can be seen when proclitics are added to the noun. Proclitics generally do not have tone, as (11) shows.

(11) *Luganda toneless proclitics*
    a. byaa=ba=mulondo   '(it's) those of the Mulondos'
    b. na=ku=byaa=ba=mulondo 'and on those of the Mulondos'

Luganda UTP 'sees' over these as well, no matter how many are stacked on to the noun, as in (12).

(12) *Luganda verb + noun sequences with proclitics*
    a. /tw-áa-láb-a byaa=walúsimbi/ → tw-áá-láb-á byáá-wálúsimbi

      'we saw those of Walusimbi'
    b. /tw-áa-génd-a na=byaa=ba=walúsimbi/

          → tw-áá-génd-á ná=byáá=bá=wálúsimbi

      'we went with those of Walusimbi'

These examples show UTP operating over a six-TBU span of toneless TBUs created by prefixation, illustrating triggers separated from their targets by five TBUs on each side. Thus Luganda UTP satisfies both criterion (6c), with respect to unboundedness, and criterion (6b), applying across five TBUs.
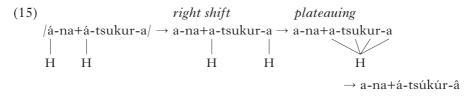
2.2.2 *Digo*. Verbs in Digo (Bantu, Kenya and Tanzania; Kisseberth 1984) show complex interactions between underlying H tones, including UTP. Underlyingly, Digo has a privative H/Ø system. Both verb roots and affixes may carry a H tone, although this is not obligatory. A single H tone in a verbal H shifts to the end of the word, surfacing as a rising/ falling pattern on the final two TBUs. This is illustrated with the addition of the 3rd person plural object prefix /á/ to the toneless root /tsukur/ 'take' in (13b) (the underlying position of the H is underlined).

(13) *Digo*
    a. ni-na+tsukur-a   'I am taking'
    b. ni-na+<u>a</u>-tsukǔr-â 'I am taking them'
    c. <u>a</u>-na+<u>á</u>-tsúkúr-â   's/he is taking them'

An autosegmental analysis of (13b) is given in (14).

(14) /ni-na+á-tsukur-a/ → ni-na+a-tsukŭr-â
        | \/
        H   H

The forms in (13a, b) have the toneless 1st person prefix /ni/; if this is substituted for the H-toned 3rd person singular subject prefix /á/, a plateau occurs from the object prefix to the end of the root, as in (13c). Again, we see the presence of two H tones creating a long-distance plateau across the length of the root. The Digo example is complicated by tone shifting; Kisseberth analyses it as a two-step process, in which the first H shifts to the initial vowel of a 'verbal complex' (marked with '+'), the second H shifts to the end of the word (as in (14)) and then the first H triggers a plateau across the root between them. This analysis is illustrated with the derivation in (15).

(15)   *right shift*   *plateauing*
   /á-na+á-tsukur-a/ → a-na+a-tsukur-a → a-na+a-tsukur-a
    | | | | \|/
    H   H   H   H   H

                                              → a-na+á-tsúkúr-â

The domain for both right shift and plateauing is larger than the verb – they also apply to verb + noun constructions. (16a) shows right shift applying when there is only one H in the verb, and (b) shows plateauing over the phrase resulting from two Hs associated with the verb.
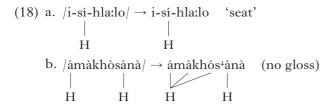
(16) a. ku+<u>a</u>fŭnâ   'to chew'
       naz<u>i</u>   'coconut'
       ku+<u>a</u>fun-a nǎzî   'to chew a coconut'
   b. a-k<u>a</u>+ts<u>ú</u>kú ts-â   'he has cleaned'
       chi-ronda   'wound'
       a-k<u>a</u>+ts<u>ú</u>kú ts-á chí-róndâ 'he has cleaned a wound'

The phrase in the final form in (16b) results in a plateau over six TBUs, which shows target TBUs three TBUs away from their triggers, and thus satisfies criterion (6b) for unboundedness. It also contains a plateau created over a syntactic domain, satisfying criterion (6c).

2.2.3 *Other languages.* Kisseberth & Odden (2003) cite a plateauing process in Xhosa which is similar to that in Digo. In Xhosa, underlying H tones shift to the antepenult. A phrase with two H tones shows a plateau between the first and the second H, and shifted H on the antepenult, as in (17) (data from Kisseberth & Odden (2003: 67–68).

(17) <u>u</u>-ku-qononóndiis-a   'to emphasise'
    ndi-f<u>ú</u>n-a   'I want'
    ndi-f<u>ú</u>n <u>ú</u>-kú-qónónóndiis-a 'I want to emphasise'

The second H shifts to the third /o/ in /ú-ku-qononondis-a/, and becomes the right trigger for UTP; this shows triggers on both sides affecting targets four TBUs away. In Zulu, related to Xhosa (Laughren 1984, Cassimjee & Kisseberth 2001, Downing 2001), a single H shifts to the antepenult (if it originates on a prefix) or penult (if it originates on a stem). In forms with two Hs, a plateau forms between the Hs. In Zulu, the two H tones do not fuse; instead, as shown in (18), the first spreads up to the second, creating downstep.

(18) a. /i-sí-hlaːlo/ → i-sí-hlaːlo   'seat'
             | |
            H   H

      b. /ámàkhòsánà/ → ámákhós⁺ánà   (no gloss)
          | | | |
         H   H   H   H

Laughren (1984: 221) states that 'the rule only applies to a H which is followed by a LH tonal sequence' (Yip 2002 represents the LH as a single H, which is then downstepped), and, while she only gives examples of the plateau operating over two TBUs, analyses the process as operating over an arbitrary number of TBUs.

Other examples of UTP can be found throughout Bantu. In Yaka (Kidima 1991), 'all toneless syllables flanked by Hs become H by right-ward spreading of the H to the left of the domain' (Kidima 1991: 44). (19), from Kidima (1991: 180), shows plateauing alternations. Kidima gives Yaka tonal assignment a complex accentual analysis; the underlying Hs marked in the following data result from a tonal assignment rule.[3] In the following examples, *bakhoko* is not assigned a tone, prompting the alternations.

(19) a. bakhoko ba ngwaási   'Uncle's chickens'
      b. bakhoko ba kabeénga   'the red chickens'
      c. bakhoko ba kabeéngá bá ngwáási Málóóngi 'Uncle Maloongi's
                                                    red chickens'

(19c) shows two plateaus, each with their target TBUs separated from their triggers by two TBUs.

UTP also occurs outside of Bantu. In Saramaccan, a creole spoken in Suriname (Roundtree 1972, Good 2004, McWhorter & Good 2012), a phrasal version of UTP occurs across words in certain syntactic configurations. This analysis follows Good (2004), who posits an underlying H/L/Ø distinction, as in the nouns in (20).

---

[3] A surface distinction between regular H and raised H – the latter occurring on accented syllables with an associated H – is ignored in these transcriptions.

(20)　/wómi/　　wómì　　'man'
　　　/mujέɛ/　　mùjέɛ̀　　'woman'
　　　/wajamáka/　wàjàmákà　'iguana'
　　　/sὲmbὲ/　　sὲmbὲ　　'person'
　　　/amɛɛká/　　àmὲɛ̀ká　　'American'

According to Roundtree (1972: 314), 'all changeable low tones [Good's Ø TBUs] between the highs in successive morphs in certain syntactic positions are changed to high'. One such syntactic position is an adjective–noun sequence. In the following, the final /o/ in /hánso/ is realised as low [ó] before /sὲmbὲ/ (21a), but as high before /wómi/ (21b) and /mujέɛ/ (21c). Relevant vowels are underlined.

(21)　a.　/dí hánso sὲmbὲ/　　→ dí háns<u>ò</u> sὲmbὲ
　　　　　'the handsome person'　　　　　　　(Roundtree 1972: 315)
　　　b.　/dí hánso sὲmbὲ wómi/　→ dí háns<u>ó</u> wómì
　　　　　'the handsome man'　　　　　　　　(Roundtree 1972: 324)
　　　c.　/dí hánso sὲmbὲ mujέɛ/　→ dí háns<u>ó</u> m<u>ú</u>jέɛ̀
　　　　　'the handsome woman'　　　　　　　(Roundtree 1972: 316)
　　　d.　/dí wajamáka=dé á óbo/　→ dí w<u>á</u>j<u>á</u>máká=dé á óbo
　　　　　'the iguana there has eggs'　　　　　　(Good 2004: 28)
　　　e.　/dí taánga amɛɛká wómi/ → dí tàáng<u>á</u> <u>á</u>m<u>ɛ́ɛ́</u>ká wómì
　　　　　'the strong American man'　　　(McWhorter & Good 2012: 48)

Note that in (21b) the initial /u/ of /mujέɛ/ also surfaces as H, illustrating plateauing over two TBUs. In (e) a plateau occurs over four TBUs, the final /a/ of /taánga/ and the first three vowels of /amɛɛká/; Saramaccan UTP thus satisfies (6b), with its targets and triggers separated by three TBUs on each side.

Similar to the Saramaccan system, as noted by Good (2004), is the intonational phonology of the Uto-Aztecan language Papago (Hale & Selkirk 1987: 152), in which H tones are associated 'to each stressed vowel and to all vowels in between'; they list examples of plateaus created in between stressed vowels three TBUs apart. Finally, Do & Kenstowicz (2011) discuss plateauing between Hs in certain intonational phrases in South Kyungsang Korean, giving a spectrogram showing a plateau between the two H tones in [seccók khaliphonía] 'Western California' (2011: 3, 11), which shows unspecified TBUs affected by triggers two TBUs away (on either side).

2.3　Unbounded circumambient processes in segmental phonology

Turning to segmental phonology, there are the only two potential attestations of unbounded circumambient processes of which I am aware. One is Sanskrit *n*-retroflexion. The other is 'plateauing harmony' in Yaka, which also has UTP, as explained in §2.2. Both satisfy the unboundedness criterion in (6b), but neither satisfy (6c).

2.3.1 *Sanskrit*. In Sanskrit *n*-retroflexion (Whitney 1889, Macdonell 1910, Schein & Steriade 1986, Hansson 2001, Graf 2010, Ryan forthcoming), an underlying alveolar /n/ becomes retroflex [ɳ] after retroflex /r ʂ/, which can appear far to the left of the target /n/. Both trigger and target in (22), from Hansson (2011: 225), are underlined.

(22) a. /iʂ-n̠aː-/   iʂ-ɳaː-   'seek (PRES STEM)'
     b. /cakʂ-aːn̠a-/   cakʂ-aːɳa-   'see (MIDDLE PART)'
     c. /kr̥p-a-maːn̠a-/   kr̥p-a-maːɳa-   'lament (MIDDLE PART)'

One restriction on *n*-retroflexion potentially gives the process an unbounded circumambient quality. Hansson (2001: 230) states that retroflexion fails 'when there is also an /ʂ/ or /r/ *later* in the word'.[4] He cites the data in (23), from Macdonell (1910) and Monier-Williams (1899).[5]

(23)   *attested*   *unattested*
     a. pr̠aː-nr̥tyat   *praː-ɳr̥tyat   *from* -nrt- 'dance'
     b. par̠i-nakʂati   *pari-ɳakʂati   'encompasses'
     c. -n̠iʂt̠ʰaː-   *-ɳiʂtʰaː-   'eminent'
     d. -n̠iʂʂidʰ-   *-ɳiʂʂidʰ-   'gift'
     e. -n̠irɳiéa-   *-ɳirɳiɟa-   'adornment'
     f. -n̠rmɳa-   *-ɳrmɳaː-   'manhood'
     g. pr̠a-naŋkʂyati   *pra-ɳaŋkʂyati   'causes to disappear'

Given the data in (22), one would expect, for example, that the underlying /n/ in [pari-nakʂati] would surface as [ɳ], as in (23b), because it follows a trigger /r/ for n- retroflexion. However, it instead surfaces as [n], which is attributed to the 'blocking' /ʂ/ three segments to the right. Thus Sanskrit *n*-retroflexion fits the definition for circumambience in a slightly different way than UTP: the crucial context includes the presence of a trigger on the left side and a blocker on the right. The evidence in the examples here also fit criterion (6b) for unboundedness; (23g) shows a blocker three segments to the right of the target, and (22c) shows a trigger five segments to the left. Thus, by criterion (6b), Sanskrit *n*-retroflexion is an unbounded circumambient process. However, Ryan's (forthcoming) study of several Sanskrit corpora finds no examples of retroflexion being blocked when a long vowel or multiple syllables intervene between the target and blocker. Thus, whether or not it is truly unbounded is the subject of some doubt. Furthermore, there is no evidence satisfying (6c), i.e. that the distance this second blocker can be from the target may be extended by a morphological or syntactic process.

  [4] There are additional restrictions; see Ryan (forthcoming) for a detailed discussion of Sanskrit retroflexion.
  [5] Many thanks to Kevin Ryan for drawing my attention to the final form in (23b) and for information about Sanskrit in general.

2.3.2 *Yaka vowel harmony.*   The other attested unbounded circumambient segmental process is vowel-height harmony in Yaka (Hyman 1998, 2011). Hyman (2011) cites Yaka as a unique example of vowel 'plateauing'. In Yaka, the initial vowel of the perfective suffix /ile/ lowers to mid [e] when the vowel in the stem is also mid (24c, d). Otherwise, a progressive harmony converts the final /e/ to [i] ([l] and [d] alternate depending on the following vowel, and with [n] following a root nasal; hence the allomorph [ene] in (25c) below). The lowering of this middle /i to [e] does not happen in the applicative suffix /ila/, which does not end in a mid vowel. The examples in (24), from Hyman (2011: 501), show this alternation.

(24)   root   *applic* /ila/   *perf* /ile/
   a. /kik/ 'obstruct'   kik-ila   kik-idi
   b. /kas/ 'bind'   kas-ila   kas-idi
   c. /keb/ 'pay attention'   keb-ila   keb-ele
   d. /sol/ 'clear brush'   sol-ila   sol-ele

That this is part of a more general process lowering high vowels to mid if and only if they are between two mid vowels can be seen in (25), from Hyman (1998: 19). This process can take place over at least three vowels.

(25)   stem   *FV* /a/   *perf* /ile/
   a. /hit-ik/   'send'   hit-ik-a   hit-ik-idi
   b. /bet-ilik/ 'lower'   bet-idik-a   bet-elek-ele
   c. /kel-umuk/ 'do an about-face'   kel-umuk-a   kel-omok-ene

As in UTP, the mid-vowel plateau shows triggers on both sides of the target vowels. (25b) and (c) thus show triggers two vowels from the right and two vowels from the left from their targets. As such, Yaka vowel harmony satisfies criterion (6b) for an unbounded process. However, according to Hyman (1998), plateauing alternations can only be seen with verb roots and the perfective /-ile/; it thus does not satisfy (6c).

2.4 Empirical summary

The preceding sections presented eight attestations of UTP, an unbounded circumambient process in tonal phonology, and two examples of separate unbounded circumambient processes in segmental phonology. All satisfied criterion (6b), with the greatest distance between target and trigger being five TBUs. Seven examples of UTP were shown to operate over domains extended by morphology or syntax, thus satisfying criterion (6c). Neither segmental process satisfied (6c), and in Yaka, the distribution of the process was quite limited. Both satisfied (6b), although in Yaka the greatest attested distance between trigger and target was three vowels, and in Sanskrit there was no evidence for blocking beyond the syllable following the target.

Most importantly, to the best of my knowledge, these are the only examples of such processes in segmental phonology. This is notable, given the

wide attestation of long-distance segmental processes in general, as documented in the comprehensive surveys on feature-spreading harmony (Rose & Walker 2004), vowel harmony (Baković 2000, Nevins 2010), consonant harmony (Rose & Walker 2004, Hansson 2001, 2010) and consonant disharmony (Suzuki 1998, Bennett 2013). In fact, Wilson (2003, 2006b), reviewing the typologies of nasal, emphasis and vowel harmonies, characterises segmental spreading processes as 'myopic'. This means that even segmental spreading processes which affect multiple segments proceed in a local fashion, never 'looking ahead' beyond immediately adjacent segments. As unbounded circumambient processes by definition depend on information unboundedly far away from the target, any myopic process is not an unbounded circumambient process, and any unbounded circumambient spreading process is necessarily not myopic. Yaka vowel harmony and Sanskrit *n*-retroflexion are thus exceptions to the myopic spreading generalisation; this highlights how atypical these two processes are.[6]

There is thus a typological asymmetry between tonal and segmental processes: unbounded circumambient patterns are extremely rare in segmental processes, but widely attested in tonal phonology, as evidenced by the variants of UTP discussed. A comparison between the proportions of circumambient unbounded processes found in typological surveys (such as those just mentioned) of segmental and tonal processes would be ideal, but comparable surveys of tonal processes, or even particular kinds of processes, do not exist (to the best of my knowledge). Regardless, the evidence reviewed in this paper clearly shows an asymmetry, despite the absence of such surveys for tone.

It should be noted that bidirectional spreading processes, in which a feature spreads outwards from a single trigger in two directions, are common in segmental harmony. An example is Arabic emphasis spreading, in which an emphatic gesture spreads in an unbounded fashion in both directions from an underlying emphatic segment. In the following example from Southern Palestinian Arabic (Al Khatib 2008: 2), emphasis spreads to the left (26a), to the right (b) and to both the left and right (c) (spreading is blocked by high front segments, such as /j/).[7]

(26) a. /balːaːsˤ/   bˤaˤlˤːaˤːsˤ   'thief'
     b. /sˤajːaːd/   sˤaˤjːaːd   'hunter'
     c. /ʔatˤfaːl/   ʔaˤtˤfˤaˤːlˤ   'children'

Other bidirectional processes are nasal spread in Capanahua and Southern Castillian (Safir 1982) and cases of stem control in vowel harmony (Baković 2000) and consonant harmony (Hansson 2001, 2010). While these processes apply in an unbounded fashion, and operate in two directions, they only have one trigger, and are thus not circumambient.

---

[6] See Gafos (1999) and Hansson (2001) for arguments that *n*-retroflexion is best analysed as spreading.

[7] Thanks to an anonymous reviewer for providing this example.

2.5 The sour-grapes pattern in tonal phonology

One final piece of evidence for the unbounded circumambient asymmetry involves the unattested 'sour-grapes' vowel-harmony pattern (Baković 2000, Wilson 2003, McCarthy 2010, Heinz & Lai 2013).[8] The sour-grapes pattern is a non-myopic harmony pattern predicted to exist by ranking permutations of classic OT with AGREE constraints, but not attested in segmental harmony. The sour-grapes pattern is also unbounded circumambient, and, as will be discussed momentarily, a sour-grapes-like process appears in tonal phonology.

Sour-grapes harmony works as follows. Given a spreading [+F] feature (which targets underlying [−F] segments) and a blocking feature [!F], [−F] segments become [+F] after another [+F], provided that there is no blocking feature following in the word, as shown in (27).

(27) a. $[-F]^n \rightarrow [-F]^n$ no trigger, no blocker → no harmony
       b. $\ldots [+F][-F]^n \rightarrow \ldots [+F][+F]^n$ trigger, no blocker → harmony
       c. $\ldots [+F][-F]^n[!F] \ldots \rightarrow \ldots [+F][-F]^n[!F] \ldots$
                                             trigger, blocker → no harmony

In other words, [+F] spreads if and only if it can spread all the way.[9] Note that sour-grapes spreading is not myopic: the spread of the [+F] has to 'look ahead' to see if there is a blocker before it can apply. As such, it is also an unbounded circumambient pattern, because the presence or absence of both [+F] and [!F] segments, which can be any distance apart, bears on the realisation of [−F] segments in between (this blocking aspect makes it circumambient in a similar way to Sanskrit *n*-retroflexion).

However, a sour-grapes-like pattern does exist in tonal phonology. In Copperbelt Bemba (Bickmore & Kula 2013, Kula & Bickmore 2015), underlying H tones undergo one of two spreading processes, bounded ternary spreading or unbounded spreading. The latter is blocked by the presence of another H tone, as shown in (28). In phrase-final forms, unbounded spreading applies to the rightmost H.

(28) a. /u-ku-tul-a/   ù-kù-tùl-à
                                    'to pierce'
       b. /bá-ka-fik-a/   bá-ká-fíká
                                     'they will arrive'
       c. /bá-ka-mu-londolol-a/   bá-ká-mú-lóóndólól-á
                                     'they will introduce him/her'
       d. /tu-ka-páapaatik-a/   tù-kà-páápáátík-á
                                     'we flatten'

---

[8] The term 'sour grapes', originally due to Padgett (1995), refers to behaviours under certain formulations of OT in which either all features in a particular domain assimilate or none do. See also McCarthy (2010).

[9] This characterisation is thanks to an anonymous reviewer.

Bounded spreading occurs when another H appears to the right, as in (29). Bounded spreading obeys the Obligatory Contour Principle; it will spread up to two additional TBUs, maintaining at least one L TBU before the second H. All other intervening TBUs surface with a L tone.

(29) a. /bá-ka-pat-a kó/   bá-ká-pát-à kó
                                 'they will hate'

    b. /bá-ka-londolol-a kó/   bá-ká-lóòndòlòl-à kó
                                 'they will introduce them'

    c. /tu-ka-béleeŋ-el-an-a kó/   tù-kà-bélééŋ-él-àn-à kó
                                 'we will read for each other'

    d. /tu-ka-lás-a kapembuá/   tù-kà-lás-á kápèèmbwá
                                 'we will hit Kapembwa'

    The formalisations in (30) summarise the facts. When no Hs are present, as in (a), all TBUs surface as L (cf. (28a): [ù-kù-tùl-à]). When one H is present, it spreads to all remaining TBUs in the domain (and the rest surface as L, as in (28d): [tù-kà-páápáátík-á]). When two Hs are present, the first only spreads to the next two TBUs (cf. (29c): [tù-kà-bélééng-él-àn-à kó]).

(30) a. $\mu^n \rightarrow \mu^n$ (cf. (28a))   b. $\mu^m\mu\mu^n \rightarrow \mu^m\mu\mu^n$ (cf. (28d))
         | | | ⫽
        L   H   L H

    c. $\mu^m\mu\mu^n\mu \rightarrow \mu^m\mu\mu\mu\mu^{n-2}\mu$ (cf. (29c))
       | | | ⫽ | |
      H H   L H L H

    Note that, modulo the bounded spreading, the formalisations in (30) are almost identical to the sour-grapes generalisations in (27). In other words, in Copperbelt Bemba the second H can be seen as a blocker for unbounded spread. This makes it an unbounded circumambient process, because the realisation of unspecified TBUs depends on the presence or absence of Hs on both sides which can be arbitrarily far away. As a tonal process, it does not seem particularly aberrant, in contrast to Sanskrit *n*-retroflexion and Yaka vowel harmony. Thus it provides more evidence for the unbounded circumambient asymmetry.

3 The computational complexity of phonological maps

A relevant measure of computational complexity can be found in formal language theory (FLT). This section introduces FLT, shows how it relates to phonology, and introduces properties of finite-state transducers that allow us to measure the relative complexity of phonological processes.

3.1 Formal language complexity and cognitive complexity

A FORMAL LANGUAGE is a set of strings, and FLT studies the relationships between formal languages and the expressive power of grammars that

describe them. FLT characterisations of natural language patterns have been argued to reflect domain-specific cognitive biases for and against patterns of a certain level of complexity. For example, the REGULAR class of formal languages is insufficient to describe English syntax, and English syntax is at least CONTEXT-FREE (Chomsky 1956). Phonology, on the other hand, appears to be at most regular (Johnson 1972, Kaplan & Kay 1994).[10] This notion of complexity has been explicitly linked to cognitive complexity (Folia *et al.* 2010, Rogers & Hauser 2010, Rogers & Pullum 2011), and results from artificial language learning experiments provide evidence in support of the psychological reality of the context-free/regular division between syntax and phonology (Lai 2015).

The primary goal of this paper is to use this notion of complexity to characterise the unbounded circumambient asymmetry discussed in §2. Instead of formal languages, however, the following sections study the relationships between regular maps, where a MAP is a relation between strings in which an input string is paired with at most one output string.[11] Analogous to the regular/context-free distinction for formal languages, maps can be classified according to their complexity. The distinctions important for this paper centre around the property of SUBSEQUENTIALITY (Mohri 1997), which can be defined in terms of FINITE-STATE TRANSDUCERS (FSTs). FSTs are idealised machines that match pairs of strings; they are described in more detail below. While FSTs in general can describe any regular map, subsequential FSTs describe a more restricted set of maps.
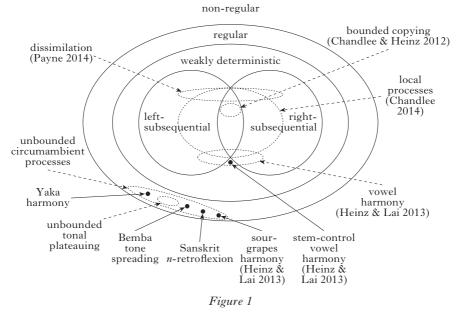
3.2 Overview: formal language complexity and phonology

The classes of interest in this paper are LEFT- and RIGHT-SUBSEQUENTIAL maps, WEAKLY DETERMINISTIC maps and REGULAR maps. Figure 1 depicts the relationships between these classes as a nested hierarchy, in which more complex classes properly include lesser ones. Left- and right-subsequential maps are provably less complex than weakly deterministic maps (Heinz & Lai 2013) and regular maps (Mohri 1997), and weakly deterministic maps have been conjectured to be less complex than regular maps (Heinz & Lai 2013). Assuming this conjecture to be true (see further below), this means that all weakly deterministic maps are also regular maps, but not all regular maps are weakly deterministic. Regular maps conjectured not to be weakly deterministic can be referred to as FULLY REGULAR.

Figure 1 also depicts where typological research examining segmental processes as maps place these processes in the complexity hierarchy. This work has found all of them to be within the weakly deterministic

---

[10] For more on the regular/non-regular split between phonology and syntax, see Heinz (2011) and Heinz & Idsardi (2011, 2013).

[11] Cases of free variation are thus not maps, as one UR can be paired with multiple SRs. The focus of the present study is on subclasses of maps, and so such cases will not be considered, but (finite) free variation can be studies in a similar way with the p-subsequential transducers of Mohri (1997) or the semi-deterministic transducers of Beros & de la Higuera (2014).

Figure 1

Phonological processes in the subregular hierarchy.

class, with most in the less complex left- and right-subsequential classes. This has led to the hypothesis that that the weakly deterministic class forms a bound on the complexity of phonology. As Heinz & Lai (2013) discuss, this hypothesis is supported by the absence of sour-grapes patterns, which they prove to be neither left- nor right-subsequential, and conjecture to not be weakly deterministic.

In addition, Fig. 1 shows the other unbounded circumambient processes discussed in this paper belonging to the fully regular region. As will be explained in detail in §4, this is because UTP and other unbounded circumambient processes are, like sour-grapes harmony, neither left- nor right-subsequential, and there is no known weakly deterministic characterisation of them. By revising the above hypothesis, then, we have a characterisation for the typological asymmetry established earlier: segmental phonology, but not tonal phonology, is weakly deterministic. The existence of Sanskrit *n*-retroflexion and Yaka harmony are admittedly exceptions, but as already discussed in §2, they are not only rare, but also exceptions to Wilson's myopia generalisation. These exceptions will be discussed in §5.3.

3.3 Finite-state transducers, subsequentiality and determinism

This subsection gives an informal introduction to FSTs, and discusses subsequential transducers, a type of FST which is strictly less expressive than the fully regular non-deterministic transducers (Mohri 1997). As an
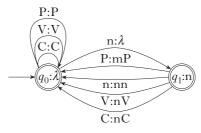
Figure 2

A deterministic finite-state transducer for nasal place assimilation.

example, consider the regressive assimilation generalisation 'nasals become labial before a labial consonant'. A rule for this generalisation is given in (31).

(31) [+nasal] → [labial] / __ [labial]

To simplify the notation, let us abstract away from place features other than [labial] as follows. 'C' represents non-labial, non-nasal consonants and 'P' labial consonants. A FST corresponding to the generalisation in (31) is given in Fig. 2. It comprises a set of STATES (the circles labelled with a q) and TRANSITIONS between the states (the labelled arrows). The FST can be interpreted as reading an input and writing an output as follows. Beginning from the START STATE (marked with the unlabelled arrow; q_0 in Fig. 2), it traverses the transition for each symbol (indicated to the left of the colon in a transition label) in the input string, each time writing out the corresponding output (indicated to the right of each label).[12] At the end of the input string, it appends to the output the output for the current state (indicated to the right of the colon on the state label). The EMPTY STRING λ indicates that there is no output. In this way, the transitions and states define which input/output string pairs are accepted by a machine. As an example, the derivation in (32) shows how an input CVnC corresponds to an output CVnC (i.e. no change occurs). The derivation highlights each state, input symbol and output as the machine reads CVnC.

(32) *A derivation for* CVnC → CVnC *in the finite-state transducer in Fig. 2*

input	C		V		n		C		
state	q_0	→	q_0	→	q_0	→	q_1	→	q_0
output	C		V					nC	

Note that, while on the initial C and V the FST simply loops on state q_0, outputting C and V respectively, on the following 'n' it makes a transition

---

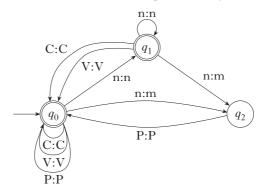[12] For a more detailed introduction to FSTs see Beesley & Karttunen (2003).

Figure 3
A non-deterministic finite-state transducer for nasal place assimilation.

to state q_1, outputting nothing. This is because the machine cannot yet 'decide' the output for this 'n', as it could be 'n' or 'm', depending on whether or not the following input symbol is a P. Let us thus call state q_1 a 'wait' state. As the next symbol is a C, the machine takes the C:nC transition from q_1 to q_0, outputting nC. If the input is instead CVnP, the derivation is similar, but in the last step the machine takes the P:mP transition from state q_1, as in (33).

(33) *A derivation for* CVnP → CVmP *in the finite-state transducer in Fig. 2*

input		C		V		n		P	
state	q_0	$\to$	q_0	$\to$	q_0	$\to$	q_1	$\to$	q_0
output		C		V				mP	

In this way, state q_1 allows the machine to 'look ahead' one symbol in the input. The reader can verify that this machine will take any permutation of C, V, n and P as an input, outputting 'n' as 'm' only before a P. Thus Fig. 2 describes exactly the map as (31).[13]

For all states in Fig. 2, there is only one possible transition, given a particular input symbol. This means the FST in Fig. 2 is deterministic. Not all FSTs are deterministic; it is possible to write a FST such that a state may have two separate transitions on a particular output. Figure 3 is such a machine. Here, there are no outputs specified for the states, and q_2 is a 'non-accepting' state (indicated by the single, rather than double, circle), meaning that the machine rejects an input/output pair for which the input ends on that state.

---

[13] This abstracts away from other constraints and phonological processes (for example, those preventing an output 'CCCC'). This is no different from modelling a specific phonological generalisation with a single *SPE*-style rewrite rule or partial OT constraint ranking.

Non-deterministic FSTs are more powerful than deterministic FSTs: any map describable by a deterministic FST can be described by a non-deterministic FST, but not *vice versa* (Mohri 1997). The ability to be described by a deterministic FST is central to the definition of the classes of left- and right-subsequential maps (Mohri 1997). The regular class of maps are those maps which can be described by any non-deterministic FST. Thus the left- and right-subsequential maps form a strict subset of the regular class.

UTP is a concrete example of a map which cannot be described with a deterministic FST, and will be described in detail in §4. For now, the restriction which determinism places on the subsequential maps can be understood in terms of 'wait' states, as in Fig. 2, which allow deterministic FSTs to look ahead in the input. Because the number of states is finite, there can only ever be a finite number of wait states, and so a deterministic FST has BOUNDED look-ahead. Non-determinism, in contrast, allows a FST to 'postpone' a decision about a particular input symbol indefinitely.

3.4 Subsequentiality and segmental phonology

Computational analyses of typologies of segmental processes have shown that they are describable with deterministic FSTs, with some variation regarding directionality. Mohri (1997) described two kinds of subsequential maps: right-subsequential and left-subsequential, which are describable by deterministic FSTs reading an input string left-to-right and right-to-left respectively. Collectively, they can be referred to as subsequential maps.

The discussion in the previous section suggests how local processes are subsequential. For a thorough survey of the subsequentiality of local processes, including epenthesis, deletion, metathesis, substitution and partial reduplication, see Chandlee (2014). Work on long-distance segmental processes such as vowel harmony (Gainor *et al.* 2012, Heinz & Lai 2013) and dissimilation (Payne 2014) has also found them to be largely left- or right-subsequential. To see how, imagine a long-distance progressive consonantal harmony process in which a feature [−F] becomes [+F] after some other consonant specified [+F], no matter how early in the word this consonant appeared (an example is /l/ in Yaka becoming [n] in suffixes attaching to a root containing a nasal). Such a map, given in (34), is describable with the deterministic FST in Fig. 4 (the tailed arrow indicates an input/output pair that belongs to a map).

(34) ...[+F]...[−F]... ↦ ...[+F]...[+F]...

This is impossible for a regressive harmony process, in which [−F] becomes [+F] before some [+F] segment an unspecified distance later in the word, as in (35).

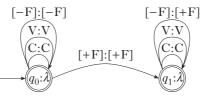(35) ...[−F]...[+F]... ↦ ...[+F]...[+F]...

Figure 4

A deterministic finite-state transducer for progressive harmony
(C = consonant, V = vowel segments not participating in harmony).

This requires unbounded look-ahead from left to right; in order to determine the output for a target $[-F]$ in the input, a FST reading from left to right would have to wait indefinitely to see if a trigger $[+F]$ appeared later in the string. However, this map is right-subsequential, because reading the input from right to left requires no look-ahead. Reading from right to left can be thought of as reversing the input, feeding it into the FST and then reversing the output (Heinz & Lai 2013). If we reverse the strings in the map in (35), we get the same map as in (34). We can thus describe the reverse of (35) with a deterministic FST, so it is right-subsequential.

Thus local processes and unidirectional long-distance processes are either left- or right-subsequential. However, there are two classes of phonological processes which are neither left- nor right-subsequential: unbounded circumambient processes and bidirectional spreading processes. These processes are the focus of the following two sections.

4 The non-subsequentiality of unbounded circumambient processes

This section provides proof that UTP is neither left- nor right-subsequential. This is compared to a similar result obtained for sour-grapes processes by Heinz & Lai (2013), and is then generalised to the class of unbounded circumambient processes. Additionally, §4.4 defends the use of a linear representation of UTP.

4.1 Unbounded tonal plateauing as a map

To analyse UTP using the computational framework for studying string maps outlined above, this section uses a string-based representation which marks associations to H tones on each TBU. The use of this representation will be discussed in §4.4. For now, we can say that if UTP is viewed with this kind of string-based representation, then it is neither left- nor right-subsequential.

The unbounded tonal plateauing generalisation was formalised in (7) in §2. Its string-based counterpart is given in (36), where H represents a TBU associated to a H tone, and Ø represents an unspecified TBU. Superscript m, n and p represent any natural number.
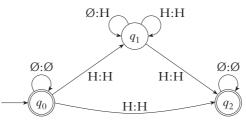
Figure 5

A non-deterministic finite-state transducer for unbounded tonal plateauing.

(36) a. $\varnothing^n \qquad\qquad \mapsto \varnothing^n$
　　b. $\varnothing^m H \varnothing^n \qquad\qquad \mapsto \varnothing^m H \varnothing^n$
　　c. $\varnothing^m H \{\varnothing, H\}^n H \varnothing^p \mapsto \varnothing^m H H^n H \varnothing^p$

The linear map in (36) makes explicit every possible situation in UTP. In (a) and (b), in which there are fewer than two H-toned TBUs in the underlying form, no plateauing occurs. Plateauing occurs instead when there are two or more Hs in the input (e.g. (10b) above from Luganda). This is seen in (36c): all TBUs between the first and last H-toned TBUs surface as H. The notation $\{\varnothing, H\}^n$ denotes a string of n TBUs, either H or $\varnothing$.

4.2 The non-subsequentiality of unbounded tonal plateauing

The UTP map in (36) is neither left- nor right-subsequential. As §5.3 will discuss in detail, this holds for any unbounded circumambient process. This section presents an informal illustration of the proof; for the full proof, see the online appendix.

　　The UTP map in (36) is regular, as it can be modelled with a non-deterministic FST. This FST is given in Fig. 5, where underlying $\varnothing$ TBUs are H in the output in state q_1, from which a final state can only be reached if there is another input H following; one can think of q_1 as the 'plateau' state.

　　This non-determinism is necessary, as there is no way to capture this map with a deterministic FST. Recall that a deterministic FST must have at most one transition per input symbol at every state. We cannot determinise the FST in Fig. 5. To see why not, let us adopt the 'wait' strategy employed in the FST in Fig. 2 in §3.3. There are many ways to try this; this discussion follows one. The proof in the appendix ensures that all will fail, but the discussion here is intended to give an intuition as to why.

　　In the deterministic FST in Fig. 6, q_2 is a wait state representing the knowledge that a sequence H$\varnothing$ – which may be a plateauing environment – has been seen in the input. Here, it 'waits' one symbol to see whether the next symbol in the input is a H or a $\varnothing$. It will thus correctly transform $\varnothing$H$\varnothing$H, with only one intervening $\varnothing$ TBU, to $\varnothing$HHH. However it incorrectly maps inputs for which the second H is further away; for example, the input $\varnothing$H$\varnothing\varnothing$H would
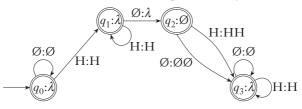
Figure 6

First attempt at a deterministic finite-state
transducer for unbounded tonal plateauing.

be mapped to ØHØØH (itself). Thus, because there is only one wait state, the machine in Fig. 6 can only describe plateaus of at most three Hs.

The machine in Fig. 7 adds an additional wait state to try to remedy this situation. Here the FST is much like the one in Fig. 6, except that it has two wait states, q_2 and q_3, which represent the look-ahead necessary to capture the behaviour of one and two Øs following an input H. Thus, Fig. 7 correctly maps ØHØH ↦ ØHHH and ØHØØH ↦ ØHHHH. However, it incorrectly maps inputs like ØHØØØH, where two Hs are separated by three or more Øs, to themselves. By now it is perhaps obvious that we are on a wild-goose chase; any 'wait n symbol' strategy will fail for any mapping in the UTP relation whose input string includes the sequence $HØ^{n+1}H$. However, given the restriction of determinism, 'wait n symbols' is the best we can do. Simply reversing the string, in an attempt to create a right-subsequential transducer, will not help; the position of the first triggering H is just as arbitrarily far to the right as the second is to the left. Thus a deterministic FST representation of the UTP map is impossible, and so it lies outside both the left- and right-subsequential classes.
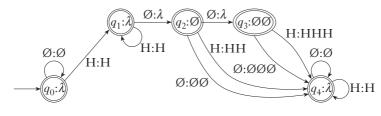

Figure 7

Second attempt at a deterministic finite-state
transducer for unbounded tonal plateauing.

4.3 Unbounded circumambient processes in general

That UTP is neither left- nor right-subsequential follows from its unbounded circumambient nature: as triggers may lie any distance away on

either side of a given target, a FST describing the map requires unbounded look-ahead in both directions.

Heinz & Lai (2013) prove this is also true for sour-grapes harmony, for the same reasons. For these processes, the fate of an input segment that can potentially assimilate rests on whether or not a trigger appears to one side *and* whether or not a blocker appears to the other. They show that this means it cannot be described by a deterministic FST reading in either direction. We can then see why no unbounded circumambient process can be subsequential. For unbounded circumambient processes, it is *definitional* that crucial information may appear on either side of the target, unboundedly far away. This means any unbounded circumambient process will require unbounded look-ahead in both directions, and will not be subsequential. This result is then key to characterising the unbounded circumambient asymmetry in terms of computational complexity, as discussed further in §5.

As stated at the outset, this result depends on a particular kind of string-based representation. Before we compare unbounded circumambient processes with bidirectional spreading in §5, it is therefore necessary to address the potential objection that these results are invalid because they hinge on this string-based representation for UTP.

4.4 Autosegmental representations and linear representations

Tone has been widely analysed in terms of autosegmental representations (though for alternatives see Cassimjee & Kisseberth 2001 and Shih & Inkelas 2014), while the notion of subsequentiality is defined in terms of strings. The preceding analysis used a string-based representation of UTP, even though it is a tonal process. This section explains and defends two interrelated facts about such string-based representations, formulated in (37).

(37) a. Each symbol in a string represents associations to a timing tier unit in any corresponding autosegmental representation.
   b. 'Unbounded look-ahead' is therefore determined by what in autosegmental terms is distance on the timing tier, not the melody tier.

(37a) is a common assumption about the relationship between string-based representations and autosegmental representations, and is the one implicit in the string-based representations used in the computational literature cited in §3.4. As for (37b), measuring distance on the timing tier as opposed to the melody tier is a representational choice. Indeed, Kornai (1995) compares different ways of encoding string-based representations of autosegmental representations. The following sections explain and justify string-based representations based on (37).

4.4.1 *Translating between string and autosegmental representations.*   Let us compare string and autosegmental representations of the sour-grapes map in (27). We have a number of options for autosegmentally representing the

featural contrasts given by the symbols in (27). Some possible representations for the underlying form $[+F][-F]^n[!F]$ in (27c) are given in (38).

(38) a.
$$
\begin{array}{cccccc}
\text{V} & \text{V}_1 & \ldots & \text{V}_n & \text{V} \\
| & & \searrow\!\!\downarrow\!\!\swarrow & & | \\
{+}\text{F} & & {-}\text{F} & & !\text{F}
\end{array}
$$

b.
$$
\begin{array}{cccccc}
\text{V} & \text{V}_1 & \ldots & \text{V}_n & \text{V} \\
| & | & & | & | \\
{+}\text{F} & {-}\text{F} & \ldots & {-}\text{F} & !\text{F}
\end{array}
$$

c.
$$
\begin{array}{ccccc}
\text{V} & \text{V}_1 & \ldots & \text{V}_n & \text{V} \\
| & & & & | \\
{+}\text{F} & & & & {-}\text{F}
\end{array}
$$

In (a) and (b), the target $[-F]$ vowels in the string are represented as underlyingly associated to $[-F]$ features in the autosegmental diagram. In (a), a single $[-F]$ feature is associated to multiple vowels, whereas in (b), each $[-F]$ vowel is associated to its own $[-F]$ feature (in violation of the Obligatory Contour Principle; Leben 1973, McCarthy 1986). As above, $[!F]$ is used as shorthand for some vowel which is also associated to some other feature $[+G]$ which prevents $[+F]$ from spreading to it (e.g. as in $[+low]$ vowels in Akan (Clements 1976), which block the spreading of a $[+ATR]$ feature). In (c), the targets are analysed as underspecified on the $[F]$ tier, and the blocker is analysed as underlyingly specified as $[-F]$ (as in Clements' analysis of $[+low]$ vowels in Akan as underlyingly $[-ATR]$).

One property that all possible autosegmental analyses share is that each symbol in a string in (27a) corresponds to the featural associations of a particular timing tier unit in (38). As an explicit example of this, (39) is a translation between symbols in UR strings in (27) and the autosegmental information at each timing tier unit in (38a). The translation for each string symbol $[+F]$, $[-F]$ and $[+F]$ is given in (39a), with an example of correspondence between the string $[+F][-F][-F][!F]$ and an autosegmental representation in (b).

(39) a.
$$
\begin{array}{ccc}
\text{V} \Leftrightarrow [+\text{F}] & \text{V} \Leftrightarrow [-\text{F}] & \text{V} \Leftrightarrow [!\text{F}] \\
| & | & | \\
{+}\text{F} & {-}\text{F} & !\text{F}
\end{array}
$$

b.
$$
\begin{array}{cccc}
[+\text{F}] & [-\text{F}] & [-\text{F}] & [!\text{F}] \\
\updownarrow & \updownarrow & \updownarrow & \updownarrow \\
\text{V} & \text{V} & \text{V} & \text{V} \\
| & & \searrow\!\!\swarrow & | \\
{+}\text{F} & & {-}\text{F} & !\text{F}
\end{array}
$$

While simplified in the sense that it focuses on one feature, translations like in (39) are what phonologists commonly, if implicitly, use when moving back and forth between linear strings of feature bundles and autosegmental representations. Again, what is important is that each symbol in the string corresponds to a timing tier unit in the autosegmental representation – the property highlighted in (37a).

As (39) makes clear, information about the melody (or featural) tier units in the autosegmental representation is obscured in the string-based representation. In (39b), the string does not encode the multiple associations of the [−F] autosegment, and it also does not show that [+F] and [!F] are only separated by a single [−F] autosegment. This ambiguity is apparent in (38), in which different autosegmental interpretations of the same string-based representation have different information on their featural tiers.

4.4.2 *Representational assumptions and look-ahead*.   As a result, when a FST reads a symbol in a string-based representation as in (27), it can be thought of in autosegmental terms as reading the featural information of a particular timing tier unit. Thus, because subsequentiality depends on look-ahead in terms of a FST reading such a string-based representation, the results regarding subsequentiality in the work of Heinz & Lai (2013) and others are based on what translates in autosegmental terms to look-ahead on the timing tier, not the melody tier.

There are a number of arguments for this assumption. One is that locality on the melody tier depends on certain representational assumptions such as underspecification and the Obligatory Contour Principle, both of which have been argued against as phonological universals (see Inkelas 1995 for the former, and Odden 1986 for the latter). The string-based representation is agnostic about such assumptions.

More importantly, as detailed in this paper, it is measuring unboundedness on the timing tier which distinguishes UTP and sour-grapes harmony from common segmental processes via their non-subsequentiality. Thus, while this choice is an assumption on which the results in this and the literature cited in §3.4 is based, it appears to be correct, in that it helps us to characterise the unbounded circumambient asymmetry. How this relates to different representational assumptions will be discussed in a moment.

Returning to UTP, it is important to establish that the linearisation used in the previous section is indeed comparable to that of previous studies of subsequentiality, such as for the sour-grapes process (Heinz & Lai 2013), in that it also measures look-ahead in terms of the timing tier. The linearisation in (35) can be made explicit with a map along the lines of the one in (39), as in (40a).

(40) a. $\mu \Leftrightarrow \mathrm{H}$ $\mu \Leftrightarrow \varnothing$ b. $\mu \quad \mu \quad \mu \quad \mu \Leftrightarrow \varnothing \mathrm{H} \varnothing \varnothing$
          | |
          H   H

As can be seen in (40b), the string symbols H and Ø can encode both contrasts in association in the underlying representation and changes in these associations in the surface, as claimed in (37a). Additionally, the string-based representation in (40b) preserves the linear order of the TBU tier of the autosegmental representation, just as in (39). Look-ahead is thus measured in terms of the timing tier, as in (37b). Note again that, as they simply encode the associations to each timing tier

unit, linearisations like in (39) and (40) are very general, and can be applied to any set of autosegmental representations for a particular process.

4.4.3 *Representation and subsequentiality.* This concludes the arguments for why this particular kind of representation is used in this paper. However, there remains an important open question: what if we don't use this particular representation? This is a valid direction for research, but any answers to this question will not change the result argued for in this paper. Formally, representation is related to expressive power (Medvedev 1964, Rogers *et al*. 2013). Work studying the computational properties of autosegmental representations (Bird & Ellison 1994, Kornai 1995, Yli-Jyrä 2013) has not yet offered a hierarchy of complexity like the one presented in §3.4. Future research may use this work in computational autosegmental phonology as a starting point to study the relationship between representation and computational complexity of phonological processes.

It may even be possible to derive non-subsequentiality of unbounded circumambient processes from some aspect of representation (such as look-ahead on the melody tier). This requires, however, that if such representations were translated into the string-based representations according to (37), non-subsequentiality in the string-based representations would somehow emerge. Thus any such explanation based on representation would duplicate the computational results outlined here – it would not refute them. That is, it would have to maintain the result that much of segmental phonology, when viewed with string-based representations that obey the criteria in (37), is subsequential, while UTP and sour-grapes harmony are not.

5 Unbounded circumambient processes and weak determinism

Having shown that unbounded circumambient processes are not left- or right-subsequential, we must make one final distinction. This section argues that unbounded circumambient processes are computationally distinct from the bidirectional spreading processes introduced in §2.5. This leads to the characterisation of the unbounded circumambient asymmetry in terms of a weakly deterministic complexity bound on segmental phonology which is absent in tonal phonology.

5.1 Bidirectional spreading and the weakly deterministic class

The class of unbounded circumambient processes is not the only one which is not left- or right-subsequential. In patterns of stem-control harmony, or cases of bidirectional spreading, such as the Arabic emphasis spreading discussed in §2.5, a feature spreads outward both to the right and the left, as in (41).

(41) …[−F]…[+F]…[−F]… ↦ …[+F]…[+F]…[+F]…

Such a map requires unbounded look-ahead in either direction: a target may follow or precede the trigger, in any direction. Thus, as Heinz & Lai (2013) also show, such cases (which will henceforth be referred to under the umbrella term bidirectional spreading) are not subsequential. However, there is a crucial difference between bidirectional spreading and unbounded circumambient processes. Bidirectional spreading hinges on a single trigger, whose influence spreads outward. In contrast, unbounded circumambient processes hinge on two triggers/blockers, whose targets lie between.

Heinz & Lai (2013) observe that a bidirectional spreading process is essentially the same unidirectional map applied left-to-right and then right-to-left. They propose a superclass of the subsequential maps, called weakly deterministic maps, which includes this kind of process. A weakly deterministic map can be decomposed into a left- and right-subsequential map, such that the left-subsequential map is not allowed to change the alphabet or increase the length of the string.

The process in (41) can be decomposed into two left- and right-subsequential maps, describable by the consonant harmony FST in Fig. 4. First, the input string is read by the FST left-to-right (applying the left-subsequential map), then the resulting output is fed back into the FST right-to-left (applying the right-subsequential map to the output). This decomposition is schematised in (42).

(42) a. *left-subsequential* $[-F]\ldots[+F]\ldots[-F] \mapsto [-F]\ldots[+F]\ldots[+F]$
     b. *right-subsequential* $[-F]\ldots[+F]\ldots[+F] \mapsto [+F]\ldots[+F]\ldots[+F]$

This composition of the two maps is special, because it does not change the alphabet or increase the length of the string. Intuitively, this is because these bidirectional processes can be thought of as one unidirectional process operating in two directions. This is highlighted by the fact that both sub-maps use the same FST. Thus, bidirectional spreading is weakly deterministic, as first seen in Fig. 1.

The weakly deterministic class is defined as such by Heinz & Lai (2013) as a restriction on Elgot & Mezei's (1965) result that any regular map can be decomposed into a left-subsequential and right-subsequential map, as long as the left-subsequential map is allowed to enlarge the alphabet. This result holds because the left-subsequential map can 'mark up' the string with extra symbols in the first map, and then erase them in the second.[14] However, no attested segmental maps studied in the literature cited above require such a mark-up.

5.2 Unbounded circumambient processes and the weakly deterministic class

In contrast to attested segmental processes, Heinz & Lai (2013) argue that sour-grapes harmony requires such a mark-up, and thus is not weakly

[14] This is not unlike the use of abstract intermediate forms in early derivational phonology; see e.g. Clements (1977).

deterministic. This is because it is not simply the application of the same process in two directions. Because the sour-grapes pattern is unbounded circumambient, for any decomposition into two sub-maps the left-subsequential process must somehow encode whether or not a [+F] has been seen to the right of the [−F] targets. It is difficult to see how this can be done without intermediate mark-up, and thus Heinz & Lai (2013) conjecture that sour-grapes harmony is not weakly deterministic. The exact same arguments apply to UTP.

UTP can be decomposed into two subsequential maps with an augmented alphabet, as in (43). First, the left-subsequential process marks all Ø following a H as ?, and then the right-subsequential process changes all ? preceding a H to H (and all ? not preceding a H to Ø).

(43) *input*   ØØØH   HØØØ   HØØH
    left-subsequential   ØØØH   H???   H??H
    right-subsequential   ØØØH   HØØØ   HHHH
    output   ØØØH   HØØØ   HHHH

Crucially, this decomposition relies on the intermediate ? symbols to 'carry forward' the information that a H appears to the left in the string. This allows the right-subsequential map to correctly apply without any unbounded look-ahead. However, as in the case of sour-grapes processes, it is hard to see how there could be a similar decomposition which uses only H and Ø, maintains string length in the left-subsequential sub-map and obtains the exact same map. This is because using the same alphabet to create an encoding will always lead to distinctions between input strings being lost, and so such an encoding is bound to fail.

Again, this is based on the unbounded circumambient nature of the process: because there is crucial information on either side of the targets, it is necessary to mark targets to the right of the left trigger in order for the right-subsequential function to correctly process them without unbounded look-ahead. Thus, Heinz & Lai (2013)'s conjecture applies not just to UTP, but also to any unbounded circumambient process.

An important implication of Heinz & Lai (2013)'s conjecture is that there exist fully regular maps outside of the weakly deterministic class, and thus that the weakly deterministic class of maps is a proper subclass of the regular class of maps, as depicted in Fig. 1. This section has argued that unbounded circumambient processes fall into this fully regular class, as in Fig. 1, and that this distinguishes them in complexity from bidirectional spreading.

5.3 The weakly deterministic hypothesis

The conjecture that the weakly deterministic class is a proper subclass of the regular class provides a way to capture the unbounded circumambient asymmetry. Heinz & Lai (2013) propose that phonology is weakly deterministic. This hypothesis correctly predicts the absence of unbounded

circumambient processes like the sour-grapes pattern in the typology of vowel harmony. However, this paper has shown that unbounded circumambient processes are widely attested in tonal phonology, and has proposed that the weakly deterministic bound only applies to segmental phonology. This accurately predicts that fully regular maps such as UTP and Copperbelt Bemba H-tone spreading exist in tonal phonology. Furthermore, the proposal that tone is more computationally complex than segmental phonology is in line with the assertion of Hyman (2011) and others that tone can 'do more' than phonology.

An explanation for how the weakly deterministic bound is manifested in the phonological system will be left for future work, although it is possible to speculate on a few points. For example, formal language complexity correlates with an increase in the computational resources necessary for parsing and generation. It could be that tone has access to such resources because prosodic information more commonly interacts with syntax (see e.g. Hyman & Katamba 2010), and thus requires more powerful computation. This issue of computational power can also be directly related to learning, as empirical work suggests that formal complexity constrains phonological learning (Moreton & Pater 2012, McMullin & Hansson 2014, Lai 2015).

5.4 Explaining the exceptions to the weakly deterministic hypothesis

The weakly deterministic hypothesis in its strongest form cannot account for the rare cases of unbounded circumambient segmental processes discussed in this paper, i.e. Sanskrit *n*-retroflexion and Yaka vowel harmony. There are a few ways to reconcile these exceptions with the weakly deterministic hypothesis.

One possibility is that the accounts of these processes in the literature are incorrect in classifying them as unbounded. As mentioned in §2.3, Ryan (forthcoming) observes that blocking of Sanskrit *n*-retroflexion may be bound to the adjacent syllable. In Yaka harmony, the greatest attested distance is only three vowels, and is restricted to a particular morphological context.

Another explanation comes from potentially interfering factors. For example, the presence of an unbounded circumambient process in the tonal phonology may license an unbounded circumambient process in segmental phonology. This may be the case for Yaka, which as noted in §2.2, also has UTP. It could be that Yaka speakers, having first internalised the plateauing process in tonal phonology, may then be able to generalise it to their segmental phonology.

Finally, it may simply be that the constraint against fully regular maps in segmental phonology is not categorical, but in some sense gradient, and thus can admit exceptions. One way in which this may be manifested is through a learning bias in which individuals are more receptive to some patterns than others (Wilson 2006a, Moreton 2008). Children may show

a strong preference for weakly deterministic maps when learning segmental phonology, but change to a more general learner in the face of sufficient data. As Staubs (2014) shows, gradient typological generalisations may also result from the transmission of patterns among multiple learning agents. It is likely that non-weakly deterministic patterns require more kinds of data to learn, and perhaps this data is only available in tonal phonology, which is known to operate over much larger domains than segmental phonology.

Hence there are a number of reasons why the cases of Sanskrit *n*-retroflexion and Yaka vowel harmony do not immediately invalidate the proposal offered here. Even if they cannot ultimately be explained away, they are exceptional with respect to other characterisations of segmental phonology, such as Wilson's (2003, 2006b) myopia generalisation. Finally, Sanskrit *n*-retroflexion and Yaka vowel harmony appear to be the only such cases – Hyman (2011: 218) states, for example, that Yaka harmony is the 'only one example' of such a process of which he is aware. As such, the development of the potential explanations above will be left to future research.

In sum, this paper has characterised the unbounded circumambient asymmetry in terms of computational complexity, and has identified the weakly deterministic boundary as the relevant difference between tonal and segmental phonology which captures this asymmetry.

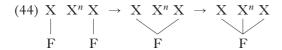# 6 The unbounded circumambient asymmetry in Optimality Theory

This paper has laid out a computational explanation for the unbounded circumambient asymmetry. Here I briefly outline why this explanation is superior to any possible explanation using current theories based on Optimality Theory (Prince & Smolensky 1993) or its variants. This is because current theories of OT do not provide a unified characterisation of the unbounded circumambient processes, meaning that making OT empirically adequate with regards to the asymmetry runs into a 'duplication of effort' problem. Banning the particular non-local effects that generate segmental unbounded circumambient processes requires changes both to how OT manipulates segments autosegmentally and to how it compares candidates. Thus OT does not provide a unified characterisation of the asymmetry comparable to the one based on computational complexity put forward in the preceding sections.

Sour-grapes vowel harmony is an appropriate place to begin. It is unattested in segmental phonology, but predicted by parallel OT and local AGREE constraints checking the agreement between features of adjacent segments. Given such constraints, candidates in which spreading is blocked still violate AGREE (as the last segment to which the feature has spread still disagrees with the blocker) and are thus harmonically bounded by candidates in which no spreading occurs, as they also violate

AGREE but do not register any faithfulness violations incurred by the spreading feature (see McCarthy 2010 for how the sour-grapes pattern is generated in OT). Thus the nature of optimisation, which involves comparison of candidates with non-local changes to those with local ones, allows for unbounded behaviour with local constraints. In fact, this is exactly this type of behaviour that Hyman (1998) harnesses to describe vowel harmony in Yaka.

One proposal for dealing with the sour-grapes problem is McCarthy's (2010) Harmonic Serialism (HS) analysis of spreading, which restricts GEN to the production of candidates with a single change (with winning candidates fed back into the grammar until no changes produce a more optimal candidate). As only candidates with single changes are compared, global comparison cannot take place, and thus patterns such as sour-grapes harmony can no longer be generated with local AGREE constraints. Of course, assuming a HS framework provides no explanation for the presence of a sour-grapes-like tone pattern in Copperbelt Bemba.

Furthermore, this solution is specific to sour-grapes vowel harmony resulting from AGREE, and does not extend to autosegmental analyses of plateauing. Plateauing, tonal or otherwise, can be seen as two identical features adjacent on the tier merging to satisfy the OCP, with the intervening anchor units associating to the resulting fused feature in order to avoid a gapped structure. This can be schematised derivationally as in (44).

(44) $\text{X} \quad \text{X}^n \text{X} \quad \rightarrow \quad \text{X} \quad \text{X}^n \text{X} \quad \rightarrow \quad \text{X} \quad \text{X}^n \text{X}$
         $\quad | \qquad \quad | \qquad\qquad\qquad \vee \qquad\qquad\qquad\qquad \vee$
         $\;\; \text{F} \qquad \text{F} \qquad\qquad\quad \text{F} \qquad\qquad\qquad\quad \text{F}$

While (44) is a derivational sketch of the process, the intuition can be implemented in either serial or parallel versions of OT, given basic constraints governing the behaviour of autosegments proposed in the OT literature. Fusion of identical adjacent autosegments can be motivated by an OCP markedness constraint (Leben 1973, McCarthy 1986, Yip 2002) ranking above a UNIFORMITY faithfulness constraint militating against the fusion of autosegments (Myers 1997, Pater 2004). The 'filling in' of the associations to the intervening anchor units can then be motivated by ranking NOGAP (Itô *et al*. 1995, Yip 2002), a constraint against structures in which associations skip over potential anchors, above faithfulness constraints prohibiting the addition of association lines. Given such a ranking in classical OT, only a candidate featuring both fusion and filling in will satisfy both the OCP and NOGAP markedness constraints. In HS, a candidate with fusion will represent an increase in harmony (as it satisfies the OCP), and then successive candidates representing a gradual filling in of the associations in the intervening anchor units will represent an increase in harmony with respect to NOGAP. Thus, both fusion and filling in can be achieved straightforwardly either in parallel OT or in HS.

We therefore have a duplication of effort problem: in order to categorically remove segmental unbounded circumambient processes from the

typology predicted by OT theories of phonology, it is necessary to adopt at least one of the changes proposed for sour-grapes harmony and to remove the segmental versions of the constraints above that achieve autosegmental plateauing. Thus, while it appears technically possible to have a theory of OT that is empirically adequate with regards to the unbounded circumambient asymmetry, it does not provide a unified explanation. In contrast, the FLT-based account presented here characterises the difference: segmental phonology is weakly deterministic, but tonal phonology is not.

7 Conclusion

This paper has made three contributions. First, it has documented the asymmetry in the attestation of unbounded circumambient processes in tonal phonology and segmental phonology. Second, it has shown that UTP is similar to sour-grapes harmony, in that they are both unbounded circumambient processes, and that this similarity has formal consequences. Third, it has characterised the asymmetry between segmental and tonal phonology by arguing that unbounded circumambient processes are fully regular, and thus more computationally complex than processes which do not require unbounded look-ahead in two directions. This was shown to be a superior characterisation than that offered in OT, which cannot account for the asymmetry in a unified way.

The conclusions in this paper raise a number of interesting questions for future research. What computational constraints are there on tone? In other words, is there a subregular class of maps which includes (or corresponds to) unbounded circumambient processes? How does this relate to Elgot & Mezei's (1965) result using intermediate mark-up to generate any regular map from two subsequential maps? As Heinz & Lai (2013) point out, this idea may be brought to bear on the question of how abstract intermediate representations are in phonology. The questions of representation raised at the end of §4.4 can be approached in a similar way – how do changes in representation correlate with changes in generative capacity? Finally, it remains to be seen how the FLT insight presented here can be incorporated into traditional phonological theory. One potential approach is to appeal to learning, as considered in §5.3.

Unfortunately, going into these concerns in detail is beyond the scope of this paper. Instead, its goal is similar to that of Kisseberth (1970)'s prophetic work on conspiracies in Yawelmani. Kisseberth writes that he is not 'principally interested in proposing detailed formalism; instead I would like to encourage phonologists to *look at* the phonological component of a grammar in a particular way' (1970: 293). This work, too, aims to encourage phonologists to look at phonology in a new way. Regardless of how it is incorporated into our previous understanding of phonology, the unbounded circumambient asymmetry between tonal and segmental phonology is robust, and the best available characterisation of this generalisation comes from computational complexity.

REFERENCES

Al Khatib, Sam (2008). On the directionality of emphasis spread. In Susie Jones (ed.) *2008 Canadian Linguistic Association Conference Proceedings*. Available (May 2016) at http://cla-acl.ca/actes-2008-proceedings.

Baković, Eric (2000). *Harmony, dominance, and control*. PhD dissertation, Rutgers University.

Beesley, Kenneth R. & Lauri Karttunen (2003). *Finite state morphology*. Stanford: CSLI.

Bennett, William (2013). *Dissimilation, consonant harmony, and surface correspondence*. PhD dissertation, Rutgers University.

Beros, Achilles & Colin de la Higuera (2014). A canonical semi-deterministic transducer. In Alexander Clark, Makoto Kanazawa & Ryo Yoshinaka (eds.) *JMLR Workshop and Conference Proceedings*. Vol. 34: *Proceedings of the 12th International Conference on Grammatical Inference*. 33–48. Available (May 2016) at http://jmlr.org/proceedings/papers/v34.

Bickmore, Lee S. & Nancy C. Kula (2013). Ternary spreading and the OCP in Copperbelt Bemba. *Studies in African Linguistics* **42**. 101–132.

Bird, Steven & T. Mark Ellison (1994). One-level phonology: autosegmental representations and rules as finite automata. *Computational Linguistics* **20**. 55–90.

Cassimjee, Farida & Charles W. Kisseberth (2001). Zulu tonology and its relationship to other Nguni languages. In Shigeki Kaji (ed.) *Cross-linguistic studies of tonal phenomena: tonogenesis, Japanese accentology, and other topics*. Tokyo: ILCAA. 327–359.

Chandlee, Jane (2014). *Strictly local phonological processes*. PhD dissertation, University of Delaware.

Chandlee, Jane, Angeliki Athanasopoulou & Jeffrey Heinz (2012). Evidence for classifying metathesis patterns as subsequential. *WCCFL* **29**. 303–309.

Chomsky, Noam (1956). Three models for the description of language. *IRE Transactions on Information Theory* **2**. 113–124.

Clements, G. N. (1976). *Vowel harmony in nonlinear generative phonology: an autosegmental model*. Distributed 1980, Indiana University Linguistics Club.

Clements, G. N. (1977). Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. *NELS* **7**. 49–64.

Clements, G. N. & John A. Goldsmith (eds.) (1984). *Autosegmental studies in Bantu tone*. Dordrecht: Foris.

Do, Young Ah & Michael Kenstowicz (2011). A note on phonological phrasing in South Kyungsang. Ms, MIT. Available (May 2016) at http://web.mit.edu/linguistics/people/faculty/kenstowicz/do-kens.pdf.

Downing, Laura J. (2001). How ambiguity of analysis motivates stem tone change in Durban Zulu. *UBC Working Papers in Linguistics* **4**. 39–55.

Elgot, C. C. & J. E. Mezei (1965). On relations defined by generalized finite automata. *IBM Journal of Research and Development* **9**. 47–68.

Folia, Vasiliki, Julia Uddén, Meinou de Vries, Christian Forkstam & Karl Magnus Petersson (2010). Artificial language learning in adults and children. *Language Learning* **60**. Suppl. 2. 188–220.

Gafos, Adamantios I. (1999). *The articulatory basis of locality in phonology*. New York: Garland.

Gainor, Brian, Regine Lai & Jeffrey Heinz (2012). Computational characterizations of vowel harmony patterns and pathologies. *WCCFL* **29**. 63–71.

Goldsmith, John A. (1990). *Autosegmental and metrical phonology*. Oxford & Cambridge, Mass.: Blackwell.

Good, Jeff (2004). Tone and accent in Saramaccan: charting a deep split in the phonology of a language. *Lingua* **114**. 575–619.

Graf, Thomas (2010). *Logics of phonological reasoning*. MA dissertation, University of California, Los Angeles.

Hale, Kenneth & Elisabeth Selkirk (1987). Government and tonal phrasing in Papago. *Phonology Yearbook* **4**. 151–183.

Hansson, Gunnar Ólafur (2001). *Theoretical and typological issues in consonant harmony*. PhD dissertation, University of California, Berkeley.

Hansson, Gunnar Ólafur (2010). *Consonant harmony: long-distance interaction in phonology*. Berkeley: University of California Press.

Heinz, Jeffrey (2011). Computational phonology. Part 1: Foundations. *Language and Linguistics Compass* **5**. 140–152.

Heinz, Jeffrey & William J. Idsardi (2011). Sentence and word complexity. *Science* **333**. 295–297.

Heinz, Jeffrey & William J. Idsardi (2013). What complexity differences reveal about domains in language. *Topics in Cognitive Science* **5**. 111–131.

Heinz, Jeffrey & Regine Lai (2013). Vowel harmony and subsequentiality. In András Kornai & Marco Kuhlmann (eds.) *Proceedings of the 13th Meeting on the Mathematics of Language*. Sofia: Association for Computational Linguistics. 52–63.

Hyman, Larry M. (1998). Positional prominence and the 'prosodic trough' in Yaka. *Phonology* **15**. 41–75.

Hyman, Larry M. (2011). Tone: is it different? In John A. Goldsmith, Jason Riggle & Alan C. L. Yu (eds.) *The handbook of phonological theory*. 2nd edn. Oxford: Wiley-Blackwell. 197–239.

Hyman, Larry M. & Francis Katamba (2010). Tone, syntax and prosodic domains in Luganda. *ZAS Papers in Linguistics* **53**. 69–98.

Hyman, Larry M., Francis Katamba & Livingstone Walusimbi (1987). Luganda and the strict layer hypothesis. *Phonology Yearbook* **4**. 87–108.

Inkelas, Sharon (1995). The consequences of optimization for underspecification. *NELS* **25:1**. 287–302.

Itô, Junko, Armin Mester & Jaye Padgett (1995). Licensing and underspecification in Optimality Theory. *LI* **26**. 571–613.

Johnson, C. Douglas (1972). *Formal aspects of phonological description*. The Hague: Mouton.

Kaplan, Ronald & Martin Kay (1994). Regular models of phonological rule systems. *Computational Linguistics* **20**. 331–378.

Kenstowicz, Michael (1994). *Phonology in generative grammar*. Cambridge, Mass. & Oxford: Blackwell.

Kidima, Lukowa (1991). *Tone and accent in Kiyaka*. PhD dissertation, University of California, Los Angeles.

Kisseberth, Charles W. (1970). On the functional unity of phonological rules. *LI* **1**. 291–306.

Kisseberth, Charles W. (1984). Digo tonology. In Clements & Goldsmith (1984). 105–182.

Kisseberth, Charles W. & David Odden (2003). Tone. In Derek Nurse & Gérard Philippson (eds.) *The Bantu languages*. London & New York: Routledge. 59–70.

Kornai, András (1995). *Formal phonology*. New York: Garland.

Kula, Nancy & Lee Bickmore (2015). Phrasal phonology in Copperbelt Bemba. *Phonology* **32**. 147–176.

Lai, Regine (2015). Learnable vs. unlearnable harmony patterns. *LI* **46**. 425–451.

Laughren, Mary (1984). Tone in Zulu nouns. In Clements & Goldsmith (1984). 183–234.

Leben, William R. (1973). Suprasegmental phonology. PhD dissertation, MIT.

McCarthy, John J. (1986). OCP effects: gemination and antigemination. *LI* **17**. 207–263.

McCarthy, John J. (2010). Autosegmental spreading in Optimality Theory. In John A. Goldsmith, Elizabeth Hume & W. Leo Wetzels (eds.) *Tones and features*. Berlin & Boston: De Gruyter Mouton. 195–222.

Macdonell, A. A. (1910). *Vedic grammar*. Strasbourg: Trübner.

McMullin, Kevin & Gunnar Ólafur Hansson (2014). Locality in long-distance phono-tactics: evidence for modular learning. *NELS* **44:2**. 1–14.

McWhorter, John H. & Jeff Good (2012). *A grammar of Saramaccan Creole*. Berlin & Boston: De Gruyter Mouton.

Medvedev, Yu. T. (1964). On the class of events representable in a finite automaton. In Edward F. Moore (ed.) *Sequential machines: selected papers*. Reading, Mass.: Addison-Wesley. 215–227.

Mohri, Mehryar (1997). Finite-state transducers in language and speech processing. *Computational Linguistics* **23**. 269–311.

Monier-Williams, Monier (1899). *A Sanskrit–English dictionary*. Oxford: Clarendon. Maintained at http://www.sanskrit-lexicon.uni-koeln.de.

Moreton, Elliott (2008). Analytic bias and phonological typology. *Phonology* **25**. 83–127.

Moreton, Elliott & Joe Pater (2012). Structure and substance in artificial-phonology learning. Part 1: Structure. *Language and Linguistics Compass* **6**. 686–701.

Myers, Scott (1997). OCP effects in Optimality Theory. *NLLT* **15**. 847–892.

Nevins, Andrew (2010). *Locality in vowel harmony*. Cambridge, Mass.: MIT Press.

Odden, David (1986). On the role of the Obligatory Contour Principle in phonology. *Lg* **62**. 353–383.

Padgett, Jaye (1995). Partial class behavior and nasal place assimilation. In Keiichiro Suzuki & Dirk Elzinga (eds.) *Proceedings of the 1995 Southwestern Workshop on Optimality Theory* (*SWOT*). Tucson: Department of Linguistics, University of Arizona. 145–183.

Pater, Joe (2004). Austronesian nasal substitution and other NÇ effects. In John J. McCarthy (ed.) *Optimality Theory in phonology: a reader*. Malden, Mass.: Blackwell. 271–289.

Payne, Amanda (2014). Dissimilation as a subsequential process. *NELS* **44:2**. 79–90.

Prince, Alan & Paul Smolensky (1993). *Optimality Theory: constraint interaction in generative grammar*. Ms, Rutgers University & University of Colorado, Boulder. Published 2004, Malden, Mass. & Oxford: Blackwell.

Rogers, James & Marc D. Hauser (2010). The use of formal language theory in studies of artificial language learning: a proposal for distinguishing the differences between human and nonhuman animal learners. In Harry van der Hulst (ed.) *Recursion and human language*. Berlin & New York: De Gruyter Mouton. 213–231.

Rogers, James, Jeffrey Heinz, Margaret Fero, Jeremy Hurst, Dakotah Lambert & Sean Wibel (2013). Cognitive and sub-regular complexity. In Glyn Morrill & Mark-Jan Nederhof (eds.) *Formal grammar*. Berlin & Heidelberg: Springer. 90–108.

Rogers, James & Geoffrey K. Pullum (2011). Aural pattern recognition experiments and the subregular hierarchy. *Journal of Logic, Language and Information* **20**. 329–342.

Rose, Sharon & Rachel Walker (2004). A typology of consonant agreement as correspondence. *Lg* **80**. 475–531.

Roundtree, S. Catherine (1972). Saramaccan tone in relation to intonation and grammar. *Lingua* **29**. 308–325.

Ryan, Kevin M. (forthcoming). Attenuated spreading in Sanskrit retroflex harmony. *LI*.

Safir, Ken (1982). Nasal spreading in Capanahua. *LI* **13**. 689–694.

Schein, Barry & Donca Steriade (1986). On geminates. *LI* **17**. 691–744.

Shih, Stephanie & Sharon Inkelas (2014). A subsegmental correspondence approach to contour tone (dis)harmony patterns. In John Kingston, Claire Moore-Cantwell, Joe Pater & Robert Staubs (eds.) *Proceedings of the 2013 Meeting on Phonology*. Available (May 2016) at http://journals.linguisticsociety.org/proceedings/index.php/amphonology.

Staubs, Robert (2014). *Computational modeling of learning biases in stress typology*. PhD dissertation, University of Massachusetts Amherst.

Suzuki, Keiichiro (1998). *A typological investigation of dissimilation*. PhD dissertation, University of Arizona.

Whitney, William Dwight (1889). *A Sanskrit grammar*. Cambridge, Mass.: Harvard University Press.

Wilson, Colin (2003). Analyzing unbounded spreading with constraints: marks, targets, and derivations. Ms, University of California, Los Angeles.

Wilson, Colin (2006a). Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cognitive Science* **30**. 945–982.

Wilson, Colin (2006b). Unbounded spreading is myopic. Paper presented at the Phonology Fest Workshop on Current Perspectives on Phonology, Indiana University.

Yip, Moira (2002). *Tone*. Cambridge: Cambridge University Press.

Yli-Jyrä, Anssi (2013). On finite-state tonology with autosegmental representations. In *Proceedings of the 11th International Conference on Finite State Methods and Natural Language Processing*. Stroudsburg, PA: Association for Computational Linguistics. 90–98.