

Decomposing phonological transformations in serial derivations

Andrew Lamont

University of Massachusetts, Amherst
alamont@linguist.umass.edu

Abstract

While most phonological transformations have been shown to be subsequential, there are tonal processes that do not belong to any subregular class, thereby making it difficult to identify a tighter bound on the complexity of phonological processes than the regular languages. This paper argues that a tighter bound obtains from examining the way transformations are computed: when derived in serial, phonological processes can be decomposed into iterated subsequential maps.

1 Introduction

Phonological transformations map underlying representations (UR) onto surface forms (SF). The maps between UR and SF are known to be REGULAR (Johnson, 1972; Kaplan and Kay, 1994), meaning they can be modeled with finite state transducers (FST). This generalization is stated as the Regular Hypothesis (1).

- (1) **Regular Hypothesis:** Phonological transformations are regular.

The Regular Hypothesis is not strong enough. There are many regular maps that are phonologically implausible, and most UR→SF maps belong to the SUBREGULAR classes shown in Figure 1. The majority are in the SUBSEQUENTIAL classes in gray. Bidirectional long-distance processes like stem-controlled vowel harmony belong to the more powerful WEAKLY DETERMINISTIC class (Heinz and Lai, 2013). Only two tonal processes, unbounded tonal plateauing and conditional rightward

spreading (in bold), have been shown to not belong to any subregular class (Jardine, 2016a).

Because of their wide empirical coverage and computational properties, the union of the subsequential classes was an early candidate for a tighter bound on the complexity of phonological processes than the regular class (Chandlee and Heinz, 2012; Gainor et al., 2012). Heinz (forthcoming) states this as the Subsequential Hypothesis (2). The Subsequential Hypothesis is stronger than the Regular Hypothesis, while maintaining its uniform generalization over all phonological transformations.

- (2) **Subsequential Hypothesis:** Phonological transformations are left- or right-subsequential.

In light of the weakly deterministic and regular processes, the Subsequential Hypothesis is too strong. Because there are phonological transformations that are not subregular, there is not a uniform revision of the Subsequential Hypothesis stronger than the Regular Hypothesis. Jardine (2016a) argues that only tonal processes exceed the weakly deterministic class, so a possible revision states that segmental processes are weakly deterministic and tonal processes are regular.¹ In short, from examining the UR→SF maps on their own, there is no subregular class that subsumes all phonological transformations.

This paper argues that a uniform revision of the Subsequential Hypothesis obtains by examining not only the UR→SF maps, but also how their derivations are computed. There is an open question in

¹Tutrugbu vowel harmony challenges this generalization (McCollum et al., 2017).

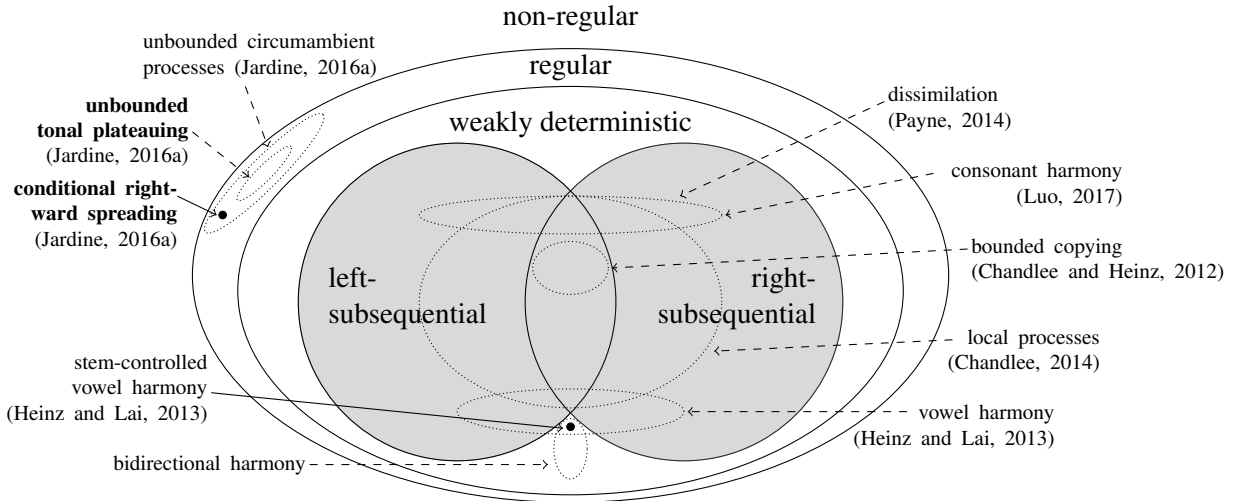


Figure 1: Subregular hierarchy of phonological transformations (based on Jardine’s Figure 1 (2016, 263))

phonological theory whether $UR \mapsto SF$ maps are derived in one fell swoop, or whether they are broken down into sub-derivations. For example, consider the sibilant harmony process that transforms the UR /sasasaf/ into the SF [ʃaʃaʃaʃ] in Figure 2. The dashed line directly from the UR to the SF shows the PARALLEL derivation, where every /s/ changes at the same time. The solid lines from UR to SF via two intermediate forms show the SERIAL derivation, where only one /s/ changes at a time. Each line represents one computation made by the phonology. Both derivations yield the same SF, the parallel derivation in one step and the serial in three.




Figure 2: Serial and parallel sibilant harmony

In a parallel derivation, the SF is derived directly from the UR, so the derivation is exactly the $UR \mapsto SF$ map. Because they are identical, parallel derivations have the same computational complexity as $UR \mapsto SF$ maps. This paper argues that in a serial derivation, where the SF is derived gradually over a number of steps, each step is subsequential. This is stated as the Serial Subsequential Hypothesis (3). Restricting each step to making a single change requires iterating processes. The solid lines in Figure 2, represent a process that changes one /s/, which applies three times to gradually yield the

SF. As Section 4 argues, this restriction also predicts that some regular maps are not possible phonological processes.

- (3) **Serial Subsequential Hypothesis:** Phonological transformations are decomposable into iterated left- or right-subsequential maps.

The paper is organized as follows. Section 2 reviews the characterization of the classes in Figure 1 in terms of FSTs, providing empirical examples, and discusses the serial counterparts of the subregular classes. Section 3 demonstrates that the regular tonal processes can be broken down into subsequential steps in a serial derivation. Sections 4 and 5 discuss the predictions of the Serial Subsequential Hypothesis and conclude.

2 Phonological transformations and FSTs

2.1 Subsequential transformations

Subsequential transformations include local processes like place assimilation, and unidirectional long-distance processes like regressive sibilant harmony. They can be computed by SUBSEQUENTIAL FSTs (Mohri, 1997). These are deterministic FSTs where every state is accepting. When the end of the input is reached, an additional string is appended to the output, determined by the current state the machine is in. This can be thought of as standing in for transitions on boundary symbols (Chomsky and Halle, 1968). LEFT-SUBSEQUENTIAL

FSTs (L-SFST) read inputs left-to-right; RIGHT-SUBSEQUENTIAL FSTs (R-SFST) read inputs right-to-left.

Example subsequential FSTs are given in Figures 3 and 4 for a toy sibilant harmony system. The machines remember the identity of the first sibilant they read. If the first sibilant is /s/, they transition to state q_1 ; if /f/, state q_2 . That information and the current position in the input determine what is written to the output. In the diagrams, the end-of-input string is shown after states' labels. λ stands for the empty string, so neither machine appends to the output when the input is exhausted.

The L-SFST in Figure 3 reads inputs left-to-right. The leftmost sibilant controls harmony, yielding progressive sibilant harmony, such as that in Aari (Hansson, 2010, 51). Table 1 gives sample derivations; the I row gives the current symbol in the input read by the FST, Q the state the machine is in, and O the string written to the output.




Figure 3: L-SFST for progressive sibilant harmony

/safafa/ → [sasasa]							
I	s	a	f	a	f	a	
Q	q_0	$\rightarrow q_1$	$\rightarrow q_1$	$\rightarrow q_1$	$\rightarrow q_1$	$\rightarrow q_1$	$\rightarrow q_1$
O	s	a	s	a	s	a	λ
/fasasa/ → [fafafa]							
I	f	a	s	a	s	a	
Q	q_0	$\rightarrow q_2$	$\rightarrow q_2$	$\rightarrow q_2$	$\rightarrow q_2$	$\rightarrow q_2$	$\rightarrow q_2$
O	f	a	f	a	f	a	λ

Table 1: Sample derivations for the L-SFST in Figure 3

The R-SFST in Figure 4 is the mirror image of the L-SFST in Figure 3. It reads inputs right-to-left, so the rightmost sibilant controls harmony, yielding regressive harmony, such as that in Navajo (Hansson, 2010, 43). For the URs in Table 1, this machine pro-

duces SFs with the opposite direction of harmony: /safafa/ → [fafafa]; /fasasa/ → [sasasa].




Figure 4: R-SFST for regressive sibilant harmony

The direction in which subsequential FSTs read inputs determines whether a long-distance process is regressive or progressive. That is, R-SFSTs model regressive harmony, but L-SFSTs cannot. The R-SFST in Figure 4 first identifies a trigger and remembers its identity. This is enough information to write the correct output for every target in the input. A L-SFST would read the targets first, and face the insurmountable problem of anticipating the identity of the trigger. Until it finds the trigger, a L-SFST does not have enough information to write the correct output for a target. Because the trigger may be arbitrarily far away, the L-SFST would have to wait until the end of the input to correctly output the targets. Because FSTs cannot remember arbitrarily long sequences, this strategy fails.

The FSTs in Figures 3 and 4 compute the UR → SF maps and, equivalently, the parallel derivations of these processes. Once the machines transition into a harmonizing state, q_1 or q_2 , they remain in that state until the input is exhausted, because the only transitions from these states are self-loops. The machines therefore apply harmony to every focus in an UR.²

In the corresponding serial derivation, each computation applies harmony to only one sibilant in the input. Compare the L-SFST in Figure 3 to its serial counterpart in Figure 5. In the latter, the unfaithful transitions, i.e. the arcs leaving q_1 on f and q_2 on s, lead to a state q_3 , where the input is copied faithfully to the output. Thus, once this machine makes a sin-

²In a phonological rule of the form $A \rightarrow B / C_D$, A is called the *focus*, B the *structural change*, C_D the *context*, and CAD the *structural description*.

gle change, it transitions to a state where it is unable to make any further changes. Restricted to making one change at a time, inputs with multiple foci must pass through the machine a number of times before the final SF is computed.




Figure 5: L-SFST for serial progressive sibilant harmony

Subsequential FSTs have enough memory to compute long-distance processes like sibilant harmony. To satisfy the restriction that they make only one change, they simply have to remember whether they have already made a change. In the L-SFST in Figure 5, this is implemented by transitioning into the faithful state q_3 on the unfaithful arcs leaving q_1 and q_2 . This modification does not require any additional computational power, so the serial FST is still left-subsequential.

2.2 Weakly deterministic transformations

Weakly deterministic maps are defined as length- and alphabet-preserving³ compositions of a left-subsequential and a right-subsequential map (Heinz and Lai, 2013), and include long-distance bidirectional processes like stem-controlled vowel harmony. Because SFSTs are limited to unidirectional long-distance processes, they are not powerful enough to compute these maps on their own. Characterizing these maps in terms of independent unidirectional processes is empirically sound, as blocking and other restrictions can vary with directionality (Rose and Walker, 2011).

Emphasis spreading in South Palestinian Arabic is an illustrative example. Emphasis spreads bidirectionally from a pharyngealized segment. Examples are given in (4–6) (Al Khatib, 2008; Jardine,

³These restrictions are necessary to define a subregular class. Without them, composing a left- and right-subsequential map can produce any regular map (Elgot and Mezei, 1965).

2016a); targeted segments are underlined in the SFs. Regressive spreading is unrestricted; the final obstruent in (4) triggers pharyngealization of the entire word. Progressive spreading is blocked by /i, j, ʒ, dʒ/; pharyngealization in (5) affects one vowel and is stopped by the /j/. In (6), there are no blockers, and the medial stop triggers pharyngealization of the entire word.

- (4) $/\chi a j : a : t^f/ \mapsto [\chi^f a^f j^f : a^f : t^f]$ ‘tailor’
 (5) $/s^f a j : a : d/ \mapsto [s^f \underline{a}^f j : a : d]$ ‘hunter’
 (6) $/\text{ʔ} a t^f f a : l/ \mapsto [\text{ʔ}^f \underline{a}^f t^f f^f a^f : l^f]$ ‘children’

While this process is beyond the capability of a subsequential FST, it can be computed by feeding a UR into a left-subsequential FST and its output into a right-subsequential FST. Table 2 makes this explicit, showing emphasis spreading as the outcome of ordering progressive spreading before regressive spreading. Though the SFSTs are not shown here, these processes can be computed by a left- and right-subsequential FST, respectively. As with sibilant harmony, these SFSTs can be restricted to making only one change without affecting their computational complexity. Thus, because weakly deterministic maps can be decomposed into subsequential maps, they can be further decomposed into iterated subsequential maps in a serial derivation.

UR	$/\chi a j : a : t^f/$	$/s^f a j : a : d/$	$/\text{ʔ} a t^f f a : l/$
L→R	—	$s^f \underline{a}^f j : a : d$	$\text{ʔ} a t^f f^f a^f : l^f$
R→L	$\underline{\chi}^f a^f j^f : a^f : t^f$	—	$\underline{\text{ʔ}}^f a^f t^f f^f a^f : l^f$
SF	$[\underline{\chi}^f a^f j^f : a^f : t^f]$	$[s^f \underline{a}^f j : a : d]$	$[\text{ʔ}^f \underline{a}^f t^f f^f a^f : l^f]$

Table 2: Emphasis spreading as two directional processes

2.3 Regular transformations

Weakly deterministic maps can be decomposed into two unidirectional processes because each process has a single trigger. Thus, even though crucial information may be at a distance from a target, a subsequential FST can identify the trigger before it encounters a target. This is not the case with regular transformations that are UNBOUNDED CIRCUMAMBIENT, meaning crucial information lies both to the left and to the right of a target and may be arbitrarily far away in both directions (Jardine, 2016a). This

property means that a subsequential FST will not have enough information when it reaches a target to write the correct output (see Jardine (2016a) for a formal account). In the attested unbounded circumambient processes, the crucial information consists of a trigger and a blocker, as in conditional rightward spreading, or two triggers, as in unbounded tonal plateauing.

Conditional rightward spreading (CRS) is a process in Copperbelt Bemba that exemplifies the combination of a trigger and a blocker. Tone bearing units (TBU) surface with high tones if there is a high tone to the left and there is no high tone at the right edge of the prosodic word; examples are given in (7-8) (Bickmore and Kula, 2013; Kula and Bickmore, 2015; Jardine, 2016a). High tones are indicated with an acute accent (´) and low tones with a grave accent (`). When no underlying high tones are present, words surface with all low tones (7a). In words with high tones and underlyingly toneless final vowels, the rightmost high tone spreads all the way to the right edge; in (7b-c), the subject marker /bá-/ provides the high tone. Word-final high-tones block conditional rightward spreading, and high tones only spread across two TBUs; in (8), the locative enclitic /=kó/ provides the blocker.

- (7) a. /u-ku-tul-a/ ↦ [ùkùtùlà] ‘to pierce’
 b. /bá-ka-fik-a/ ↦ [bákáfíká] ‘they will arrive’
 c. /bá-ka-mu-londolol-a/ ↦ [bákámúlóóndólólá] ‘they will introduce him/her’
- (8) a. /bá-ka-pat-a=kó/ ↦ [bákápátákó] ‘they will hate’
 b. /bá-mu-luk-il-a=kó/ ↦ [bámúlúkìlákó] ‘they will plait a bit for him’
 c. /bá-ka-londolol-a=kó/ ↦ [bákálóóndòlòlákó] ‘they will introduce’

CRS cannot be decomposed into two unidirectional processes because a subsequential FST will not have enough information once it encounters a target TBU. Reading inputs left-to-right, a L-SFST remembers whether a triggering high tone is present, but cannot anticipate whether a blocking high tone is

present word-finally. Likewise, reading right-to-left, a R-SFST remembers whether a blocker is present, but cannot anticipate the presence of a trigger.

Unbounded tonal plateauing (UTP) exemplifies the combination of two triggers. TBUs surface with high tones if there is a high tone to the left and a high tone to the right; examples from Luganda are given in (9-10) (Hyman and Katamba, 2010; Jardine, 2016a). Underlyingly toneless TBUs in phrases without high tones (9a) and in phrases with only one high tone span (9b) surface with low tones. In words with multiple high tone spans, underlyingly toneless TBUs flanked by two high tones surface with high tones (10).

- (9) a. /mu-tund-a/ ↦ [mùtùndà] ‘seller’
 b. /mu-tém-a/ ↦ [mùtémà] ‘chopper’
- (10) a. /mu-tém-a-bi-sikí/ ↦ [mùtémábísíkí] ‘log-chopper’
 b. /tw-áa-láb-w-a walúsimbi/ ↦ [twáálábwá wálúsimbì] ‘we were seen by Walusimbi’
 c. /tw-áa-génd-a na=byaa=ba=walúsimbi/ ↦ [twáágéndá nábyáábáwálúsimbì] ‘we went with those of Walusimbi’

Like CRS, UTP cannot be decomposed into two unidirectional processes. A SFST reading the input in either direction will have only seen one trigger when it identifies a target. Because both triggers must be present for an underlyingly toneless TBU to surface with high tone, the SFST will not have enough information to write the correct output.

Jardine (2016a) argues that CRS and UTP must be modeled by non-deterministic FSTs. Unlike finite state acceptors, non-deterministic FSTs cannot in general be determinized (Lothaire, 2013), and can compute regular maps that deterministic FSTs cannot. The FST for UTP is given in Figure 6 (Jardine, 2016a, 268). Following Jardine (2016a), tonal FSTs read inputs in TBU-sized chunks; U indicates a TBU unspecified for tone and H a high-toned TBU. Anticipating Section 3, while the circumambient processes cannot be decomposed into two subsequential maps like weakly deterministic processes, they can be decomposed into arbitrarily many iterated subsequential maps in a serial derivation.

Non-determinism allows the FST to anticipate the presence of a second trigger, effectively granting it




Figure 6: Non-deterministic FST for UTP (based on Jardine’s Figure 5 (2016, 268))

unbounded lookahead. The FST in Figure 6 accepts an input with exactly one high tone only if it takes the lower path, transitioning to q_2 on its high tone. Inputs with more than one high tone must take the upper path, transitioning to q_1 on the first high tone. In this state, high tones are written for unspecified TBUs. Reading inputs without high tones, the FST does not transition out of q_0 , and faithfully maps the input string. Table 3 gives sample derivations.

/UUHUU/ → [UUHUU]						
I		U	U	H	U	U
Q	q_0	→ q_0	→ q_0	→ q_2	→ q_2	→ q_2
O		U	U	H	U	U
/HUUUH/ → [HHHHH]						
I		H	U	U	U	H
Q	q_0	→ q_1	→ q_1	→ q_1	→ q_1	→ q_2
O		H	H	H	H	H

Table 3: Sample derivations for the FST in Figure 6

Unlike weakly deterministic maps, which can be decomposed into two independently attested subsequential processes, Jardine (2016a) argues that unbounded circumambient processes are different. Taking a derivational perspective, Section 3 argues that in a serial derivation, the regular tonal maps can be decomposed into empirically-motivated subsequential processes that are iterated.

3 Decomposing regular transformations in serial derivations

This section presents the main contribution of this paper, that, under a serial derivation, the computation of the unbounded circumambient processes can be decomposed into iterated subsequential maps which are empirically motivated. In a serial derivation, the left high tone trigger is always adjacent to the target, so no step depends on non-local informa-

tion on both sides. By exploiting this locality, FSTs in serial do not require unbounded lookahead, obviating nondeterminism. In parallel, this connection is not obvious, and the computation requires non-determinism just like the UR → SF map.

The tonal processes discussed in §2.3 have been characterized as spreading (Kisseberth and Odden, 2003; Kula and Bickmore, 2015). Spreading involves associating a tone onto a TBU adjacent to a TBU already associated with that tone. This is represented visually in Figure 7 for UTP; solid lines indicate associations between tones and TBUs (here syllables) in the UR and dashed lines indicate new associations. The high tones of the UR /HUUUH/ trigger spreading, which is assumed to be progressive (Hyman, 2011), yielding the SF [HHHHH].




Figure 7: UTP as high tone spreading

Evidence supporting a spreading analysis comes from inhibitory effects by blockers, which produce partial spreading. A clear example of this is found in Digo, where voiced obstruents impede high tone spreading; examples are given in (11-12) (Kisseberth, 1984). In these examples, UTP is fed by a process that displaces a high tone to the word-final vowel, which realizes as a final rise-fall contour ($\check{V}\hat{V}$). The displaced high tone of the subject prefix /á-/ realizes as this final rise-fall in isolation (11a). In (11b-c), the object prefix /á-/ provides a second high tone, creating the context for UTP. Verb stems with initial voiceless obstruents show rightward spreading (11b), but stems with initial voiced obstruents do not (11c).

The words in (12) have verb stems with initial high tones that interact with the displaced high tone of the tense/aspect prefix /ká-/, creating the context for UTP. The high tone on the verb stem spreads rightwards until it reaches a voiced obstruent. In (12a), rightward spreading is blocked entirely, because the voiced obstruent is adjacent to the left high tone. In (12b), the voiced obstruent is further away, so the high tone spreads, but only across one syllable before it is blocked.

- (11) a. /á-na-tʃukur-a/ ↦ [ànàtsùkùrà] ‘he/she is taking’
 b. /á-na-á-tʃukur-a/ ↦ [ànàátsúkúrâ] ‘he/she is taking them’
 c. /á-na-á-demurir-a/ ↦ [ànàádèmùrìrà] ‘he/she is scolding them’
- (12) a. /a-ká-ézeker-a/ ↦ [àkàézèkèrà] ‘he/she has thatched with/for’
 b. /a-ká-súrúbik-a/ ↦ [àkàsúrúbìk-â] ‘he/she is strong/firm’

The pattern in Digo reveals a local relation between the target and left high tone trigger in UTP. TBUs to the right of a voiced obstruent do not surface with high tones because the left high tone cannot spread across voiced obstruents to establish adjacency. These facts receive a natural explanation under a spreading analysis.

Further evidence for UTP being sensitive to locality comes from Saramaccan Creole. Saramaccan Creole has an underlying three-way contrast between high-toned TBUs, low-toned TBUs, and TBUs unspecified for tone. UTP only targets unspecified TBUs. When a low tone intervenes between either high tone and a span of toneless TBUs, plateauing is blocked; examples are given in (13-14) (Good, 2004; McWhorter and Good, 2012). When toneless TBU spans contact only one high tone, they surface with default low tone (13a). When toneless TBUs are flanked by two high tones, as with the subject and verb in (13b), they surface with high tones. The high tones must be adjacent to the toneless span as the examples in (14) show. In (14a), two low tones intervene between the left high tone and the toneless span, and in (14b), the intervention is between the toneless span and the right high tone. In both cases, spreading is blocked and the unspecified TBUs surface with low tones.

- (13) a. /páúlu lègèdè/ ↦ [páúlu lègèdè] ‘Paul lies’
 b. /dí wómi kulé alá/ ↦ [dí wómí kulé àlá] ‘the man runs there’
- (14) a. /dí káimà kulé alá/ ↦ [dí káimà kulé àlá] ‘the alligator runs there’
 b. /dí wómi bà wáta/ ↦ [dí wómì bà wátà] ‘the man carried water’ (Good, p.c.)

The Saramaccan Creole pattern supports the generalization drawn from Digo that the left high tone trigger must be adjacent to a target TBU. Further, it also shows that the right high tone trigger must be adjacent to the span of toneless TBU targets. This indicates that the structural description is subject to locality constraints defined over the tonal tier rather than being unbounded over the timing tier (Jardine, 2016b; Jardine, 2017).

The locality between the left high tone trigger and the target may be implicit in the UR→SF map, but it is made explicit in a serial derivation. Consider Figure 8, which gives the serial derivation of the UR→SF map in Figure 7. In each step, the target is adjacent to the left high tone trigger, but may be at a distance from the right trigger. Because only the triggering high tone to the right is ever at a distance, no single step is unbounded circumambient.




Figure 8: UTP as high tone spreading in a serial derivation

Each step in the serial derivation of UTP is right-subsequential and can be computed by the R-SFST in Figure 9. Reading an input without a high tone, the R-SFST does not transition out of q_0 , and faithfully outputs the input. Likewise, inputs with one high tone span do not trigger any changes (e.g. /UUHUU/ in Table 4). On inputs with U...H sequences, the R-SFST transitions to q_2 , where it waits to identify a HU string. If it does, it writes two high tones to the output and makes no more changes (e.g. /HUUUH/ in Table 4). If the input ends before another high tone trigger is identified, the end-of-input function writes a toneless TBU to the output.




Figure 9: R-SFST for serial UTP

Looking for a string of a bounded length obviates the need for non-determinism. The R-SFST only needs to lookahead one segment at a time once it

finds the right high tone trigger and a toneless TBU. State q_2 enables lookahead. Transitioning to q_2 , the R-SFST writes nothing onto the output; instead it remembers that it has seen one U. Rather than having to memorize an arbitrarily long sequence of toneless TBUs, the self-loop on q_2 allows it to keep one in memory. If the input ends in q_2 , the memorized U is written to the output. Otherwise, if a second high tone is found, the memorized U is forgotten, and a second high tone is written to the output.

/UUHUU/↦[UUHUU]						
U	U	H	U	U	U	I
q_2	← q_2	← q_1	← q_0	← q_0	← q_0	Q
U	U	λ	H	U	U	O
/HUUUH/↦[HHUUH]						
H	U	U	U	H	H	I
q_3	← q_2	← q_2	← q_2	← q_1	← q_0	Q
λ	HH	U	U	λ	H	O

Table 4: Sample derivations for the FST in Figure 9

CRS can be similarly broken down. The R-SFST in Figure 10 computes each step in a serial derivation, using exactly the same lookahead strategy as the R-SFST for UTP in Figure 9. The only difference is that it checks for a blocker rather than a trigger. Reading an input right-to-left, if the first TBU is high-toned, the machine transitions directly from the starting state to the faithful state q_2 (bounded spreading is assumed to be a separate process). Otherwise, it looks for a HU sequence and spreads the high tone.




Figure 10: R-SFST for serial CRS

4 Discussion

Phonologists have long debated whether to describe phonological processes as applying in parallel or in serial. In rule-based models, the debate is whether a rule applies simultaneously to a string with multiple foci (Chomsky and Halle, 1968; Anderson, 1974) or whether rules apply to each focus one-by-

one (Howard, 1972; Johnson, 1972; Lightner, 1972; Kenstowicz and Kisseberth, 1977). In constraint-based models, the debate is whether competing possible outputs may differ from inputs in unlimited ways, as in parallel Optimality Theory (Prince and Smolensky, 1993/2004), or whether they can only differ from the input by only one change, as in Harmonic Serialism (McCarthy, 2000, et seq.), where the derivation iterates until converging. This is a fundamental question in phonological theory.

This paper has argued that serial phonological models, where only one focus is changed at a time, are advantageous to characterizing the class of $UR \mapsto SF$ maps. Parallel models of phonology, where derivations are identical to $UR \mapsto SF$ maps, do not offer stronger generalizations than the Regular Hypothesis. Serial models, on the other hand, allow for the decomposition of $UR \mapsto SF$ maps, yielding the Serial Subsequential Hypothesis. The Serial Subsequential Hypothesis is stated as a uniform generalization over phonological transformations and does not need to distinguish between segmental and tonal processes. This is desirable as a general, restrictive characterization.

As a general characterization, it can reduce the computational differences between related phonological processes. Consider the case of UTP, which avoids sequences of $H \dots U \dots H$ from surfacing (Yip, 2002). Cross-linguistically, this sequence is also avoided by deleting the second high tone, giving maps like $/HUH/ \mapsto [HU\underline{U}]$. Examples of this progressive lowering process in Barasana are given in (15-16) (Gomez-Imbert and Kenstowicz, 2000; Gomez-Imbert, 2001; Hyman, 2010); the tilde \sim marks nasalized stems. This is also attested in Yongning Na (Michaud, 2017). The high tone of the diminutive suffix $/-áka/$ surfaces when attached to stems with all high-toned TBUs (15a) and stems with UH contours (15b). Attaching the suffix to stems with HU contours creates a $H \dots U \dots H$ sequence, triggering the diminutive high tone to lower (16). Progressive lowering also targets the underlying high tone on the suffix $/-íri/$ (16c). This process is left-subsequential; the L-SFST in Figure 11 computes the $UR \mapsto SF$ map.

- (15) a. $/\sim kúbú-áka/ \mapsto [\sim kúbúákà]$ ‘small shaman’

- b. /gohé-áka/ \mapsto [gòhéákà] ‘small hole’
- (16) a. / \sim céda-áka/ \mapsto [\sim cédààkà] ‘a bit of pineapple’
- b. / \sim céda-a-áka/ \mapsto [\sim cédààkà] ‘a small pineapple’
- c. / \sim céda-a-íri-áka-re/ \mapsto [\sim cédààriàkàrè] ‘small pineapples-OBJ’

Progressive lowering and UTP target the same marked structure. In a parallel derivation, languages like Luganda use a regular process to repair H...U...H sequences, while languages like Barasana use a subsequential repair. This computational gap disappears in a serial derivation, where both repairs comprise iterated subsequential maps.




Figure 11: L-SFST for progressive lowering

As a restrictive characterization, the Serial Subsequential Hypothesis predicts that certain regular maps are not possible phonological transformations. For example, consider a variant of Saramaccan Creole, Saramaccan', in which low tones do not act as blockers. In Saramaccan', arbitrarily many low tones may intervene between a toneless TBU and the two triggering high tones, yielding maps such as /HL^mULⁿH/ \mapsto [HL^mHLⁿH], where L indicates a low-toned TBU. This cannot be modeled by an iterated subsequential map, as there is no guarantee that a SFST will identify the second high tone trigger within a bounded distance from the target. Besides its larger alphabet, the non-deterministic FST for Saramaccan' in Figure 12 is indistinguishable from that for UTP in Figure 6. Without considering the derivation, it is not clear how to predict the non-existence of Saramaccan', or for that matter, any regular map in a principled way.

While the Serial Subsequential Hypothesis excludes some regular maps like Saramaccan', it also overgenerates. Consider the phonological rule $\emptyset \rightarrow ab / a_b$. Kaplan and Kay (1994) demonstrate that iteratively applying this rule $n \geq 1$ times to an input /ab/ produces the context-free string set $a^n b^n$. Such a derivation is within the scope of the Serial Sub-




Figure 12: Non-deterministic FST for Saramaccan'

quential Hypothesis, because the rule can be computed subsequentially. The overgeneration caused by this particular rule stems from its ability to apply without motivation. Moreton (2004) argues that constraint-based frameworks like Optimality Theory, which limit processes to apply only if motivated by marked structures, avoid circular processes and infinite augmentation. In the case of inserting [ab] into the string, there is no clear improvement in these terms. Imposing restrictions of this sort from formal phonological analyses is a promising direction to take to characterize the class of serial SFSTs.

5 Conclusion

This paper argued that a phonologically uniform characterization of UR \mapsto SF maps obtains under a serial model of phonology. Phonological transformations can be decomposed into iterated left- or right-subsequential maps, when computed in serial. Combined with restrictions already imposed on phonological computations, the Serial Subsequential Hypothesis was argued to make restrictive typological predictions. This demands a precise characterization of the class of maps produced by iterating subsequential maps, and is left to future work.

The result of this paper draws on the interaction between formal language theory and traditional phonology. The decision to examine whether UR \mapsto SF maps are computed in parallel or serial is not arbitrary, but has been an important question in phonological research for decades. A formal language understanding of phonological transformations is enriched by an appreciation for the models that compute them. Likewise, the models should be restricted by a knowledge of the formal language landscape. The interface between these two disciplines can produce interesting generalizations and should continue to be probed.

Acknowledgments

This work has greatly benefited from discussions with Carolyn Anderson, Thomas Graf, Jeff Heinz, Adam Jardine, Gaja Jarosz, John McCarthy, Joe Pater, Brandon Prickett, Kristine Yu, participants in the Phonology Reading Group and Sound Workshop at the University of Massachusetts, Amherst, and the audience at NECPHON 11, as well as comments from three anonymous reviewers for SCiL 2018. This work was supported by the National Science Foundation through grant BCS-424077. All remaining errors are of course my own.

References

- Sam Al Khatib. 2008. On the directionality of emphasis spread. In Susie Jones, editor, *Proceedings of the 2008 annual conference of the Canadian Linguistic Association*.
- Stephen R. Anderson. 1974. *The Organization of Phonology*. Academic Press, New York.
- Lee S. Bickmore and Nancy C. Kula. 2013. Ternary spreading and the OCP in Copperbelt Bemba. *Studies in African Linguistics*, 42:101–132.
- Jane Chandlee and Jeffrey Heinz. 2012. Bounded copying is subsequential: Implications for metathesis and reduplication. In *Proceedings of SIGMORPHON 2012*, pages 42–51.
- Jane Chandlee. 2014. *Strictly local phonological processes*. Ph.D. thesis, University of Delaware.
- Noam Chomsky and Morris Halle. 1968. *The Sound Pattern of English*. Harper & Row, New York.
- C. C. Elgot and J. E. Mezei. 1965. On relations defined by generalized finite automata. *IBM Journal of Research and Development*, 9:47–68.
- Brian Gainor, Regine Lai, and Jeffrey Heinz. 2012. Computational characterizations of vowel harmony patterns and pathologies. In *Proceedings of the 29th West Coast Conference on Formal Linguistics*, pages 63–71.
- Elsa Gomez-Imbert and Michael Kenstowicz. 2000. Barasana tone and accent. *International Journal of American Linguistics*, 66:419–463.
- Elsa Gomez-Imbert. 2001. More on the tone versus pitch accent typology: Evidence from Barasana and other Eastern Tukanooan languages. In Shigeki Kaji, editor, *Proceedings of Symposium Cross-Linguistic Studies of Tonal Phenomena: Tonogenesis, Japanese Accentology, and Other Topics*, pages 369–412.
- Jeff Good. 2004. Tone and accent in Saramaccan: charting a deep split in the phonology of a language. *Lingua*, 114:575–619.
- Gunnar Ólafur Hansson. 2010. *Consonant Harmony: Long-Distance Interactions in Phonology*. University of California Press, Berkeley, CA.
- Jeffrey Heinz and Regine Lai. 2013. Vowel harmony and subsequentiality. In *Proceedings of the 13th Meeting on the Mathematics of Language*, pages 52–63.
- Jeffrey Heinz. forthcoming. The computational nature of phonological generalizations. In Larry Hyman and Frans Plank, editors, *Phonological Typology*. Mouton.
- Irwin Howard. 1972. *A Directional Theory of Rule Application*. Ph.D. thesis, Massachusetts Institute of Technology.
- Larry M. Hyman and Francis Katamba. 2010. Tone, syntax and prosodic domains in Luganda. *ZAS Papers in Linguistics*, 53:69–98.
- Larry M. Hyman. 2010. Amazonia and the typology of tone systems. Technical report, UC Berkeley Phonology Lab Annual Reports.
- Larry M. Hyman. 2011. Tone: Is it different? In John Goldsmith, Jason Riggle, and Alan C. L. Yu, editors, *The Handbook of Phonological Theory*, pages 197–239. Blackwell Publishing, Malden, MA, 2 edition.
- Adam Jardine. 2016a. Computationally, tone is different. *Phonology*, 33:247–283.
- Adam Jardine. 2016b. *Locality and Non-Linear Representations in Tonal Phonology*. Ph.D. thesis, University of Delaware.
- Adam Jardine. 2017. The expressivity of autosegmental grammars. Unpublished manuscript.
- C. Douglas Johnson. 1972. *Formal Aspects of Phonological Description*. Mouton, The Hague.
- Ronald Kaplan and Martin Kay. 1994. Regular models of phonological rule systems. *Computational Linguistics*, 20:331–378.
- Michael Kenstowicz and Charles Kisseberth. 1977. *Topics in Phonological Theory*. Academic Press, New York.
- Charles W. Kisseberth and David Odden. 2003. Tone. In Derek Nurse and Gérard Philippson, editors, *The Bantu languages*, pages 59–70. Routledge, London.
- Charles W. Kisseberth. 1984. Digo tonology. In George N. Clements and John Goldsmith, editors, *Autosegmental Studies in Bantu Tone*, pages 105–182. Foris Publications, Dordrecht.
- Nancy C. Kula and Lee S. Bickmore. 2015. Phrasal phonology in Copperbelt Bemba. *Phonology*, 32:147–176.
- Theodore M. Lightner. 1972. *Problems in the Theory of Phonology*. Linguistic Research, Inc., Edmonton.
- M. Lothaire. 2013. *Applied Combinatorics on Words*. Cambridge University Press, Cambridge.
- Huan Luo. 2017. Long-distance consonant agreement and subsequentiality. *Glossa*, 2(52).

- John J. McCarthy. 2000. Harmonic serialism and parallelism. In Masako Hirotsu, Andries Coetzee, Nancy Hall, and Ji-Yung Kim, editors, *Proceedings of the North East Linguistics Society 30*, pages 501–524. GLSA, Amherst, MA.
- Adam McCollum, Eric Baković, Anna Mai, and Eric Meinhardt. 2017. Conditional blocking in Tugrugbu requires non-determinism: implications for the subregular hypothesis. Presentation at NELS 48.
- John H. McWhorter and Jeff Good. 2012. *A Grammar of Saramaccan Creole*. De Gruyter Mouton, Berlin.
- Alexis Michaud. 2017. *Tone in Yongning Na*. Language Science Press, Berlin.
- Mehryar Mohri. 1997. Finite-state transducers in language and speech processing. *Computational Linguistics*, 23:269–311.
- Elliott Moreton. 2004. Non-computable functions in Optimality Theory. In John J. McCarthy, editor, *Optimality Theory in Phonology: A Reader*, pages 141–163. Blackwell Publishing, Malden, MA.
- Amanda Payne. 2014. Dissimilation as a subsequential process. In Jyoti Iyer and Leland Kusmer, editors, *NELS 44*, volume 2, pages 79–90.
- Alan Prince and Paul Smolensky. 1993/2004. *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell Publishing, Malden, MA.
- Sharon Rose and Rachel Walker. 2011. Harmony systems. In John Goldsmith, Jason Riggle, and Alan C. L. Yu, editors, *The Handbook of Phonological Theory*, pages 240–290. Wiley-Blackwell, Malden, MA, 2 edition.
- Maira Yip. 2002. *Tone*. Cambridge University Press, Cambridge.