From Agglutination to Fusion
Coalescence in the Athabaskan Verb

Alessandro Jaker, Stanford University
Qualifying Paper #2
Final Version

Acknowledgements

I wish to thank my consultants, Mary Louise Drygeese and Michel Paper, the staff of the Goyatiko Language Center, Mary Rose Sundberg and Betty Harnum, and the people of the Yellowknives Dene First Nation in Yellowknife, NWT for all their help and support, without which this work would not be possible.

I also wish to thank my committee, Will Leben, Paul Kiparsky, Arto Anttila, Peter Sells, and Joan Bresnan, as well as Keren Rice and participants of WAIL (Workshop on American Indigenous Languages) in May 2007 for comments on previous versions of this paper. All remaining errors are my own.

Table of Contents

Agglutination to Fusion

1.0 Introduction.

Coalescence is a process by which, synchronically or diachronically, two constituents merge into a single constituent. Coalescence may happen at any level: two segments may fuse together to become a single segment (often combining properties of both of the original segments), two syllables may become one syllable, two moras may become a single mora (as I shall argue), and, ultimately, two morphemes may become a single morpheme. At the same time, coalescence is still relatively poorly understood. This is because, by its very nature, coalescence means that the units of representation in the output are different from those in the input, which makes it difficult to evaluate input-output mappings. For example, if a sequence of two vowels /a + i/ in the input fully coalesces into a single long vowel [ee] in the output, how is this to be evaluated? Does the output contain the same two segments as in the input, which have merely assimilated to each other? Or is there now only a single, new segment, which merely preserves features from both underlying segments? What about the morpheme boundary—is it still present in the output? To address these questions, in this paper I will present a formal model for the proper representation of coalescence, as well as undertake a detailed case study of Dogrib, an Athabaskan language of northern Canada, whose phonology essentially revolves around coalescence.

1.1 Agglutination to Fusion in Four Stages.

Diachronically, the ultimate end-point for phonological coalescence is morphological fusion: for some pair of morphemes, the underlying representation is re-structured such that instead of two separate morphemes, we have a single, separate lexically listed form, which carries all of the morphosyntactic features of the original separate morphemes. Synchronically, morphological fusion may be thought of as the most extreme form of coalescence, a state which diachronically arises through morphological reanalysis (cf. section 4.0). The representation of morpheme boundaries plays a crucial role in our understanding of coalescence. In order to account for cases of partial fusion, such as Northern Athabaskan ablaut (cf. section 3.1), I argue that it is necessary that both phonological constraints and phonological representations make reference to morpheme boundaries. I posit an independent morphological level of representation with which both segments and suprasegmentals are associated. It then becomes possible, e.g. to speak of a Low tone violating an alignment constraint, if it crosses a morpheme boundary, or of faithfulness to the underlying form of a morpheme (both of which have typological implications). With morphemes reified on an independent level of representation, it also becomes possible to speak of a progression from agglutination to fusion in four stages, as presented in Figure 1.1 below.

Figure 1.1: Four Stages towards Morphological Fusion[1].

Stage 1: Crisp Agglutination	Stage 2: Harmonic Agglutination	Stage 3: Partial Fusion	Stage 4: Total Fusion

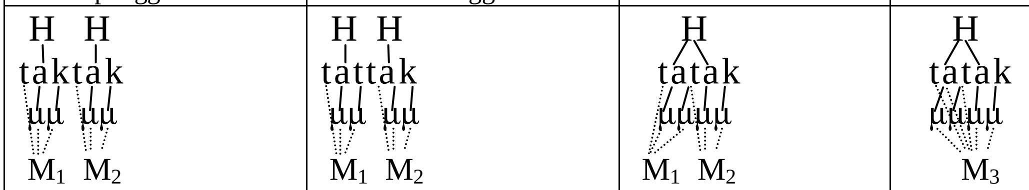

Figure 1.1 presents a schematic representation of four diachronic stages from agglutination to fusion. These stages are characterized by different autosegmental and morphological representations, each of which are, in turn, derived from different constraint-rankings in an Optimality-Theoretic grammar (cf. Figure 1.2 below). I refer to stage 1 as “crisp agglutination”. In this stage, when two morphemes are concatenated, the overall mora count for the string remains unaffected, and there are no segmental or tonal changes. Crucially, it is also the case that there are no units of representation—tones, segments, or moras—which cross a morpheme boundary. Stage 1 is a highly idealized state of affairs which seldom obtains in practice, i.e., a language without morphophonemic alternations; nevertheless, some languages approach this state of affairs, and so it must be describable by the theory. Stage 2, “harmonic agglutination,” is a more common state of affairs, the state of affairs, e.g., which is found in vowel harmony. At this stage, surface morphs remain discrete (i.e. segments, tones, and moras do not cross morphological boundaries), but morphophonemic alternations occur, motivated by Markedness. For example, in Figure 1.1, /*taktak*/ has assimilated to [*tattak*]; nevertheless the number of segments is the same. Jumping ahead, the final stage, “total fusion,” or stage 4, describes a state in which a single morpheme carries all

[1] Some association lines have been omitted for presentational clarity.

the morphosyntactic features of what were originally two separate morphemes. As the result consists only of a single morpheme, with no internal structure, this stage contains no adjacent identical segments, in accordance with the OCP (Leben 1973, Keer 1999). Morphophonologically, stage 4 is in fact identical to stage 1 (or 2); it differs only in the number of morphosyntactic features that the morpheme carries. The existence of stages 1, 2, and 4 is uncontroversial; what is unique to the current proposal is stage 3, which I call "partial fusion," which is defined as a sharing of phonological structure across a morpheme boundary. In stage 3 in Figure 1.1, we see a form [*tattak*], in which the geminate *t* consists not of two adjacent identical segments, but rather a single segment associated with a mora. This representation of geminates is the one preferred by the OCP, and, indeed is predicted to be the only possible representation morpheme-internally (Keer 1999). Across a morpheme boundary however, sequences of adjacent identical segments (or tones) which arise through concatenation might not be simplified by the OCP due to the role of Alignment, which militates against phonological constituents crossing morpheme boundaries (cf. section 2.3; McCarthy & Prince 1995, Keer 1999). To summarize, the progression from agglutination to fusion proceeds as follows: In stage 1, morphemes are concatenated without affecting each other in any way. In stage 2, adjacent morphs influence each other so that they agree in certain phonological features, resulting in adjacent identical elements. In stage 3, adjacent identical elements are simplified into a single element shared across a morpheme boundary, in order to satisfy the OCP. Finally, in stage 4, the morpheme boundary itself is erased, resulting in a simpler representation. Each of these four stages is derived from a different constraint-ranking in OT, as illustrated in Figure 1.2 below.

Figure 1.2: Constraint Rankings Deriving Each Stage:

Stage(s)	**Constraint Ranking**
Crisp Agglutination, Harmonic Agglutination	NO-ALLOMORPHY, ALIGN-R(MORPH, X) >> OCP(X)
Partial Fusion	NO-ALLOMORPHY, OCP(X) >> ALIGN-R(MORPH, X)
Total Fusion	OCP(X), ALIGN-R(MORPH, X) >> NO-ALLOMORPHY

Stages 1 and 2, Crisp Agglutination and Harmonic Agglutination, are characterized by violations of the OCP. For example in the form *tat-tak*, as represented in stage 2 in Figure 1.1, there are two adjacent, idential segments: *t-t*. In the Harmonic Agglutination stage, it may be that there are some forms which do not violate the OCP, such as *tak-tak* in stage 1 in Figure 1.1; nevertheless, the key point is that, in stages 1 and 2, representations in which a single phonological constituent crosses a morpheme boundary are not tolerated. If adjacent identical elements should arise through morphological concatenation, they must be tolerated in spite of the OCP violation, since it is more important to maintain crisp edges, as demanded by Alignment.

In Stage 3, the Partial Fusion stage, the OCP is promoted over Alignment. Recall that in stage 3 in Figure 1.1, the form *tattak* is represented as having a single segment, *t*, linked to a mora and shared across two surface morphs, and likewise a single High tone is shared across two vowels, straddling the morpheme boundary. Both of these constitute an Alignment violation: specifically, the constraint Align-R(Morph, X) assigns one violation mark to each phonological constituent *X* for each morpheme boundary that it crosses (cf. section 2.3.2 for a formal definition). In other words, in stage 3, the OCP causes phonological structure to be shared across morpheme boundaries, in violation of alignment.

While the difference between stages 2 and 3 is a trade-off between the OCP and Alignment, in Stage 4 the language seeks to have things both ways: Alignment is once again promoted, so that both Alignment and the OCP are undominated. How is this possible? Recall that the OCP is a constraint against adjacent identical elements, while Alignment is a constraint against constituents which cross morpheme boundaries. There is a very straightforward way to satisfy both of these, namely, *eliminate the morpheme boundary*. In one sense, this morphological reanalysis happens in the course of every derivation, as the result of bracket-erasure after each cycle (cf. section 1.2). This type of synchronic fusion poses essentially no problem, as it is follows as an automatic consequence of cyclicity. More problematic, however, is *diachronic* fusion. The constraint violated, NO-ALLOMORPHY, is in fact not a phonological constraint at all, but a constraint in the morphology, which in effect penalizes morphemes which carry multiple morphosyntactic features. Impressionistically, diachronic fusion may be described as a situation in which some surface string which previously was phonologically derived by the concatenation of two or more morphemes comes to be lexically listed as-is as a single morphological constituent in the underlying representation. Phonologically, this has several advantages: OCP violations are eliminated, there are no Alignment violations (since there are no longer any morpheme boundaries to be crossed), and some Faithfulness violations can be elminated as well, insofar as some phonological processes which resulted from morphological concatenation are eliminated as the underlying representation is re-structured to more closely match the surface form. Morphologically, however, the language becomes more complex (in a sense) insofar as there is now a greater inventory of morphemes which must be stored in the lexicon. To summarize, Total Fusion amounts to a trading of phonological complexity for morphological complexity. The formal implementation of this process is problematic, yet see section 4.0 for extended discussion.

1.2 Morphological Domains in the Athabaskan Verb.

All four stages presented in Figure 1.1—Crisp Agglutination, Harmonic Agglutination, Partial Fusion, and Total Fusion, exist synchronically within the verb-complex of Northern Athabaskan languages, and are associated with different lexical strata, within the framework of Lexical Phonology (Kiparsky 1982, to appear). These stages also result in different morphophonemic effects, which will be the topic of an in-depth case-study in section 3.0. The starting point for this analysis is my claim that, in Northern Athabaskan, there are not 3 but 4 lexical strata, as illustrated in Figure 1.3: the Root Level, Stem Level, Word Level, and Postlexical level.

Figure 1.3: Morphological Domains and Bracket Erasure.

Stratum	**Structure**
Root Level	[Root + A_1, A_2 +....A_n]$_{Stem}$
Stem Level	[C_1 + C_2 + C_3 + ...C_n + Stem]$_{ConjunctDomain}$
Word Level	[D_1 + D_2 + D_3 +....D_n +ConjunctDomain]$_{ProsodicWord}$
Postlexical	[Clitics + PrWd + PrWd + PrWd + ...PrWd]$_{ProsodicPhrase}$

The input to the Root Level phonology consists of the verbal root itself, plus one or more ablaut-inducing suffixes (cf. section 3.1), which, when combined, form the stem. At the Stem Level, the stem-internal brackets are erased (i.e. between the root and ablauting suffixes), and the stem is combined with the conjunct prefixes. At the Word Level, the morphological brackets between the stem and conjunct prefixes are erased, and the entire conjunct domain as a whole is combined with the disjunct prefixes. Finally, at the Postlexical level, the entire prosodic word (which now contains no internal morphological structure) is combined with clitics and other prosodic words, to participate in phrasal phonology (e.g. intonation) as part of a prosodic phrase. An example may suffice to illustrate—consider the form *bòkàweht'è*, which is the 1st person singular optative of the verb 'cook' (cf. section 3.2), which is parsed broken up into morphemes as *bò-kà-we-h-t'e-è*. A cyclic derivation of this form is given in Figure 1.4 below.

Figure 1.4: Derivation of *bòkàweht'è* in cycles.

Stratum	Structure
Root Level	bò + kà + we + h + [t'e + è]$_{Stem}$
Stem Level	bò + kà + [we + h + t'è]$_{ConjunctDomain}$
Word Level	[bò + kà + weht'è]$_{ProsodicWord}$
Postlexical	[bòkàweht'è]$_{ProsodicPhrase}$

While all of the input morphemes are, in an abstract sense, present throughout the derivation, at the Root Level, the grammar operates only on the root domain, which consists of the root plus the ablauting suffix. Thus, *bò, kà, we,* and *h* are not visible to the grammar at this stratum. At the Root Level, *t'e* and *è* combine to yield an output *t'è* (cf. section 3.1), which is a stem. The stem *t'è*, which is the output of the Root Level, becomes the input to the stem level. Before the Stem Level phonology can perform any operations on the stem, all internal morpheme boundaries are erased; therefore, *t'è*, as an unanalysable unit, is combined with all of the conjunct prefixes (in this case *we* and *h*) to form the conjunct domain. Similarly, at the Word Level, the entire conjunct domain is combined with the disjunct prefixes, and so forth.

I claim that in Dogrib, the Root Level is characterized by partial fusion (cf. section 3.1). Thus, while the input representation for the root for 'cook' plus its suffix is /t'e + è/, the output is represented as [t'[è]]: the High toned vowel of the root and the Low toned vowel of the suffix have coalesced into a single, Low toned vowel, which is shared between both morphemes. This situation, characteristic of partial fusion, incurs a violation of Alignment, since a segment extends across a morpheme boundary (cf. section 2.3.2), as well as violation of MAX(H) from the perspective of the root (cf. section 2.2). When *t'è* becomes the input to the Stem Level, however, its internal brackets are erased, and it is combined, as a unit with *we* and *h*. The Stem Level, unlike the Root Level, is characterized by harmonic agglutination. While not evident in the particular form *bòkàweht'è*, the Stem Level is characterized by a complex set of phonological processes resulting from the concatenation of the conjunct prefixes with the stem—these are oral vowel coalescence (cf section 3.2.2.1), nasal coalescence (cf. section 3.2.2.2), vowel deletion (cf. section 3.2.2.3) and sonorant deletion (cf. section 3.2.2.4). The Stem Level processes themselves are interesting in that these are all segmental processes conditioned by tone, a state of affairs which is typologically unusual and has in fact been claimed to be impossible (Blumenfeld 2006). Regarding the larger typology of agglutination to fusion, the relevant observation is that these coalescence processes are represented as being assimilatory in nature. For example, for a case of nasal coalescence such as *bòkàwhį̨t'e* 'you have cooked,' is represented as /bò + kà + whe + ne + t'e/ in the input to the Stem Level phonology, and [bò][kà][whį̨][į̨][t'e] in the output (where the square brackets signify morphological constituency). Finally, it was historically the case, until recently, that the Word Level was characterized by crisp agglutination. That is, phonological processes that normally occurred with conjunct prefixes were blocked in disjunct prefixes (Ackroyd 1982, Marinakis 2002). This is illustrated in Figure 1.5.

Figure 1.5: Conjunct vs. Disjunct Prefixes at the Word Level.

	Conjunct Prefix	Disjunct Prefix
Conservative	/bò + kà + whį̨t'e/ → [bò][kà][**whį̨į̨**t'e]	/nà + nezè/ → [ną̀][nezè]
Innovative	/bò + kà + whį̨t'e/ → [bò][kà][**whį̨į̨**t'e]	/nà + nezè/ → [**ną̀**][**ą**zè]

Figure 1.5 illustrates the process of Nasal Coalescence (cf. section 3.2.2.2), which is highlighted in bold in Figure 1.5, in the cases where it has applied. In terms of Stratal OT, we may say that formerly nasal coalescence was exclusively a Stem Level process, which in innovative speech has been extended to the Word Level. Diachronically, these facts further illustrate the general progression from agglutination to fusion. The Athabaskan verb complex arose historically by adding more and more prefixes to the left of the stem: the conjunct prefixes, being diachronically older, are consistent in their linear order across the family, and engage in many coalescence processes; the disjunct prefixes, having been added more recently, show more variation in their ordering and enter into fewer morphophonemic alternations (Rice 2000: 79). Thus Figure 1.5 illustrates yet another step from agglutination towards fusion, which is currently in progress, as this varies according to conservative or innovative speech.

Finally, total fusion is, in the present model, by definition beyond the scope of phonology proper. Totally fused morphemes are found in Northern Athabaskan, though their distribution is governed by morphological, not phonological, principles (Jaker

2006). Nevertheless, in section 4.0 I will discuss some of the mechanisms which may induce the final transition from partial to total fusion.

2.0 Theoretical Background.

As mentioned in section 1.0, coalescence is theoretically problematic because the units of representation involved—segments, moras, syllables, tones, and even morphemes—are not held constant between input and output. In this section, I outline my assumptions regarding morphological structure and autosegmental representations, to enable a precise formal treatment of coalescence in section 3.0.

2.1 Autosegmental Tiers: what tier are morphemes on?

There has been no general consensus in the literature regarding how autosegmental tiers should be arranged, or whether any such arrangement should hold cross-linguistically. Schematically, three types of tier structure are shown in Figure 2.1, of which numerous permutations are possible: I call these the Multidirectional Model, the Mediation Model, and the Hierarchical Model.

Figure 2.1: Three Models of Associations between Autosegmental Levels.

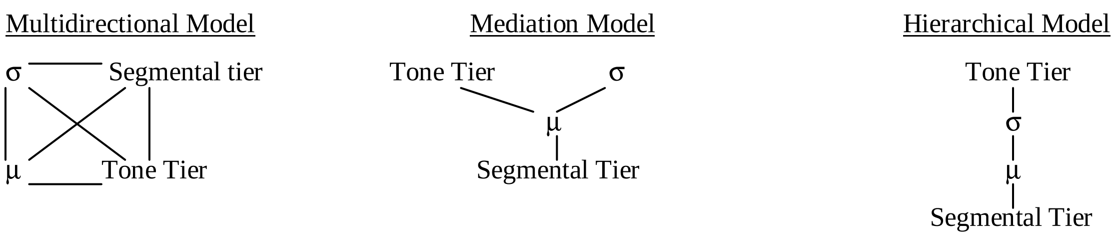

A hierarchical model (though not of the type shown in Figure 2.1) was proposed by Mester (1988). In this model, the segmental tier is subdivided into head tiers and dependent tiers, according to a theory of feature geometry. For example, in Ngbaka, high front and high back vowels do not cooccur (*uCi), and non-high front and non-high back vowels do not cooccur (*oCe). In Mester's model, this is derived from a tier structure in which backness is dependent upon height (1988: 9-13). A version of the mediation model is presented in Hyman (2003 [1986]). In Hyman's model, segments by default linked to a Weight Unit (i.e. mora), unless de-linked by an Onset Creation Rule or Margin Creation Rule, to derive non-moraic onsets and codas, respectively (1986: 15-21). These weight units are then linked to tones and/or syllables. Finally, the multidirectional model has not been explicitly argued for in the literature, to my knowledge. This represents a sort of null hypothesis: any type of phonological interaction is possible in such a framework. It seems, in fact, that every type of interaction predicted by a multidirectional model is in fact attested (*contra* Blumenfeld 2006). For example, glottalic segments may affect tone (Krauss 2005, Kingston 2005) just as tone may trigger gemination (cf. section 3.2.3.4) and vowel deletion (cf. section 3.2.2.3). Contour tones may be restricted by moras, as in Northern Athabaskan languages, or by syllable count, as in ancient Greek (Kiparsky, to appear). It is beyond the scope of this paper to explore every typological possibility predicted by the multidirectional model; rather, I will simply assume it, as in the constraint MAXASSOC(V, TONE), "every vowel must remain associated with its tone," (cf. section 3.2.2.3) and illustrate those interactions between autosegmental tiers which are relevent in Athabaskan languages.

Assuming that any tier may interact with any other tier, the question arises, on what tier are morphemes located? Strictly speaking, morphemes are not located on any tier, but are off in another dimension, as illustrated in Figure 2.2.

Figure 2.2: Morphemes are not on a 'tier'.

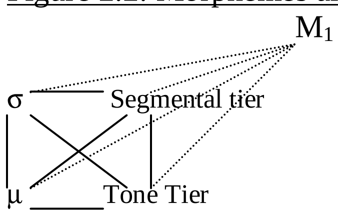

That is, a morpheme has no inherent linear properties, but is an abstraction regarding the way in which morphosyntactic features are associated with segments, tones, etc., which are arranged in a linear order on some tier. Such a conception poses problems for the use of Alignment constraints, however: in what way does it make sense, e.g., to align a tone with the edge of a morpheme, if a morpheme cannot be identified as a linear string? In practice, alignment with a morphological category is almost always interpreted as alignment with the edgemost *segment* associated with that morphological category; this applies both to Alignment proper (McCarthy & Prince 1993) and the closely related constraint family CRISPEDGE (Ito & Mester 1999, Pater 2001). I will use the same convention in this paper: thus, while strictly speaking, it is not correct to speak of alignment with a *morpheme*, a *morph* can be said to have left and right edges, which are defined as the leftmost and rightmost elements on the segmental tier associated with some set of morphosyntactic features. Thus, in summary, while any tier may interact with any other tier, misalignment is understood as misalignment with some segment (cf. section 2.3.2 for further discussion).

2.2 Morpheme-Based Faithfulness.

In the formal analysis beginning in section 3.0, I will use a definition of Faithfulness which is fundamentally different from the standard Correspondence Theory definition (McCarthy & Prince 1995). I call this *Morpheme-Based Faithfulness*, and the formulation of the faithfulness constraint itself relies crucially on morphological bracketing and level ordering. For illustrative purposes, the full formal definition of MAX is given below.

Figure 2.3: Definition of MAX:

·Let M be a minimal morphological constituent at level L.
·Given a candidate ([o], <i>, <o>, ℛ), where:
  ·<o> is a concatenative decomposition of the output [o], ***such that every element in <o> is associated with M***, and
  ·<i> is the underlying representation of M, and
  · ℛ is a total bijective function from <i> to <o>,
for every string κ in <i> where ℛ (κ) = #
  for every element in κ
    assign a violation mark.

Informally: "At every level, the output domain of a morph must contain everything in the input of that morph."

At first glance, what may seem unfamiliar about the above definition is the use of the theory of *string-based correspondence* (McCarthy & Wolf 2005). This notion of correspondence differs from the classical notion of segment-based or element-based correspondence in that the correspondence relation is viewed as *categorical* between input and output (formally, a "total bijective function"). Under element-based correspondence, if a segment from the input was deleted in the output, it was assumed that the correspondence relation did not hold (since it is not possible to be in correspondence with nothing). Under string-based correspondence, however, the correspondence relation is assumed to always hold. This requires the introduction of null elements into the phonology: thus deletion is represented as a correspondence relation between some segment in the input and "#" in the output. Likewise epenthesis is a relation between some segment in the output and "#" in the input. The only exceptions to this involve violations of, in McCarthy's model, the single constraint MPARSE, or, in my model, NO-ALLOMORPHY (cf. section 4.0).

For the purposes of the present analysis, however, what is crucial to the definition in Figure 2.3 is not string-based correspondence per se, but the association clause, highlighted in bold: every string in the output must be associated with the morpheme, *M*, in order to count for satisfying MAX. Put another way, it is not sufficient that some segment from the input be present somewhere in the linear string of the output; rather, it must be present in the output and morphologically affiliated with the same morpheme in the input.

An example illustrating the importance of morphological affiliation is given in Figure 2.4 below.

Figure 2.4: Output Moras under Partial Fusion.

/ta + ak/ μ μμ	MAX(μ)	DEP(μ)
a. all moras retained tak μμμ M_1 M_2		
b. all moras retained tak μμμ M_1 M_2		*(M_1)
c. loss of one mora tak μμ M_1 M_2		
d. loss of one mora tak μμ M_1 M_2		*(M_1)

In Figure 2.4, we see an input string /ta + ak/, which contains three moras: one mora from the first morpheme, and two moras from the second morpheme. The candidates, (a)-(d), vary not only in the overall number of morphemes in the entire string, but also in the way these are affiliated morphologically, represented as a dotted line linking each of the moras to the morphemes M_1 and/or M_2. Candidate (a) is a faithful parse, in which all moras from the input string are retained (a total of 3), and their morphological affiliation does not change between input and output. However, note that candidate (c) is *also* a faithful parse, even though one mora has been deleted: it incurs no violations of either MAX(μ) or DEP(μ). How is this possible? Recall that faithfulness is defined at the level of morphemes. In Figure 2.4, M_1 was associated with one mora in the input, while M_2 was associated with 2 moras in the input. In candidate (c), this holds in the output as well, by virtue of the fact that the first mora is doubly associated. In a similar fashion, candidate (b) incurs a violation of DEP(μ), even though the overall number of moras in the output is held constant, since the 2^{nd} mora is doubly associated, and, in candidate (d), a DEP(μ) violation is incurred even though the overall number of moras in the output is reduced, because both moras are doubly affiliated morphologically.

That phonological units may change their morphological affiliation between input and output has not, to my knowledge, been proposed in the OT literature to date. Two objections might be raised against this proposal: one typological, the other formal. Typologically, it may seem that the morpheme-based faithfulness proposal allows for unconstrained re-affiliation between morphemes and phonological units. This is not the case. In the vast majority of cases, it is in fact Faithfulness which prevents morphological re-bracketing. Take for example an input sequence /pat + ka/. If the overall string is held constant, to maintain the same morphological bracketing ([pat][ka]) incurs no faithfulness violations, to move the bracket over one segment incurs one violation of MAX and one of DEP ([patk][a]), and to move the bracket over two segments incurs two violations of each (i.e. [patka][]). Because, in all of these cases, the entire string is the same, all of these candidates fare equally well with respect to markedness; thus those candidates which exhibit re-bracketing incur faithfulness violations gratuitously, and are therefore harmonically bounded. It is only in those cases in which, across a morpheme boundary, there are two phonological units which are adjacent and identical (or at least identical in some features) that re-bracketing is desirable, on account of the OCP (cf. section 2.3.1). In this way, the formal system presented here is able to straightforwardly derive the typology sketched out in Figure 1.1, as well as its directionality: the reason that harmonic agglutination (i.e. featural identity across a morpheme boundary) must precede partial fusion is that, formally, adjacent identical elements are necessary to trigger an OCP violation, which in turn triggers re-bracketing.

Formally, there is a potential problem regarding the correspondence status of segments (or other phonological units) whose morphological affiliation changes between input and output. Imagine a situation in which there are two adjacent identical segments across a morpheme boundary, e.g. [pat][ta]. Under element-based correspondence, if one of these two segments were shared across the morpheme boundary, the output candidate would incur one violation of MAX and one violation of DEP. Thus in the case of $/p_1a_2t_3 + t_4a_5/ \rightarrow [p_1a_2[t_4]a_5]$, although the segment labeled t_4 is affiliated with the first morpheme, it counts as an epenthetic segment and thus incurs a DEP violation, just as a MAX violation is incurred for the missing segment t_3. This situation would severely compromise the typological predictions of the model, in that there would be at best a weak connection between the featural identity of adjacent segments across a morpheme boundary and their propensity to coalesce: the mapping $/p_1a_2t_3 + t_4a_5/ \rightarrow [p_1a_2[t_4]a_5]$ would incur exactly the same Faithfulness violations as $/p_1a_2k_3 + t_4a_5/ \rightarrow [p_1a_2[t_4]a_5]$ (although only the former would be motivated by an OCP violation). Under string-based correspondence, however, this problem disappears. Whether two segments are in correspondence or not is not decided at the level of individual segments, but rather at the level of an entire output string (in the present model, at the level of the morpheme). So long as MPARSE is satisfied (in the present model, NO-ALLOMORPHY), then it is assumed that all phonological units within some morphological domain are in a correspondence relation. Thus for a mapping /pat + ta/ → [pa[t]a], there are numerous possible input-output mappings for the domain of the first morpheme, e.g. <p, a, t> → <p, a, t> (faithful parse), <p, a, t, #> → <p, a, #, t> (deletion and epenthesis), <p, a, t, #, #, #> → <p, a, #, #, #, t> (multiple deletion and epenthesis), etc., although since all fare equally well with respect to markedness, all except the faithful parse are harmonically bounded. Therefore, phonological structure-sharing is possible under string-based correspondence and poses no problem for the typology presented in Figure 1.1.

2.3 Domain-Based Markedness.

I will assume that, unlike Faithfulness, the domain of Markedness constraints is the *maximal* domain of some stratum, unless otherwise specified (in effect, this is the same as standard practice). Two types of markedness constraints are of particular interest to the present analysis: the Obligatory Contour Principle (OCP) and Alignment, which will be explored in the following sections.

2.3.1 The Obligatory Contour Principle (OCP).

In its simplest formulation, the Obligatory Contour Principle, or OCP (Leben 1973), states that two adjacent, identical elements are prohibited. In practice, this generalization is complicated by what counts as "adjacent" or "identical". Since the advent of Autosegmental Phonology (Goldsmith 1976), adjacency has been understood as relativized to some autosegmental tier. Thus in Arabic, in a form such as *katab* 'he wrote,' *t* and *b* are adjacent on their own tier, the vowels being off on a separate tier (thus not "adjacent" to any of the consonants), at least a some relevant point during the derivation (cf. McCarthy 1986). Identity, in turn, depends on a combination of one's assumptions about tier structure and feature geometry. For example, a form such as *ampa*, in which *m* and *p* are two separate segments (i.e. independent at the root node), might nevertheless be represented either as having two separate instantiations of the feature [+labial], or else a single instance of [+labial] shared across a syllable boundary. The consequences of this are explored in Figure 2.5 below.

Figure 2.5: Featural representation and the OCP.

	OCP	ALIGNMENT
a. Identical features. [+lab][+lab] | | a m p a	?	✓
b. Feature spreading. [+lab] / \\ a m p a	✓	?

The question now becomes, if two adjacent segments are identical in some feature, with each feature represented independently, as in candidate (a), is that a violation of the OCP? Conversely, if two segments share some feature by feature-spreading, as in candidate (b), would this be a violation of Alignment (cf. Ito & Mester 1999 for discussion)? In this paper, I will make the simplifying assumption that Alignment and the OCP apply only at the level of segments—formally, the root node on the feature hierarchy. This assumption may need to be refined in future work on syllable structure and permissible codas in Dogrib, e.g., but is sufficient for the purposes of the present paper, to investigate prosody and relatively high-level morphophonemics.

2.3.2 Alignment.

In this paper I will be using the notion of Generalized Alignment (McCarthy & Prince 1993), which I assume to be gradient (contra McCarthy 2003). While Generalized Alignment provides a framework in which generalizations about edges may be stated, as mentioned previously, certain complications arise when talking about alignment with a morpheme boundary, as this is an abstract notion which does not directly correspond to the edge of any particular tier. For the purposes of the present paper, however, I will in effect equate "morpheme boundary" with some edgemost element on the segmental tier affiliated with some morpheme. An example of such an alignment constraint is given below.

ALIGN(MORPH, R, X, R):
·Let X be an autosegmental tier consisting of a set of ordered elements x_1, x_2, x_3,....x_n.
·Let S be the segmental tier consisting of a set of ordered elements s_1, s_2, s_3,....s_n.
·Let M_1, M_2, M_3,....M_n be a sequence of morphemes.

> For all s such that s is the rightmost element on the segmental tier affiliated with morpheme M_α, there exists an x such that x is the rightmost element on autosegmental tier X affiliated with M and x is associated with s.
>
> For each M such that $M \neq M_\alpha$, and M is affiliated with some element of S which is associated with the rightmost element of X which is affiliated with M_α, assign a violation mark.

Informally: "An element of type X cannot extend rightwards across a morpheme boundary; assign a violation mark for each morpheme boundary it crosses."

The definition given above is complex as it requires reference to three levels of representation simultaneously: the segmental tier, an autosegmental tier, and an abstract sequence of morphemes with which they are both affiliated. The definition is stated in terms of the number of morphemes (except for M_α) which are affiliated with the segmental material which intervenes between the rightmost segment affiliated with M_α

and the rightmost element on some autosegmental tier X affiliated with M_α. The reason for this seemingly roundabout formulation is that, as explained previously, morphemes do not exist directly on any particular tier, but are an abstraction referring to the relationship between a set of morphosyntactic features and a set of phonological entitites. The constraint as formulated above is gradient, as it assigns an increasing number of violation marks based on the number of morpheme boundaries crossed by some autosegment X. Other formulations of Alignment are possible, such as violations assigned by the number of segments, or syllables, or feet, all of which make different typological predictions in cases of long-distance spreading, although a full survey of these possibilities is beyond the scope of this paper.

A key point of this paper is that the difference between Harmonic Agglutination and Partial Fusion is derived from different rankings of Alignment and the OCP. The interaction of Alignment and the OCP, for both heteromorphemic and tautomorphemic sequences, is illustrated in Figures 2.6 and 2.7 below.

Figure 2.6: Alignment and OCP Heteromorphemically.

H H /whe-ne/	ALIGN-R(MORPH, TONE)	OCP
☞ a. H H whe-ne		*
b. H whe-ne	*!	

In Figure 2.6, we are presented with an input which contains two adjacent High tones, thus a sequence which would incur an OCP violation. In candidate (a), this surfaces faithfully as such, while in candidate (b), the first High tone spreads rightwards, across a morpheme boundary, and incurs one violation of ALIGN-R(MORPH, TONE). Figure 2.6 therefore predicts, by factorial typology, that, across a morpheme boundary, there will be some languages which prefer to maintain two adjacent identical elements (Harmonic Agglutination), while other languages prefer to spread an element across a morpheme boundary, to satisfy the OCP (Partial Fusion).

On the other hand, tautomorphemically, a simplified representation is always preferred, as is illustrated in Figure 2.7 below.

Figure 2.7: Alignment and OCP Tautomorphemically.

H H /łuwe/	ALIGN-R(MORPH, TONE)	OCP
a. H H łuwe		*!
☞b. H łuwe		

In Figure 2.7, it can be seen that both candidates (a) and (b) satisfy Alignment, since there are no morpheme boundaries to be crossed. Candidate (a), however, violates the OCP. The present analysis, therefore, makes the same prediction as Keer (1999), namely, that adjacent, identical elements may be preferred across a morpheme boundary, but never tautomorphemically.

3.0 Coalescence and the Lexical Phonology of Dogrib.

3.1 Root Level: Ablaut as Partial Fusion.

In the Northern Athabaskan literature, "ablaut" refers to a phenomenon in which the vowel of the verbal root undergoes some morphologically conditioned changes in vowel quality, tone, and/or nasality. Historically, these vowel changes were conditioned by verb-final suffixes, and indeed this is still the case in some more conservative Athabaskan languages (Rice 2000, Marinakis 2002). In Dogrib, however, the phonological conditioning environment has been lost entirely, such that the ablaut trigger is (or appears to be) entirely morphological. For instance, many stems undergo a change in vowel quality and/or nasality in the perfective aspect, and many other stems undergo a tone change in the imperfective and/or optative aspect, yet all of these type of ablaut are blocked if the verb is iterative (Ackroyd 1982: 79). It should be noted that ablaut is not a phenomenon which affects all verbs: some stems are invariant, while others exhibit full suppletion of the root, e.g. *tį* 'eat, sg/du subject' ~ *zhe* 'eat, pl. subject'. A full account of the various factors which may trigger ablaut is beyond the scope of this paper; however, ablaut seems likely to result from the interaction of the historical phonological shape of the root (Leer 1979) and lexical semantic properties of the root, in particular situation aspect (Rice 2000: 284).

In analyzing Dogrib ablaut synchronically, there are, broadly speaking, three possible analyses: Total Fusion, Harmonic Agglutination, and Partial Fusion. The Total Fusion analysis is essentially the account given by Ackroyd 1982, in which, for an ablauting stem, each form of the root is separately listed lexically. Thus for the verb 'handle 3-dimensional object' in Figure 3.1 below, the imperfective, perfective, and optative forms *ʔà*, *ʔǫ*, and *ʔa*, respectively, would each be given a separate lexical entry. The Total Fusion account therefore claims that ablaut is, de facto, no different from total stem suppletion of the *tį* ~ *zhe* variety mentioned above. It is of little consequence that *ʔà*, *ʔǫ*, and *ʔa* happen to resemble each other in form: any surface resemblance owes its explanation to diachrony; they are unrelated from the point of view of the synchronic phonology.

There are at least two arguments against this analysis. First, and most obviously, if one supposes that stem forms which have been subjected to ablaut are unrelated, it immediately becomes impossible to state any constraints upon what types of ablaut are allowed or disallowed in Dogrib (for example, ablaut never alters the vowel length of a root, nor does it ever create a contour tone). Secondly, true stem suppletion in Dogrib is almost exclusively conditioned by the grammatical number of the subject or object;[2] ablaut is conditioned only by aspect. Therefore, although a Total Fusion analysis of ablaut in Dogrib is compatible with the data, it leads to a loss of generalizations at both the morphological and phonological levels.

A second possible type of analysis is the Harmonic Agglutination analysis, which claims that ablaut is reducible to simple affixation, modulo some fairly predictable morphophonological processes. Such an analysis of ablaut was proposed by Stonham for Sanskrit, where ablaut is characterized as "the linking of certain phonological features to the base vowel of the root" (1994: 116). Here again, such analysis is not incompatible with the facts per se; rather, it is incompatible with Optimality Theory (Prince & Smolensky 1993). In Dogrib, the Harmonic Agglutination analysis would require positing fairly abstract entities as affixes, i.e. floating tones and floating vowel features. To see why, let us first look at my own proposal, the Partial Fusion analysis, in Figure 3.1.

---

[2] A single counterexample is given by Ackroyd (1982: 81), the verb 'say': *ʔareh**sį*** 'I say' vs. *ʔanen**dı*** 'you (sg) say' [boldface added]. This appears to be a case in which, historically, the 1sg subject marker *s* was re-analyzed as part of the root, and later the new 1sg subject morpheme *h* was restored by analogy.

Figure 3.1: Proposed Analysis of Ablaut Patterns.

	surface form			analysis		
root meaning	imp.	perf.	opt.	imp.	perf.	opt.
1. handle plural objects	le	la	le	√le μ	√le + a μ μ	√le μ
2. go out, die (fire)	kwi	kwo	kwi	√kwi μ	√kwi + o μ μ	√kwi μ
3. win (conclusive	nè	no	nè	H L √ne + e μ μ	H H √ne + o μ μ	H L √ne + e μ μ
4. handle three-dimensional object	ʔà	ʔǫ	ʔa	H L √ʔa + a μ μ	H √ʔa + n μ μ	H √ʔa μ
5. handle animate object	tè	tį	tè	H L √te + e μ μ	H √te + n μ μ	H L √te + e μ μ
6. cook[3]	t'è	t'e	t'è	H L √t'e + e μ μ	H √t'e μ	H L √t'e + e μ μ

In my proposal, each morpheme, both the root and the ablauting affix, is fully specified for moraicity, vowel quality, and tone, and indeed in many cases this information is represented redundantly (although I have omitted the tone tier in (1) and (2), for expository purposes). In a hypothetical Stonham-style affixation analysis, the affixes would consist of, e.g., a feature [+low] in the perfective in (1), a feature [+nas] in the perfective of (5), and a floating low tone in the imperfective and optative of (6).

While such an analysis may seem intuitively appealing, it is problematic from an OT perspective. Richness of the Base requires that no matter what one posits as the input, the winning output candidate always be a form which is grammatical in the language. For the affixation analysis, this is true so long as the only affixes that exist at the root level are floating tones and abstract features. But what would happen if one were to try to affix to the root something which actually looks like an affix? The affixation analysis immediately collapses. Take, for example, the perfective form *tį* in (5). What if, instead of merely affixing a feature [+nas] to the root *√te*, one affixed an actual nasal coda-consonant, i.e. */√te + n/*? All other things being equal, the result would be a long nasal vowel, i.e. *tįį*. Such a situation could be easily avoided if there were a constraint

[3] (6) is based on my own field notes; (1)-(5) are based on Ackroyd 1982: 76-79.

against long vowels in roots, but this is not the case: there *are* roots with underlying long vowels, e.g. √tłòo, and these surface faithfully as long. In other words, it is not the case that Dogrib disallows long vowels in roots; rather *long vowels are never created through affixation.* Short-voweled roots remain short no matter how many ablauting suffixes (if any) are added, and long-voweled roots remain long. A Harmonic Agglutination analysis cannot account for this except by stipulation.

A second argument against the Harmonic Agglutination analysis pertains to Lexicon Optimization. Lexicon Optimization states that, although any input must surface as something grammatical in the language, the *actual* input for a particular form should deviate from the winning candidate as little as possible. With respect to the Harmonic Agglutination hypothesis, this puts us in a rather strange situation, in that the forms which we are forced to posit as underlying for the various ablaut morphemes (despite Richness of the Base) are in fact forms which could *never* be grammatical in the language. Dogrib does not have non-moraic coda consonants or floating tones—on the surface, all nuclear segments bear a mora (WEIGHTBYPOSITION, cf. Borrelli 2000, Morén 2001), all tones are linked to vowels, and all vowels are linked to tones. So to sum up, under a Harmonic Agglutination analysis, contrary to Richness of the Base, only a small class of inputs will yield a grammatical output for ablauting roots, and, contrary to Lexicon Optimization, every member of this class is a phonological form which is ill-formed in the language (if not universally).

The Partial Fusion analysis avoids both of these difficulties. In Figure 3.1, every morpheme posited as an ablauting suffix *could* be a morpheme in Dogrib. For example, in (5), the affix /è/ has a low tone and bears one mora; there is, e.g., an affix /é/, which has a high tone and bears a mora, which is an aspectual marker, cf. bòkà**e**ht'è 'I am cooking'. Furthermore, there is nothing in the Partial Fusion analysis which *requires* that morphemes be fully specified. A morpheme can have any amount of degree of specification underlyingly; it will nevertheless surface as grammatical in the language. A summary of the different predictions of the Hamonic Agglutination vs. Partial Fusion analyses is given in Figure 3.2.

Figure 3.2: Predictions of the Harmonic Agglutination vs. Partial Fusion analyses.

input	**predicted result**	
	agglutination	**partial fusion**
H a. /√kwi + [+back]/ μ	H kwo μ	H kwo μ
H b. /√kwi + o/ μ	H kwo μ	H kwo μ
H H c. /√kwi + o/ μ μ	H kwo μμ	H kwo μ
H H H H d. /√kwi + o + o + o/ μ μ μ μ	H kwo μμμμ	H kwo μ

In Figure 3.2, I compare the predicted outputs of four hypothetical inputs, under a harmonic agglutination and a partial fusion analysis. Under a partial fusion analysis, any degree of underlying featural specification for the ablauting affix, indeed, even multiple copies of the same affix, still yields the same result: *kwo*, with a single mora and a high tone. A harmonic agglutination analysis certainly could derive the correct result for inputs (a) and (b), since the input string already contains the number of moras desired in the output (i.e. one), the feature [+back] can cause underlying /i/ to lower and backen in accordance with the vowel hierarchy (cf. section 3.1.1), and high tone is the default tone, regardless of where or how it is specified. Inputs (c) and (d) are more problematic, however. Multiple copies of the same tone, or of the same segment, will be default coalesce into a single (multiply linked) tone or segment, in accordance with the OCP (Leben 1973, Keer 1999). Yet even under the most generous interpretation, it is not clear how any harmonic agglutination analysis can deal with multiple copies of the same mora. This is because all presently existing definitions of faithfulness constraints (cf. McCarthy & Prince 1995, McCarthy & Wolf 2005) make no reference to morphological structure, only the output string as a whole. It is possible to say that, if forced to choose, a grammar will preserve a mora from the root rather than from an affix, although this is empirically meaningless, since all moras are identical. It is also possible to say that a mora can be added to an affix, but not to a root, although it is impossible to enforce such a distinction in the cases presented in Figure 3.2, since positional faithfulness (Beckman 1995) presupposes that morphological constituents are coindexed with discrete, non-intersecting phonological domains; when two morphemes share phonological structure, it becomes impossible to determine the locus of violation, rendering positional faithfulness, as well, empirically vacuous in this case.

The goal of an OT grammar is that, for any input, the grammar will generate a structure which is grammatical in the language. An OT account of ablaut, therefore, should ensure that, no matter what one affixes to the root, the result is some stem alternant which is possible in the language. The partial fusion account does just this, to which we turn in the next section.

3.1.1 Vowel Ablaut.

We now turn to the formal treatment of ablaut in Dogrib. In all of the following tableaux, I will be assuming that all inputs and outputs, both roots and affixes, are fully specified for tone and moraic content, although I will sometimes omit the tone tier or moraic tier for expository purposes. Nothing hinges upon the assumption that inputs are fully specified; indeed, this makes the analysis more difficult, by creating more feature clashes than would arise if some part of the input were underspecified.

In general, as the root level is characterized by partial fusion, OCP(X) >> Align(X) (cf. section 2.3.2, Figures 2.6 & 2.7), it is preferred to share a feature (be it a segment, tone, or mora) across a morpheme boundary than to maintain morpheme boundaries with crisp edges. In the case of moras, this leads to the loss of exactly one mora across the root-affix boundary. In the tableaux which follow, I have only considered candidates in which both morphemes coalesce into a single syllable in the output. While there is no *a priori* reason to assume this, formally it follows from the ranking ONSET >> ALIGN(MORPH, σ) at the root level.[4] Therefore, the outputs are subject to constraints on multiply linked features within the domain of a single syllable. In the case of vowels, this means that they are subject to the family of *DIP constraints (Casali 1997), which militate against diphthongs. The most severe of these constraints is the simplex *DIP constraint which forbids any more than a single vocalic segment linked to a syllable node, and which I claim is undominated at the root level. The result is then determined by a language-specific[5] vowel hierarchy (cf. Rice 1989, Marinakis 2002), which, in Dogrib, is expressed as Max(a) >> Max(o) >> Max(i) >> Max(e). This is illustrated in Figure 3.3.

Figure 3.3 considers seven possible output candidates from the concatenation of the root /√le/ 'handle plural objects' and the perfective ablauting suffix /a/. These candidates can be divided into three groups. The first group, candidates (a), (b), and (c), are perfectly aligned, that is, they show no sharing of segments or moras across the root-affix boundary—thus, these are all agglutinating candidates. Candidate (a) represents crisp agglutination: each segment from the input remains unchanged, and is linked to its

[4] See Marinakis 2002 for an extended discussion of syllabification in Dogrib.
[5] See Casali 1998 for discussion; the a > o > i > e hierarchy is derived from universal hierarchies of place and sonority, though the manner of their conjunction is language-specific.

Figure 3.3: Perfective of √*le*, ‘handle plural objects’.

/le + a/ μ μ	*DIP	OCP(SEG)	OCP(μ)	ALIGN-R (MORPH, SEG)	ALIGN-R (MORPH, μ)	MAX(a)	MAX(e)
a. crisp agglutination [le][a] [μ][μ]	*!		*				
b. assimilation [le][e] [μ][μ]		*!	*			*	
c. assimilation [la][a] [μ][μ]		*!	*				*
d. coalescence [l[e]] [μ][μ]			*!	*		*	
e. coalescence [l[a]] [μ][μ]			*!	*			*
f. coalescence, shortening [l[e]] [[μ]]				*	*	*!	
☞ g. coalescence, shortening [l[a]] [[μ]]				*	*		*

own mora and affiliated with its own morpheme, even though they are now part of the same syllable. Candidate (a) fatally violates *DIP, however, since this constraint forbids two non-identical vocalic segments within a syllable nucleus. Candidates (b) and (c) satisfy, *DIP, since, in these candidates, one of the vowels has assimilated to the other yielding two adjacent, identical vowels, but, in doing so, these candidates violate OCP(SEG). The first three candidates also violate OCP(μ), since they contain more than one mora within the domain of evaluation (cf. Section 2.3.1). The next two candidates, (d) and (e), are partially fusional on the segmental tier, but still agglutinating on the moraic tier. That is, these candidates satisfy alignment and violate the OCP with respect to moras, thus yielding a long vowel, but violate the alignment and satisfy the OCP at the segmental level, since, in these candidates, a single vocalic segment is shared across the morpheme boundary. Finally, candidates (f) and (g) exhibit partial fusion on both the segmental and moraic tiers; therefore, these candidates satisfy the OCP and violate alignment for both segments and moras. The choice between these two candidates comes down to the vowel hierarchy in Dogrib. Candidate (f) deletes *a*, which is the most sonorous vowel, and highest on the hierarchy, and thus fatally violates MAX(a), whereby candidate (g) emerges as the winner.

3.1.2 Tone Ablaut.

Within the root domain, tones behave in a way very much analogous to vowels in the previous example. As before, any candidates containing more than one mora are eliminated due to the constraint-ranking OCP(μ) >> ALIGN-R(MORPH, μ), as in candidate (a) below. The remaining candidates therefore only have a single, short vowel. Next, candidates containing contour tones are eliminated. I am assuming two different constraints referring to contour tones, which refer to different tiers to which the tones are linked (cf. section 2.1, Figure 2.1). *CONTOUR(σ) is violated by two or more tones linked to a single syllable node. Since Dogrib contains numerous examples of this (i.e. on long vowels), this constraint is low-ranked and inactive in the language. A second type of *CONTOUR constraint, *CONTOUR(μ), is relevant here. This constraint forbids

Figure 3.4: Imperfective/Optative of √*t'e* 'cook'.

H L μ μ /t'e+e/	*CONTOUR(μ)	OCP(μ)	ALIGN-R(MORPH, μ)	MAX(L)	MAX(H)
a. contour, long vowel [H] [L] [μ] [μ] [t'[e]]		*!			
b. contour, short vowel [H] [L] [[μ]] [t'[e]]	*!		*		
c. high tone, short vowel [[H]] [[μ]] [t'[e]]			*	*!	
☞d. low tone, short vowel [[L]] [[μ]] [t'[e]]			*		*

multiple tones linked to a single mora—this constraint is undominated in the language. Since vowels must be short, and short vowels can only contain a single tone, by necessity one of the input tones must be eliminated. The choice remains between candidates (c), and (d), which is decided by the tone hierarchy in Dogrib. That is, since Dogrib is a

"Low-marked" language, it is more important to preserve a Low tone from the input than a High tone. Candidate (c) fatally violates MAX(L) and candidate (d) is therefore chosen as the winner.

3.1.3 Vowel and Tone Ablaut.

Section 3.1.1 dealt with the root /√le/ , which takes an affix /o/ in the perfective; section 3.1.2 dealt with the root /√t'e/, which takes a Low tone affix in the imperfective and optative. There is nothing to prevent a root from taking both of these suffixes, and indeed, the root /√ne/ is just such a root: it takes a Low tone affix in the imperfective and optative, and the affix /o/ in the perfective. This is illustrated in Figure 3.5.

Figure 3.5: Perfective and Imperfective/Optative of *√ne* 'win (conclusive)'

H L /ne + e/ μ μ	*DIP	*CONTOUR(μ)	MAX(L)	MAX(o)	MAX(H)	MAX(e)
a. contour [H] [L] [n[e]]		*!				
b. high tone [[H]] [n[e]]			*!			
☞c. low tone [[L]] [n[e]]					*	
H H /ne + o/ μ μ						
d. diphthong [[H]] [ne][o]	*!					
e. root vowel [[H]] [n[e]]				*!		
☞f. ablaut vowel [[H]] [n[o]]						*

Candidate (a), in the imperfective/optative, and candidate (d), in the perfective, are eliminated because they violate the constraints on diphthongs and contour tones, respectively. The remaining candidates are evaluated based on the hierarchies of the language. Candidate (c) in the imperfective/optative is preferred over candidate (b) due to the ranking MAX(L) >> MAX(H), while candidate (f) is preferred over candidate (e) in the perfective due to the ranking MAX(o) >> MAX(e).

3.1.4 Vowel, Tone, and Nasal Ablaut.

One problematic case for the analysis sketched out so far is the root /√ʔa/, ‘handle 3-dimensional object’, which exhibits the pattern *ʔà ~ ʔǫ~ ʔa* in the imperfective, perfective, and optative, respectively (Ackroyd 1982: 77). At first glance, this pattern would seem to require analyzing the root as underlyingly *ʔo*, a form which never surfaces as such, and positing: (1) a low-tone affix in the imperfective only, (2) a nasal affix in the perfective only, and (3) an affix /a/ in both the imperfective and optative. (2) is uncontroversial, and, although the most common pattern is for a tone suffix to be present in both the imperfective and optative (cf. Rice 2000: 284), though several other patterns are attested (Ackroyd 1982: 77-79). Problematic, however, is the alleged /a/ suffix in the imperfective and optative, which is otherwise unattested, and, indeed, exactly the reverse of the normal pattern. Furthermore, to posit *ʔo* as the underlying form of the root is a violation of the Alternation Condition (Kiparsky 1973).

I propose instead that the underlying form of the root is *√ʔa*, which takes only two affixes: a low-tone suffix in the imperfective, and a nasal suffix in the perfective, which triggers a process which I refer to as *nasal raising*. An analysis of the imperfective is presented in Figure 3.6-a, below.

Figure 3.6-a: Imperfective of *√ʔa* ‘handle 3-dimensional object’.

H L /ʔa + a/ μ μ	NOCODA	MAX(±NAS)	OCP(μ)	RAISENASAL	MAX(L)	MAX(a)	MAX(H)	DEP(±NAS)
a. high tone [[H]] [ʔ[a]]					*!			
b. low tone [[L]] [ʔ[a]]							*	

In Figure 3.6-a, I consider two candidates, both with a single, short vowel. Candidate (b) is preferred over candidate (a) on account of Dogrib's tone hierarchy, MAX(L) >> MAX(H). In Figure 3.6-b below, I consider six possible outcomes from adding a nasal affix in the perfective.

Figure 3.6-b: Perfective of √*ʔo* 'handle 3-dimensional object'.

H /ʔa + n/ μ μ	CODACON	MAX(±NAS)	OCP(μ)	RAISENASAL	MAX(L)	MAX(a)	MAX(H)	DEP(±NAS)
a. nasal coda [ʔa][n] [μ][μ]	*!		*					
b. short oral vowel [ʔa][] [[μ]]		*!						
c. short nasal vowel [ʔ[ą]] [[μ]]				*!				*
d. long nasal vowel [ʔ[ą]] [μ][μ]			*!	*				*
☞e. short raised nasal vowel [ʔ[ǫ]] [[μ]]						*		*
f. long raised nasal vowel [ʔ[ǫ]] [μ][μ]			*!			*		*

Candidate (a) is excluded for violating CODACON, which states that the only codas allowed in the language are *h* and geminates.

3.1.5 Contrastive Vowel Length in Roots.

As predicted by the partial fusion analysis, under the same constraint ranking used in the preceding tableaux, a root with an underlyingly long vowel remains long in the output. This is illustrated in Figure 3.7 below.

Figure 3.7: Root with Underlying Long Vowel.

L H μ μ /tło/	MAX(μ)	*CONTOUR(μ)	OCP(μ)	MAX(L)	MAX(H)
a. faithful parse L H μ μ tło			*		
b. shortening, contour L H μ tło	*!	*			
c. shortening, high tone H μ tło	*!			*	
d. shortening, low tone L μ tło	*!				*

3.1.6 Summary of Root Level Phonology.

Tier 1: MAX(μ), MAX(±NAS), RAISENASAL >> *DIP, *CONTOUR(μ) >> NOCODA, OCP(SEG), OCP(μ), OCP(TONE), >>

Tier 2: ALIGN-R(MORPH, SEG), ALIGN-R(MORPH, μ), ALIGN-R(MORPH, TONE) >>

Tier 3: MAX(L), MAX(a) >> MAX(H), MAX(o) >> MAX (i) >> MAX(e) >>

Tier 4: Dep(μ), DEP(TONE), DEP(SEG), DEP(±NAS), ORALVOWEL...etc.

3.2 Stem Level and Word Level: Tonal Feet at work.

At the Stem Level and Word Level, the driving force behind coalescence in Dogrib is a prosodic unit which I refer to as the *tonal foot*. Feet in Dogrib are essentially moraic trochees, i.e. stress feet, but tone is relevant both in how feet are constructed, and in conditioning phonological processes to improve foot structure. Specifically, tonal feet are required to be *level*, as required by the constraint LEVELFOOT. A typology of tonal foot types in Dogrib is given in Figure 3.8.

Figure 3.8: Typology of Tonal Feet.

tonal foot type	***good or bad***	***example***
a. high-high trochee. H (σ σ) ta ta	good	L H L dze(k'oo)(lane) 'wild rose'
b. low-low trochee. L (σ σ) ta ta	good	L L L bo(kawi)t'e '2 of us will cook'
c. high-low trochee. H L (σ σ) ta ta	bad	L H L H *(boka)(whewi)t'e '2 of us have cooked'
d. low-high trochee. L H (σ σ) ta ta	bad	L H H *(shets'e)zhe 'we eat'
e. low-mid trochee. L M (σ σ) ta ta	acceptable as last-resort	L M H (shets'e)zhe 'we eat'

Example (a) is well-formed underlyingly and surfaces faithfully: all feet are level. Example (b) is derived from an underlyingly non-level sequence by deletion: /bòkàwewìdt'è/. This is because an /L-L-H-L-L/ cannot be parsed as level feet, so the H-syllable is deleted. ·Example (c), an unattested form, would be the faithful parse of '2 of us have cooked'; it surfaces instead as *bò(kàwhì)t'e*, again with the High toned vowel deleted. Finally, example (d) is ungrammatical, but, on account of a generalization known as Schneider's Law (Lipscomb 1992, Dresher & Johns 1995; cf. section 3.3.2), is unfixable by deletion or re-footing; instead, H is lowered to M, as in (e). Although LEVELFOOT would seem to rule out contour tones *a fortiori* (since a non-level syllable is also a non-level foot), in practice contour tones are common. Mostly, this is because un-level sequences are fixed by shifting foot-boundaries, and a foot-boundary cannot occur in the middle of a syllable, on account of Strict Layering (Kiparsky 2003).

This conception of the tonal foot differs from other conceptions of tonal foot proposed in the literature. DeLacy (2002) discusses a case of tone-to-stress interaction in Ayutla Mixtec. In this system, although stress by default falls on the leftmost syllable, stress is also attracted to higher-toned syllables: the language prefers to stress High toned syllables, followed by Mid tones, and Low tones are least preferred. While Ayutla Mixtec is similar to Dogrib in that tone may cause stress to shift, it is not the case in Dogrib that High tones attract stress over low tones: LEVELFOOT requires that tone remain level within a foot; *Low-Low* or *High-High* sequences fare equally well according to this constraint. Zec (1999) discusses a case of bidirectional foot-tone interaction in Neo-Stokavian (NS) Serbo-Croatian. For example, in NS, a pitch accent seeks to align itself with the right edge of certain derivational suffixes (1999: 229), which in turn attracts footing by a TONE-TO-FT constraint (1999: 234). In Dogrib, however, the interaction between tone and metrical structure is unidirectional: if a faithful parse of the input would yield a non-level tone sequence within a foot, one or more repair strategies will be invoked to repair the violation: segmental deletion, gemination, and/or H to M lowering, as will be illustrated in the following sections.

3.2.1 Tonal Feet in Dogrib.

In Figure 3.9, I give the surface footing for the entire paradigm of the verb bòkà√t'è, 'cook', in Dogrib. There are some cells missing from the paradigm; I am still unsure whether this represents my own failure to elicit these forms, or whether the optative paradigm is in fact defective, as it is in Hupa (Golla 1970). At first, the question of footing might appear fairly trivial: in all of the examples in Figure 3.9, all of the feet consist of either a single heavy syllable (either a long vowel, or a coda consonant, or both) or two light syllables, and, with the exception of the 2nd person singular optative, the final syllable is always extrametrical. It would appear, then, that foot-assignment in

Dogrib is a relatively simple matter of parsing moraic trochees from right to left, with the final syllable extrametrical. The problem with this, however, is that in fact the surface forms given in Figure 3.9 have, in some cases, been modified quite radically from their underlying representations, and, furthermore, they have been modified in such a way as to optimize foot structure. A second issue is the role of *tone* in tonal feet. Notice that the majority of feet in Figure 3.9 are level in tone, i.e. (LL) as in (bòkà) or (HH)/(H) as in (whene)/(whęę). This, one might argue, is not evidence for anything: it could be a statistical accident, or it could follow from independent principles having nothing to do with feet. Yet, as I will demonstrate in section 3.2.2, the phonological processes of Dogrib actively conspire to create feet that are level (i.e. no high-low or low-high sequences within a foot).

Figure 3.9: paradigm for bòkà√t'è[6] 'cook', surface forms, with footing.

Imperfective:

	singular	dual	plural
1st person	(bòkà)(eh)t'è[7]	bò(kàì)t'è	(bòkà)(ts'ee)t'è
2nd person	bò(kàį)t'è[8]	bò(kàah)t'è	bò(kàah)t'è
3rd person	(bòkà)(et)t'è	(bòkà)(gee)t'è	(bòkà)(gee)t'è

Perfective:

	singular	dual	plural
1st person	(bòkà)(whıh)t'e	bò(kàwhì)t'e[9]	(bòkà)(ts'ıh)t'e
2nd person	(bòkà)(whęę)t'e	(bòkà)(whah)t'e	(bòkà)(whah)t'e
3rd person	(bòkà)(whet)t'e	(bòkà)(geh)t'e	(bòkà)(geh)t'e[10]

Optative:

	singular	dual	plural
1st person	(bòkà)(weh)t'è	bò(kàwì)t'è	(bòkà)(ts'ıh)t'è
2nd person	(bòkà)(wįt'è)		
3rd person	(bòkà)(wet)t'è		

[6] bò can be prenazalized throughout the paradigm: mbò
[7] The sequence àe is pronounced as two syllables.
[8] The sequence àį is pronounced as a single syllable.
[9] Also: bòkàdıìht'e
[10] These forms are the perfective punctual, "they cooked for a short time". bòkàgįht'e means "they cooked, and it went on for a while." Similarly with bòkàts'įht'e.

In demonstrating the relevance of tone to footing, it is first necessary to explore what happens if we construct feet without reference to tone. An example of this is shown in Figure 3.10 below. In Figure 3.10, the input /bòkàgeet'è/ is parsed into moraic trochees, using four constraints that are standard in the literature: NONFINALITY, FOOTBINARITY, PARSESYLLABLE, and ALIGNRIGHT. Candidate (a) violates non-finality since the final syllable in the word, *t'è*, is footed. FOOTBINARITY, which requires that a foot have exactly two moras, is violated by candidates (a) and (c), and candidate (c) is excluded for this reason. Finally, candidate (d) is eliminated on account of the constraint PARSE(σ). PARSE(σ) comes into direct conflict with other constraints in the system. It conflicts directly with NONFIN(FT), which requires that the final syllable in a prosodic word *not* be parsed into a foot. In the case of candidate (d), candidate (d) is identical to candidate (b), the winning candidate, in the position of feet, except that it has one fewer foot. This is an improvement according to the constraint ALIGN-R(FT, PRWD), since the foot constructed on (bòkà), whose right edge would be two syllables removed from the right edge of the prosodic word, is not present. However, under the ranking given below, PARSE(σ) outranks ALIGN-R(FT, PRWD), and so candidate (b) emerges as the winner.

Figure 3.10: Feet assigned atonally.

/bò-kà-ge-e-t'è/	NONFIN(FT)	FTBIN	PARSE(σ)	ALIGN-R(FT, PRWD)
a. (bò.kà)(gee.t'è)	*!	*		**
☞b. (bò.kà)(gee)t'è			*	***
c. bò(kà.gee)t'è		*!	**	*
d. bòkà(gee)t'è			***!	*

There are other potentially relevant metrical constraints which I have omitted, for example, RHTYPE=TROCHEE, which states that trochees are preferred over iambs. However, since I have seen no evidence of iambs, I will simply assume that this constraint is undominated in the constraint hierarchy. Also relevant is a more precise definition of FOOTBINARITY. FOOTBINARITY (FTBIN) can be broken down along two dimensions: whether binarity is defined in terms of moras or syllables, and whether violations are incurred by feet having less than two syllables/moras (so-called "degenerate feet", which violate FTBIN(MIN)) or more than two syllables/moras (so-called "ternary feet," which violate FTBIN(MAX)). Finally, I am also presupposing that the rightmost foot is the "head foot" (this is the official term, not an oxymoron). That is, the foot that is rightmost in the prosodic word receives main stress. This can be derived straightforwardly by using the constraint ALIGN-R(HDFT, PRWD), “the Head Foot is rightmost in the Prosodic Word),” which dominates NONFIN(HDFT, PRWD) “the Head Foot is not final in the prosodic word”. However, since, again, I have not seen any examples in which the Head Foot is not final in the Prosodic Word, I will simply assume that ALIGN-R(HDFT, PRWD) is undominated, and it will not feature in any subsequent tableaux.

The preceding discussion, and the example shown in Figure 3.10, might give one the impression that Dogrib is basically a garden-variety moraic trochee language à la Hayes 1995. This is far from true, and in fact the only cases in which this would seem to be true are cases in which the underlying representation is already optimal according to the strictly metrical constraints (as given in Figure 3.10) and the constraints SWP and LEVELFOOT, to be introduced in section 3.2.2. In the majority of cases, the output has been modified in some significant way in order to improve tonal foot structure. This is the case, for example, in (bòkà)(et)t’è, whose underlying form is /bòkàet’è/. Thus, between the input and the output, gemination has occurred. The gemination, as I shall argue, is prosodically motivated, but it is unclear why this gemination should have occurred, since either *bòkàet’è* or *bòkàett’è* can be parsed equally well into moraic trochees.[11] A more extreme example is *bòkàwìt’è*, whose underlying form is /bòkàwewìdt’è/. There is nothing in the constraints in Figure 3.10 which would explain why an entire syllable should delete, let alone the syllable (*we*) which would receive main stress by simple right-to-left parsing of moraic trochees. The answer is that, in Dogrib, the metrical is tightly intertwined with the tonal. Metrical structure also refers to tone, which is why I, as I claim, this structure is best referred to as “tonal feet”. The precise nature of the “tonality” of these feet, and its consequences, is explored in the next section.

[11] The actual output, *bòkàett’è* satisfies PARSE(σ), while the ungeminated form incurs one violation. However, this would be a very typologically unusual reason to geminate, and in any case would not apply to other examples.

3.2.2 Stem-Level Processes.

The stem-level phonology of Dogrib may be described succinctly as follows. All phonological processes that occur at this level are prosodically motivated, both by the classical metrical constraints as given in Figure 3.10, and by the constraints SWP and LEVELFOOT. SWP is the stress-to-weight principle, "if stressed, then heavy," which dislikes stressed, light syllables (Borrelli 2000, Morén 2001). Thus, where moraic trochees are concerned, if the foot bears primary lexical stress, then SWP prefers a single heavy syllable over two light syllables. LEVELFOOT prefers feet in which there is no disagreement in tone, that is, no high-low or low-high sequences, whether these be within a single syllable (i.e. contour tones) or across two syllables. These preferences are illustrated in Figure 3.11.

Figure 3.11: Preferred configurations for SWP and LevelFoot.

	SWP	LEVELFOOT
a. (tátá)	*	✓
b. (tàtà)	*	✓
c. (tátà)	*	*
d. (tàtá)	*	*
e. (táá)	✓	✓
f. (tàà)	✓	✓
g. (táà)	✓	*
h. (tàá)	✓	*

There are eight logically possible types of moraic trochees in a tone language, derived from four possible tone patterns (high-high, high-low, low-low, low-high) times two syllable patterns (light-light versus heavy), assuming, of course, that heavy syllables may bear a contour tone, while light syllables cannot. Of these, as shown above, all of the light-light feet are dispreferred by SWP, while all of the low-high and high-low feet are dispreferred by LEVELFOOT. Note, in Figure 3.11, that there are two foot types which satisfy both constraints, candidates (e) and (f). These are the ideal foot-types towards which the language aspires; conversely, there are two extremely bad foot types according to this scheme, represented by (c) and (d). It is in order to avoid these most despised foot patterns that the language resorts to its most radical repair strategy, namely, to delete an entire syllable, as will be illustrated in sections 3.2.2.3 and 3.2.2.4. Other, less extreme repairs include simple oral vowel coalescence and nasal coalescence. These will be illustrated in the following sub-sections.

3.2.2.1 Oral vowel coalescence.

Figure 3.12 below shows that, in the case of words such as *bòkàwhaht'e* "you (du/pl) have cooked," the force which drives coalescence is SWP. That is, in the faithful parse, candidate (a), the head foot is a light-light moraic trochee, in which primary lexical stress falls on a light syllable. This fatally violates the SWP. In all of the remaining candidates, (b)-(e), some form of coalescence occurs, in order to create a heavy syllable. In the case of (b) and (c), these violate a constraint which forbids trimoraic (i.e. superheavy) syllables.[12] This requires some form of vowel shortening, which violates MAX(μ), as in candidates (d) and (e). The choice between (d) and (e) comes down to the vowel hierarchy in Dogrib: it is more important to preserve an [a] than an [e], and so candidate (e) is chosen as the winner.

Figure 3.12: Oral vowel coalescence, assuming harmonic agglutination.

/bò-kà-whe-ah-t'e/	SWP	$*\mu\mu\mu]_\sigma$	MAX(a)	MAX(μ)	MAX(e)
a. bò-kà-(whe.ah)t'e	*!				
b. bò-kà-(wheeh)t'e		*!	*		
c. bò-kà-(whaah)t'e		*!			*
d. bò-kà-(wheh)t'e			*!	*	
☞e. bò-kà-(whah)t'e				*	*

In saying that candidate (e) is chosen over candidate (d) on account of the constraint MAX(a), the question arises as to the exact autosegmental representation of thse candidates. That is, do candidates (d) and (e) represent a case of outright *deletion* of a segment, or is it a case of assimilation accompanied by shortening, or even partial fusion as in section 3.1? Theory-internal considerations dictate that this must be considered deletion. That is, a partial fusion analysis would be inconsistent with my claim that partial fusion happens only at the root level in Dogrib. As for assimilation accompanied by shortening, this amounts to the same thing as deletion, as far as I can tell: the vowel [e] on the segmental tier deletes, and its mora fails to reassociate, since

[12] Trimoraic syllables do arise in Dogrib, and to account for them probably requires a revision of this constraint; for now I will make the simplifying assumption that they are not allowed. They are, in any case, dispreferred.

moras do not reassociate in Dogrib (cf. section 3.2.2.3). The question then becomes, why is it that all cases of assimilation are not accompanied by shortening (i.e. the loss of a mora)? The problem is illustrated in Figure 3.13 below.

Figure 3.13: Problematic nature of vowel assimilation representation.

μ μ e a	*DIP	MAXASSOC(μ, V)	MAX(a)	MAX(e)
a. faithful parse μ μ e a	*!			
b. reassociation of mora μ μ a		*!		*
c. vowel assimilation μ μ e e			*! (?)	
d. vowel assimilation μ μ a a				* (?)

In Figure 3.13, the input is a sequence of vowels, /ea/, which, if parsed faithfully, would result in a diphthong, as in candidate (a). This violates the constraint *DIP. The remaining candidates, (b)-(d), present some form of monophthongization. The question is, if two vowels monophthongize, but length is retained (i.e. no moras are lost), how is this to be represented? In candidate (b), the vowel /e/ is literally deleted, and its mora reassociates to [a]. In candidates (c) and (d), both vowels are retained, but one assimilates to the other. Now if, as I argue (and evidence for this will be given in section 3.2.2.3), moras do not reassociate in Dogrib, a representation such as candidate (b) is excluded by MAXASSOC(μ, V), which says, "for every mora, it must remain associated to its vowel". The only remaining viable candidates then are (c) and (d), which are distinguished based on the vowel hierarchy of Dogrib: it is more important to preserve *a* than to preserve *e*. The vowel hierarchy is expressed formally as MAX(a) >> MAX(e). Then question then arises regarding in what sense candidate (d) is superior in satisfying MAX(a), since, in candidate (c), no vocalic segment has actually been deleted. It seems, therefore, that the use of the constraints MAX(a), MAX(e), etc. is, in some sense a misnomer. That is, these "MAX" constraints do not really refer to retention of the vowels themselves (formally, their slot on the segmental tier, and the root node on the feature hierarchy), but rather certain *features* of the vowels. This is in fact the original sense in which vowel hierarchies were presented in Casali (1997): there are universal hierarchies of MAX/DEP[13] feature constraints, and these are combined in different ways to yield a variety of vowel hierarchies cross-linguistically.

3.2.2.2 Nasal coalescence.

Another way in which tonal foot structure is optimized in Dogrib is a process which I refer to as "nasal coalescence". In nasal coalescence, a nasal consonant situated between two short oral vowels is deleted, leaving a single long nasal vowel. This process is in fact variable; thus from the input /bòkàwhenet'e/, we find both *bòkàwhenet'e* and *bòkàwhęęt'e*. In addition, nasal coalescence interacts with nasal raising, so that *bòkàwhįįt'e* is also attested. This is illustrated in Figure 3.14 below.

Figure 3.14: Nasal coalescence and variability.

/bò-kà-whe-ne-t'e/	MAX[±Nas]	MAX(μ)	SWP	IDENT(V)	MAX(son)	RAISENASAL
(☞) a. bò-kà-(whene)t'e			*			
b. bò-kà-(whını)t'e			*	*!		
☞c. bò-kà-(whęę)t'e					*	*
(☞) d. bò-kà-(whįį)t'e				*	*	
e. bò-kà-(whę)t'e		*!	*		*	*
f. bò-kà-(whį)t'e		*!	*	*	*	
g. bò-kà-(whee)t'e	*!	*			*	

The two highest-ranked constraints in Figure 3.14 above are MAX[±Nas] and MAX(μ). That is, it is necessary to retain the nasal feature, whether by preserving the nasal consonant itself, or by nasalizing an adjacent vowel, and it is important to preserve all moras. This excludes candidates (e), (f), and (g). Of the remaining candidates, (a), (c), and (d) are all attested in Dogrib, although (c) is the most frequent (I do not have any data on usage frequency, at the present time). The variation between these candidates arises because the four lower-ranked constraints are all unranked. The faithful parse in candidate (a), *bòkàwhenet'e*, violates SWP, because, in it, the main-stressed syllable is light. Candidate (b) is harmonically bounded by candidate (a), since it incurs a

[13] Actually, in his model, PARSE/FILL.

superfluous violation of IDENT(V). That is, although vowel raising also occurs in candidate (d), this is in order to avoid the violation of RAISENASAL that occurs in candidate (c); the vowel raising in *bòkàwhinit'e* is unmotivated. Finally, note that candidates (c) and (d), which exhibit nasal coalescence, violate the constraint MAX(son), on account of the nasal segment /n/ which was deleted, but satisfy MAX[±Nas], since the nasal feature is preserved on the adjacent vowels. The constraint Max(son) is somewhat vague and probably inadequate. That is, in addition to *n*, *r*, *d*, and *w* may delete intervocalically, while *m*, *t*, *l*, and *wh* may not. In this sense, MAX(son), in Figure 3.14, functions as a sort of place-holder for a class of phonemes which is theoretically problematic.

3.2.2.3 Tone-conditioned vowel deletion.

In this section I come to the most dramatic and typologically unusual manifestation of tonal feet, tone-conditioned vowel deletion. The driving force behind this process is the constraint LEVELFOOT, which forbids trochees whose tone pattern is either high-low or low-high. Based on this, one might expect that the language would simply adjust the tone pattern, that is, make a high tone low or a low tone high to fix the problem. Instead, the result is deletion, on account of the MAXASSOC constraints, as shown below.

Figure 3.15: Tone-conditioned vowel deletion.

/bò-kà-e-wìd-t'è/	CODACON	LEVELFOOT	MAXASSOC (V, TONE)	MAXASSOC (μ, V)	MAX(L)	FTBIN	MAX(μ)	MAX(H)
a. bò-kà-(ewìt)t'è	*!	*				*		
b. bò-kà-(ewì)t'è		*!						
c. bò-kà-(èwì)t'è			*!					*
d. bò-kà-(ewı)t'è			*!		*			
e. bò-kà-(èwı)t'è		*!	**					
f. bò-kà-(e)t'è					*!	*	*	
☞g. bò-kà-(wì)t'è						*	*	*
h. bò-kà-(wìì)t'è				*!				*
ı. bò-kà-(ee)t'è				*!	*			

Candidate (a) violates CODACON on account of the geminate consonant in *bòkàewìtt'è*. Since, as I argue elsewhere, Dogrib does in fact have geminates, this requires some explanation. I am assuming that gemination is strictly a *word* level process. Thus, at the stem level, the only acceptable coda consonant is *h*, while at the word level both *h* and geminates (and, in some cases, coda nasals) are allowed. Thus, CODACON is to be understood as a sort of place-holder for whatever coda consonants are allowed at a given level.

Candidates (b) and (e) are in direct violation of the constraint LEVELFOOT, since they contain feet with high-low and low-high sequences, respectively. The issue then becomes, how are these violations of LEVELFOOT to be repaired? Herein comes the most potentially controversial part of my proposal, the constraint MaxAssoc(V, TONE), "for every vowel, it must remain associated with its tone". This constraint directly militates against vowels being de-linked from the tones with which they are associated underlyingly but, crucially, is satisfied vacuously if the vowel itself is deleted. Therefore, if some tonal configuration becomes unacceptable for independent reasons, and MaxAssoc(V, TONE) is highly ranked, then the vowel in question will delete. This is what rules out candidates (c) and (d) in Figure 3.15. In candidate (c), a high tone was lowered to make a level foot *bòkà(èwì)t'è*, while in candidate (d) a low tone was raised *bòkà(ewi)t'e*.

Having established that one of the vowels is to be deleted, the question arises as to which one. Here the choice comes down to the vowel hierarchy in Dogrib. Candidate (f) violates MAX(L), which is ranked higher than MAX(H), which is violated by candidate (g); thus, candidate (g) emerges as the winner. Finally, there is the issue of compensatory lengthening. That is, even though we had to delete a high tone to satisfy LevelFoot, and even though we had to delete its vowel along with it to satisfy MAXASSOC(V, TONE), might we not still spare its mora, and reassociate it to another vowel? Here again, this is not allowed, due to the constraint MAXASSOC(μ, V), "for every mora, it must remain associated with its vowel." In a parallel fashion, this constraint is satisfied if the mora itself is deleted: just as a captain goes down with his ship, so too a mora must go down with its vowel. We are left, it seems, with a rather curious situation: the autosegmental phonology of Dogrib doesn't seem very autosegmental at all, and both tones and moras behave more like properties of segments rather than independent elements on their own tier. Historically, there may be an explanation for this: both tone (Krauss 2005) and weight-sensitivity (Marinakis 2002) are relatively recent developments in Dogrib. Tone arose, historically, from glottal stops, which are segments, and likewise phonemic vowel length arose as the result of the deletion of intervocalic consonants (i.e. segments). Thus it is possible that, when a tone language or quantity-sensitive language is still in its early stages, tones and moras still behave as if they were on the segmental tier (which, historically, they were). This hypothesis is only speculative, but it may ultimately help us to reconcile the facts of Dogrib with well-known typological generalizations about the stability of tones and moras.

Finally, a residual issue is that of floating tones and floating moras. While I am agnostic as to the existence of invisible floating entities, I am assuming that, in Dogrib at least, if a tone or mora is not associated with an element on the segmental tier, it is eliminated by stray-erasure at the end of each cycle.

3.2.2.4 Tone-conditioned vowel and glide deletion.

A similar situation to that described in section 3.2.2.3 is found in the case of the underlying form /bò-kà-whe-wìd-te/, which surfaces ultimately as *bò(kàwhì)t'e*. In this case, both the offending high-toned vowel *e* and the following glide *w* delete. This situation is slightly unusual in that *e*and *w* are neither part of the same syllable nor the same morpheme. This is illustrated below.

Figure 3.16: Tone-conditioned vowel and glide deletion.

/bò-kà-whe-wìd-t'e/	LEVELFOOT	MAXASSOC(V, TONE)	MAX(obs)	MAX(son)
a. bò-kà-(whewì)t'e	*!			
b. bò-kà-(whèwì)t'e		*!		
c. bò-kà-(wì)t'e			*!	
☞d. bò-kà-(whì)t'e				*

Candidate (a), which contains a high-low trochee, violates LEVELFOOT, just as before. Candidate (b) is representative of a class of candidates which somehow re-adjust the offending tone pattern, but in doing so violate MAXASSOC(V, TONE). The choice then comes down to candidates (c) and (d). At this point, I can only describe what is going on at a somewhat informal level. I am presupposing, first of all, that one cannot have two adjacent glides which disagree in voicing, e.g. **bòk(à**whw**ì)t'e*. This could fairly easily be accomodated by some sort of AGREE(VOICE) constraint. The problem, then is why *w* is deletable while *wh* is not. Here, the generalization is, informally, that there is a class of unmarked and highly sonorous phonemes which may be deleted intervocalically in Dogrib. This includes *d, r, n,* and, crucially, *w*, but not *wh*. In Figure 3.16, I have stated this generalization in terms of sonority, although strictly speaking this does not work: *w* and *wh* are both sonorants. In addition, *d* deletes even though it is not a sonorant, and *l* is a sonorant but never deletes. It appears that Dogrib is a language in which, in general, unmarked segments delete but marked segments don't (DeLacy 2006), but markedness in this sense is neither a question solely of place of articulation nor solely of sonority but rather, like the vowel hierarchy and the tone hierarchy, are derived from multiple, intersecting universal hierarchies which interact in language-specific ways. I will therefore use "sonority" as a sort of placeholder for the time being, until a more fine-grained understanding of consonantal markedness in Dogrib can be established.

3.2.2.5 Summary of Stem-Level processes.

To summarize, the phonological processes which occur at the stem level are vowel deletion, glide-deletion, oral vowel coalescence, and nasal coalescence. All of these processes are prosodically driven. Ideally, the language seeks to create, at the stem level, a moraic trochee consisting of a single, heavy syllable with a level tone. However, the level-ness of tone is more important than metrical foot structure *per se*, and so in some cases, the output is actually a degenerate foot, as in 3.2.2.3 and 3.2.2.4. Finally, while deleting intervocalic consonants is one strategy to improve foot structure, only certain consonants are deletable and not others, and this notion remains yet to be defined precisely.

3.2.3 Word Level Processes.

Phonological processes at the word level in Dogrib in many ways overlap with those at the stem level. Deletion of unmarked segments intervocalically, both *r* and *n*, also occurs at the word level, as does nasal raising (which occurs at the root, stem, and word levels, but not postlexically). Deletion of vowels for tonal reasons does not occur at the word level as it does at the stem level. In the domain of metrical structure, there is somewhat of a shifting of priorities: violations of LEVELFOOT are now tolerated (that is, high-low or low-high trochees are either passively tolerated or even actively produced by phonological processes), but degenerate feet are no longer tolerated. This amounts, in effect to a re-ranking of LEVELFOOT and FTBIN at the word level, as compared to the stem level. These phonological processes are explored in greater depth in the following sub-sections.

3.2.3.1 [r]-deletion.

r-deletion is a variable process in Dogrib. Retention of *r* is a characteristic of more conservative or formal speech, while deletion of *r* characterizes more informal or innovative speech. A formal characterization of r-deletion is given in Figure 3.17 below, with data taken from Ackroyd (1982: 29).

Figure 3.17: Illustration of *r*-deletion.

/(gere)ko/	FTBIN	DEP(μ)	SWP	MAX(son)
a. (geere)ko	*!	*		
b. ge(ree)ko		*!		
c. (gere)ko			*	
d. (gee)ko				*

In Figure 3.17, the driving force behind *r*-deletion is hypothesized to be SWP, that is, the main-stressed syllable of a prosodic word should be a heavy syllable. I show footing as already present in the input, since metrical constraints apply at the stem, word, and postlexical levels, and *(gere)ko* is the footing which one expects as the output of the stem level. However, as far as I can tell, footing is transparent at each level and so the analysis would work just as well if the input were unfooted, (i.e. by ROTB). Candidates (a) and (b) satisfy SWP by adding a mora, and in doing so violate DEP(μ). In addition, candidate (a) has a trimoraic foot, which also violates foot binarity. The choice remains between candidates (c) and (d), both of which are attested in Dogrib. Candidate (c) is the fully faithful candidate, which violates SWP but satisfies foot binarity and all faithfulness constraints. Candidate (d) deltes *r* intervocalically, leaving a single, heavy syllable to constitute the main-stressed foot.

One limitation of the analysis in Figure 3.17 is that it is based entirely on data from Ackroyd, who, unfortunately, makes no mention of geminates and provides no phonetic data on consonant length. Thus, although I assume that both of the attested candidates have a singleton *k* as the stem-initial syllable, in fact I do not know this. If, e.g., the actual attested forms are *gerekko* and *geeko*, this would change the analysis considerably, as the variation would no longer be a question of the satisfaction/violation of the SWP (since, in both cases, the stressed syllable would be heavy, in one case because of a long vowel and in the other case because of a geminate), but rather it would amount to competition between MAX(son) and an antigemination constraint. More data are necessary to settle this issue.

Finally, there is, once again, the nature of the constraint MAX(son). It does so happen that *r* is a sonorant, though, again, I am using the term loosely, as a placeholder for a class of unmarked consonants which are deletable intervocalically in Dogrib.

3.2.3.2 Nasal Coalescence.

Nasal coalescence occurs at the word level just as it does at the stem level, although here again, it operates variably (I do not have any frequency data at present). Two items noteworthy of nasal coalescence at the word level are that (1) it operates in spite of non-level feet, and (2) it is preferred over gemination. With regards to (1), although it is not immediately evident from this tableau in particular, arguably LEVELFOOT should be placed between FTBIN and *Cμ. This is because, although constructing binary feet (as in candidates (b) through (e) in Figure 3.18 below) is more important than constructing level feet, nevertheless the construction of level feet plays a key role in conditioning gemination; cf. section 3.2.3.4 on gemination below.

Figure 3.18: Illustration of Nasal Coalescence.

/nà-(ne)zè/	FTBIN	MAX[±nas]	*Cμ	SWP	MAX(son)
a. nà(ne)zè	*!			*	
b. (nàa)zè		*!			*
c. nà(nez)zè			*!		
☞ d. (nàne)zè				*	
☞ e. (nąą)zè					*

Candidate (a) is the faithful parse, which maintains the same prosodic structure as in the input (although, as I have mentioned previously, I have not yet found any evidence of faithfulness to prosodic structure per se). The problem is that candidate (a) contains a degenerate foot, and so fatally violates foot binarity. Candidate (b) simply deletes the intervocalic nasal *n* with no trace, which satisfies both foot binarity and SWP, but fatally violates Max[±nas]. Candidate (c) satisfies both foot binarity and SWP by geminating the stem-initial consonant *z*. In general, cross-linguistically, the type of constraint which militates against geminates is *Cμ (Borrelli 2000), which may be broken down as to which particular types of consonants are more or less resistant to bearing a mora (Zec 1988, 1995). Here one runs into a potential difficulty, in that gemination *is* widespread in Dogrib, so one must rule out gemination here while simultaneously allowing it in examples such as those in 3.2.3.4. The general line I take is that gemination does *not* occur in order to satisfy SWP, but it does occur in order to satisfy foot binarity, i.e. to repair degenerate feet. Thus for example, a sequence such as *(tata)tà*, which violates SWP, will not be repaired by gemination, but a word such as *(ta)tà*, which contains a degenerate foot, most likely will. This then raises the issue of why such degenerate feet should arise in the first place. There are two possible scenarios. In one scenario, there are simply not enough syllables in the word, as in the form *nàzè*, 'he hunts', footed presumably *(nà)zè*, which I have heard the *z* pronounced both geminated and un-geminated. In the other scenario (to be explored in detail in section 3.2.3.4), there is actually no shortage of syllables with which to construct a binary foot, but a foot is displaced leftward one syllable, due to the constraint LEVELFOOT. In any case, candidate (c) in Figure 3.18 above is excluded because, under the constraint-ranking *Cμ >> SWP, which states that it is better to have a light syllable which is stressed, as in candidate (d), than allow a consonant to bear a mora, as in candidate (c). Finally, both candidates (d) and (e) are attested forms in Dogrib, candidate (d) being characteristic of older speakers and more careful speech, (e) representing more informal and innovative speech. The difference comes down to whether one prefers to have a stressed syllable which is heavy, as in (e), or to preserve the sonorant *n*, as in (d). Thus, SWP and Max(son) are shown unranked in Figure 3.18, to represent this variation.

3.2.3.3 Nasal Raising and Tonal Footing.

Here I come to perhaps the most complicated phonological process in Dogrib morphophonology, namely, the interaction of tonal feet and nasal raising. To remind the reader, nasal raising is the process by which nasal vowels gain height along either the front or back dimension, depending on whether they were front or back to begin with: thus *ą* goes to *ǫ*, *ę* goes to *į*, and *ǫ* and *į* remain as they are. Recall that, at the root level, nasal raising is obligatory (cf. section 3.1), while at the stem level, nasal raising is variable.[14] In the case of forms such as /nà-whe-ne-zè/, this variablity presents no problem, as *nàwhenezè, nàwhęęzè,* and *nàwhįįzè* are all well-formed outputs and surface completely transparently. In the case of /bò-kà-e-ne-t'è/, however, which surfaces, obligatorily, as *bòkàįt'è*, the optionality of nasal raising at the stem level is problematic. This is because, while the output of the stem-level phonology could be either /bò-kà-(ęę)t'è/ or /bò-kà-(įį)t'è/, when these are run through the grammar at the word level, the predicted outputs are *(bòkà)(ęę)t'è* and *bò(kàį)t'è* respectively, only the latter of which is a grammatical output for this combination of morphemes. Formally, this poses a problem, since it is not possible to stipulate that one of the two possible inputs to the word level will simply crash the derivation, and even if this were possible, it would be against Richness of the Base. Therefore, we must find an analysis in which both possible inputs to the word level phonology, /bò-kà-(ęę)t'è/ and /bò-kà-(įį)t'è/, will yield a grammatical output. Since the treatment of /bò-kà-(įį)t'è/ is simple (it involves merely shortening of a long vowel, to avoid a trimoraic nucleus, and re-footing), I will treat the more difficult case, /bò-kà-(ęę)t'è/.

[14] I will argue that, at the word level and postlexically, there is no nasal raising.

Figure 3.19: Simplified attempt at nasal raising-tonal foot interaction.

/bò-kà-(ę̨ę̨)t'è/	ALIGN-R(FT, PRWD)	$*\mu\mu\mu]_{\sigma-nuc}$	MINDIP=3	Ident(V-height)
a. (bòkà)(ę̨ę̨)t'è	***!			
b. (bòkà)(į̨į̨)t'è	***!			*
c. bò(kàę̨ę̨)t'è	*	*!	*	
d. bò(kàį̨į̨)t'è	*	*!		*
e. bò(kàę̨)t'è	*		*!	
☞ f. bò(kàį̨)t'è	*			*

There are two questions to raise with respect to the input /bò-kà-(ę̨ę̨)t'è/: why should re-footing occur in the first place, so as to yield bò(kàį̨)t'è, and why should such re-footing pose any more of a problem for /bò-kà-(ę̨ę̨)t'è/ than /bò-kà-(į̨į̨)t'è/? The response to the first question is simply the constraint ALIGN-R(FT, PRWD): an output such as candidate (a) or (b) in Figure 3.19 above incurs three violations of ALIGN-R, while outputs (c) through (f), having only a single foot which is misaligned from the right edge by only one syllable, incur only a single violation. As for the second question the problem is the constraint MINDIP=3, based on Casali (1997: 58). This constraint requires that two vowels which form a diphthong differ in at least 3 features. This is satisfied by *ai*, in which the vowels *a* and *i* differ from each-other in the features [±front], [±high], and [±low], but not satisfied by *ae*, in which the vowels differ only in frontness and lowness (they are both non-high). Assuming that nasality is an independent dimension which does not count for MINDIP, this means that the sequence *a-į̨* can diphthongize by morphological concatenation, while the sequence *a-ę̨* cannot, *unless ę̨* raises to *į̨*. The solution presented in Figure 3.19 above (which I will reject) is that MINDIP=3 should simply outrank IDENT(V-HEIGHT): tonal foot structure forces us to construct a diphthong, constraints on diphthongs want *ę̨* to raise to *į̨*, therefore *aę̨* becomes *aį̨* and everything seems to work.

While the analysis in Figure 3.19 does work for the case of *bò(kàį̨)t'è*, unfortunately it makes the wrong predictions about Dogrib generally. That is, if it were always the case that reducing the number of alignment violations were always more important than preserving underlying vowel height, then forms such as *(bòkà)(eh)t'è* and *(bòkà)(et)t'è* would not exist either, but rather would coalesce into *bò(kàih)t'è* and *bò(kài)t'è*, respectively. The question then becomes, why do we get coalescence in *bò(kàį)t'è* but not *bòkàeht'è* or *bòkàett'è*? Recall that, in general, it is not possible to alter a vowel's height in Dogrib, *except* when that vowel is nasal, on account of the constraint RAISENASAL. It must be, therefore, that the RAISENASAL constraint somehow facilitates the raising of *ę* to *į* in words such as *bòkàįt'è*. This is illustrated in Figure 3.20 below.

Figure 3.20: Local conjunction of tonal feet and nasal raising.

/bò-kà-(ęę)t'è/	[[ALIGN-R(FT, PRWD)]^2 & RAISENASAL]$_{PrWd}$	*μμμ]$_{\sigma-nuc}$	MINDIP=3	IDENT-V[±high]	RAISENASAL	ALIGN-R(FT, PRWD)
a. (bòkà)(ęę)t'è	(*!)				*	***
b. (bòkà)(įį)t'è				*!		***
c. bò(kàęę)t'è		*!	*		*	*
d. bò(kàįį)t'è		*!		*		*
e. bò(kàę)t'è			*!		*	*
☞ f. bò(kàį)t'è				*		*

Figure 3.20 presents six possibilities with respect to foot structure, nuclear moras, and vowel height. Candidates (c) and (d) are excluded more or less off-hand since they contain 3 nuclear moras, which is impermissible phonotactically in Dogrib. Presumably, candidates (e) and (f) violate the constraint Max(μ), ranked somewhere below *μμμ]$_{\sigma-nuc}$, though this is not directly relevant here. Candidate (e) is likewise ill-formed on account of its violation of MINDIP=3. The three remaining viable candidates (temporarily disregarding the local constraint conjunction) are (a), (b), and (f). In order to account for other examples such as *(bòkà)(eh)t'è*, one should be able to say that, if nasality were not present, then candidate (a) would win. This obtains by use of the constraint IDENT-V[±high]. Candidates (b) and (f) *would*, if the high-ranked locally conjoined constraint were not present, fatally violate this constraint. Indeed candidate (b) is excluded for this reason, because IDENT-V[±high] outranks RAISENASAL, and so raising of *ę* to *į* as in candidate (b) incurs a fatal violation. Thus, we would expect candidate (a) to be the winner, were it not for the high-ranked locally conjoined constraint [[ALIGN-R(FT, PRWD)]^2 & RAISENASAL]$_{PrWd}$, "two violations of Align-Right and an unraised nasal are

not allowed within a prosodic word". I admit that this type of local conjunction is somewhat odd and perhaps typologically unusual, though I do not see what other conclusion one could draw. Since (as will be shown in the next tableau) nasality is crucial in ruling out candidates such as (a) in Figure 3.20, it must be that RAISENASAL is conjoined with something. The only other alternative I can think of is a local conjunction of RAISENASAL and ONSET, which would say in effect that hiatus is not tolerated where nasal vowels are involved. This may sound more plausible *prima facie*, although I am not aware of any claims that hiatus is articulatorily more difficult when the velum is lowered (some form of rhinoglottophilia, perhaps?). In any event, an appeal to ONSET would complicate the formal analysis, since ONSET per se is not the driving force behind coalescence; Alignment is. Therefore, although $[[\text{ALIGN-R(FT, PRWD)}]^2$ & $\text{RAISENASAL}]_{PrWd}$ may be somewhat problematic typologically, I see no other alternative given the facts of Dogrib.[15]

Figure 3.21 below demonstrates that, using the exact same constraint ranking as in Figure 3.20, the correct result is obtained when the input vowel is non-nasal.

Figure 3.21: Local conjunction of tonal feet and nasal raising.

/bò-kà-(eh)t'è/	$[[\text{ALIGN-R(FT, PRWD)}]^2$ & $\text{RAISENASAL}]_{PrWd}$	$*\mu\mu\mu]_{\sigma-nuc}$	MINDIP=3	IDENT-V[±high]	RAISENASAL	ALIGN-R(FT, PRWD)
☞a. (bòkà)(eh)t'è						***
b. (bòkà)(ɪh)t'è				*!		***
c. bò(kàeh)t'è			*!			*
d. bò(kàɪh)t'è				*!		*

In Figure 3.21, there are no candidates that violate the constraint RaiseNasal, since there are no nasal vowels, and so no candidate violates the high-ranked locally conjoined constraint, either. No candidate is able to licitly raise its vowel *e* to *i* either, since, IDENT-V[±high] dominates ALIGN-R, and so candidates (b) and (d) fail due to unmotivated vowel raising. Candidate (c) fatally violates MINDIP=3 for the sequence *àeh*. Thus, the correct output, *bòkàeht'è*, is selected as the winner, and the same constraint-ranking which predicted that diphthongization and raising should occur in Figure 3.20 predicts hiatus in Figure 3.21, on account of the difference in nasality.

[15] Insert references regarding typological implications of local conjunction.

3.2.3.4 Gemination.

The last phenomenon to be treated here at the word level is gemination. The main points to be made about gemination are that, firstly, although it always serves to create a heavy syllable in stressed position, it is not motivated by SWP *per se* but rather serves to repair degenerate feet created at the stem level, i.e. to satisfy FTBIN. Secondly, gemination is also sensitive to tone, that is, whether the resulting prosodic word will contain level feet is an important consideration in whether or not to geminate. An instance of gemination is analyzed in Figure 3.22.

Figure 3.22: Analysis of gemination in Dogrib.

/bò-kà-(e)t’è/	FTBIN	*LAPSE	IDENT-V[±high]	MINDIP=3	LEVELFOOT	Align-R (Ft, PrWd)	*Cμ	SWP
a. (bòkà)(e)t’è	*!					***		*
b. (bòkà)et’è		*!				**		*
c. (bòkà)(i)t’è	*!		*			***		*
d. bò(kài)t’è			*!		*	*		
e. bò(kàe)t’è				*!	*	*		
f. bò(kà.e)t’è					*!	*		*
☞g. (bòkà)(et)t’è						***	*	

Candidates (a), (b), and (c) are excluded immediately because they violate the basic requirements of metrical well-formedness in Dogrib: (a) and (c) contain degenerate feet and violate FTBIN, while (b) has two unfooted syllables in a row and violates *LAPSE. Candidate (d) is excluded because it raised the vowel *e* to *i* in order to create a well-formed diphthong, yet in doing so fatally violated IDENT-V[±high]. The interesting case is candidate (f). Although candidate (f) is superior to candidate (g) in terms of alignment violations, (f) contains a low-high trochee, and thus violates LEVELFOOT. Candidate (g) on the other hand, incurs more violations of ALIGN-R, but has all level feet. What this illustrates is that, in a sequence of all light syllables, LEVELFOOT is able to shift a foot one syllable further to the left, as in candidate (g). In order to avoid a violation of foot binarity or *LAPSE, candidate (g) also geminates the stem-initial consonant, which violates *Cμ. Note, crucially, that the gemination in candidate (g) was forced by foot binarity and LEVELFOOT, *not* by the stress-to-weight principle, which is ranked below *Cμ. The importance of the relative ranking of *Cμ and SWP is illustrated in the next example.

Figure 3.23: Analysis of gemination in Dogrib.

/bò-kà-(ì)t'è/	FTBIN	*LAPSE	MINDIP=3	LEVELFOOT	Align-R (Ft, PrWd)	*Cμ	SWP
a. (bòkà)ìt'è		*!			**		*
b. (bòkà)(ì)t'è	*!				***		*
☞c. bò(kàì)t'è					*		
d. bò(kà.ì)t'è					*		*!
e. (bòkà)(ìt)t'è					***!	*	

Figure 3.23 above differs from the previous example in that the vowel immediately preceding the stem has a low tone rather than a high tone. Just as in the previous tableau, candidates (a) and (b), which violate foot binarity and *LAPSE, respectively, are eliminated immediately. What differs in this case is that candidate (e), similar to the winning candidate in Figure 3.22, is eliminated because it fatally violates ALIGN-R, even though it incurs the same number of violations of Align-R as *bòkàett'è* in Figure 3.22. This is because, since the vowel is already low-toned, there is no issue of creating a low-high trochee, and so the winning candidate is decided by alignment instead. Finally, candidate (c) is distinguished from candidate (d) based on syllabification. In candidate (d), *a.i* is syllabified as two separate syllables, that is, a light-light trochee, which means that a light syllable receives main stress, whereas in candidate (c) they are syllabified as a diphthong, which creates a single, heavy, main-stressed syllable, satisfying SWP.

3.2.4: Summary of Stem-Level and Word-Level processes.

A summary of stem-level and word-level processes in Dogrib is given in Figure 3.24. Each level is to be read as a stage in a serial derivation, where the output of the stem level is the input to the word level, the output of the word level is the input to the postlexical level, etc. The processes which occur at each level are summarized in the leftmost column.

Figure 3.24: Sample derivation for bòkà√t'è 'cook' .

	/bò-kà-e-ne-t'è/	/bò-kà-e-wìd-t'è/	/bò-kà-whe-wìd-t'e/	/bò-kà-whe-e-ne-t'e/
Stem level nasal coalescence, vowel coalescence, nasal raising, glide deletion, vowel deletion,	bò-kà-(ę̀ę̨)t'è	bò-kà-(ì)t'è	bò-kà-(whì)t'e	bò-kà-(whę̀ę̨)t'e
Word level nasal coalescence, vowel coalescence, re-footing, gemination	bò(kàį)t'è	bò(kàì)t'è	bò(kàwhì)t'e	(bòkà)(whę̀ę̨)t'e
Postlexical tonal processes, nasalization, contrastive consonant length	bò(kàį)t'è	bò(kàì)t'è	bò(kàwhì)t'e	(bòkà)(whę̀ę̨)t'e

So far, I have not given any examples of postlexical processes. For the most part, postlexical processes involve phrasal intonation, although consonant length is also rendered opaque by nasalization, which I will describe in the next section.

3.3 Postlexical Phonology.

3.3.1 Re-footing, geminates, and nasals.

In Yellowknife Dogrib, there is a variable loss of nasal codas, accompanied by nasalization of the preceding vowel. For example, the 2nd person singular optative *bòkàwįt'è*, 'you may/will cook' varies freely with *bòkàwint'è* . Crucially, however, when the nasal coda is lost, the preceding vowel is not lengthened. In addition, this failure of compensatory lengthening results in an exceptional metrical pattern: in the output form for the nasalized variant, the rightmost foot is final in the prosodic word, i.e. (*bòkà)(wįt'è)*. This is in violation of NONFINALITY. What we might have expected instead was that gemination should apply, i.e. **(bòkà)(wįt)t'è*, but this also does not occur. How are we to explain these facts?

Figure 3.25: Nasalization and re-footing.

/(bòkà)(win)t'è/	CODACON	MAX(C)	FTBIN	MAXASSOC (μ, C)	DEPASSOC (C, μ)	NONFIN(FT, PRWD)	MAX(μ)
☞a. (bòkà)(win)t'è	*						
☞b. (bòkà)(wįt'è)		*				*	*
c. (bòkà)(wį)t'è		*	*!				
d. (bòkà)(wįį)t'è		*		*!			
e. (bòkà)(wįt)t'è		*			*!		

That nasalization itself is a variable process is indicated by the dotted line separating CODACON and MAX(C). Here, CODACON is to be understood as a sort of placeholder for a combination of several constraints which state that the only allowable codas in Dogrib are *h* and geminates, while MAX(C) seeks to retain all underlying consonants. That these are unranked with respect to each other at the top of the hierarchy means that we should expect two winning candidates to vary with each-other: one which retains the underlying nasal coda, and one which deletes it. The candidates which delete the nasal coda, (b)-(e), are distinguished according to several criteria. I am presupposing, first of all, a high-ranked constraint Max[±nas], which demands that the nasal feature be preserved even if its host segment is deleted. Candidate (c) is a representational variant of candidate (b): (c) moves the right edge of the rightmost foot one syllable further to the left. This satisfies NONFINALITY, but creates a degenerate foot, on account of which (c) is excluded. Of greater interest are candidates (d) and (e). In candidate (d), nasalization occurs with compensatory lengthening of the preceding vowel. This violates the constraint MAXASSOC(μ, C), "for every mora, it must remain associated with its consonant". Crucially, this constraint is satisfied vacuously *if the mora is deleted*. This is yet another example of the same pattern by which, as stated previously, the autosegmental phonology of Dogrib does not seem very autosegmental: moras, like tones, behave like properties of the segments with which they are associated. Ironically, the fact that nasalization of the preceding vowel occurs at all means that, implicitly, the nasal feature has reassociated. Formally, this poses no problem, since the hypothetical constraint that would cause the feature [±nasal] to behave like a property of segments, MAXASSOC([±nas], C) merely needs to be low-ranked. However, this fact does seem to preclude making any sort of across-the-board, parametric statement about whether or not segmental features can float around in some language. I leave it as a question for typological investigation whether nasality is more likely than tone or moras to re-associate cross-linguistically. Finally, candidate (e) is excluded because gemination cannot happen postlexically, as represented by the constraint DEPASSOC(C, μ), "for every consonant, do not add a link to another mora". In effect, this means that gemination is constrastive post-lexically, and geminates are a "quasi-phoneme" in Dogrib. Thus the winning candidates are (a) and (b).

3.3.2 H to M tone in Tonal Foot Domain.

In other work (Jaker, to appear), I provide F0 data to show that, under certain conditions, a phonological high tone (H) is lowered to a middle tone (M) within a tonal foot. In this section, I will explain the distribution of this middle tone and the phonological motivation for H to M lowering.

I claim that the middle tone is a last-resort repair strategy for non-level feet postlexically, that is, to satisfy the constraint LEVELFOOT. But first, we must ask, how do non-level feet ever arise at the lexical level? This is illustrated in Figure 3.26.

Figure 3.26: Creation of non-level foot lexically in trisyllabic words.

/shè-(ts’e)zhe/	FtBin	NonFin(Ft, PrWd)	$[*C_\mu]^2_{PrWd}$	Parse(σ)	LevelFoot	$*C_\mu$
a. shè(ts’e)zhe	*!			**		
b. shè(ts’ezhe)		*!		*		
c. (shèt)(ts’ez)zhe			*!	*		**
d. shè(ts’ez)zhe				**!		*
☞f. (shèts’e)zhe				*	*	

The first item of note is that non-level feet only arise, lexically, in trisyllabic words. In words with 4 syllables or more, such as *bò(kàwì)t’è,* unparsed syllables to the left of the main-stressed foot are tolerated because they save on alignment violations. That is, every foot in addition to the rightmost foot causes additional violations of the constraint Align-R(Ft, PrWd) (not shown in the tableau above). In trisyllabic words, however, alignment plays no such role. We will first of all exclude candidate (c), which satisfies LevelFoot by geminating twice. I have posited a constraint $[*C_\mu]^2_{PrWd}$ against two geminates within a prosodic word. It is likely that this is in fact a prohibition against adjacent geminates, which exists uncontroversially in Inuktitut, and is referred to as Schneider’s Law (Lipscomb 1992, Dresher & Johns 1995). The remaining candidates, (a), (b), (d), and (e) each have only one foot within the prosodic word. Candidate (a) is excluded on account of being a degenerate foot, and (b) is excluded for violating NonFinality. The choice then comes down to (d) and (e). Based on the examples seen so far, cf. section 3.2.3.4, we might expect candidate (d), with gemination, to win, parallel to quadrisyllabic words such as *(bòkà)(wit)t’è* The crucial constraint in Figure 3.26 above is Parse(σ): candidate (f) is preferred over (d) because the maximum number of syllables is parsed, while still respecting NonFinality. In other words, parsing the maximum number of syllables is, in general, more important than maintaining level feet. The reason why this does not apply in the case of *(bòkà)(wit)t’è* is that, for this lexical item, parsing is independent of level feet: if the two leftmost syllables were unparsed, e.g. **bòkà(wit)t’è,* there would still be no LevelFoot violations. In the case of *(shèt’se)zhe*, however, Parse(σ) and LevelFoot are in direct conflict, as illustrated by candidates (d) and (e), and the candidate with the non-level foot, (f), is chosen at the winner.

Given that the output of the lexical phonology contains a non-level foot, we may now examine how this is dealt with postlexically. This is illustrated in Figure 3.27.

Figure 3.27: H → M postlexically (F0 compression).

/(shèts'e)zhe/	FTBIN	DEPASSOC(C, μ)	NONFIN(FT, PRWD)	LEVELFOOT	IDENT(Tone)
a. L H H she (ts'e) zhe	*!				
b. L H H (shet)(ts'ez)zhe		**!			
c. L H she (ts'ezhe)			*!		
d. L H H (shets'e)zhe				*!	
e. L L H (shets'e)zhe					**!
f. L M H (shets'e)zhe					*

Candidate (a) in Figure 3.27 is excluded because it contains a degenerate foot, and candidate (b) is excluded because, postlexically, consonant length is contrastive and thus not available as a repair to improve foot structure. The remaining candidates, (c), (d), and (e), differ with respect to the output tone on the 2nd syllable, *t'se*. Candidate (d) has a high tone, just as in the input, candidate (e) changes this to a low tone, and candidate (f) has a middle tone. Informally, I have assumed a gradient scale by which lowering from H to M incurs one violation of IDENT(TONE), while H to L incurs two violations. I have also assumed that a L-H sequence violates LEVELFOOT, while a L-M sequence is acceptable. Whether or not this is the correct formal characterization of H to M lowering will require additional evidence; however, the general pattern is clear, which is that H to M lowering represents a sort of compromise between full lowering (H to L) and a fully faithful mapping.

4.0 Further Issues: Total Fusion and Morphological Reanalysis.

In this section I address the formal characterization of morphological reanalysis, that is, the process that transforms a partially fused representation into a fully fused representation, represented as stages 3 and 4, respectively, in Figure 1.1 at the beginning of this paper. To begin with, one must distinguish synchronic from diachronic fusion. Synchronic fusion happens to every morpheme sequence in the course of a derivation: by the time one reaches the postlexical level of representation, all morpheme boundaries have been erased by bracket-erasure (McCarthy 1986), and thus all morphemes are de facto fused postlexically. Thus, synchronically the issue is not *if* but *when* morphemes fuse together, that is, at what level of representation: root level, stem level, word level, or postlexically. Diachronically, the issue is, then, how two (or more) morphemes which previously were only fused postlexically come to be fused at an earlier level of representation. Figures 4.1a and 4.1b illustrate the root *√t'e* and the ablaut morpheme *è*, which are partially fused at the root level (as shown in 4.1a) but fully fused at the stem level (as shown in 4.1b).

Figure 4.1a below is similar to Figure 3.4 in section 3.1.2, except that now an additional constraint is shown, NO-ALLOMORPHY(ROOT). This constraint did not appear in any of the tableaux in section 3.0, as I presupposed that it was undominated. Yet it is precisely this constraint which is at issue in the transition from partial fusion to total fusion (cf. Figure 1.1). Allomorphy is defined, informally, as a situation in which two surface morphs are associated with the same set of morphosyntactic features, but are not in phonological correspondence. In Figure 4.1a, this situation is represented by candidate (e), which is identical to the winning candidate, (d), in every respect, except the lack of phonological correspondence. I am using "correspondence" in the sense of McCarthy & Wolf (2005): thus candidate (e) is completely unrelated phonologically to the input. This gives candidate (e) the advantage that it does not incur any violations of any faithfulness constraint; it only violates the single constraint NO-ALLOMORPHY.

Figure 4.1a: Partial Fusion at Root Level, Total Fusion at Stem Level.

H L μ μ /t'e+e/	NO-ALLOMORPHY(ROOT)	*CONTOUR(μ)	OCP(μ)	ALIGN-R(MORPH, μ)	MAX(L)	MAX(H)
a. contour, long vowel [H] [L] [μ] [μ] [t'[e]]			*!			
b. contour, short vowel [H] [L] [[μ]] [t'[e]]		*!		*		
c. high tone, short vowel [[H]] [[μ]] [t'[e]]				*	*!	
☞d. low tone, short vowel [[L]] [[μ]] [t'[e]]				*		*
e. fused, not in correspondence [L] [μ] [t'e]	*!					

Another way of stating the issue, informally, is that NO-ALLOMORPHY is a constraint operating at the morphological level of the grammar: once NO-ALLOMORPHY is violated, the morphology becomes more complex (since there are now, de facto, two morphemes with the same meaning) yet at the same time the phonology becomes more simple: the underlying form is the same as the surface form, and there are no faithfulness violations (all other things being equal). This situation is further illustrated in Figure 4.1b, which shows what happens to the root + ablaut complex in the stem level phonology. That is, 4.1b illustrates what happens when the output of the root level phonology becomes the input to the stem level phonology, and all brackets from the root level have been erased (that is, the root and ablaut morpheme have fully fused).

Figure 4.1b: Stem Level.

/t'$_1$è$_2$/	NO-ALLOMORPHY(STEM)	ALIGN-R(MORPH, μ)	MAX(L)	MAX(H)	NO-ALLOMORPHY(ROOT)
☞a. low tone, in correspondence t'$_1$è$_2$					*
b. high tone, in correspondence t'$_1$é$_2$			*!		*
c. low tone, not in correspondence t'$_3$è$_4$	*!				*
d. high tone, not in correspondence t'$_3$é$_4$	*!				*

The first item of note is that all four candidates incur a violation of the constraint NO-ALLOMORPHY(ROOT). The reason for this is that, at the stem level, the internal morphological structure of the root + ablaut suffix sequence is not visible, therefore the two variants *t'é* and *t'è* are treated as separately listed, atomic entities, and thus incur one violation of the NO-ALLOMORPHY(ROOT) constraint.[16] The constraint, however, refers only to the root + ablaut suffix complex, i.e. that which is the input to the stem-level phonology. From the point of view of the stem-level phonology, the root 'cook' simply has two alternate, separately listed forms which are phonologically unrelated. On the other hand, the high-ranked constraint NO-ALLOMORPHY(STEM) prohibits the creation of any *new* allomorphs at the stem level, i.e. candidates which are evaluated with respect to the input (which can be either allomorph) but are not in correspondence with it.

What we see described in Figures 4.1a and 4.1b is the canonical state of affairs in agglutinating languages: morphological fusion arises after each cycle through bracket-erasure, but during each cycle, the constraint-ranking maintains the phonological correspondence between input and output, and morpheme boundaries remain intact. We may now ask the question, under what circumstances will new allomorphs be generated at some level of representation?

[16] This would not seem to follow based on the formalism I have presented in this paper; in fact, the constraint No-Allomorphy requires paradigmatic evaluation: in order for a candidate such as *t'è* (candidate a) to incur a violation of No-Allomorphy, it must be the case that a related form *t'é* is part of the same paradigm and also visible. I will not further elaborate the mechanism by which this is evaluated, but see Jaker 2006.

Figure 4.2: Total Fusion: When Phonology Crashes.

H L μ μ /t'e+e/	*CONTOUR(μ)	OCP(μ)	ALIGN-R(MORPH, μ)	MAX(L)	MAX(H)	NO-ALLOMORPHY(ROOT)
a. contour, long vowel [H] [L] [μ] [μ] [t'[e]]		*!				
b. contour, short vowel [H] [L] [[μ]] [t'[e]]	*!		*			
c. high tone, short vowel [[H]] [[μ]] [t'[e]]			*!	*		
d. low tone, short vowel [[L]] [[μ]] [t'[e]]			*!		*	
☞e. fused, not in correspondence [L] [μ] [t'e]						*
☞f. fused, not in correspondence [H] [μ] [t'e]						*

Figure 4.2 represents a hypothetical situation in which, at the root level, markedness, faithfulness, OCP, and alignment are all high-ranked and undominated. Under such a situation, the phonology cannot generate any winning candidate. The phonology in effect crashes, and the violation is moved into the morphology. This is represented formally as a violation of the No-Allomorphy(root) constraint, and the winning candidate will be a candidate which is not in phonological correspondence with the input (and thus incurs no faithfulness violations) and also satisfies the relevant markedness constraints which caused the phonology to crash in the first place.

There is, however, a very serious problem with the analysis presented in Figure 4.2. If the winning candidate is not in correspondence with the input, and is not subject to any faithfulness constraints, the grammar has no means to distinguish between (e), the actual fully fused form, and (f), its mirror-image with a high tone. For that matter, any number of other candidates such as *ki*, *tl'a* or *whe* would fare equally well using the constraints just given.

Intuitively, the result we want is that a fully fused form from later in the derivation is "moved up" one level. In this case, that would mean that the fully fused form [t'e], which already exists at the stem level (as in Figure 4.1b) is moved up to the root level, so that the representation at the root level is now [t'e] rather than [t'[e]], as in 4.1a. But this is formally impossible given the architecture of Lexical Phonology (Kiparsky 1982, to appear): the derivation cannot "look ahead" to see what the fully fused output will look like in the next cycle, nor can a fully fused form travel "upstream" into an earlier cycle in the derivation.

There are two very different possible solutions to this problem. One would be to say that morphological reanalysis is outside a formal theory of phonology and morphology proper, and instead operates only diachronically, through language acquisition. Such an account would proceed roughly as follows: the default representation for language learners is to have no brackets: everything is fused. Morpheme boundaries (i.e. brackets) are constructed during acquisition based on similarities in morphosyntactic and phonological features. The phonological similarity between different surface morphs bearing the same morphosyntactic features is not exact: they will differ in some phonological features, but it is still possible to recognize them as manifestations of the same morpheme based on predictable phonological correspondences. Formally these predictable correspondences are represented as low-ranked faithfulness constraints, dominated by some markedness constraint. When both faithfulness and markedness are ranked too high, these correspondences break down (as in Figure 4.2). In the absence of such correspondences, future generations of language learners will be unable to construct internal morphological structure (e.g. [t'[e]]) and instead will represented it as fused ([t'e]).

There is something unsatisfying about an acquisition-based account, however. In general, diachronic changes in OT are described simply as constraint re-ranking. For example, for an input /ki/ in which IDENT(PLACE) >> PAL, the input will surface faithfully. Phonological change is described as a change in constraint-ranking: when PAL comes to outrank IDENT(PLACE), then /ki/ will surface as *ci*, and subsequently the underlying representation may or may not be restructured to /ci/ by Lexicon Optimization, depending on what other morphophonemic alternants are derived from the same input.

Why should morphology be any different? Why should change in a morphological system not be statable as a change in constraint-ranking? To achieve this, however, would require a substantial re-thinking of the foundations of Lexical Phonology. In particular, it would be necessary to assume that different strata (root level, stem level, word level, and postlexical) are not serially ordered, but parallel and mutually visible, and differ from each-other only in morphological bracketing. Since, in the analysis I have argued for thusfar, differences in bracketing have immediate consequences for the way faithfulness constraints are evaluated, it is possible that some types of phonological opacity may be attributed to conflicting bracketings. This is an area for future investigation.

5.0 Conclusion.

The preceding analysis has presented a general theory of the phonological mechanisms, especially coalescence, which may ultimately result in morphological fusion, and has sketched out an analysis of Dogrib, in which coalescence plays a major role. Several areas remain for future investigation. A formal implementation of morphological reanalysis, i.e. from Partial Fusion to Total Fusion, seems to be beyond the capabilities of LPM-OT in its current form, and may require a substantial revision of the architecture of the theory. At the micro-level, a detailed analysis of coalescence at the level of features, with reference to a more highly articulated theory of tier structure and feature geometry, awaits further investigation. Finally, a larger corpus of phonological data, including quantitative and instrumental measurements, will be necessary in order to confirm the analysis presented thusfar, which will be available following future field expeditions by the author.

Bibliography

Ackroyd, Lynda. Dogrib Grammar. Manuscript, University of Toronto, 1982.

Beckman, Jill. Shona Height Harmony: Markedness and Positional Identity." In Jill Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, eds., *University of Massachusetts Occasional Papers: Papers in Optimality Theory* 18:53-76. GLSA, Amherst, Massachusetts.

Blumenfeld, Lev. Constraints on Phonological Interactions. Doctoral Dissertation, Stanford University, 2006. Available on ROA: http://roa.rutgers.edu.

Borrelli, Doris. *Raddoppiamento Sintattico* in Italian: a Synchronic and Diachronic Cross-Dialectal Study. Doctoral Dissertation, Cornell University, 2000.

Casali 1997. Resolving Hiatus. Garland Publishing, 1997.

Coleman, Phyllis. Dogrib Phonology. Doctoral Dissertation, University of Iowa, 1976.

deLacy, Paul. The interaction of tone and stress in Optimality Theory. *Phonology* 19 (2002): 1-32.

Dresher, B. Elan and Johns, Alana. The Law of Double Consonants in Inuktitut. *Linguistica Atlantica* 17 (1995) 79-95.

Flemming 2001. Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology* 18 (2001): 7-44.

Golla, Victor. Hupa Grammar. Doctoral Dissertation, University of California, Berkeley, 1970.

Goldsmith, John. Autosegmental Phonology. Doctoral Dissertation, MIT, 1976.

Hyman, Larry and Schuh, Russel. Universals of Tone Rules: Evidence from West Africa. *Linguistic Inquiry* 5.1 (Winter 1974): 81-115.

Hyman, Larry. A Theory of Phonological Weight. CSLI Publications, 2003 [1986].

Itô, Junko and Mester, Armin. Realignment. In Kager, van der Hulst, and Zonneveld, eds., *The Prosody-Morphology Interface,* Cambridge University Press, 1999.

Jaker, Alessandro. Split Subject Agreement and Morphological Typology. Manuscript, Stanford University, 2006.

Keer, Edward. Geminates, OCP, and the nature of CON. Doctoral Dissertation, Rutgers University, available on ROA: http://roa.rutgers.edu.

Leben, William R. Suprasegmental Phonology. Doctoral Dissertation, MIT.

Leer, Jeff. Proto-Athabaskan Stem Variation, Part 1: Phonology. Alaska Native Language Center Research Papers, #1, 1979.

Lipscomb, David Robert. Differences and Similarities in the Prosodic systems of Western and Eastern Eskimo. *Acta Linguistica Hafniensia* vol 24, Copenhagen, 1992.

Kiparsky, Paul. Abstractness, Opacity, and Global Rules. In *Three Dimensions in Linguistic Theory,* ed. by Osamu Fujimura, 57-86. Tokyo: TEC 1973.

Kiparsky, Paul. Lexical Phonology and Morphology. In I.S. Yang, ed., *Linguistics in the Morning Calm,* Linguistic Society of Korea, Hanshin, Seoul, 3-91 1982.

Kiparsky, Paul. Syllables and Moras in Arabic. In *The Syllable in Optimality Theory*, Féry and van der Vijver (eds.), Cambridge University Press, 2003.

Kiparsky, Paul. Paradigms and Opacity. Manuscript, Stanford University, to appear.

Kingston, John. The Phonetics of Athabaskan Tonogenesis. In Sharon Hargus and Keren Rice, eds., *Athabaskan Prosody,* John Benjamins, 2005.

Krauss, Michael. Athabaskan Tone. In Sharon Hargus and Keren Rice, eds., *Athabaskan Prosody*, John Benjamins, 2005.

Marinakis, Aliki. Seeking Simplicity: The preference for minimal syllable structure in Dogrib. Master's Thesis, University of Victoria, 2004.

McCarthy, John. OCP Effects: Gemination and Antigemination. *Linguistic Inquiry* 17:2, 1986, 207-263.

McCarthy, John. OT Constraints are Categorical. *Phonology* 20 (2003): 75-138.

McCarthy, John and Prince, Alan. Generalized Alignment. In G.E. Booij and J. van Marle (eds.), *Yearbook of Morphology 1993*, Dordrecht: Kluwer, 79-153.

McCarthy, John and Prince, Alan. Faithfulness and Reduplicative Identity. In Jill Beckman, Laura Walsh Dickey, and Suanne Urbanczyk, eds., *University of Massachusetts Occasional Papers: Optimality Theory* 18:249-384. GLSA, Amherst, Massachusetts.

McCarthy, John and Wolf, Matthew. Less Than Zero: Correspondence and the Null Output. Availalbe on ROA: http://roa.rutgers.edu.

Mester, Armin. Studies in Tier Structure. Doctoral Dissertation, University of Massachusetts, Amherst, 1988.

Morén, Bruce. Distinctiveness, Coercion, and Sonority: A Unified Theory of Weight. Routledge, 2001.

Myers, Scott. OCP Effects in Optimality Theory. *Natural Language and Linguistic Theory* 15.4 (1997): 847-92.

Pater, Joe. Austronesian Nasal Substitution Revisited. In *Segmental Phonology in Optimality Theory*, Cambridge University Press, 2001.

Rice, Keren. Prosodic Constituency in Hare (Athapaskan): Evidence for the Foot. *Lingua* 82 (1990): 201-244.

Rice, Keren. Morpheme Order and Semantic Scope: Word Formation in the Athapaskan Verb. Cambridge University Press, 2000.

Stonham, John. Combinatorial Morphology. John Benjamins, 1994.

Zec, Draga. Footed tones and tonal feet: rhythmic constituency in a pitch-accent language. *Phonology* 16 (1999): 225-264.