☐ **ROA Version, 8/2002**. Essentially identical to the Tech Report, with new pagination (but the same footnote and example numbering); correction of typos, oversights & outright errors; improved typography; and occasional small-scale clarificatory rewordings. Citation should include reference to this version.

OPTIMALITY THEORY

Constraint Interaction in Generative Grammar

First circulated: April, 1993
RuCCS-TR-2; CU-CS-696-93: July, 1993
Minor Corrections: December, 1993
ROA Version: August, 2002

Alan Prince
Department of Linguistics
Rutgers Cognitive Science Center
Rutgers University
prince@ruccs.rutgers.edu

Paul Smolensky
Department of Cognitive Science
The Johns Hopkins University
[1993: University of Colorado at Boulder]
smolensky@cogsci.jhu.edu

Everything is possible but not
everything is permitted …
— Richard Howard, "The Victor Vanquished"

"It is demonstrated," he said, "that things cannot be otherwise: for, since everything was made for a purpose, everything is necessarily made for the best purpose."
— *Candide ou l'optimisme*. Ch. I.

Remark. The authors' names are arranged in lexicographic order.

2. Optimality in Grammar: Core Syllabification in Imdlawn Tashlhiyt Berber

Here we argue that certain grammatical processes can only be properly understood as selecting the *optimal output* from among a set of possibilities, where the notion *optimal* is defined in terms of the constraints bearing on the grammatical domain at issue.

2.1 The Heart of Dell & Elmedlaoui

The Imdlawn Tashlhiyt dialect of Berber (ITB) has been the object of a series of remarkable studies by François Dell and Mohamed Elmedlaoui (Dell & Elmedlaoui 1985, 1988, 1989). Perhaps their most suprising empirical finding is that in this language any segment — consonant or vowel, obstruent or sonorant — can form the nucleus of syllable. One regularly encounters syllables of the shape *tK, rB, xZ, wL*, for example. (Capitalization represents nucleus-hood of consonants.) The following table provides illustrative examples, with periods used to mark syllable edges:[4]

Nucleus Type	Example	Morphology	Reference
voiceless stop	.ra.t**K**.ti.	ra-t-kti	1985: 113
voiced stop	.b**D**.dL. .ma.ra.t**G**t.	bddl ma=ra-t-g-t	1988: 1 1985: 113
voiceless fricative	.t**F**.tKt. .t**X**.zNt.	t-ftk-t t-xzn-t	1985: 113 1985: 106
voiced fricative	.tx**Z**.nakk^w.	t-xzn#nakk^w	1985: 113
nasal	.tz**M**t. .t**M**.zħ.	t-zmt t-mzħ	1985: 112 1985: 112
liquid	.t**R**.g**L**t.	t-rgl-t	1985: 106
high vowel	.**i**l.d**i**. .rat.l**u**lt.	i-ldi ra-t-lul-t	1985: 106 1985: 108
low vowel	.tR.b**a**.	t-rba	1985: 106

Dell and Elmedlaoui marshall a compelling range of evidence in support of the claimed patterns of syllabification. In addition to native speaker intuition, they adduce effects from segmental phonology (emphasis spread), intonation, versification practice, and prosodic morphology, all of which agree in respecting their syllabic analysis.

---

[4] Glosses are *ratkti* 'she will remember'; *bddl* 'exchange!'; *maratgt* 'what will happen to you?'; *tftkt* 'you (2psg) suffered (pf.)a strain', *txznt* 'you stored'; *txznakk^w* 'she even stockpiled'; *tzmt* 'it(f.) is stifling; *tmzħ* 'she jested'; *trglt* 'you locked'; *ildi* 'he pulled'; *ratlult* 'you will be born'; *trba* 'she carried-on-her-back'.

The domain of syllabification is the phonological phrase. All syllables must have onsets except when they occur in absolute phrase-initial position. There, syllables may begin with vowels, either with or without glottal striction (Dell & Elmedlaoui 1985: 127 fn. 20), evidently a matter of phonetic implementation. Since any segment at all can form the nucleus of a syllable, there is massive potential ambiguity in syllabification, and even when the onset requirement is satisfied, a number of distinct syllabifications will often be potentially available. But the actual syllabification of any given string is almost always unique. Dell & Elmedlaoui discovered that assignment of nuclear status is determined by the relative sonority of the elements in the string. Thus we find the following typical contrasts:

(3) **Sonority Effects on Nuclear Status**

(a) tz**M**t — *t**Z**mt ‘*m* beats *z* as a nucleus’

(b) rat.l**u**lt — *ra.t**L**.w**L**t.‘*u* beats *l* as a nucleus’

Orthography: we write *u* for the nuclear version, *w* for the marginal version of the high back vocoid, and similarly for *i* and *y*: as with every other margin/nucleus pair, we assume featural identity.

All the structures in (3), including the ill-formed ones, are locally well-formed, composed of licit substructures. In particular, there is nothing wrong with syllables ***tZ***, ***tL***, or ***wLt*** nor with word-final sequences *mt* — but the more sonorous nucleus is chosen in each case. By examining the full range of such contrasts, Dell and Elmedlaoui establish the relevance of the following familiar kind of 8-point hierarchy:

(4) **Sonority Scale**

|Low V|>|High V|>|Liquid| >|Nasal|>|Voiced Fric.|>|Voiceless Fric.|>|Voiced Stop|>|Voiceless Stop|

We write $|\alpha|$ for the sonority or intrinsic prominence of α.

With the sonority scale in hand, Dell and Elmedlaoui then propose an iterative syllable-construction procedure that is designed to select the correct nuclei. Their algorithm can be stated in the following way, modified slightly from Dell & Elmedlaoui 1985: 111(15):

(5) **Dell−Elmedlaoui Algorithm for Core Syllabification** (DEA)

Build a core syllable (“CV”) over each substring of the form XY, where

X is any segment (except [*a*]), and

Y is a matrix of features describing a step of the sonority scale.

Start Y at the top of the sonority scale and replace it successively with the matrix of features appropriate to the next lower step of the scale.

(Iterate from Left to Right for each fixing of the nuclear variable Y.)

Like all such procedures, the DEA is subject to the Free Element Condition (FEC: Prince 1986), which holds that rules establishing a level of prosodic structure apply only to elements that are not already supplied with the relevant structure. By the FEC, the positions analyzed by the terms X,Y must be free of syllabic affiliation. Effectively, this means that any element seized as an onset is no longer eligible to be a nucleus, and that a segment recruited to nucleate a syllable is not then available to serve as an onset.

There are other syllabification phenomena in ITB that require additional rules beyond the DEA; we will abstract away from these and focus on the sense of DEA itself.[5] We will also put aside some wrinkles in the DEA which are related to parenthesized expressions in (5) — the lack of a glide counterpart for /a/, the phrase-initial loosening of the onset requirement, and the claimed left-to-rightness of the procedure.[6]

The DEA is a rule, or rather a schema for rules, of exactly the classical type A → B / C—D. Each rule generated by the schema has a Structural Description specified in featural terms and a Structural Change ('construct a core syllable'). To see how it works, consider the following derivations:

---

[5] Not the least of these is that syllables can have codas; the DEA serves essentially to locate syllable nuclei, which requires that onsets be taken into consideration. But it is not difficult to imagine plausible extensions which lead to adjunction of codas. More subtle, perhaps, are these phenomena:

a. obstruents are always nonsyllabic in the envs. #— and —#.

b. sonorant C's are optionally nonsyllabic —# under certain conditions.

c. the 1st element of a tautomorphemic geminate is never an onset.

In addition, the DEA does completely resolve sequences /~aa~/, which according to other sources, surface as *~aya~* (Guerssel 1985). The appropriate approach to epenthetic structure within OT involves the constraint FILL, which makes its appearance below in §3.1 and receives full discussion in §6.

[6] We deal with the fact that [a] cannot occupy syllable margins in §8.1.1. The commonly encountered relaxation of the onset requirement in initial position is resolved in McCarthy & Prince 1993 in terms of constraint interaction, preserving the generality of ONS. Dell & Elmedlaoui are themselves somewhat ambivalent about the need for directionality (Dell & Elmedlaoui 1985: 108); they suggest that "the requirement [of directionality] is not concerned with left to right ordering *per se*, but rather with favoring application of [the DEA] that maximize the sonority differences between [onset and nucleus]" (Dell & Elmedlaoui 1985:127, fn. 22). In addition, they note that directionality falsely predicts **.i.tBd.rin.* from /i=t-!bdri-n/ 'for the cockroaches', whereas the only licit syllabification is *.it.bD.rin*. The reason for this syllabification is not understood. A directionless theory leaves such cases open for further principles to decide.

(6) **DEA in Action**

Steps of the DEA	/ratlult/ 'you will be born'
Seek [X][+low, −cns] *& Build*	(**ra**)tlult
Seek [X][−low,−cns] *& Build*	(ra)t(**lu**)lt
Seek [X][+cns,+son,−nas]	*−blocked by FEC−*
Seek [X][+cns,+son,+nas]	—
Seek [X][−son, +cnt, +voi]	—
Seek [X][−son, +cnt, −voi]	—
Seek [X][−son,−cnt,+voi]	—
Seek [X][−son,−cnt,−voi] *& Build*	(ra)t(lu)(**lT**)[7]

(7) **DEA in Action**

Steps of the DEA	/txznt/ 'you sg.stored'
Seek [X][+low,−cns]	—
Seek [X][−low,−cns]	—
Seek [X][+cns,+son,−nas]	—
Seek [X][+cns,+son,+nas] *& Build*	tx(**zN**)t
Seek [X][−son, +cnt, +voi]	—
Seek [X][−son, +cnt, −voi] *& Build*	(**tX**)(zN)t
Seek [X][−son,−cnt,+voi]	—
Seek [X][−son,−cnt,−voi]	—

[7] We show the form predicted by the DEA. The form is actually pronounced r**a**t.l**u**lt. because obstruents cannot be nuclear next to phrase boundaries, as mentioned in fn. 5.

(8) **DEA in action**

Steps of the DEA	/txznas/ 'she stored for him'
Seek [X][+low,−cns] *& Build*	txz(**na**)s
Seek [X][−low,−cns]	—
Seek [X][+cns,+son,−nas]	—
Seek [X][+cns,+son,+nas]	*−blocked by FEC−*
Seek [X][−son, +cnt, +voi] *& Build*	t(**xZ**)(na)s
Seek [X][−son, +cnt, −voi]	*−blocked by FEC−*
Seek [X][−son,−cnt,+voi]	—
Seek [X][−son,−cnt,−voi]	*−blocked by FEC−*

The DEA provides an elegant and straightforward account of the selection of syllable nuclei in the language. But it suffers from the formal arbitrariness characteristic of re-writing rules when they are put to the task of dealing locally with problems that fall under general principles, particularly principles of output shape. (By 'formal arbitrariness', we mean that a formal system rich enough to allow expression of the desired rule will also allow expression of many undesired variations of the rule, so that the rule itself appears to be an arbitrary random choice among the universe of possibilities.) The key to the success of the DEA is the way that the variable Y scans the input, starting at the top of the sonority scale and descending it step by step as the iterative process unfolds. We must ask, why start at the top? why *descend* the scale? why not use it in some more elaborate or context-dependent fashion? why apply the scale to the nucleus rather than the onset? [8]

The answers are to be found in the theory of syllable structure markedness, which is part of Universal Grammar. The more sonorous a segment is, the more satisfactory it is as a nucleus. Conversely, a nucleus is more satisfactory to the degree that it contains a more sonorous segment. It is clear that the DEA is designed to produce syllables with optimal nuclei; to ensure that the syllables it forms are the most *harmonic* that are available, to use the term introduced in §1. Dell and Elmedlaoui clearly understand the role of sonority in choosing between competing analyses of a given input string; they write:

> When a string …PQ… could conceivably be syllabified as …Pq… or as …pQ… (i.e. when either syllabification would involve only syllable types which, when taken individually, are possible in ITB), the only syllabification allowed by ITB is the one that takes as a syllabic peak the more sonorous of the two segments.
>
> — Dell & Elmedlaoui 1985:109

[8] These are exactly the sort of questions that were fruitfully asked, for example, of the classic TG rule of Passive that moved subject and object, inserted auxiliaries, and formed a PP: why does the post-verbal NP move *up* not *down*? why does the subject NP move at all? why is by+NP a PP located in a PP position? and so on.

But if phonology is couched in re-writing rules, this insight cannot be cashed in as part of the function that assigns structural analyses. It remains formally inert. Dell and Elmedlaoui refer to it as an "empirical observation," emphasizing its extra-grammatical status.

The DEA itself makes no contact with any principles of well-formedness; it merely scans the input for certain specific configurations, and acts when it finds them. That it descends the sonority scale, for example, can have no formal explanation. But the insight behind the DEA can be made active if we re-conceive the process of syllabification as one of choosing the optimal output from among the possible analyses rather than algorithmic structure-building. Let us first suppose, with Dell and Elmedlaoui, that the process of syllabification is serial, affecting one syllable at a time (thus, that it operates like Move-α or more exactly, Move-x of grid theory). At each stage of the process, let all possible single syllabic augmentations of the input be presented for evaluation. This set of candidates is evaluated by principles of syllable well-formedness and the most harmonic structure in the set is selected as the output. We can state the process informally as follows:

(9) **Serial Harmonic Syllabification (informal)**.
Form the optimal syllable in the domain.
Iterate until nothing more can be done.

This approach depends directly on the principles of well-formedness which define the notion 'optimal'. No instructions are issued to the construction process to contemplate only one featurally-specified niche of the sonority scale. Indeed, the Harmonic syllabification algorithm has no access to any information at all about absolute sonority level or the specific featural composition of vowels, which are essential to the DEA; it needs to know whether segment α is *more* sonorous than segment β, not what their sonorities or features actually are. All possibilities are entertained simultaneously and the choice among them is made on grounds of general principle. That you start at the top of the scale, that you descend the scale rather than ascending it or touring it in some more interesting fashion, all this follows from the universal principles that define the relative Harmony of nucleus-segment pairings. The formal arbitrariness of the DEA syllable-constructing procedure disappears because the procedure itself ('make a syllable') has been stripped of intricacies.[9]

This is an instance of Harmony-increasing processing (Smolensky 1983, 1986; Goldsmith 1991, 1993). The general rubric is this:

(10) **Harmonic Processing**
Go to the most harmonic available state.

We speak not of 'relative well-formedness' but rather of *relative Harmony*: Harmony is a well-formedness scale along which a maximal-Harmony structure is well-formed and all other structures are ill-formed.

---

[9] Further development of this idea could eliminate complications at the level of the general theory; in particular, the appearance of obeying the Free Element Condition during serial building of structure could be seen to follow from the fact that disobeying it inevitably decrements the Harmony of the represention.

We conclude that the Dell-Elmedlaoui results establish clearly that harmonic processing is a grammatical mechanism, and that optimality-based analysis gives results in complex cases. Let us now establish a formal platform that can support this finding.

2.2 Optimality Theory

What, then, is the *optimal* syllable that Harmonic Syllabification seeks? In the core process that we are focusing on, two constraints are at play, one ensuring onsets, the other evaluating nuclei. The onset constraint can be stated like this (Itô 1986, 1989):

(11) **The Onset Constraint (ONS)**. Syllables must have onsets (except phrase initially).

As promised, we are not going to explicate the parenthesized caveat, which is not really part of the basic constraint (see McCarthy & Prince 1993: §4). The nuclear constraint looks like this:[10]

(12) **The Nuclear Harmony Constraint (HNUC)**. A higher sonority nucleus is more harmonic than one of lower sonority.

I.e. If $|x| > |y|$ then Nuc/x $\succ$ Nuc/y.

The formalizing restatement appended to the constraint uses some notation that will prove useful.

For 'x is more harmonic than y' we write x $\succ$ y.

For 'the intrinsic prominence of x' we write $|x|$.

'A/x' means 'x belongs to category A, x is the constituent-structure child of A'

The two kinds of order $\succ$ and $>$ are distinguished notationally to emphasize their conceptual distinctness. Segments of high sonority are not more harmonic than those of lower sonority. It is only when segments are contemplated in a structural context that the issue of well-formedness arises.

It is necessary to specify not only the relevant constraints, but also the set of candidates to be evaluated. To do this we need to spell out the function Gen that admits to candidacy a specific range of structurings or parses of the input. In the case at hand, we want something roughly like this:

(13) **Gen** ($input_i$): the set of (partial) syllabifications of $input_i$ which differ from input_i in no more than one syllabic adjunction.

For any form $input_i$ to undergo Serial Harmonic Syllabification, the candidate set Gen($input_i$) must be evaluated with respect to the constraints ONS and HNUC. There would be little to say if evaluation were simply a matter of choosing the candidate that satisfies both constraints. Crucially,

[10] It is also possible to conceive of the operative constraint in a kind of 'contrapositive' manner. Because all underlying segments of ITB are parsed, a segment is a nucleus iff it is not a member of the syllable margin. Consequently, negative constraints identifying the badness of syllable margins can have the same effect as positive constraints identifying the goodness of nuclei. We investigate this approach below in §8.1.1, §8.3.3, §8.4.2.

and typically, this straightforward approach cannot work. Conflict between the constraints ONS and HNUC is unavoidable; there are candidate sets in which no candidate satisfies both constraints.

Consider, for example, the syllabification of the form /ħaul-tn/ 'make them (m.) plentiful' (Dell & Elmedlaoui 1985:110). Both ONS and HNUC agree that the core syllable ***ħa*** should be formed: it has an onset as well as the best possible nucleus. Similarly, we must have a final syllable *t**N***. But what of the rest of the string? We have two choices for the sequence /ul/: a superior nucleus lacking an onset, as in ***u**l*; or an onsetted syllable with an inferior nucleus, as in *w**L***. This situation can be perspicuously displayed in tabular form:[11]

(14) **Constraint Inconsistency**

Candidates /ħaul-tn/	ONS	HNUC
~.w**L**.~		|l|
~.**u**l.~	*	|u|

The cells contain information about how each candidate fares on the relevant constraint. A blank cell indicates that the constraint is satisfied; a star indicates violation. (In the case of a scalar constraint like HNUC we mention the contents of the evaluated element.) The first form succeeds on ONS, while the second form violates the constraint. The relative performance is exactly the opposite on HNUC: because |u| > |l|, the second, onsetless form has the better nucleus. The actual output is, of course, *.ħa.w**L**.tN*. The onset requirement, in short, takes priority.

Such conflict is ubiquitous, and to deal with it, we propose that a relation of *domination*, or priority-ranking, can be specified to hold between constraints. When we say that one constraint *dominates* another, we mean that when they disagree on the relative status of a pair of candidates, the dominating constraint makes the decision. If the dominating constraint does not decide between the candidates — as when both satisfy or both violate the constraint equally — then the comparison is passed to the subordinate constraint. (In the case of a more extensive hierarchy, the same method of evaluation can be applied repeatedly.)

In the case at hand, it is clear that ONS must dominate HNUC. The top priority is to provide syllables with onsets; the relative Harmony of nuclei is a subordinate concern whose force is felt only when the ONS issue is out of the way. We will write this relation as ONS ≫ HNUC. Given such a hierarchy, an optimality calculation can be usefully presented in an augmented version of display (14) that we will call a *constraint tableau*:

[11] Properly speaking, if we limit our attention to the core syllable stage of the procedure, we should be comparing core *.u.* with core *.wL.* But the comparison remains valid even after coda consonants are adjoined and we wish to emphasize that the two cited analyses of /ħaul-tn/ differ only in treatment of the sequence /ul/.

(15) **Constraint Tableau** for partial comparison of candidates from /ħaultn/

Candidates	ONS	HNUC
☞ ~.w**L**.~		|l|
~.**u**l.~	* !	|u|

Constraints are arrayed across the top of the tableau in domination order. As above, constraint violations are recorded with the *mark* *, and blankness indicates total success on the constraint. These are the theoretically important conventions; in addition, there is some clarificatory typography. The symbol ☞ draws the eye to the optimal candidate; the ! marks the *crucial* failure for each suboptimal candidate, the exact point where it loses out to other candidates. Cells that do not participate in the decision are shaded. In the case at hand, the contest is decided by the dominant constraint ONS; HNUC plays no role in the comparison of *.w**L**.* and *.**u**l.* HNUC is literally irrelevant to this particular evaluation, as a consequence of its dominated position — and to emphasize this, we shade its cells. Of course, HNUC is not irrelevant to the analysis of *every* input; but a precondition for relevance is that there be a set of candidates that tie on ONS, all passing it or all failing it to the same extent.

If we were to reverse the domination ranking of the two constraints, the predicted outcome would be changed: now *.**u**l.* would be superior to *.w**L**.* by virtue of its relative success on HNUC, and the ONS criterion would be submerged. Because of this, the ranking ONS ≫ HNUC is *crucial*; it must obtain in the grammar of Berber if the actual language is to be generated.

The notion of domination shows up from time to time in one form or another in the literature, sometimes informally, sometimes as a clause clarifying how a set of constraints is to be interpreted. For example, Dell and Elmedlaoui write, "The prohibition of hiatus…*overrides*" the nuclear sonority comparison (Dell & Elmedlaoui 1985: 109, emphasis added). For them, this is an extra-grammatical observation, with the real work done by the Structural Descriptions provided by the DEA and the ordering of application of the subrules. Obviously, though, the insight is clearly present. Our claim is that the notion of domination, or 'over-riding', is the truly fundamental one. What deserves extra-grammatical status is the machinery for constructing elaborately specific Structural Descriptions and modes of rule application.

To see how Serial Harmonic Syllabification (9) proceeds, let us examine the first stage of syllabifying the input /txznt/ 'you sg. stored, pf.'. It is evident that the first syllable constructed must be *.z**N**.* — it has an onset, and has the highest sonority nucleus available, so no competing candidate can surpass or even equal it. A more discursive examination of possibilities might be valuable; the larger-scale comparisons are laid out in the constraint tableau below.

Here are (some of the) leading candidates in the first round of the process:

(16) **Constraint Tableau for Serial Syllabification** of /txznt/ (partial, first step)

Candidates	ONS	HNUC	Comments
☞ tx(z**N**)t		n	optimal: onsetted, best available nucleus
txz(**N**)t	* !	n	no onset, HNUC irrelevant
t(x**Z**)nt		z !	|z| < |n|
(t**X**)znt		x !	|x| < |n|
txz(n**T**)		t !	|t| < |n|

Syllabic parsing is conceived here as a step-by-step serial process, just as in the DEA. A candidate set is generated, each produced by a single licit change from the input; the relative status of the candidates is evaluated, yielding an optimal candidate (the output of the first step); and that output will then be subject to a variety of further single changes, generating a new candidate set to be evaluated; and so on, until there are no bettering changes to be made: the final output has then been determined.

This step-by-step Harmony evaluation is not intrinsic to the method of evaluation, though, and, in the more general context, when we discard the restricted definition of Gen in (13), it proves necessary to extend the procedure so that it is capable of evaluating entire parsed strings, and not just single (new) units of analysis. To do this, we apply the same sort of reasoning used to define domination, but *within* the constraint categories. To proceed by example, consider the analysis of /txznt/ taking for candidates all syllabified strings. We present a sampling of the candidate space.

(17) **Parallel Analysis of Complete Syllabification** of /txznt/

Candidates	ONS	HNUC	Comments
☞ .t**X**.z**N**t.		n x	optimal
.**T**x.z**N**t.		n t !	|n| = |n|, |t| < |x|
.t**X**z.n**T**.		x ! t	|x| < |n|, t irrelevant
.tx**Z**.**N**t.	* !	n z	HNUC irrelevant
.**T**.**X**.**Z**.**N**.**T**.	* ! ***	n z x t t	HNUC irrelevant

In evaluating the candidates we have kept to the specific assumptions mentioned above: the onset requirement is suspended phrase-initially, and the nonnuclear status of peripheral obstruents is, as in the DEA itself, put aside (see fn. 5).

In this tableau, all the relevant information for harmonic evaluation of the parse of the whole string is present. We start by examining the first column, corresponding to the dominant constraint ONS. Only the candidates which fare best on this constraint survive for further consideration. The first three candidates all have syllables with onsets; the last two do not (to varying degrees). Lack of onset in even a single non-initial syllable is immediately fatal, because of the competing candidates which satisfy ONS.

The remaining three parses are not distinguished by ONS, and so HNUC, the next constraint down the hierarchy, becomes relevant. These three parses are compared by HNUC as follows. The most sonorous nucleus of each parse is examined: these are the most harmonic nuclei according to HNUC. For each of the first two candidates the most sonorous nucleus is *n*. For the last candidate, the most sonorous nucleus is *x*, and it drops out of the competition since *n* is more sonorous than *x*. We are left with the first two candidates, so far tied on all comparisons. The HNUC evaluation continues now to the next-most-harmonic nuclei, where the competition is finally settled in favor of the first candidate .tX.zNt.

What we have done, in essence, is to replace the iterative procedure (act/evaluate, act/evaluate,…) with a recursive scheme: collect the results of all possible actions, then sort recursively. Rather than producing and pruning a candidate set at each step of sequential processing, striving to select at each step the action which will take us eventually to the correct output, the whole set of possible parses is defined and harmonically evaluated. The correct output is the candidate whose complete structure best satisfies the constraint hierarchy. And 'best satisfies' can be recursively defined by descending the hierarchy, discarding all but the best possibilities according to each constraint before moving on to consider lower-ranked constraints.

The great majority of analyses presented here will use the parallel method of evaluation. A distinctive prediction of the parallel approach is that there can be significant interactions of the top-down variety between aspects of structure that are present in the final parse. In §3, §4 and §7 we will see a number of cases where this is borne out, so that parallelism is demonstrably crucial; further evidence is presented in McCarthy & Prince 1993. 'Harmonic serialism' is worthy of exploration as well, and many hybrid theories can and should be imagined; but we will have little more to say about it. (But see fn. 49 below on Berber syllabification.)

The notion of parallel analysis of complete parses in the discussion of constraint tableau (17) is the crucial technical idea on which many of our arguments will rest. It is a means for determining the relative harmonies of entire candidate parses from a set of conflicting constraints. This technique has some subtleties, and is subject to a number of variant developments, so it is worth setting out with some formal precision exactly what we have in mind. A certain level of complexity arises because there are two dimensions of structure to keep track of. On the one hand, each individual constraint typically applies to several substructures in any complete parse, generating a *set* of evaluations. (ONS, for example, examines every syllable, and there are often several of them to examine.) On the other hand, every grammar has multiple constraints, generating multiple sets of evaluations. Regulating the way these two dimensions of multiplicity interact is a key theoretical commitment.

Our proposal is that evaluation proceeds by constraint. In the case of the mini-grammar of ONS and HNUC, entire syllabifications are first compared via ONS alone, which examines each syllable for an onset; should this fail to decide the matter, the entire syllabifications are compared via HNUC alone, which examines each syllable's nucleus.

Another way to use the two constraints would be to examine each (completely parsed) candidate syllable-by-syllable, assessing each syllable on the basis of the syllabic mini-grammar. The fact that ONS dominates HNUC would then manifest itself in the Harmony assessment of each individual syllable. This is also the approach most closely tied to continuous Harmony evaluation during a step-by-step constructive derivation. Here again, we do not wish to dismiss this conception, which is surely worthy of development. Crucially, however, this is not how Harmony evaluation works in the present conception (see §5.2.3.1 for further discussion).

In order to characterize harmonic comparison of candidate parses with full generality and clarity, we need to specify two things: first, a means of comparing entire candidates on the basis of a single constraint; then, a means of combining the evaluation of these constraints. The result is a general definition of *Harmonic Ordering of Forms*; this is, in its formal essence, our theory of constraint interaction in generative grammar. It is the main topic of §5.

2.3 Summary of discussion to date

The core syllabification of Imdlawn Tashlhiyt Berber provides a particularly clear case where the function assigning structural analyses must be based on the optimality of the output if it is to be properly founded on principle. Once the relevant principles have been moved into grammatical theory, it becomes possible to undertake a radical simplification of the generative procedure that admits candidate syllable structures. The focus shifts away from the effort to construct an algorithm that assembles the correct structure piece-by-piece, an effort that we believe is doomed to severe explanatory shortcomings. Linguistic theory, properly conceived, simply has little to say about such constructional algorithms, which (we claim) are no more than implementations of grammatical results in a particular computational-like framework. The main explanatory burden falls the constraints themselves, and on the apparatus that governs their interactions.

The Berber situation is particularly interesting in that core syllabification simply cannot proceed without the intervention of *two* distinct constraints. As with other forms of prosodic organization, the most common picture is one in which the structure is built (more-or-less) bottom-up, step-by-single-step, with each step falling under the constraints appropriate to it. Taking this seriously in the syllable structure domain, this would mean, following Levin [Blevins] (1985) and ultimately Kahn (1976), that you first locate the nuclei — the heads of syllables; then project higher order structure that includes the onsets; then project the structure that includes postnuclear consonantism. In ITB, however, as in many other languages, the availability of nuclei depends on the choice of onsets: an early step in the derivational constructive procedure, working on a low level in the structural hierarchy, depends on later steps that deal with the higher levels. Indeed, the higher level constraint is very much the more forceful. Technical solutions to this conundrum can be found in individual cases — Dell and Elmedlaoui's being a particularly clever one; but the theme will re-appear persistently in every domain of prosody, defying a uniform treatment in constructionist terms.

In the theory advocated here, where outputs are evaluated, we expect exactly this kind of interaction. The whole output is open to inspection; how we choose to inspect it, or how we are forced by UG to inspect it, is not determined by the course that would be taken in bottom-up construction. The potential force of a constraint is not indexed to the level of structure that it pertains to, and under certain circumstances (UG permitting or demanding), constraint domination will be invoked to give higher-level constraints absolute priority over those relevant to the design of lower structural levels.