Fred Karlsson
Department of General Linguistics
University of Helsinki

A PARADIGM-BASED MORPHOLOGICAL ANALYZER

1. Introduction

Computational morphology has advanced by leaps in the past few years. Since the pioneering work of Kay (e.g. Kay 1977), major contributions have been submitted especially by Karttunen (Karttunen & al. 1981) and Koskenniemi (1983). A common linguistic trait of this line of work has been a fairly strict adherence to the basic principles of generative phonology and morphology (especially of the IP type). The theories and models proposed have been decisively based on the notion of rules relating different levels of representation. Typically, the rules describe morphophonological alternations by which surface-level word-forms deviate from postulated lexical or underlying forms. Central concepts have also been the representation of lexicons as tree structures, minilexicons for describing morphotactic structure in terms of pointers to subsequent classes of allowed morphological categories (e.g. Karttunen & al. 1981), and the implementation of IP rules as finite-state transducers. A major achievement was Koskenniemi's (1983) truly bidirectional language-independent formalism for word-form production and recognition. Notions such as intraparadigmatic dependencies between subsets of endings and/or stems, as well as productivity and the mechanisms of lexical extension, have so far played only a minor role (however, cf. Ejerhed and Church's paper in the present volume).

This paper discusses a morphological analyzer called PARMORF that was designed for simulating not IP rules but **paradigmatic relationships.** One of the most notable recent trends in morphological theory has been the natural morphology advocated especially by Dressler, Mayerthaler, and Wurzel (e.g. Dressler 1985, Wurzel

A paradigm-based morphological analyzer
Fred Karlsson
Proceedings of NODALIDA 1985, pages 95-112

1984). One of its key concepts is the notion of **paradigmatic dependency** that has been elaborated especially by Wurzel (also cf. Bybee 1985). This body of work has provided important impetus for the present effort. In particular, it is my intention to explore how feasible a paradigm view of morphology is in building computational models of word-form recognition. Another point of interest is how easily such a model can be designed to incorporate morphological productivity and lexical extension.

An important feature of PARMORF is that it renounces the use of morphophonemic symbols on the lexical level, and also does away with the corresponding phonological rules. Diacritics are used only for the purpose of singling out members of truly non-productive and closed inflectional types. Whatever morphophonological alternations there are will be expressed by stating intra-paradigmatic dependencies between stems and ending classes.

The central property of PARMORF is that the lexicon tree operational in word-form analysis is based on **stems** that are derived by **paradigmatic pattern rules** from base forms which may be either entries in the main lexicon or new words that are about to be integrated in the lexicon. The base forms of the lexical entries as such are not directly involved in word-form recognition. The PARMORF main lexicon for Finnish thus contains i.a. the noun lexeme kauppa 'shop' (N.B. in straightforward phonological shape without morphophonemes). For this lexeme, general pattern rules determine four stems with their appropriate morphotactic information (here omitted), viz. kauppa, kaupa, kauppo, and kaupo. These stems are inserted in the tree used for word-form recognition.

It is my hypothesis that once the inflectional behavior of a word is known, recognition of individual instances of it takes place in relation to the concrete stems in the lexicon tree. No (analogues of) IP/IA rules are invoked in the actual process of word-form recognition.

PARMORF embodies the hypothesis that morphological processing in the sense of "applying rules" consists primarily in determining how words so far unknown to the language user are inflected. For any word, this piece of knowledge should be supplied by a working theory of morphological productivity (here formalized as pattern rules). Supposing that all words belonging to unproductive and closed inflectional subclasses are marked in the lexicon, the pattern rules will derive appropriate stem sets for them, and productive default stem sets for all unmarked words (whether in the lexicon or not).

This approach to morphological productivity makes the process of lexical extension fall out from entities already in the grammar. Since word-forms are recognized just by scanning concrete stems and concrete endings, PARMORF should lend itself to psycholinguistic interpretation more directly than models invoking generative rules. These models face the problem of determining how, precisely, phonological rules and their implementation as finite-state automata should be related to real behavior.

2. Lexical representations

There are at least eight ways in which the lexical forms of words may be construed:

(1) Minimal listing of the SPE-type where even distantly related word-forms are derived from a shared lexical source whose composition is claimed to be (systematic-)phonological. This underlying form lexically represents all word-forms (the whole inflectional paradigm). A central goal is to minimize the number of lexemes and to maximize the statement of morphophonolocial alternations as IP rules. Word-forms are indirectly related to the lexical representations (i.e. derived by rules).

(2) Constrained minimal listing where remotely related (especially morphophonemically irregular) word-forms are not derived from a common source. The number of lexemes postulated is therefore somewhat larger than under (1). The vast majority of words is represented by a unique lexical form as in (1). However, these base-forms as well as the rules are subject to more restricted (naturalness) conditions than are SPE-type rules. This is a modified SPE-position advocated by several variants of natural generative phonology (e.g. Hooper 1976).

(3) Unique lexical forms allowing diacritics and morphopho-

nemes. This position is embodied in most two-level implementations based on Koskenniemi's (1983) model. A lexical form may contain several morphophonemic and diacritic (e.g. juncture) symbols. Otherwise, it resembles (1,2), especially in the use of phonological rules (to be compiled as finite-state automata). Paradigms are represented by a single base-form, as in (1,2).

(4) Stems. This solution is advocated here. I regard all phonologically distinct bound variants of a base-form as separate stems. A stem-based lexicon is bound to be somewhat larger than a lexicon containing unique base forms for most words. One of the present purposes is to explore whether the amount of repetition will be prohibitively large so as to render this approach unfeasible. It deserves to be stressed that common initial substrings, meanings, category information, syntactic features, etc., in a set of stems manifesting one lexeme will not be repeated but shared in the stem tree. We are thus not heading for a theory involving whole-sale listing. - No comprehensive stem-based theory of morphology has so far been advanced, apart from the "technical stem" stem approach (5), and some general mention of (full) stems as a theoretical possibility for lexical representation (e.g. Linell 1979).

(5) Technical stems. This concept refers to the minimal invariant phonological substance occurring in all (full) stems, e.g. kaup in Fi. kauppa. Such technical stems have been used by Hellberg (1978) in his description of Swedish morphology, and by Karttunen & al. (1981) in their Finnish morphology (TEXFIN). In this approach, stem alternations are described e.g. by postulating minilexicons pointed to by the relevant technical stems.

(6) Full listing hypothesis (FLH). FLH claims that all word-forms are listed in the lexicon. This view is widely entertained in psycholinguistic research on word-form recognition (cf. Butterworth 1983). We shall discard this possibility since it leads to implausible consequences for highly inflected languages such as Finnish. Given that a Finnish verb has some 15,000 forms and an English verb less than five, FLH entails that learning Finnish verbal morphology would be thousands of times more cumbersome than learning English, and that a Finn would need much more neural space to internalize his verbs than would an Englishman.

Furthermore, according to FLH, upon learning a new verb a Finn should have to internally generate all the 15,000 forms - most of which he would never use. All this seems implausible. In face of these remarks, FLH without precisions and amendments is not acceptable as a general (psycholinguistic) theory of lexical organization. Stems provide a more uniform crosslinguistic characterization of the lexicon. E.g. English and Finnish don't differ decisively in regard to how many stems a word may have. Finnish verbs and nouns have maximally five or six stems.

(7) Semantically feasible word-forms. This would be a more realistic reduced version of FLH (to my knowledge, not yet elaborated). It would claim that the lexicon contains word-forms, but only those that are semantically feasible. Thus, the English lexicon would not (normally) contain e.g. plural forms for proper names or mass words, or personal forms for meteorological verbs.

(8) Prototypical word-forms. Given that most words, due to obvious semantic reasons, favour certain forms (e.g Fi. local nouns favour the local cases, mass nouns the partitive case, countables the nominative), it is more reasonable to suppose that the core lexicon of a language user contains the very word-forms that he/she has learnt, especially those that are in frequent active use, i.e. the prototypical ones (cf. Karlsson 1985).

All of (1-8) are not mutually exclusive. Any "realistic" model (i.e. striving not only for system description but also for isomorphy with psycholinguistic facts) must be able to account at least for frequency effects which often manifest themselves on the level of individual word-forms (cf. Garnham 1985:45 for an overview). This would presuppose special treatment (e.g. separate listing) of the most frequent and deeply engraved word-forms, regardless of whether the bulk of the lexemes are represented according to one of the alternatiuns (1-5). However, in what follows we shall only consider the feasibility of (4).

In approaches (1-3), the basic set-up of word-form processing is this:

LEXICON(S) (often compiled as tree structures)
RULES (often implemented as finite-state transducers)
SURFACE WORD-FORMS

In computational models, the (main and ending) lexicons are normally implemented as trees. These trees are direct operational analogues of the respective lexicons and are therefore the only processually relevant lexical structure. The lexicon list is an epiphenomenon helpful in inspecting the existing stock of words.

The present approach is sligthly different. I postulate a main lexicon (list) containing the stock of lexemes. Here, each lexeme is represented as a quintuple:

<base-form nextLexicon meaning syntFeatures cat>

Each lexeme has a unique base-form consisting of phonemes only. No morphological markings are needed when all stems of a base-form are predictable by general pattern rules. E.g., all Swedish nouns ending in -*el*, -*en*, -*er* lose their -*e*- in certain morphological environments and therefore no individual base-forms need diacritics. However, predicting the morphophonological behavior of the Finnish inflectional types *vesi* (nom.) : *vede+n* (gen.) and *lasi* (nom.) : *lasi+n* (gen.) presupposes that the members of the former closed, unproductive, complex class are marked (say, *vesi>*). Pattern rules tell what special stems -*si>* -nouns have. Unmarked -*si* -nouns constitute the unmarked default pattern.

The Finnish main lexicon thus contains nominal and verbal entries such as the following ones. nextLex will be specified for each stem by the pattern rules, the meaning is here just represented by a translation into English, and the syntactic features occur in bare outline.

```
(talo NIL house (Countable ...) N)
(vesi> NIL water (Mass ...) N)
(hullu NIL mad NIL A)
(suuri> NIL big NIL A)
(raskas> NIL heavy NIL A)
(kannas NIL isthmus (Concrete ...) N)
(anta NIL give (Trans AllRection) V)
(asu NIL live (Intrans IneRection) V)
...
```

Given the information supplied by each lexical entry, pattern rules compile the **stem lexicon tree** active in word-form recognition. The stem lexicon is crucially different from the main lexicon list since it contains full sets of stems. The stems of each lexeme share initial substrings, meaning, syntactic features, and part of speech, i.e. all lexical information apart from alternating stem segments is given just once. The core of PARMORF is thus:

PATTERN RULES (predicting stems)
STEM TREE
WORD-FORMS

3. Pattern rules

Pattern rules embody the predictive power of morphology. They are in active use only when **a new word is added to the stem tree.** Given appropriate information, the stems of a base-form are predicted and inserted in the stem tree. Once integrated, PARMORF presupposes no more (IP or IA type) processing for recognizing forms of a word. In many respects, this model is equally applicable to children's acquisition of morphology and to an adult's adding words to his/her lexicon. Note that this model embodies the core of FLH without endless listing of concrete word-forms, but also without rule processing.

The pattern rules also explicate one aspect of paradigm constitution. They determine what stems belong together and also what morphophonological alternations belong together. Such clusterings are at the heart of traditional paradigms.

Pattern rules are IF-THEN -rules obeying the following format where parentheses indicate elements not necessarily used in all pattern rules:

IF	THEN
base form coda	stem-coda$_1$ + nextLex$_1$
part of speech	(stem-coda$_2$ + nextLex$_2$)
(number of syllables)	(stem-coda$_n$ + nextLex$_n$)
(morphosyntactic feature(s))	

The core of the IF-part is the **base-form coda** (closely related to Bybee and Slobin's (1982) notion "schema"), i.e. the shortest segment string extracted from the **end** of the base form that suffices for predicting the stems. The coda is expressed as a sequence of phonemes (plus a diacritic, where needed). The part of speech is also needed by the IF-part. Syllable number is often required, as might be specific morphosyntactic features (e.g. Swedish stem prediction often needs gender). Apart from the number of syllables (which is determined by a separate algorithm) the IF-part information is given in the main lexicon entries.

For new words, this information must be made available by context of use. Evidently, inflectional behavior cannot be predicted without knowledge of part of speech, etc.

The THEN-part provides a set of pairs (at least one) each consisting of a **stem-coda** and a reference to the appropriate ending lexicon.

The inserted full stems are formed by appending the residue of the base form (i.e. what is to the left of the base-form coda) to each stem-coda. Typical Finnish pattern rules look as follows (by convention, names of ending trees are prefixed by a slash):

IF	THEN
kko, N, 2	kko /huppu
	ko /hupu

IF	THEN
ppa, N, 2	ppa /nom/sg/str
	pa /nom/sg/w
	ppo /j/pl
	po /nom/pl/w

IF	THEN
ppa, N, 3	ppa /nom/sg/str
	pa /nom/sg/w
	ppo /j/pl
	po /ammate

IF	THEN
ppää, V	ppää /loukkaa
	ppä /loukka
	pä /louka

A disyllabic gradable noun ending in -kko' thus has two stems and the appropriate ending trees are /huppu and /hupu, respectively. A disyllabic gradable noun in -ppa has four stems. A trisyllabic noun in -ppa has the same four stems but a difference in what endings are allowed in weak grade plurals (ulapoita vs. *kaupoita). A verb ending in -ppää has three stems, etc.

Pattern rules are normally differentiated at least for nouns and verbs, often also for nouns and adjectives (not so in Finnish). All base-form codas generated by the pattern rules for a certain part of speech are inserted into a **pattern tree.** There is one pattern tree for each distinct part of speech. The segments of each base-form coda are inserted in reverted order, prefixed by an integer indicating the number of syllables where needed. Thus, the strings inserted in the nominal pattern tree for the first three pattern rules just mentioned are okk, app, 3app.

THEN-parts are entries under the last node of each identifiable coda in this tree. Once this base-form pattern tree exists for a given part of speech, the stem set for any such base-form is found by **picking the longest match in the pattern tree** for the search key consisting of the base-form segments in reverted order.

Thus, when the stem set for the noun ulappa is to be determined, a match for the string 3appalu is sought in the pattern tree (the integer "3" having been prefixed by the syllable counting algorithm). The longest match found will be 3app and the corresponding entry is retrieved. The four stems thus determined are inserted in the stem tree and then used in the recognition process. For nongradable trisyllabic nouns in -a the longest match found will be 3a providing only two stems (-a, -o).

The base-form codas of the pattern rules are expressed as strings of phonemes and eventual inflectional diacritics. This leads to repetition especially for words subject to consonant gradation and mutation of the final vowel (both exemplified by ulappa). E.g. up to 15 individual instances of consonant gradation will be separately stated for the paradigms where they actually occur. There are thus some 15 pattern rules for disyllabic nouns ending in -o, viz. -kko, -ppo, -nto, etc.

Deviating from generative practice, I have deliberately chosen not to generalize consonant gradation, vowel mutation, and similar morphophonological alternations across paradigms. At first sight, this seems to lead to prohibitively unilluminating repetition. However, there are positive linguistic arguments in favour of this solution. **Particular paradigms might contain morphophonological gaps** that should somehow be accounted for. Thus, trisyllabic Fi. nouns allow only a few gradable stops at the final syllable boundary: pp, tt, kk, nt, nk. Nouns of the kaikki> -type disallow i.a. the gradable combinations lt, nt, rt at the syllable boundary. Paradigms like *karampa : karamman, *kanti : kannen are not just accidentally lacking but morphophonologically ungrammatical. That is, individual pattern rules explicitly state the allowed possibilities up to systematic gaps but exclude the latter, thereby accounting for systematic restrictions.

The principle of longest match used in searching the pattern tree gives a convenient and uniform way of handling exceptions. If the inverted form of a whole word is found in the pattern tree, it will by definition be the longest match. Thus, exception features for individual words are generally not needed.

The total number of pattern rules with the above concrete properties invoked in my full description of Finnish nominal and verbal morphology is some 1,130 (600 for nouns, 530 for verbs; some 250 exceptional pronominal forms are not included in the first figure). This number includes all idiosyncracies (roughly half of these rules concern one item only). Considering that the power of the pattern rule system is such as to predict the inflection of all nouns, adjectives, and verbs in the lexicon, including all exceptions, **and** the default inflection of any such word not in the lexicon, and furthermore excluding many types of impossible paradigms, we would not regard the number as "prohibitively large", especially when one takes into account that no further morphophonological rules or processing is invoked in word-form recognition. I.e., full productive mastery of Finnish morpho(no)logy presupposes learning some 1,100 concrete phoneme-level rules.

4. Ending lexicons

Similarly behaving endings are grouped into ending lexicons which are triples with the following structure:

<name, otherLex, endings>

Each ending lexicon has a name (conventionally prefixed by a slash) normally chosen so as to give a mnemonic hint of what kinds of stems or words it is normally appended to. The component "otherLex" provides a (possibly null) list of **other ending lexicons paradigmatically included in the present one.** This facility provides a convenient opportunity of stating paradigmatic relationships between distributionally related subsets of endings. Finally, the compartment "endings" is a (possibly empty) set of endings belonging to the current ending lexicon (i.e. possibly empty because an ending lexicon may consist exclusive of references to other ending lexicons under otherLex). Each ending, in turn, is a triple:

<item, nextLex, entry>

where "item" is the ending in phonemic shape, "nextLex" a reference to the next morphotactic position, and "entry" contains a list of morphological categories. Vowel harmony is an exception to the phonemic principle of item structure, i.e. suffix vowel harmony pairs are lexically represented as the archiphonemes A, O, U, which are spelled out as a-ä, o-ö, u-y when the ending lexicons are compiled into trees used in actual processing.

The endings and entries are often listed as wholes, especially in close-knit combinations of e.g. number and case for nouns. Such combinations are often subject to bidirectional dependencies that are hard to capture otherwise. The /j/pl lexicon below contains good examples of this dependence. The plural allomorph j occurs only if the following ptv. or gen. case morph starts with a vowel, and the latter occur only if pl. j precedes. Furthermore, for gradable nouns the -jA, -jen -combinations are tied to strong-grade stems only (koivikkojen vs. *koivikojen). This complex paradigmatic interdependence between a certain stem, a certain number morph, and a certain case morph has proven laborious to capture by (morpho)phonological rules. Under the present approach, it suffices to point from one stem to one lexicon.

A psycholinguistic argument for treating (some) ending sequences as wholes comes from the observation that children acquiring inflectional languages seldom make errors involving the order of morphemes in a word (cf. Bybee 1985:114ff. for an overview).

The following are typical examples of ending lexicons. The name is given on the first line, otherLex on the second, and the endings, if any, are indented.

```
(/nom/sg/str
(/clit/nom /ill/Vn /poss3)
     (A /poss4 (PTV SG))
     (nA /poss4 (ESS SG)))

(/nom/sg/w
NIL
     (n /clit (GEN SG))
     (llA /poss4 (ADE SG))
     (ltA /poss4 (ABL SG))
     (lle /poss5 (ALL SG))
     (ssA /poss4 (INE SG))
     (stA /poss4 (ELA SG))
```

-106-

```
     (ksi /clit (TRA SG))
     (kse /poss6 (TRA SG))
     (ttA /poss4 (ABE SG))
     (t /clit (NOM PL)))

(/nom/pl/w
NIL
     (illA /poss4 (ADE PL))
     (iltA /poss4 (ABL PL))
     (ille /poss5 (ALL PL))
     (issA /poss4 (INE PL))
     (istA /poss4 (ELA PL))
     (iksi /clit (TRA PL))
     (ikse /poss6 (TRA PL))
     (ittA /poss4 (ABE PL))
     (in /clit (INS PL))
     (i /poss2 (INS PL)))

(/i/pl
NIL
     (iA /poss4 (PTV PL))
     (ien /clit (GEN PL))
     (ie /poss2 (GEN PL))
     (iin /clit (ILL PL))
     (ii /poss2 (ILL PL))
     (inA /poss4 (ESS PL))
     (ine /poss6 (COM SG/PL)))

(/j/pl
NIL
     (jA /poss4 (PTV PL))
     (jen /clit (GEN PL))
     (je /poss2 (GEN PL))
     (ihin /clit (ILL PL))
     (ihi /poss2 (ILL PL))
     (inA /poss4 (ESS PL))
     (ine /poss6 (COM SG/PL)))
```

```
(/huppu
(/nom/sg/str /j/pl))

(/hupu
(/nom/sg/w /nom/pl/w))

(/nom/2s/all
(/huppu /hupu))

(/puolisko
(/nom/2s/all /itA/iden))

(/itA/iden
NIL
     (itA /poss4 (PTV PL))
     (iden /clit (GEN PL))
     (itten /clit (GEN PL))
     (ide /poss2 (GEN PL))
     (itte /poss2 (GEN PL)))
```

Endings in the same ending lexicon behave alike. An ending lexicon constitutes a kind of "paradigmatic natural class". Thus, /nom/sg/str contains endings occurring after strong-grade sg. stems of (certain) gradable nouns. These endings are ptv. -A and ess. -nA, plus certain clitics, possessives, and illatives included via the specifications in otherLex. /nom/sg/w contains the corresponding weak-grade sg. endings, /nom/pl/w the weak-grade pl. endings.

The paradigm formalism enables us to capture complex **intersecting paradigmatic networks** by way of otherLex references. Thus, the lexicon /huppu (covering strong-grade sg. and pl. stems like huppu, lakko) contains the members of /nom/sg/str and /j/pl but no endings of its own. /hupu (covering the corresponding sg. and pl. weak-grade stems) contains the members of /nom/sg/w, /nom/pl/w. Then one may continue: /nom/2s/all covers the corresponding non-gradable stems (words like talo, hullu) and is described by referring via otherLex to /huppu, /hupu. Yet another layer may be added by describing trisyllabic non-gradable nouns (e.g.puolisko) as /puolisko consisting of /nom/2s/all and /itA/iden. This captures the generalization that these nouns depart from the disyllabic ones only in having some more alternative plural endings.

In other words, references via otherLex are recursively broken down by tracing all the lexicons invoked. The hierarchical paradigmatic lexicon network may be displayed as follows:

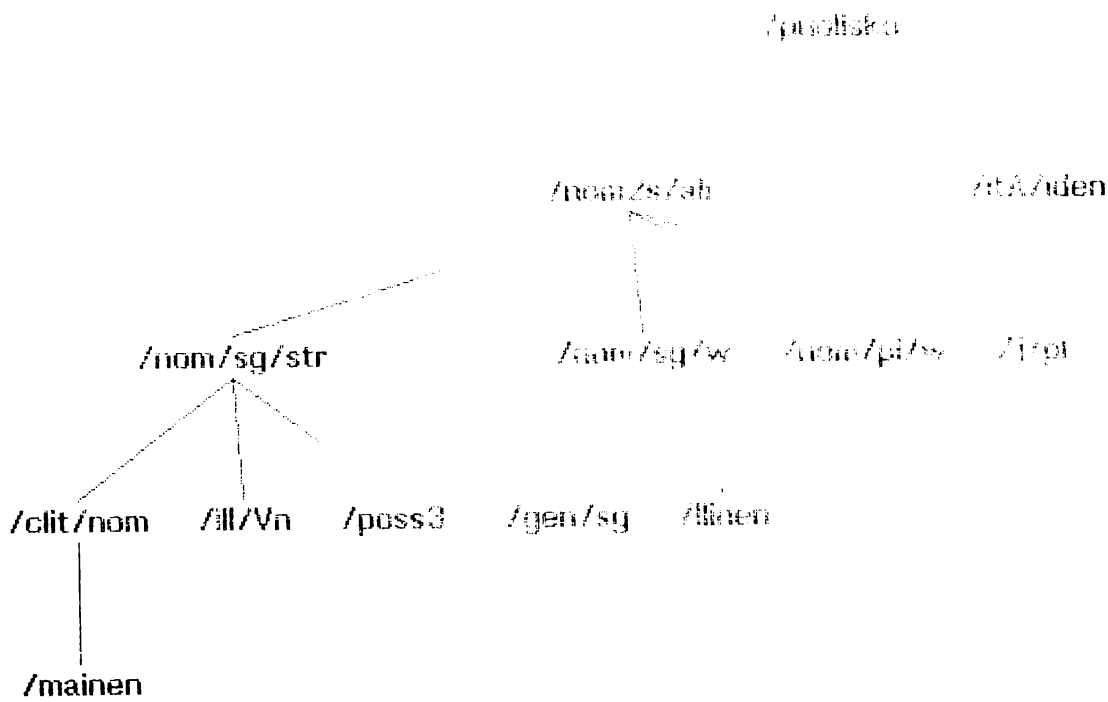

The full description of Finnish contains 134 ending lexicons. At run-time, two options are available for compiling ending lexicons to trees. In the minimal version, each ending tree contains only the endings listed in the respective lexicon, and when a word-form is to be analyzed, eventual otherLex references are all checked separately and recursively by jumping from tree to tree. E.g. when the /puolisko tree is consulted, all 13 trees in the display above are run through. The 134 minimal ending trees require some 1,000 nodes. The maximal option lumps together into one tree the endings of the current lexicon plus all endings found by recursively checking the otherLex references (e.g. all 13 trees under /puolisko). In this mode, the lexicon trees require some 8,000 nodes. Of course, using maximal ending trees speeds up the recognition process (roughly by a factor of three).

This kind of paradigmatic description does capture significant generalizations. It also makes interesting predictions, e.g. that paradigm levelling or extension is likely to concern **all** the members of a given ending lexicons (in due course).

5. Implementation and evaluation

The formalism for expressing pattern rules, stems, and ending lexicons is language-independent. The pattern rules must, of course, be determined by the linguist before they can be read by the program, i.e. before the pattern tree is constructed. The program reads lexical entries of the specified type upon constructing the stem tree.

So far, I have only tested the model on Finnish. The current size of the Finnish main lexicon is roughly 9,000 items (of which 4,300 are nouns and 2,000 verbs). On the average, a Fi. noun has 2,5 stems and a verb 3,2 stems (in the sense of phonologically distinct from the base-form). When all stems of these 9,000 items are compiled into the stem tree, its size is roughly 41,000 nodes. A rough comparison to Koskenniemi's (1983; personal communication) Fi. lexicon shows that a full stem-approach **less than doubles the number of nodes in the main lexicon tree.** I find this rough ratio interesting as it proves that a stem-based lexicon is **not** prohibitively much larger than a lexicon based on unique lexical forms. For IE languages stem-based lexicons would be even more manageable than in Finnish.

The "cost" of the stem-based approach is thus a doubling of the number of main lexicon nodes. This is counterbalanced by an elimination of morphophonological processing at run-time which of course streamlines and speeds up the actual process of word-form recognition. Using maximal ending trees, word-form recognition over the 9,000-item stem tree takes 30 ms on the average (including multiple analyses of homonyms). Short unambiguous words are analyzed in 10-15 ms.

The program provides for productive morphological analysis of any compound just by turning a switch. In normal mode, all analyses are produced. Another switch constrains the analyzer to producing one analysis only. The given efficiency figures pertain to this non-compound mode.

Acknowledgement

I am indebted to Kimmo Koskenniemi and Martti Nyman for insightful comments.

References

Butterworth, Brian 1983. "Lexical Representation". In B. Butterworth, ed., Language Production, Vol. 2, New York: Academic Press, 257-294.

Bybee, J.L. 1985. Morphology. A Study of the Relation between Meaning and Form. Amsterdam: Benjamins.

Bybee, J.L. and Slobin, D.I. 1982. "Rules and schemas in the development and use of the English past tense". Language 58, 265-289.

Dressler, Wolfgang U. 1985. Morphonology. Ann Arbor: Karoma Publishers.

Garnham, Alan 1985. Psycholinguistics. Central Topics. Methuen: London and New York.

Hellberg, Staffan 1978. The Morphology of Present-Day Swedish. Stockholm: Almqvist & Wiksell International.

Hooper, J.B. 1976. An Introduction to Natural Generative Phonology. New York: Academic Press.

Karlsson, Fred 1985. "Paradigms and Word-Forms". Studia gramatyczne 7, 135-154.

Karlsson, Fred & Koskenniemi, Kimmo 1985. "A Process Model of Morphology and Lexicon". Folia Linguistica 19:1/2, 207-229.

Karttunen, Lauri, Root, Rebecca, & Uszkoreit, Hans 1981. "TEXFIN: Morphological Analysis of Finnish by Computer". Paper read at the 71st Annual Meeting of the SASS, Albuquerque, New Mexico.

Kay, Martin 1977. "Morphological and Syntactic Analysis". In A. Zampolli, ed., Linguistic Structures Processing, Amsterdam: North-Holland, 131-234.

Koskenniemi, Kimmo 1983. Two-level Morphology: A Theory for Automatic Word-Form Recognition and Production. University of Helsinki, Department of General Linguistics, Publications No. 11.

Linell, Per 1979. Psychological Reality in Phonology. Cambridge University Press.

Mayerthaler, Willi 1981. Morphologische Natürlichkeit. Wiesbaden: Athenaion.

Wurzel, W.U. 1984. Flexionsmorphologie und Natürlichkeit. Berlin: Akademie-Verlag.