

# Reference-Set Constraints as Linear Tree Transductions via Controlled Optimality Systems

Thomas Graf

Department of Linguistics  
University of California, Los Angeles  
[tgraf@ucla.edu](mailto:tgraf@ucla.edu)

**Abstract.** Reference-set constraints are a special class of constraints used in Minimalist syntax. They extend the notion of well-formedness beyond the level of single trees: When presented with some phrase structure tree, they compute its set of competing output candidates and determine the optimal output(s) according to some economy metric. Doubts have frequently been raised in the literature whether such constraints are computationally tractable [6]. I define a subclass of Optimality Systems (OSs) that is sufficiently powerful to accommodate a wide range of reference-set constraints and show that these OSs are globally optimal [7], a prerequisite for them being computable by linear tree transducers. As regular and linear context-free tree languages are closed under linear tree transductions, this marks an important step towards showing that the expressivity of various syntactic formalisms is not increased by adding reference-set constraints. In the second half of the paper, I demonstrate the feasibility of the OS-approach by exhibiting an efficiently computable OS for a prominent reference-set constraint, Focus Economy [13].

**Key words:** Reference-Set Constraints, Transderivationality, Optimality Systems, Tree Transducers, Modelling, Finite-State Methods

## Introduction

Out of all the items in a syntactician's toolbox, reference-set constraints are probably the most peculiar one. When handed some syntactic tree, a reference-set constraint does not determine its well-formedness from inspection of the tree itself. Instead, it constructs a *reference set* — a set containing a number of trees competing against each other — and chooses the optimal candidate from said set.

Consider *Fewest Steps* [1]. The reference set that this constraint constructs for any given tree  $t$  consists of  $t$  itself and all the trees that were assembled from the same lexical items as  $t$ . All the trees in the reference set are then ranked by the number of movement steps that occurred during their assembly (this is usually identical to the number of traces they contain), and the tree(s) with the fewest instances of movement is (are) chosen as the winner. All other trees are flagged as ungrammatical, including  $t$  if it did not emerge as a winner.

Another reference-set constraint is *Focus Economy* [13, 18], which accounts for the empirical fact that neutral stress is compatible with more discourse situations than

shifted stress. Take a look at the utterances in (1), where main stress is indicated by **bold face**. Example (1a) can serve as an answer to various questions, among others “What’s going on?” and “What did your neighbor buy?”. Yet the virtually identical (1b), in which the main stress falls on the subject rather than the object, is compatible only with the question “Who bought a book?”. These contrasts indicate a difference as to which constituents may be *focused*, i.e. can be interpreted as providing new information.

- (1) a. My neighbor bought a **book**.  
b. My **neighbor** bought a book.

Focus Economy derives the relevant contrast by stipulating that first, any constituent containing the node carrying the sentential main stress can be focused, and second, in a tree in which stress was shifted from the neutral position, a constituent may be focused only if it cannot be focused in the original tree with unshifted stress. In (1a), the object, the VP and the entire sentence can be focused, since these are the constituents containing the main stress carrier. In (1b), the main stress is contained by the subject and the entire sentence, however, only the former may be focused because focusing of the latter is already a licit option in the neutral stress counterpart (1a).

This esoteric behavior of reference-set constraints coupled with a distinct lack of formal work on their properties has led to various conjectures that they are computationally intractable [6]. In this paper, I refute these claims by showing how reference-set constraints can be emulated by a new variant of Optimality Systems (OSs), and I contend that this route paves the way for reference-set constraints to be implemented as finite-state devices; linear bottom-up tree transducers (lbutts), to be precise. Lbutts are of interest for theoretical as well as practical purposes because both regular and linear context-free tree languages are known to be closed under linear transductions, so applying a linear transducer to a regular/linear context-free tree language yields a regular/linear context-free tree language again. On a theoretical level, this provides us with new insights into the nature of reference-set constraints, while on a practical level, it ensures that adding reference-set constraints to a grammar does not jeopardize its computability. I support my claim by exhibiting a formal model of Focus Economy as an lbuttt. My results shed new light on reference-set computation as well as on Optimality Systems and should be of interest to readers from various formal backgrounds, foremost computational phonology and Minimalist grammars; moreover, when viewed as tree transducers, reference-set constraints also exhibit some previously overlooked connections to synchronous TAG [15, 16], the exploration of which I have to leave for future research, unfortunately.

The paper is layed out as follows: After the preliminaries section, which due to space restrictions has to be shorter than is befitting, I give a brief introduction to OSs before introducing my own variant, controlled OSs, in Sec. 3. The mathematical core results of this section are a new characterization of the important property of global optimality and a simplification of Jäger’s theorem [7] regarding the properties of an OS that jointly ensure that it does not exceed the power of linear tree transducers. In the last section, I show how to model Focus Economy as such a restricted OS.

## 1 Preliminaries and Notation

Let me introduce some notational conventions first. Given a relation  $R$ , its *domain* is denoted by  $\text{dom}(R)$ , its *range* by  $\text{ran}(R)$ . For any  $a \in \text{dom}(R)$ , we let  $aR := \{b \mid \langle a, b \rangle \in R\}$ , unless  $R$  is a function, in which case  $aR = R(a)$ . The *composition* of two relations  $R$  and  $S$  is  $R \circ S := \{\langle a, c \rangle \mid \langle a, b \rangle \in R, \langle b, c \rangle \in S\}$ . The *diagonal* of some set  $A$  is  $\text{id}(A) := \{\langle a, a \rangle \mid a \in A\}$ .

Tree languages and tree transductions form an integral part of this paper, however, the technical machinery is mostly hidden behind the optimality-theoretic front-end so that only a cursory familiarity with the subject matter is required. Nevertheless the reader is advised to consult [4] and [9] for further details. I also assume that the reader is knowledgeable about string languages and generalized sequential machines [5].

**Definition 1.** A context-free tree grammar (CFTG) is a 4-tuple  $\mathcal{G} := \langle \Sigma, F, S, \Delta \rangle$ , where  $\Sigma$  and  $F$  are disjoint, finite, ranked alphabets of terminals and non-terminals, respectively,  $S \in F$  is the start symbol, and  $\Delta$  is a finite set of productions of the form  $F(x_1, \dots, x_n) \rightarrow t$ , where  $F$  is of rank  $n$ , and  $t$  is a tree with the node labels drawn from  $\Sigma \cup F \cup \{x_1, \dots, x_n\}$ .

A production is linear if each variable in its left-hand side occurs at most once in its right-hand side. A CFTG is *linear* if each production is linear. A CFTG is a *regular* tree grammar (RTG) if all non-terminals are of rank 0. A tree language is *regular* iff it is generated by an RTG, and every regular tree language has a context-free language as its string yield.

**Definition 2.** A bottom-up tree transducer is a 5-tuple  $\mathcal{A} := \langle \Sigma, \Omega, Q, Q', \Delta \rangle$ , where  $\Sigma$  and  $\Omega$  are finite ranked alphabets,  $Q$  is a finite set of states,  $Q' \subseteq Q$  the set of final states, and  $\Delta$  is a set of productions of the form  $f(q_1(x_1), \dots, q_n(x_n)) \rightarrow q(t(x_1, \dots, x_n))$ , where  $f \in \Sigma$  is of rank  $n$ ,  $q_1, \dots, q_n, q \in Q$ , and  $t(x_1, \dots, x_n)$  is a tree with the node labels drawn from  $\Omega \cup \{x_1, \dots, x_n\}$ .

**Definition 3.** A top-down tree transducer is 5-tuple  $\mathcal{A} := \langle \Sigma, \Omega, Q, Q', \Delta \rangle$ , where  $\Sigma$ ,  $\Omega$  and  $Q$  are as before,  $Q' \subseteq Q$  is the set of initial states, and all productions in  $\Delta$  are of the form  $q(f(x_1, \dots, x_n)) \rightarrow t$ , where  $f \in \Sigma$  is of rank  $n$ ,  $q \in Q$  and  $t$  is a tree with the node labels drawn from  $\Omega \cup \{q(x) \mid q \in Q, x \in \{x_1, \dots, x_n\}\}$ .

As with CFTGs, a production is linear if each variable in its left-hand side occurs at most once in its right-hand side. A transducer is *linear* if each production is linear. I denote a linear bottom-up/top-down tree transducer by lbutt/ltddt. The class of ltddts is properly contained in the class of lbutts, which in turn is closed under union and composition. The domain and the range of an lbutt are both recognizable, i.e. regular tree languages. The relation  $\tau$  induced by a (linear) tree transducer is called a (linear) *tree transduction*. For a bottom-up tree transducer, the graph of  $\tau$  consists of pairs  $\langle s, t \rangle$  such that  $s$  and  $t$  are  $\Sigma$ - and  $\Omega$ -labeled trees, respectively, and for some  $q \in Q'$ ,  $q(t)$  can be obtained from  $s$  by finitely many applications of productions  $\delta \in \Delta$ . The definition is almost unchanged for top-down tree transducers, except that we require that  $t$  can be obtained from  $q(s)$ . We call a relation *rational* iff it is a finite-state string transduction or a linear tree transduction. For any recognizable tree language  $L$ ,  $\text{id}(A)$  is a rational

relation. Furthermore, both regular string/tree languages and linear context-free tree languages are closed under rational relations.

At the end of Sec. 4.2, I make good use of  $\mathcal{L}_{K,P}^2$  [14], an incarnation of monadic second-order logic (MSO) specifically designed for linguistic purposes. MSO is the extension of first-order logic with monadic second-order variables and predicates as well as quantification over them such that the first-order variables represent nodes in the tree and the second-order variables and predicates sets of nodes. A set of finite strings/trees is definable in MSO iff it is regular. Specifics of  $\mathcal{L}_{K,P}^2$  will be briefly introduced in the relevant section. See [19] for further background material on MSO and [14] for  $\mathcal{L}_{K,P}^2$  in particular.

## 2 Traditional Perspective on Optimality Systems

OSs were introduced independently by [3] and [8] as a formalization of Optimality Theory (OT) [12], which has been the dominant theory of phonology in mainstream linguistics for the last fifteen years. In OT, well-formed expressions are no longer derived from underlying representations through iterated applications of string rewrite rules, as was the case with SPE [2]. Instead, underlying representations — which are usually referred to as *inputs* — are assigned a set of *output candidates* by a relation called *generator*, abbreviated GEN. This set is subsequently narrowed down by a sequence of constraints  $c_1, \dots, c_n$  until only the *optimal* output candidates remain. This narrowing-down process proceeds in a fashion such that only the candidates that incurred the least number of violations of constraint  $c_i$  are taken into account for the evaluation of  $c_{i+1}$ . Thus every constraint acts as a filter on the set of output candidates, with the important addendum that the order in which the filters are applied is crucial in determining optimality.

Consider the example in Fig. 1 on the facing page, which depicts an OT evaluation of output candidates using the tableau notation. Here some input  $i$  is assigned three output candidates  $o_1, o_2$  and  $o_3$ . The OT grammar uses only three constraints  $c_1, c_2$  and  $c_3$ , with each  $c_i$  preceding  $c_{i+1}$ ,  $1 \leq i < 2$ . Constraint  $c_1$  is applied first. Candidates  $o_2$  and  $o_3$  each violate it once, however,  $o_1$  violates it twice and there are no other output candidates, thus  $o_2$  and  $o_3$  are the output candidates incurring the least number of violations of the constraint and are allowed to proceed to the next round of the evaluation. Candidate  $o_1$ , on the other hand, is ruled out and does not participate in further evaluations. Neither  $o_2$  nor  $o_3$  violate  $c_2$  (nor does  $o_1$ , but this is immaterial since it has been previously ruled out), so neither is filtered out. In the third round,  $o_2$  and  $o_3$  are evaluated with respect to  $c_3$ . Each of them violates the constraint once, but since there is no candidate that fares better than them (again,  $o_1$  is not taken into consideration anymore), they also survive this round of the evaluation. Thus,  $o_2$  and  $o_3$  are the optimal output candidates for  $i$ . If  $c_3$  had been applied before  $c_1$ , on the other hand,  $o_2$  and  $o_3$  would lose out against  $o_1$ .

With this intuitive understanding of OT grammars under our belt, the formal definitions of OSs and their *output language* (not to be confused with the *candidate language*  $\text{ran}(\text{GEN})$ ) are straight-forward.

**Definition 4 (Optimality system).** An optimality system over languages  $L, L'$  is a pair  $O := \langle \text{GEN}, C \rangle$  with  $\text{GEN} \subseteq L \times L'$  and  $C := \langle c_1, \dots, c_n \rangle$  a linearly ordered sequence

$i$	$c_1$	$c_2$	$c_3$
$o_1$	2	0	0
$o_2$	1	0	1
$o_3$	1	0	1

**Fig. 1.** Example of an OT evaluation in tableau notation

of functions  $c_i: \text{GEN} \rightarrow \mathbb{N}$ . For  $a, b \in \text{GEN}$ ,  $a <_O b$  iff there is an  $1 \leq k \leq n$  such that  $c_k(a) < c_k(b)$  and for all  $j < k$ ,  $c_j(a) = c_j(b)$ .

**Definition 5 (Optimality).** Given an optimality system  $O := \langle \text{GEN}, C \rangle$ ,  $\langle i, o \rangle$  is optimal with respect to  $O$  iff both  $\langle i, o \rangle \in \text{GEN}$  and there is no  $o'$  such that  $\langle i, o' \rangle \in \text{GEN}$  and  $\langle i, o' \rangle <_O \langle i, o \rangle$ . The output language of  $O$  is  $\text{ran}(\{\langle i, o \rangle \mid \langle i, o \rangle \text{ is optimal with respect to } O\})$ .

The important insight of [3] as well as [8], which was later improved upon by [7, 9, 20], is that an OS as defined above can be understood to define a transduction from a set of inputs to its set of optimal output candidates. Moreover, if the OS is suitably restricted, it is guaranteed to define a rational relation, which implies its efficient computability.

**Theorem 6.** Let  $O := \langle \text{GEN}, C \rangle$  be an OS such that

- $\text{dom}(\text{GEN})$  is a regular string language, or a regular/linear context-free tree language, and
- $\text{GEN}$  is a rational relation, and
- all  $c \in C$  are output-markedness constraints, and
- each  $c \in C$  defines a rational relation on  $\text{ran}(\text{GEN})$ , and
- $O$  is globally optimal.

Then the transduction  $\tau$  induced by the OS is a rational relation and  $\text{ran}(\tau)$  belongs to the same formal language class as  $\text{dom}(\tau)$ .

As the reader might have guessed, the use of  $\tau$  here is carried over straight-forwardly from tree transducers, meaning that  $\tau := \{\langle i, o \rangle \mid \langle i, o \rangle \text{ is optimal with respect to } O\}$ . The theorem also makes reference to two notions we have not encountered yet at all, output-markedness and global optimality. The former is easily defined.

**Definition 7 (Output-Markedness Constraints).** Given an OS  $O := \langle \text{GEN}, C \rangle$ ,  $c \in C$  is an output-markedness constraint iff  $c(\langle i, o \rangle) = c(\langle i', o \rangle)$  for all  $\langle i, o \rangle, \langle i', o \rangle \in \text{GEN}$ .

Global optimality, on the other hand, requires a lot of finite-state machinery to be in place before it can be made formally precise, which would lead us off track here. For our purposes it is sufficient to know that an OS is globally optimal iff for every optimal output candidate  $o$  it holds that there is no input  $i$  such that  $o$  is an output candidate for  $i$  but not an optimal one. The curious reader is referred to [7] for a more rigorous definition.

### 3 Controlled Optimality Systems

OSs are perfectly capable of modelling reference-set constraints. The reference set for any input  $i$  is defined by  $i\text{GEN}$ , and the evaluation metric can straightforwardly be implemented as a sequence of constraints. Fewest Steps, for instance, can be viewed as an OS in which a tree  $t$  is related by GEN to all the trees that were constructed from the same lexical items as  $t$ , including  $t$  itself. Besides that, the OS has only one constraint \*TRACE, which punishes traces. As a consequence, only the trees with the least number of traces will be preserved, and these are the optimal output candidates for  $t$ . Focus Economy is slightly more involved and allows for different models. In one of them (which isn't the one I will formalize in Sec. 4.2), GEN relates a tree  $t$  without stress annotation but a focus-marked constituent to its analogs with neutral and shifted stress. The OS then uses two constraints, the first of which weeds out trees in which the focus-marked constituent does not contain the main stress carrier, while the second penalizes derivation length (the assumption here is that derivations for trees with shifted stress are longer than those for trees with neutral stress). Then a tree derived by stress shift will be optimal just in case the equivalent tree with neutral stress has already been ruled out by the first constraint.

While these two examples show that OSs certainly can get the job done, the way output candidates are specified for reference-set constraints actually relies on additional structure — the reference sets — that is only indirectly represented by GEN. In the following I introduce controlled OSs as a variant of standard OSs that is closer to reference-set computation by making reference sets prime citizens of OSs and demoting GEN to an ancillary relation that is directly computed from them. Note, though, that in the case of a constraint like Focus Economy, where input and candidate language are disjoint, distinct inputs might be assigned the same reference set. In such a configuration, it makes sense to map entire sets of inputs to reference sets, rather than the individual inputs themselves. Let us call such a set of inputs a *reference type*. An OS can then be defined by reference types and a function mapping them to reference sets:

**Definition 8 (Controlled Optimality Systems).** An  $\mathcal{F}$ -controlled optimality system over languages  $L$ ,  $L'$  is a 4-tuple  $O[\mathcal{F}] := \langle \text{GEN}, C, \mathcal{F}, \gamma \rangle$ , where

- GEN and  $C$  are defined as usual,
- $\mathcal{F}$  is a family of non-empty subsets of  $L$ , each of which we call a reference type,
- the control map  $\gamma : \mathcal{F} \rightarrow \wp(L') \setminus \{\emptyset\}$  associates every reference type with a reference set, i.e. a set of output candidates,
- the following conditions are satisfied
  - exhaustivity:  $\bigcup_{X \in \mathcal{F}} X = L$
  - bootstrapping:  $x\text{GEN} = \bigcup_{X \in \mathcal{F}} X\gamma$

Every controlled OS can be translated into a canonical OS by discarding its third and fourth component (i.e.  $\mathcal{F}$  and the control map  $\gamma$ ). In the other direction, a controlled OS can be obtained from every canonical OS by setting  $\mathcal{F} := \{\{i\} \mid i \in L\}$  and  $\gamma : \{i\} \mapsto i\text{GEN}$ . So the only difference between the two is that controlled OSs modularize GEN by specifying it through reference types and the control map.

Now that OSs operate (at least partially) at the level of sets, it will often be interesting to talk about the set of optimal output candidates assigned to some reference type, rather than one particular input. But whereas the set of optimal output candidates is always well-defined for inputs — recall that for any input  $i$  this set is given by  $i\tau$  — we have to be more careful when lifting it to reference types, because distinct inputs that belong to the same reference type may not necessarily be assigned the same optimal output candidates. Such a situation may arise in OSs with faithfulness or input-markedness constraints, which are sensitive to properties of the input, or when two inputs  $i$  and  $j$  are of reference type  $X$ , but in addition  $j$  is also of reference type  $Y$ . Given this ambiguity, one has to distinguish between the set of output candidates that are optimal for at least one member of reference type  $X$ , and the set of output candidates that are optimal for all members of reference type  $X$ . The former is the *up-transduction*  $X\tau^\uparrow := \bigcup_{x \in X} x\tau$ , the latter the *down-transduction*  $X\tau^\downarrow := \bigcap_{x \in X} x\tau$ .

At this point it might be worthwhile to work through a simple example. Fig. 2 on the next page depicts a controlled OS and the distinct steps of its computation. We are given a collection of reference types consisting of RED :=  $\{i_1, i_2, i_3, i_4, i_5, i_6\}$ , SIENNA :=  $\{i_4\}$ , TEAL :=  $\{i_5, i_6, i_7\}$ , PURPLE :=  $\{i_8\}$ , and LIME :=  $\{i_7, i_9, i_{10}\}$ . The reference sets are BLUE :=  $\{o_1, o_2, o_3\}$ , ORANGE :=  $\{o_3, o_4, o_5, o_6, o_7\}$ , GREEN :=  $\{o_6, o_7\}$ , and BROWN :=  $\{o_8, o_9\}$ . Finally, the graph of  $\gamma$  consists of the pairs  $\langle \text{RED}, \text{BLUE} \rangle$ ,  $\langle \text{SIENNA}, \text{BROWN} \rangle$ ,  $\langle \text{TEAL}, \text{GREEN} \rangle$ ,  $\langle \text{PURPLE}, \text{BROWN} \rangle$ , and  $\langle \text{LIME}, \text{ORANGE} \rangle$ . Note that a reference type may overlap with another reference type or even be a proper subset of it, and the same holds for reference sets. This means that an input can belong to several reference types at once. Consequently,  $x\text{GEN}$  may be a superset of  $X\gamma$  for every reference type  $X$  that contains  $x$ , as is the case for  $i_4$ , say, but not for  $i_7$ , even though both are assigned exactly two reference types. Input  $i_4$  is related by GEN to all outputs contained in  $\text{RED}\gamma \cup \text{SIENNA}\gamma = \text{BLUE} \cup \text{BROWN} = \{o_1, o_2, o_3, o_8, o_9\}$ , whereas  $i_7$  is related to  $\text{LIME}\gamma \cup \text{TEAL}\gamma = \text{ORANGE} \cup \text{GREEN} = \text{ORANGE} = \{o_3, o_4, o_5, o_6, o_7\}$ . AS SOON AS GEN has been determined from the reference types and the control map, the computation proceeds as usual with the constraints of the OS filtering out all suboptimal candidates.

Interestingly, almost all reference-set constraints fall into two classes with respect to how reference types and reference sets are distributed (see Fig. 3 on the following page). In the case of Fewest Steps, where the input language is also the candidate language, each reference type is mapped to itself, that is to say, there is no distinction between reference types and reference sets. A constraint like Focus Economy, on the other hand, requires not only the input language and the candidate language to be disjoint, but also all reference sets and reference types.

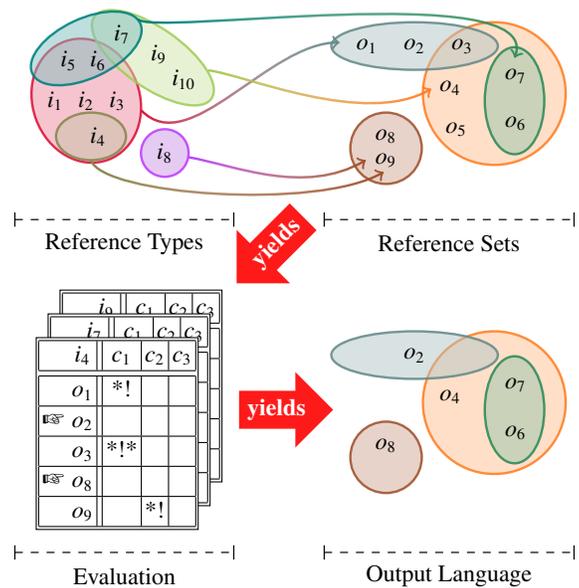
The Fewest Steps-type class of OSs is best captured by the notion of *endocentricity*.

**Definition 9.** An  $\mathcal{F}$ -controlled OS is endocentric iff  $X\gamma = X$  for all  $X \in \mathcal{F}$ .

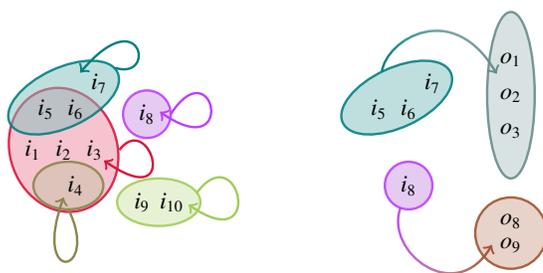
Regarding constraints like Focus Economy, a natural candidate might be the class of *output-partitioned* and *output-segregated* OSs.

**Definition 10.** An  $\mathcal{F}$ -controlled OS is

- output-partitioned iff for all distinct  $X, Y \in \mathcal{F}$ ,  $X\gamma \neq Y\gamma$  implies  $X\gamma \cap Y\gamma = \emptyset$ .



**Fig. 2.** Example of a controlled OS.  $\text{GEN}$  is defined in a modular fashion using reference types, reference sets, and the control map  $\gamma$  from reference types to reference sets.



**Fig. 3.** Almost all instances of reference-set computation in the literature use one of the two configurations above, both of which are output joint preserving.

- output-segregated iff for all distinct  $X, Y \in \mathcal{F}$ , neither  $X\gamma \subseteq Y\gamma$  nor  $Y\gamma \subseteq X\gamma$ .

While this class certainly fits Focus Economy and its ilk pretty well, it does not extend naturally to endocentric OSs, which means that we would have to study those two classes independently from each other. In particular, endocentric OSs aren't necessarily output-partitioned, as they do allow for overlapping reference-sets. Nor is output-segregation a meaningful restriction on endocentric OSs: Given some  $\mathcal{F}$ -controlled OS, assume  $x \in X, y \in Y, X \subseteq Y$ . Then  $x$  is also a member of  $Y$ , whence both  $\langle x, y \rangle$  and  $\langle y, x \rangle$  are in (the graph of) GEN. The reference type  $X$  is immaterial for computing GEN and can be removed from  $\mathcal{F}$  without consequences. Given these discrepancies, we should not be too sure that results pertaining to one class can easily be carried over to the other. As it turns out, though, a natural unification of the two subclasses is available in the form of *output joint preservation*.

**Definition 11.** An  $\mathcal{F}$ -controlled optimality system is output joint preserving iff for all distinct  $X, Y \in \mathcal{F}$ ,  $X\gamma \cap Y\gamma \neq \emptyset \rightarrow X \cap Y \neq \emptyset$ .

The OS depicted in Fig. 2 on the preceding page fails output joint preservation. It is clearly violated by SIENNA and PURPLE, which are disjoint yet mapped to the same reference set, BROWN. It isn't respected by RED and LIME, either, which are mapped to BLUE and ORANGE, respectively, the intersection of which is non-empty even though RED and LIME are disjoint. Moving on to Fig. 3 on the facing page, we observe that every endocentric OS is output joint preserving. In addition, all output-partitioned OSs trivially satisfy output joint preservation, too, because there are no joints to preserve in these OSs. These inclusions tell us that output joint preservation is certainly general enough a property to encompass the kinds of controlled OSs we are interested in. In the following, I show that it is also sufficiently restrictive to establish a link to global optimality.

**Definition 12.** An  $\mathcal{F}$ -controlled OS is type-level optimal iff  $X\tau^\uparrow \upharpoonright X\gamma = X\tau^\downarrow \upharpoonright X\gamma$  for all  $X \in \mathcal{F}$ .

**Lemma 13.** Let  $O[\mathcal{F}]$  an  $\mathcal{F}$ -controlled OS. Then  $O[\mathcal{F}]$  is type-level optimal only if it is globally optimal.

*Proof.* We prove the contrapositive. If  $O[\mathcal{F}]$  is not type-level optimal, then it holds for some  $X \in \mathcal{F}$  that  $X\tau^\uparrow \upharpoonright X\gamma \neq X\tau^\downarrow \upharpoonright X\gamma$ . But this implies that there are  $x, y \in X$  and  $z \in X\gamma$  such that  $x\tau \ni z \notin y\tau$ , which is an unequivocal violation of global optimality.

**Theorem 14.** Every output joint preserving OS is type-level optimal iff it is globally optimal.

*Proof.* The right-to-left direction follows from Lemma 13. We prove the contrapositive of the other direction. If  $O[\mathcal{F}]$  fails global optimality, then there are  $x, y \in L$  and  $z \in L'$  such that  $\langle x, z \rangle, \langle y, z \rangle \in \text{GEN}$  yet  $x\tau \ni z \notin y\tau$ . W.l.o.g. let  $x \in X$  and  $y \in Y, X, Y \in \mathcal{F}$ , whence  $z \in X\gamma \cap Y\gamma$ . As  $O[\mathcal{F}]$  is output joint preserving,  $X\gamma \cap Y\gamma \neq \emptyset$  entails  $X \cap Y \neq \emptyset$ . Pick some  $p \in X \cap Y$ . Now if  $O[\mathcal{F}]$  is type-level optimal, then it holds that  $X\tau^\uparrow \upharpoonright X\gamma = X\tau^\downarrow \upharpoonright X\gamma$  and  $Y\tau^\uparrow \upharpoonright Y\gamma = Y\tau^\downarrow \upharpoonright Y\gamma$ , so  $z \in x\tau$  implies  $z \in p\tau$ , whereas  $z \notin y\tau$  implies  $z \notin p\tau$ . Contradiction. It follows that  $O[\mathcal{F}]$  is not type-level optimal.

Intuitively, type-level optimality ensures that optimality is fixed for entire reference types, so the individual inputs can be ignored for determining optimality. However, it is too weak a restriction to rule out disagreement between reference types that are mapped to overlapping reference sets, so output joint preservation has to step in; it guarantees that if two reference types  $X$  and  $Y$  share at least one output candidates, there exists some input  $p$  belonging to both  $X$  and  $Y$  that will be faced by conflicting requirements if  $X$  and  $Y$  disagree with respect to which candidates in  $X\gamma \cap Y\gamma$  they deem optimal (since the OS is type-level optimal, optimality can be specified for entire reference types rather than their members). It should be easy to see that the conditions jointly imply global optimality.

Given our interest in using controlled OS to investigate the computability of reference-set constraints, it would be advantageous if we could read off the constraints right away whether they are type-level optimal. This is indeed very easy to do thanks to the following entailment.

**Lemma 15.** *Let  $O[\mathcal{F}] := \langle \text{GEN}, C, \mathcal{F}, \gamma \rangle$  an  $\mathcal{F}$ -controlled OS such that every  $c \in C$  is an output-markedness constraint. Then  $O[\mathcal{F}]$  is type-level optimal.*

*Proof.* Assume the opposite. Then for some  $X \in \mathcal{F}$ ,  $X\tau^\uparrow \uparrow X\gamma \neq X\tau^\downarrow \uparrow X\gamma$ , whence there are  $x, y \in X$  and  $z \in X\gamma$  with  $x\tau \ni z \notin y\tau$ . But this is the case only if there is some  $c \in C$  such that  $c(\langle x, z \rangle) \neq c(\langle y, z \rangle)$ , i.e.  $c$  isn't an output-markedness constraint.

**Corollary 16.** *Let  $O[\mathcal{F}] := \langle \text{GEN}, C, \mathcal{F}, \gamma \rangle$  an output joint preserving OS such that every  $c \in C$  is an output-markedness constraint. Then  $O[\mathcal{F}]$  is globally optimal.*

Combining these results, we arrive at the equivalent of Thm. 6 for  $\mathcal{F}$ -controlled OSs.

**Corollary 17.** *Let  $O[\mathcal{F}] := \langle \text{GEN}, C, \mathcal{F}, \gamma \rangle$  an  $\mathcal{F}$ -controlled OS such that*

- $\text{dom}(\text{GEN})$  is a regular string language, or a regular/linear context-free tree language, and
- $\text{GEN}$  is a rational relation, and
- all  $c \in C$  are output-markedness constraints, and
- each  $c \in C$  defines a rational relation on  $\text{ran}(\text{GEN})$ , and
- $O[\mathcal{F}]$  is output joint preserving.

*Then the transduction  $\tau$  induced by the OS is a rational relation and  $\text{ran}(\tau)$  belongs to the same formal language class as  $\text{dom}(\tau)$ .*

In sum, then, we see that not only do output joint preserving OSs look like a solid base for modelling reference-set constraints, they also have the neat property that the global optimality check is redundant, thanks to Lem. 15. As it is pretty easy to determine for any given reference-set constraint whether it can be modelled by output-markedness constraints alone, the only stumbling block in implementing Fewest Steps-type constraints, i.e. constraints that can be modelled by endocentric OSs, is the transducers for the constraints and  $\text{GEN}$ . If those transducers each define a rational relation, so does the entire optimality system.

## 4 Application to Reference-Set Computation

### 4.1 Focus Economy Explained

Focus Economy [13, 18] was briefly discussed in the introduction. It is invoked in order to account for the fact that sentences such as (2a), (2b) and (2c) below differ with respect to what is given and what is new information. Once again main stress is marked by **boldface**.

- (2) a. My friend Paul bought a new red **car**.
- b. My friend Paul **bought** a new red car.
- c. My friend Paul bought a new **red** car.

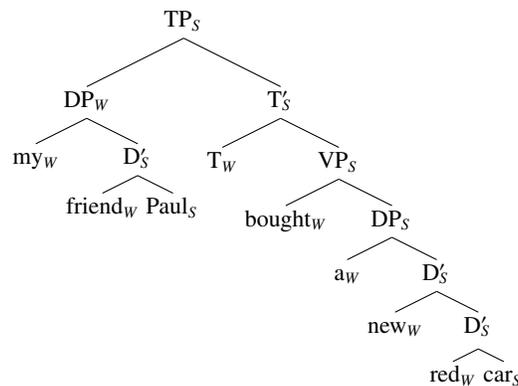
That these utterances are associated to different information structures is witnessed by the following (in)felicity judgments. For the reader's convenience, the focus, i.e. the new discourse material introduced by each answer, is put in square brackets.

- (3) What happened?
  - a. [<sub>F</sub>My friend Paul bought a new red **car**.]
  - b. # [<sub>F</sub>My friend Paul **bought** a new red car.]
  - c. # [<sub>F</sub>My friend Paul bought a new **red** car.]
- (4) What did your friend Paul do?
  - a. He [<sub>F</sub> bought a new red **car**].
  - b. # He [<sub>F</sub> **bought** a new red car].
  - c. # He [<sub>F</sub> bought a new **red** car].
- (5) What did your friend Paul buy?
  - a. He bought [<sub>F</sub> a new red **car**].
  - b. # He **bought** [<sub>F</sub> a new red car].
  - c. # He bought [<sub>F</sub> a new **red** car].
- (6) Did your friend Paul sell a new red car?
  - a. # No, he [<sub>F</sub> bought] a new red **car**.
  - b. No, he [<sub>F</sub> **bought**] a new red car.
  - c. # No, he [<sub>F</sub> bought] a new **red** car.
- (7) Did your friend Paul buy a green car?
  - a. # No, he bought a [<sub>F</sub> new red] **car**.
  - b. # No, he **bought** a [<sub>F</sub> new red] car.
  - c. He bought a [<sub>F</sub> new **red**] car.

Restricting our attention to the a-sentences only, we might conclude that a constituent can be focused just in case one of its subconstituents carries sentential main stress. A short glimpse at the b- and c-utterances falsifies this conjecture, though. Perhaps, then, main stress has to fall on the subconstituent at the right edge of the focused constituent? This is easily shown to be wrong, too. In (8) below, the stressed constituent isn't located at either edge of the focused constituent.

- (8) a. What happened to Mary?  
 b. [<sub>F</sub> John **killed** her.]

The full-blown Focus Economy system (rather than the simplified sketch given in the introduction) accounts for the data as follows. First, the *Main Stress Rule* demands that in every pair of sister nodes, the “syntactically more embedded” node [13, p.133] is assigned strong stress, its sister weak stress (marked in the phrase structure tree by subscripted S and W, respectively). If a node has no sister, it is always assigned strong stress (in Minimalist syntax, this will be the case only for the root node, as all Minimalist trees are strictly binary branching). Main stress then falls on the unique leaf node that is connected to the root node by a path of nodes that have an S-subscript. See Fig. 4 for an example.



**Fig. 4.** The stress-annotated phrase structure tree for (2a).

The notion of being syntactically more embedded isn't explicitly defined in the literature. It is stated in passing, though, that "... main stress falls on the most embedded constituent on the recursive side of the tree" [13, p.133]. While this is rather vague, it is presumably meant to convey that, at least for English, in which complements follow the heads they are introduced by, the right sister node is assigned strong stress as long as it isn't an adjunct. This interpretation seems to be in line with the empirical facts.

The second integral part of the proposal is the operation *Stress Shift*, which shifts the main stress to some leaf node *n* by assigning all nodes on the path from *n* to the root strong stress and demoting the sisters of these nodes to weakly stressed nodes. For instance, the tree for "My **friend** Paul bought a new red car" is obtained from the tree in Fig. 4 by changing friend<sub>W</sub> and Paul<sub>S</sub> to friend<sub>S</sub> and Paul<sub>W</sub>, respectively, and DP<sub>W</sub> and T'<sub>S</sub> to DP<sub>S</sub> and T'<sub>W</sub>, respectively.

While Stress Shift could be invoked to move stress from anaphoric elements to their left sister as in (8), this burden is put on a separate rule, for independent reasons. The rule in question is called *Anaphoric Destressing* and obligatorily assigns weak stress to anaphoric nodes, where a node is anaphoric iff it is "...D[iscourse]-linked to an

accessible discourse entity” [13, p.147]. Thus Anaphoric Destressing not only accounts for the unstressed anaphor in (8), but also for the special behavior of stress in cases of parallelism.

- (9) First Paul bought a new **car**.
- a. Then **John** bought one.
  - b. \* Then John bought **one**.

The overall system now works as follows. Given a phrase structure tree that has not been annotated for stress yet, one first applies Anaphoric Destressing to make sure that all d-linked constituents are assigned weak stress and thus cannot carry main stress. Next the Main Stress Rule is invoked to assign every node in the tree either W or S. Note that the Main Stress Rule cannot overwrite previously assigned labels, so if some node  $n$  has been labeled W by Anaphoric Destressing, the Main Stress Rule has to assign S to the sister of  $n$ . Now that the tree is fully annotated, we compute its *focus set*, the set of constituents that may be focused.

(10) *Focus Projection*

Given some stress-annotated tree  $t$ , its focus set is the set of nodes reflexively dominating the node carrying main stress.

The focus set of “Paul bought a new **car**”, for instance, contains [car], [.D’ new car], [.DP a new car], [.VP bought a new car] and [.TP Paul bought a new car] (equivalently, we could associate every node in the tree with a unique address and simply use these addresses in the focus set). For “Then **John** bought one”, on the other hand, it consists only of [John] and [.TP Then John bought one].

At this point, Stress Shift may optionally take place. After the main stress has been shifted, however, the focus set has to be computed all over again, and this time the procedure involves reference-set computation.

(11) *Focus Projection Redux*

Given some stress-annotated tree  $t'$  that was obtained from tree  $t$  by Stress Shift, the focus set of  $t'$  contains all the nodes reflexively dominating the node carrying main stress which aren’t already contained in the focus set of  $t$ .

So if “Then **John** bought one” had been obtained by Stress Shift from [.TP Then John bought one] rather than Anaphoric Destressing, its focus set would have contained only [John], because [.TP Then John bought one] already belongs to the focus set of “Then John bought **one**”. As an easy exercise, the reader may want to draw annotated trees for the examples in (2) and compute their focus sets.

## 4.2 A Model of Focus Economy

After this general overview, the time has come to formalize Focus Economy. In order to precisely model Focus Economy, though, I have to make some simplifying assumptions, for reasons that are entirely independent from the restrictions of OSs. First, I stipulate that adjuncts are explicitly marked as such by a subscript A on their label. This is simply

a matter of convenience, as it reduces the complexity of the transducers and makes my model independent from the theoretical status of adjuncts in syntax.

Second, I decided to take movement out of the picture, because the interaction of focus and movement is not touched upon in [13], so there is no original material to formalize. Incidentally, movement seems to introduce several interesting complications, as illustrated by sentences involving topicalization, where no other assignment of focus and main stress is grammatical.

- (12) a.  $_F$  **John**<sub>*i*</sub> Paul likes  $t_i$ .  
 b. \*  $\text{John}_i$  [ $_F$  **Paul**] likes  $t_i$ .  
 c. \*  $\text{John}_i$  [ $_F$  Paul **likes**  $t_i$ ].

The end of this section contains a brief discussion of whether my model can be extended to capture theories involving movement.

The last simplification concerns Anaphoric Distressing itself. While the core of Anaphoric Distressing, the distressing of pronominal (and possibly covert) elements, is easy to accommodate, the more general notion of d-linking is impossible to capture in any model that operates on isolated syntactic trees. Devising a working model of discourse structure vastly exceeds the scope of this contribution. Also, the role of d-linking in anaphoric distressing is of little importance to this paper, which focuses on the reference-set computational aspects of Focus Economy. Thus my implementation will allow almost any constituent to be anaphorically distressed and leave the task of matching trees to appropriate discourse contexts to an external theory of d-linking that remains to be specified.

With these provisions made explicit, the formalization of Focus Economy as a controlled OS can commence. The input language is supposedly derived by some movement-free Minimalist grammar  $\mathcal{E}$  for English [17] in which interior nodes are given explicit category labels (once more for the sake of convenience). As Minimalist grammars without remnant movement generate regular tree languages [10], it is safe to assume that Minimalist grammars without any kind of movement do so, too.

Next I define GEN as the composition of four linear transducers corresponding to Anaphoric Distressing, the Main Stress Rule, Stress Shift, and Focus Projection, respectively. Given a tree  $t$  derived by  $\mathcal{E}$ , the transducer cascade computes all logically possible variants of  $t$  with respect to stress assignment and then computes the focus in a local way. This means that GEN actually overgenerates with respect to focus, a problem that we have to take care of at a later step.

Anaphoric Distressing is modelled by a non-deterministic ldttd that may randomly add a subscript D to a node's label in order to mark it as anaphoric. The only condition is that if a node is labeled as anaphoric, all the nodes it properly dominates must be marked as such, too.

**Definition 18.** *Let  $\Sigma := \Sigma_L \cup \Sigma_A$  be the vocabulary of the Minimalist grammar  $\mathcal{E}$  that generated the input language, where  $\Sigma_L$  contains all lexical items and category labels and  $\Sigma_A$  their counterparts explicitly labeled as adjuncts. Anaphoric Distressing is the ldttd  $\mathcal{D}$  where  $\Sigma_{\mathcal{D}} := \Sigma$ ,  $\Omega_{\mathcal{D}}$  is the union of  $\Sigma$  and  $\Sigma_D := \{\sigma_D \mid \sigma \in \Sigma\}$ ,  $Q := \{q_i, q_d\}$ ,  $Q' := \{q_i\}$ , and  $\Delta_{\mathcal{D}}$  contains the rules below, with  $\sigma \in \Sigma$  and  $\sigma_D \in \Sigma_D$  and  $\alpha_{\{x,y\}}$  to be*

read as “ $\alpha_x$  or  $\alpha_y$ ”:

$$\begin{array}{ll} q_i(\sigma(x, y)) \rightarrow \sigma(q_i(x), q_i(y)) & q_i(\sigma) \rightarrow \sigma \\ q_{\{i,d\}}(\sigma(x, y)) \rightarrow \sigma_D(q_d(x), q_d(y)) & q_{\{i,d\}}(\sigma) \rightarrow \sigma_D \end{array}$$

The transducer for the Main Stress Rule is non-deterministic, too, but it proceeds in a bottom-up manner. It does not alter nodes subscripted by A or D, but if it encounters a leaf node without a subscript, it randomly adds the subscript S or W to its label. However, W is allowed to occur only on left sisters, whereas S is mostly restricted to right sisters and may surface on a left sister just in case the right sister is already marked by A or D. Note that we could easily define a different stress pattern, maybe even parametrized with respect to category labels, to incorporate stress assignment rules from other languages.

**Definition 19.** Main Stress is the lbuttt  $\mathcal{M}$  where  $\Sigma_{\mathcal{M}} := \Omega_D$ ,  $\Omega_{\mathcal{M}}$  is the union of  $\Sigma$ ,  $\Sigma_D$  and  $\Sigma_* := \{\sigma_S, \sigma_W \mid \sigma \in \Sigma\}$ ,  $\mathcal{Q} := \{q_s, q_u, q_w\}$ ,  $\mathcal{Q}' := \{q_s\}$  and  $\Delta_{\mathcal{M}}$  contains the following rules, with  $\sigma \in \Sigma$ ,  $\sigma_A \in \Sigma_A$ ,  $\sigma_x \in \{\sigma_x \mid \sigma \in \Sigma\}$  for  $x \in \{D, S, W\}$ , and  $\alpha_{a\dots z}(\beta_{a'\dots z'}, \dots, \zeta_{a''\dots z''})$  to be read as “ $\alpha_a(\beta_{a'}, \dots, \zeta_{a''})$  or ... or  $\alpha_z(\beta_{z'}, \dots, \zeta_{z''})$ ”:

$$\begin{array}{ll} \sigma_A \rightarrow q_u(\sigma_A) & \sigma_A(q_u(x), q_u(y)) \rightarrow q_u(\sigma_A(x, y)) \\ \sigma_D \rightarrow q_u(\sigma_D) & \sigma_D(q_u(x), q_u(y)) \rightarrow q_u(\sigma_D(x, y)) \\ \sigma \rightarrow q_{sw}(\sigma_{SW}) & \sigma(q_{\{u,w\}}(x), q_s(y)) \rightarrow q_{sw}(\sigma_{SW}(x, y)) \\ & \sigma(q_s(x), q_u(y)) \rightarrow q_{sw}(\sigma_{SW}(x, y)) \end{array}$$

Stress Shift is best implemented as a non-deterministic ldttt that may randomly switch the subscripts of two S/W-annotated sisters.

**Definition 20.** Stress Shift is the ldttt  $\mathcal{S}$  where  $\Sigma_{\mathcal{S}} = \Omega_S = \Omega_{\mathcal{M}}$ ,  $\mathcal{Q} := \{q_i, q_s, q_w\}$ ,  $\mathcal{Q}' := \{q_s\}$ , and  $\Delta_{\mathcal{S}}$  contains the rules below, with  $\sigma \in \Sigma_S$  and  $\sigma_* \in \Sigma_*$ :

$$\begin{array}{ll} q_s(\sigma_*(x, y)) \rightarrow \sigma_{SS}(q_{isw}(x), q_{iws}(y)) & q_s(\sigma_*) \rightarrow \sigma_S \\ q_w(\sigma_*(x, y)) \rightarrow \sigma_W(q_i(x), q_i(y)) & q_w(\sigma_*) \rightarrow \sigma_W \\ q_i(\sigma(x, y)) \rightarrow \sigma(q_i(x), q_i(y)) & q_i(\sigma) \rightarrow \sigma \end{array}$$

The last component is Focus Projection, which is formalized as a non-deterministic ldttt with two states,  $q_f$  and  $q_g$ . The transducer starts at the root in  $q_f$ . Whenever a node  $n$  is subscripted by W, the transducer switches into  $q_g$  at this node and stays in the state for all nodes dominated by  $n$ . As long as the transducer is in  $q_f$ , it may randomly add a superscript F to a label to indicate that it is focused. Right afterwards, it changes into  $q_g$  and never leaves this state again. Rather than associating a stress-annotated tree with a set of constituents that can be focused, Focus Projection now generates multiple trees that differ only with respect to which constituent along the path of S-labeled nodes is focus-marked.

**Definition 21.** Focus Projection is the ldttt  $\mathcal{F}$ , where  $\Sigma_{\mathcal{F}} = \Omega_S$ ,  $\Omega_{\mathcal{F}}$  is the union of  $\Omega_S$  and  $\Omega_S^F := \{\omega^F \mid \omega \in \Omega_S\}$ ,  $\mathcal{Q} := \{q_f, q_g\}$ ,  $\mathcal{Q}' := \{q_f\}$ , and  $\Delta_{\mathcal{F}}$  contains the rules below,

with  $\sigma \in \Sigma_{\mathcal{F}}$  and  $\sigma_{\bar{S}} \in \Sigma_{\mathcal{F}} \setminus \{\sigma_S \mid \sigma \in \Sigma\}$ :

$$\begin{array}{ll}
q_f(\sigma_S(x, y)) \rightarrow \sigma_S(q_f(x), q_f(y)) & \\
q_f(\sigma_S(x, y)) \rightarrow \sigma_S^F(q_g(x), q_g(x)) & q_f(\sigma_S) \rightarrow \sigma_S^F \\
q_f(\sigma_{\bar{S}}(x, y)) \rightarrow \sigma_{\bar{S}}(q_g(x), q_g(x)) & q_f(\sigma_{\bar{S}}) \rightarrow \sigma_{\bar{S}} \\
q_g(\sigma(x, y)) \rightarrow \sigma(q_g(x), q_g(y)) & q_g(\sigma) \rightarrow \sigma
\end{array}$$

All four transducers are linear, whence they can be composed into a single linear transducer modelling GEN. Expanding on what was said above about the inner workings of GEN, we now see that for any tree  $t$  in the input language,  $t_{\text{GEN}}$  is the set of stress-annotated trees in which, first, some subtrees may be marked as adjuncts or anaphorical material (or both) and thus do not carry stress information, second, there is exactly one path from the root to some leaf such that every node in the path is labeled by S, and third, exactly one node belonging to this path is marked as focused. The reader should have no problem verifying that in terms of controlled OSs, all reference types are singleton and their reference-sets do not overlap, i.e. output joint preservation and type-level optimality are satisfied.

Now it only remains for us to implement *Focus Projection Redux*. In the original account, Focus Projection Redux applied directly to the output of Stress Shift, i.e. trees without focus information, and the task at hand was to assign the correct focus. In my system, on the other hand, every tree is fed into Focus Projection and marked accordingly for focus. This leads to overgeneration for trees in which Stress Shift has taken place — a node may carry focus even if it could also do so in the tree without shifted main stress. Consequently, the focus set of “**John** died”, for instance, turns out to contain both [John] and [.TP John died] rather than just the former. Under my proposal, then, Focus Projection Redux is faced with the burden of filtering out focus information instead of assigning it. In other words, Focus Projection Redux is a constraint.

This is accomplished by defining a regular tree language  $L_c$  such that when GEN is composed with the diagonal of  $L_c$  (which is guaranteed to be a linear transduction), only trees with licit focus marking are preserved. Said regular language is easily specified in the monadic second-order logic  $\mathcal{L}_{K,P}^2$  [14]. First one defines two predicates, *StressPath* and *FocusPath*. The former picks out the path from the root to the leaf carrying main stress, whereas the latter refers to the path from the root to the leaf that would carry main stress in the absence of stress shift. This implies that *FocusPath* replicates some of the information that is already encoded in the Main Stress transducer. Note that in the definitions below,  $A(x)$ ,  $D(x)$  and  $S(x)$  are predicates picking out all nodes with subscript A, D, S, respectively,  $x \triangleleft y$  denotes “ $x$  is the parent of  $y$ ”,  $x < y$  “ $x$  is the left sibling of  $y$ ”, and  $\triangleleft^*$  the reflexive transitive closure of  $\triangleleft$ .

$$\begin{aligned}
\text{Path}(X) \leftrightarrow \exists x [X(x) \wedge \neg \exists y [y \triangleleft x]] \wedge \exists ! x [X(x) \wedge \neg \exists y [x \triangleleft y]] \wedge \\
\forall x, y, z [(X(x) \wedge X(y) \rightarrow x \triangleleft^* y \vee y \triangleleft^* x) \wedge (X(x) \wedge \neg X(z) \rightarrow \neg(z \triangleleft^* x))]
\end{aligned}$$

$$\text{StressPath}(X) \leftrightarrow \text{Path}(X) \wedge \forall x [X(x) \rightarrow S(x)]$$

$$\text{FocusPath}(X) \leftrightarrow \text{Path}(X) \wedge \forall x, y, z \left[ X(x) \wedge x \triangleleft y \wedge x \triangleleft z \rightarrow \right. \\ \left. ((A(y) \vee D(y)) \rightarrow X(z)) \wedge (\neg A(y) \wedge \neg D(y) \wedge y < z \rightarrow X(z)) \right]$$

In a tree where no stress shift has taken place, StressPath and FocusPath are true of the same subsets and any node contained by them may be focused. After an application of the Stress Shift rule, however, the two paths are no longer identical, although their intersection is never empty (it has to contain at least the root node). In this case, then, the only valid targets for focus are those nodes of the StressPath that aren't contained in the FocusPath. This is formally expressed by the  $\mathcal{L}_{K,P}^2$  sentence  $\phi$  below. Just like  $A(x)$ ,  $D(x)$  and  $S(x)$  before,  $F(x)$  is a predicate defining a particular set of nodes, this time the set of nodes labeled by some  $\omega^F \in \Omega_S^F$ . I furthermore use  $X \approx Y$  as a shorthand for  $\forall x[X(x) \leftrightarrow Y(x)]$ .

$$\phi := \forall x, X, Y[F(x) \wedge X(x) \wedge \text{StressPath}(X) \wedge \text{FocusPath}(Y) \rightarrow (Y(x) \rightarrow X \approx Y)]$$

Note that  $\phi$  by itself does not properly restrict the distribution of focus. First of all, there is no requirement that exactly one node must be focused. Second, nodes outside StressPath may carry focus, in which case no restrictions apply to them at all. Finally, StressPath and FocusPath may be empty, because we have not made any assumptions about the distribution of labels. As long as our monadic second-order theory is consistent, though, these cases of overgeneration are of no concern to us. Provided  $L_c$  isn't a superset of the candidate language, it can be used as intended to filter out trees with illicit foci, because  $\phi$  behaves as expected over the trees in the candidate language. As alluded to above, this filtering takes the form of taking the diagonal of  $L_c$  and composing it with GEN. Since both are linear transductions (recall that the diagonal over a regular language is a linear transduction), their composition is, too. Returning to the question of consistency, if we assume the general  $\mathcal{L}_{K,P}^2$  axioms from [14], the MSO-theory  $\{\phi\}$  is consistent and yields the desired  $L_c$ , as the reader may want to verify for themselves. This establishes the computational feasibility of Focus Economy when the input language is a regular tree language.

So far I have left open the question, though, how movement fits into the picture. First of all, it cannot be ruled out *a priori* that the interaction of movement and focus are so intricate on a linguistic level that significant modifications have to be made to the original version of Focus Economy. On a formal level, this would mean that the transduction itself would have to be changed. In this case, it makes little sense to speculate how my model could be extended to accommodate movement, so let's instead assume that Focus Economy can remain virtually unaltered and it is only the input language that has to be modified. In my model, the input language is a regular tree language by virtue of being generated by a Minimalist grammar without movement. But note that Minimalist grammars with movement generate regular tree languages, too, in the presence of a ban against more exotic kinds of movement such as remnant movement or head movement [10]. Now keep in mind that regular tree languages yield context-free string languages, which are generally assumed to be powerful enough for the greatest part of natural language syntax. Thus the restriction to regular tree languages itself does not preclude us from accommodating most instances of movement.

Things get more interesting, however, if we want the full expressive power of Minimalist grammars. Then the input language might not even be a linear context-free tree

grammar, so we aren't guaranteed that adding a reference-set constraint on top of the Minimalist grammar won't lead to an increase in generative capacity. At least for Focus Economy, though, closure under linear transductions may be more than we actually need. First, notice that the transducers the composition of which makes up GEN are very simple nondeterministic finite-state relabelings. Now we shouldn't expect the tree languages of Minimalist grammars to be closed under finite-state relabelings, despite their simplicity (the IO-context-free tree languages, for instance, aren't closed under these relabelings). However, when we consider the form of the trees in the range of GEN, it seems fairly unlikely that it couldn't be obtained directly by a Minimalist grammar. The only conditions are that there is a unique path of S-labeled nodes and that one node in this path is also marked for focus. Things are complicated slightly by the special status of adjuncts and anaphorically distressed nodes, but overall we are dealing with very simple conditions that a Minimalist grammar should have no problem with. The truly problematic question, then, is whether the tree languages of Minimalist grammars are closed under intersection with a regular language, i.e. whether we can apply the filtering step to tame the overgeneration inherent to GEN. To my knowledge, this is an open problem, although the answer might already be implicit in Kobele et al.'s recent automata-theoretic perspective on Minimalist grammars [11].

## Conclusion

I have shown that despite claims to the contrary, reference-set constraints aren't computationally intractable. Controlled OSs were introduced as a formal model for reference-set constraints. I gave a new characterization of global optimality for the subclass of output joint preserving OSs, which is general enough to accommodate most reference-set constraints. The shift in perspective induced by controlled OSs made it apparent that out of the five conditions that jointly guarantee for an OS to stay within the limits of linear tree transductions, three are almost trivially satisfied by reference-set constraints, with the only problematic areas being the power of GEN and the rankings induced by the constraints on the range of GEN. A model of a prominent reference-set constraint, Focus Economy, showed that at least for some constraints these points aren't problematic, either. These new results suggest that reference-set constraints are significantly better behaved than is usually believed.

## Acknowledgements

I am greatly indebted to Ed Stabler and Uwe Mönnich as well as the two anonymous reviewers for their motivational comments and helpful criticism. The research reported herein was supported by a DOC-fellowship of the Austrian Academy of Sciences.

## References

- [1] Chomsky, N.: The Minimalist Program. MIT Press, Cambridge, Mass. (1995)

- [2] Chomsky, N., Halle, M.: *The Sound Pattern of English*. Evanston, New York (1968)
- [3] Frank, R., Satta, G.: Optimality theory and the generative complexity of constraint violability. *Computational Linguistics* 24, 307–315 (1998)
- [4] Gécseg, F., Steinby, M.: *Tree Automata*. Akademiai Kiadó, Budapest (1984)
- [5] Hopcroft, J.E., Ullman, J.D.: *Introduction to Automata Theory, Languages, and Computation*. Addison Wesley, Reading, Mass. (1979)
- [6] Johnson, D., Lappin, S.: *Local Constraints vs. Economy*. CSLI, Stanford (1999)
- [7] Jäger, G.: Gradient constraints in finite state OT: The unidirectional and the bidirectional case. In: Kaufmann, I., Stiebels, B. (eds.) *More than Words. A Festschrift for Dieter Wunderlich*, pp. 299–325. Akademie Verlag, Berlin (2002)
- [8] Karttunen, L.: *The proper treatment of optimality in computational phonology* (1998), manuscript, Xerox Research Center Europe
- [9] Kepser, S., Mönnich, U.: Closure properties of linear context-free tree languages with an application to optimality theory. *Theoretical Computer Science* 354, 82–97 (2006)
- [10] Kobele, G.M.: Without remnant movement, MGs are context-free. In: Ebert, C., Jäger, G., Michaelis, J. (eds.) *MOL 10/11. Lecture Notes in Computer Science*, vol. 6149, pp. 160–173 (2010)
- [11] Kobele, G.M., Retoré, C., Salvati, S.: An automata-theoretic approach to minimalism. In: Rogers, J., Kepser, S. (eds.) *Model Theoretic Syntax at 10*. pp. 71–80 (2007), workshop organized as part of the European Summer School on Logic, Language and Information, ESSLLI 2007, 6-17 August 2007, Dublin, Ireland
- [12] Prince, A., Smolensky, P.: *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell, Oxford (2004)
- [13] Reinhart, T.: *Interface Strategies: Optimal and Costly Computations*. MIT Press, Cambridge, Mass. (2006)
- [14] Rogers, J.: *A Descriptive Approach to Language-Theoretic Complexity*. CSLI, Stanford (1998)
- [15] Shieber, S.M.: Unifying synchronous tree adjoining grammars and tree transducers via bimorphisms. In: *Proceedings of the 11<sup>th</sup> Conference of the European Chapter of the Association for Computational Linguistics (EACL-2006)*. pp. 377–384 (2006)
- [16] Shieber, S.M., Schabes, Y.: Synchronous tree adjoining grammars. In: *Proceedings of the 13<sup>th</sup> International Conference on Computational Linguistics*. pp. 253–258 (1990)
- [17] Stabler, E.P., Keenan, E.: Structural similarity. *Theoretical Computer Science* 293, 345–363 (2003)
- [18] Szendrői, K.: *Focus and the Syntax-Phonology Interface*. Ph.D. thesis, University College London (2001)
- [19] Thomas, W.: Languages, automata and logic. In: Rozenberg, G., Salomaa, A. (eds.) *Handbook of Formal Languages*, vol. 3, pp. 389–455. Springer, New York (1997)
- [20] Wartena, C.: A note on the complexity of optimality systems. In: Blutner, R., Jäger, G. (eds.) *Studies in Optimality Theory*, pp. 64–72. University of Potsdam, Potsdam, Germany (2000)