

Relativizations for the logic-automata connection*

Nils Klarlund
AT&T Labs–Research

*Dedicated to the memory of Bob Paige
and his contributions to automata algorithms*

August 28, 2003

Abstract

BDDs and their algorithms implement a decision procedure for Quantified Propositional Logic. BDDs are a kind of acyclic automata. But unrestricted automata (recognizing unbounded strings of bit vectors) can be used to decide monadic second-order logics, which are more expressive. Prime examples are WS1S, a number-theoretic logic, or the string-based logical notation of introductory texts. One problem is that it is not clear which one is to be preferred in practice. For example, it is not known whether these two logics are computationally equivalent to within a linear factor, that is, whether a formula ϕ of one logic can be transformed to a formula ϕ' of the other such that ϕ' is true if and only if ϕ is and such that ϕ' is decided in time linear in that of the time for ϕ .

Another problem is that first-order variables in either version are given automata-theoretic semantics according to relativizations, which are syntactic means of restricting the domain of quantification of a variable. Such relativizations lead to technical arbitrations that may involve normalizing each subformula in an asymmetric manner or may introduce spurious state space explosions.

In this paper, we investigate these problems through studies of congruences on strings. This algebraic framework is adapted to language-theoretic relativizations, where regular languages are intersected with restrictions. The restrictions are also regular languages. We introduce ternary and sextartite characterizations of relativized regular languages. From properties of the resulting congruences, we are able to carry out detailed state space analyses that allow us to address the two problems.

We report briefly on practical experiments that support our results. We conclude that WS1S with first-order variables can be robustly implemented in a way that efficiently subsumes string-based notations.

*Some of the material in this paper appeared in *Computer Aided Verification, CAV '99*, LNCS 1633, 1999, under the title "A theory of restrictions for logics and automata."

1 Motivation

The relationship between automata and logic has been very successfully exploited through Binary Decision Diagrams[4]. This technique allows formulas of propositional logic to be decided through the use of automata representations for sets of strings of bounded length. But, a more general *logic-automata* connection exists: Büchi[5], Elgot[6], and Trakhtenbrot[17] argued forty years ago that a logical notation, now called the Weak Second-order theory of 1 Successor or WS1S, would be a more natural alternative to what already was known as regular expressions. WS1S has an extremely simple syntax and semantics: it is a variation of predicate logic with first-order variables that denote natural numbers and second-order variables that denote finite sets of natural numbers; it has a single function symbol, which denotes the successor function, and has usual binary operators such as \leq , $=$, \in and \supseteq .

Büchi, Elgot, and Trakhtenbrot showed that a decision procedure exists for this logic. The idea is to view interpretations as finite strings over bit vectors and then to show by explicit constructions of automata that the set of satisfying interpretations for any subformula is a regular language. In this way, an automaton becomes an object that represents the logic semantics of a formula, and it makes sense to talk about automata-theoretic semantics, which characterize the computational approach to the logic. As with Binary Decision Diagrams, the idea behind the decision procedure is to construct inductively a deterministic automaton for each subformula. This method, which we shall review in detail, handles each connective in the logic through an automata-theoretic operation, such as product or subset construction. To *decide* a sentence is then under this view the process of building the automaton inductively.

A main motivation of this article is to make the decision procedure feasible in practice. Since 1994, we have explored the practicality of the logic-automata connection in the *Mona* project, first described in [7]. The *Mona* tool has been used for a variety of tasks, for example in linguistics[18], pointer verification[13], protocol verification[14], and hardware verification[2]. Among other implementation challenges[12], we discovered spurious state space explosions in intermediate automata. Sometimes, we would be prevented from solving even trivial problems due to phenomena that we will uncover and overcome in the present article.

1.1 The string-theoretic formulation and relativizations

The problems we encountered are in part linked to the existence of two approaches to monadic second-order logic.

The logic WS1S, the first approach, is a very natural notation that satisfies elementary expectations. For example, a sentence in WS1S is either true or false and a formula with k free first-order variables defines a relation that is simply a subset of \mathbb{N}^k , where \mathbb{N} is the set of non-negative integers. Thus, we call this approach to the logic-automata connection *number-theoretic*. The automaton for a sentence is a simple one: it has one state (if minimized)! If it is accepting, the formula is true; if rejecting, the formula is false.

The second approach, the one emphasized in presentations of the logic-automata connection (such as in [15, 16]) is more complicated to explain, at least when it comes

to the semantics, which are tied to a parameterized, finite, unbounded domain represented by a number $n \geq 0$. This number defines a set $\{0, \dots, n - 1\}$ of *positions*. Under this view, the truth status of a sentence depends on n . For example, a sentence may be defined that is true if and only if n is even. A formula with free first-order variables now defines a parameterized family of finite relations, where the n th relation is a subset of $\{0, \dots, n - 1\}^k$. We call this view the *string-theoretic* approach, since the semantics can be explained advantageously in terms of strings. The ideas of the automaton-theoretic decision procedure of WS1S still apply; in fact, the algorithm becomes much simpler, which is the reason that this approach is often preferred in introductions to the logic-automata connection. The approach is also intrinsically appealing for certain applications, for example in the description and verification of parameterized hardware[2]. Among other names, these logics have been called MSO(S)[16], SOM[+][15], and M2L(Str)[7, 9]. They vary slightly, but we will identify them as M2L(Str) in this paper.

There are two reasons for preferring the number-theoretic approach. First, its logical semantics are simpler as just argued—it can be explained to people with little mathematical background. Second, WS1S appears to be the stronger logic in the following sense: there is apparently no known polynomial time reduction f from sentences interpreted under the WS1S view to sentences interpreted under the M2L(Str) view such that $f(\phi)$ is true for all n if and only if ϕ is true. In contrast, there is a rather obvious, linear translation in the other direction: given a sentence ϕ under the string-view, we turn it into a formula ϕ' with one free variable $\$$ such that ϕ holds for n if and only if ϕ' holds under the interpretation $\$ \mapsto n$. Specifically, ϕ' is obtained from ϕ by restricting quantified variables to the domain $\{0, \dots, \$ - 1\}$. Syntactically, the restriction to $\{0, \dots, \$ - 1\}$ for a variable can be expressed as a formula that is conjoined to the formula of the existential quantifier introducing it. Such syntactic constructs are well-known in logic; they are called *relativizations*. (As usually formulated, a relativization of the formula imply that every variable quantified is restricted, see[3]. In our case, we relativize only the second-order variables that represent first-order variables.) We call the formula ϕ' a *WS1S representation* of the M2L(Str) formula ϕ .

A main focus of the present work is to show how relativizations can be guaranteed to work in practice. In fact, even though the syntactic translation is linear, there is no guarantee that the computations involved in calculating the automaton for ϕ' (under the WS1S view) are not asymptotically more involved than those involved in ϕ (under the M2L(Str) view).

Indeed, during early experiments with this procedure, our problem was that seemingly innocuous formulas would yield enormous automata after the conversion into WS1S. So, we say that an *efficient translation algorithm* is one that in linear time transforms any sentence ϕ in M2L(Str) to a representation ϕ' such that ϕ' is decided in time that is linear in the time to decide ϕ . Let us call the question of finding such an algorithm the *translation problem*.

1.2 Handling of first-order variables

Another computational problem we encountered with monadic second-order logics stems from the way that first-order variables and terms are handled. Through for-

mula rewritings, they are transformed into second-order variables. The second-order variables are subjected to relativizations that restrict them to singleton sets. Consequently, automata corresponding to subformulas are not simply determined by the logical semantics, but also according to how relativizations are formulated. So, to make automata for formulas canonically determined, extra automata product operations are used to *conjunctively normalize* these intermediate automata as we shall see. (The canonicity of representations is essential to the success of automata-based methods such as BDDs—it guarantees that intermediate results are always pruned to their minimum size.) The *first-order semantics problem* is to find an automaton representation that is no bigger than the conjunctively normalized representation, while not requiring such explicit normalization steps. This is important in practice, since we want to minimize the amount of computational work.

We note that there are other ways of deciding WS1S, for example through Ehrenfeucht-Fraïssé games[15] or through bounded-model techniques[1].

1.3 Contributions of this paper

In this paper, we propose solutions to the translation problem and the first-order semantics problem. We do so by studying relativizations in an algebraic framework. We proceed as follows.

We formulate a syntax for WS1S, where relativizations are made explicit, and we provide initially three different automata-theoretic semantics: (1) the *ad hoc* semantics that correspond to the strategy we first used in the Mona implementation for first-order variables, (2) the *conjunctive normalized semantic*, where *all* the intermediate automata are conjoined with relativizations, and (3) the *ternary semantics*, which are based on valuations that identifies the membership status (0 or 1) for each string in the restriction and assigns \perp to each string not in the restriction. We explain why the ad hoc semantics are unsuitable, and why the conjunctive normalized semantics, in addition to being asymmetric, would slow down the decision procedure. We show that the ternary semantics make most normalizations unnecessary, since they inherently propagate through the automata-theoretic constructions. Also, we indicate how the ternary semantics can be implemented based on the standard WS1S decision procedure.

To study the question of automata sizes, we give a detailed congruence-theoretic analysis of regular languages under relativizations, that is, under intersections with other regular sets that act as restrictions. We introduce a notion of a *thin* language, and we show that the relativizations occurring in the treatment of first-order variables and in the translation problem are thin. We prove that languages under thin relativizations make comparisons of the conjunctive normalized semantics and the ternary semantics easy: the latter are the same as the former except for some extra equivalence classes that we characterize. We show that if the automata of restrictions are bounded, then the sizes of intermediate automata occurring under the ternary semantics are to within this additive bound the same as the sizes of automata of the conjunctive normalized semantics.

We strengthen this result by exhibiting congruences based on a *sexpartite semantics* that are no bigger than those of the conjunctive normalized semantics. The sexpartite valuations are based on the ternary ones, but for certain strings they do not yield an

exact answer to what the ternary evaluation is. We are able to show that operations corresponding to logical connectives may be formulated directly on automata representing these congruences, under the further assumption of *crispness*, a language-theoretic property that the restrictions of the first-order semantics problem and the translation problem enjoy.

Our main result is that the resulting decision procedure, while requiring only few normalizations, involves intermediate automata that are at most the same in size (to within a linear factor) as the ones occurring under the conjunctive normalized semantics, but sometimes only logarithmic in size compared to the conjunctive normalized semantics. Thus, we have found a symmetric and efficient representation of formulas under restrictions.

We conclude that WS1S, and not a string-oriented logic, is the best interface to the logic-automata connection, since in practice the string-theoretic view is effectively subsumed by the number-theoretic view through the techniques developed in this article.

1.4 Organization

In Section 2, we review WS1S and its decision procedure. We provide further motivation for why normalizations are necessary, both in the case of first-order variables (Section 2.3) and in the case of the translation problem (Section 2.4). For the latter, we discuss an example in detail.

In Section 3, we formalize the classic semantics, the conjunctive normalized one, and the ternary one. We also discuss the relationship between a syntax that explicitly accommodates relativizations and the automata-theoretic semantics.

In Section 4, we develop an understanding of restrictions imposed on regular languages through ternary valuations. In particular, Theorem 1 relates sizes of automata under the conjunctive normalized semantics and under the ternary semantics.

In Section 5, we present state space engineering techniques that remove information from ternary valuations. The resulting sexpartite valuations are specialized ternary valuations. We discuss ways of calculating sexpartite valuations and how to reason about them.

In Section 6, we study the problem of calculating the sexpartite representation of intersections directly from the sexpartite representation of the sets involved.

In Section 7, we show how the sexpartite semantics of WS1S can be reformulated in ways that will support better algorithms; in particular, we try to avoid normalizations.

In Section 8, we present the algorithms that may be used in a decision procedure for WS1S based on techniques from the Sections 6 and 7. In particular, we formulate algorithms that allow us to formulate our second theorem: WS1S under restrictions of the kinds we are interested in can be decided in a way that is up to exponentially faster than using the conjunctive normalization technique.

In Section 9, we summarize and provide some hints about the practical performance of the Mona tool when equipped with techniques of this article.

2 WS1S: review and issues

We need to fix a syntax for WS1S. We follow the concrete, ASCII-based syntax of Mona, but we keep only a small set of primitives. *Nutshell WS1S* can be presented as follows. A formula ϕ is *composite* and of the form $\sim \phi'$, $\phi' \ \& \ \phi''$, or $\text{ex2 } P^i : \phi'$, or it is *atomic* and of the form $P^i \text{ sub } P^j$, $P^i < P^j$, $P^i = P^j \setminus P^k$, or $P^i = P^j + 1$. Here, we have assumed that variables are all second-order and named P^i , where $i \geq 1$. Other comparison operators, second-order terms with set-theoretic operators, and Boolean connectives can be introduced by trivial syntactic abbreviations, see [11, 16]. The treatment of first-order terms is discussed later.

2.1 Logic Semantics of WS1S

A decision procedure takes as input a formula ϕ_0 , called the main formula, whose truth status is to be investigated. Following standard practice, we sometimes regard the main formula as an abstract syntax or parse tree (with its root facing up—the usual convention). We define its *logic semantics* (or just *semantics*) inductively relative to a string w over the alphabet $\Sigma = \mathbb{B}^k$, where $\mathbb{B} = \{0, 1\}$ and k is the number of variables in ϕ_0 . We assume that ϕ_0 is closed and that each variable is bound in at most one occurrence of an existential quantifier. Generally, we consider only formulas that are subformulas of ϕ_0 , since the semantics are formulated in terms of strings over a bounded alphabet—something that prevents us from giving semantics to all possible formulas given a value of k . We now regard a string $w = a_0 \cdots a_{\ell-1}$, where $\ell = |w|$ is the length of w , to be of the form:

$$\begin{array}{l} P^1 \\ \dots \\ P^k \end{array} \quad \left(\begin{array}{c} a_0^1 \\ \dots \\ a_0^k \end{array} \right) \cdots \left(\begin{array}{c} a_{\ell-1}^1 \\ \dots \\ a_{\ell-1}^k \end{array} \right)$$

where we have indicated (left) that if the string is viewed as a matrix, then row i is called the P^i -*track*. Each letter a is sometimes written in a transposed notation as the vector $(a^1, \dots, a^k)^t$. The interpretation $w(P^i) \subseteq \mathbb{N}$ of P^i defined by w is the finite set $\{m \mid \text{the } m\text{th bit in the } P^i\text{-track is } 1\}$. Note that suffixing w with a *null-extension*, a string of the form $\mathbf{0} \cdots \mathbf{0}$ with $\mathbf{0} = (0, \dots, 0)^t$, does not change the interpretation of any variable.

The semantics of a formula ϕ can now be defined inductively relative to an interpretation w . We use the notation $w \models \phi$ (which is read: w satisfies ϕ) if the interpretation defined by w makes ϕ true:

$$\begin{array}{ll} w \models \sim \phi' & \text{iff } w \not\models \phi' \\ w \models \phi' \ \& \ \phi'' & \text{iff } w \models \phi' \ \text{and } w \models \phi'' \\ w \models \text{ex2 } P^i : \phi' & \text{iff } \exists \text{ finite } M \subseteq \mathbb{N} : w[P^i \mapsto M] \models \phi' \\ w \models P^i \text{ sub } P^j & \text{iff } w(P^i) \subseteq w(P^j) \\ w \models P^i < P^j & \text{iff } \forall h \in w(P^i) : \forall k \in w(P^j) : h < k \\ w \models P^i = P^j \setminus P^k & \text{iff } w(P^i) = w(P^j) \setminus w(P^k) \\ w \models P^i = P^j + 1 & \text{iff } w(P^i) = \{m + 1 \mid m \in w(P^j)\} \end{array}$$

where we use the notation $w[P^i \mapsto M]$ for the minimal string w' , called the *witness string*, that is at least as long as w and interprets all variables P^j , $j \neq i$, as w does, but interprets P^i as M . Note that the truth status of a subformula ϕ that contains an existential formula of the form $w \models \text{ex}2 P^i : \phi'$ does not depend on the P^i -track, since we assume that each P^i is bound by at most one existential quantifier, and consequently there can be no free occurrence of P^i in ϕ . We could have chosen a model, where these unnecessary tracks are removed—but that gives us the complication of working with different alphabets for each subformula.

To fully explain the decision procedure for WS1S, we need to look at the situation for quantification in more detail. This discussion will be rather technical; it is the reason why introductory texts concentrate on M2L(Str) for which existential quantification is straightforward. Looking at the witness string $w[P^i \mapsto M]$ again, we observe that it may be longer than w , since the greatest element in M may be equal to or greater than the length of w . In this case, the extension consists of all zeros except for the P^i -track, which describes the elements of M greater than $w - 1$.

To characterize such extensions more formally, we let Σ_0^i be the subset of all letters of the form $(0, \dots, 0, X, 0, \dots, 0)^t$ (where the X means that the value of the i th component is either 0 for 1). An extension z^i in Σ_0^{i*} is then called a *null-but- i extension*. Thus, the witness string is of the form $w' \cdot z^i$, where w' is the same string as w except possibly for the P^i -track.

Generally, we may decompose any w into parts \tilde{w} and z^i such that z^i contains as many letters from the end of the string as possible that are all 0 outside the P^i -track; more precisely, we let $w = \tilde{w} \cdot z^i$, where z^i is a maximal, null-but- i extension of some prefix of w . Thus, \tilde{w} is empty or at least one non- P^i -track in \tilde{w} ends with a 1, i.e. the last letter in \tilde{w} is not in Σ_0^i . We say that \tilde{w} is *but- i -minimal*.

As defined, the witness string may not be shorter than w . For example, if w is all zeros outside the P^i -track, with the P^i -track consisting of all ones and $M = \emptyset$, then the witness string is $\mathbf{0} \cdots \mathbf{0}$ of length $|w|$. We have introduced this requirement to avoid other unpleasant technicalities later.

For any formula ϕ , we associate the *language* $L_\phi = \{w \mid w \models \phi\}$. The interpretation of ϕ_0 is independent of w , since by assumption it is a closed formula. Thus, ϕ_0 is either true, when $L_\phi = \Sigma^*$, and we write $\models \phi_0$, or ϕ_0 is false, when $L_\phi = \emptyset$, and we write $\not\models \phi_0$.

Proposition 1 Words w and w' that interpret all variables in the same way satisfy the same set of formulas: if for all i , $1 \leq i \leq k$, $w(P^i) = w'(P^i)$, then for all ϕ , $w \models \phi$ if and only if $w' \models \phi$. Thus, the interpretation is invariant under null-extensions.

Proof (Idea) By a simple induction on formulas ϕ , we may show that for words w and w' such that w' is a null-extension of w the following holds: either both satisfy ϕ or both do not satisfy ϕ . For existential quantification, we note that if w' is a null-extension of w , then $w'[P^i \mapsto M]$ is a null-extension of $w[P^i \mapsto M]$. \square




Figure 1: Automaton that accepts satisfying interpretations for $P^1 \text{sub} P^2$.

2.2 Automata-theoretic semantics

The automata-theoretic semantics define a decision procedure that associates to each ϕ the deterministic, minimal automaton A_ϕ accepting the language L_ϕ .

2.2.1 Automata preliminaries

We recall that an automaton $A = (\Sigma, Q, Q_0, \rightarrow, Q^F)$ consists of an alphabet Σ , which we here assume to be \mathbb{B}^k , a finite set of states Q , a set of initial states Q_0 , a transition relation $\rightarrow \subseteq Q \times \Sigma \times Q$, and a set of final states Q^F . A *run* $(q_m)_{m \leq \ell}$ over a word $w = a_0 \cdots a_{\ell-1} \in \Sigma^*$ is a sequence of states q_0, \dots, q_ℓ such that $q_0 \in Q_0$ and for all m , $0 \leq m < \ell$, $(q_m, a_m, q_{m+1}) \in \rightarrow$. The *length* of the run is ℓ . The run is *accepting* if $q_\ell \in Q^F$, and the language accepted by A is the set of w that allows some accepting run. The automaton A is deterministic if Q_0 is a singleton and if for all $q \in Q$ and $a \in \Sigma$ there is exactly one q' such that $(q, a, q') \in \rightarrow$. Our automata are assumed deterministic if not otherwise indicated. For a deterministic automaton, any word w allows exactly one run; if its length is ℓ , then last state q_ℓ of this run is denoted $\text{LAST } A(w)$ and we say that w *brings* A to state q_ℓ .

The *size* $|A|$ of an automaton is its number of states.

2.2.2 Constructing A_ϕ

For atomic formulas ϕ , a small, deterministic, minimal automaton can be directly constructed that accepts the language L_ϕ . For example, for the formula $P^1 \text{sub} P^2$, a two-state automaton exists that accepts exactly the set of w for which $w(P^1) \subseteq w(P^2)$; see Figure 1, where we have followed usual conventions: states are circles, the initial state is denoted by an arrow pointing to it, transitions are denoted by arrows marked with letters for which they apply, and the final states are designated by an inner circle.

Other atomic formulas are treated similarly, and for composite formulas we proceed by induction.

For a formula ϕ of the form $\sim \phi'$, the automaton A_ϕ is taken to be the complement of the automaton $A_{\phi'}$ calculated by induction. This automaton will be minimal by construction since it is obtained by simply reversing final and non-final states in the automaton $A_{\phi'}$. The case of conjunction is handled by an automata-theoretic product construction: given $A_{\phi'}$ accepting $L_{\phi'}$ and given $A_{\phi''}$ accepting $L_{\phi''}$, we construct the

minimized product automaton of $A_{\phi'}$ and $A_{\phi''}$; this automaton accepts the language $L_{\phi'} \cap L_{\phi''}$.

The case of quantification is more complicated. Consider $\phi = \text{ex}2 P^i : \phi'$. Define the projection operator PROJ^i so that

$$\text{PROJ}^i A(u) = \begin{cases} 1 & \text{if } \exists M : A(u[P^i \mapsto M]) = 1 \\ 0 & \text{if } \forall M : A(u[P^i \mapsto M]) = 0 \end{cases} \quad (1)$$

Then, A_ϕ is the automaton $\text{PROJ}^i A_{\phi'}$. The quantification over unbounded (but finite) sets seems to contradict our use of finite strings. We describe how to calculate PROJ^i in two steps. In the first step, knowledge about possible future extensions that only interpret the P^i -track is used to change the labeling of the $A_{\phi'}$ automaton. In the second step, a conventional projection construction is carried out.

From an automaton A , we construct the *futurization* of A , denoted $\text{FUT}^i A$ according to null-but- i -extensions:

$$\text{FUT}^i A(u) = \begin{cases} 1 & \text{if } \exists v \in \Sigma_0^{i*} : A(u \cdot v) = 1 \\ 0 & \text{if } \forall v \in \Sigma_0^{i*} : A(u \cdot v) \neq 1 \end{cases}$$

This description can clearly be implemented as a linear time automata algorithm by changing the labeling of the states of A appropriately.

We define a substitution operator $[P^i := M]$ that overwrites the P^i track without changing the length of the string: for any finite set M with $\max M < |u|$, the string $v = u[P^i := M]$ interprets all P^j , $i \neq j$, the same way as u , but $v(P^i)$ is M . Define the *bounded projection* operator BPROJ^i such that

$$\text{BPROJ}^i A(u) = \begin{cases} 1 & \text{if } \exists M \text{ with } \max M < |u| : A(u[P^i := M]) = 1 \\ 0 & \text{if } \forall M \text{ with } \max M < |u| : A(u[P^i := M]) = 0 \end{cases}$$

This operator describes the traditional projection operation on the P^i track. It can be formulated as an algorithm that yields a deterministic automaton via the subset construction. The algorithm may run in exponential time, since the resulting automaton may be exponentially bigger.

Proposition 2

$$\text{PROJ}^i A = \text{BPROJ}^i(\text{FUT}^i A)$$

Proof We have

$$\begin{aligned} \text{PROJ}^i A(u) = 1 & \text{ iff} \\ \exists M : A(u[P^i \mapsto M]) = 1 & \text{ iff} \\ \exists \hat{M}, z^i \text{ with } \max \hat{M} < |u| \text{ and } z^i \in \Sigma_0^{i*} : A(u[P^i := \hat{M}] \cdot z^i) = 1 & \text{ iff} \\ \text{BPROJ}^i(\text{FUT}^i A)(u) = 1 & \end{aligned}$$

where the second biimplication is valid because $|u[P^i \mapsto M]| \geq |u|$ holds thanks to the definition of $u[P^i \mapsto M]$. \square

2.3 Automata-theoretic semantics of first-order variables

Let us look at the first-order semantics problem. Adding first-order variables to nutshell WS1S can easily be done as follows: a first-order variable p is regarded as a second-order term P that is restricted to take on values that are singleton sets, whose sole element denotes the value of p , see [10, 15, 16]. This relativization is imposed syntactically by conjoining a singleton predicate $\text{singleton}(P)$ to the formula where P is quantified; the symbol P is a meta-variable that stands for one of the P^i . The singleton predicate can be expressed as a formula in nutshell logic for any particular P^i . Thus, we view it as a macro, a textual expansion mechanism, not an additional primitive predicate that must be interpreted.

This *classic relativization strategy* means that the automata-semantics of a formula containing p are not entirely robust. The meaning of interpretations w not fulfilling $\text{singleton}(P)$ is not well-defined for subformulas in the parse tree below the point of relativization. We will illustrate this phenomenon with an example that assumes that we have already added the ability to express both first-order and second-order constants to the nutshell language. The formula $\phi = p=0$, where p is first-order, is then reasonably represented as $\{0\} \text{sub } P$, since this formula has the right truth value relative to singleton interpretations of P . Similarly, we would reasonably translate the formula $\phi' = p < 1$ into $P < \{1\}$. But $\{0\} \text{sub } P$ and $P < \{1\}$ are not equivalent, except when restricted to singleton interpretations. This phenomenon leads to the following problem: automata corresponding to subformulas may have many states that describe spurious and irrelevant truth functions outside the restrictions. Of course, these states corresponding to interpretations outside the restriction are eventually pruned thanks to the conjunction of the restriction in the relativized quantified formula.

In order to fix the automata semantics of subformulas, we could arbitrate as follows: conjoin the restriction to every subformula ϕ in a procedure we call *normalization*. Then, we would have a firm automata-theoretic explanation of the language $L(\phi)$ under what we call the *normalize strategy*. Unfortunately, this solution is an asymmetric one.

The practical problem with a conjunctively normalized semantics is that additional product and minimization calculations would be necessary: for each automaton A representing a subformula ϕ and each free variable P^i , the automaton representing the singleton property for P^i must be conjoined to A . Such extra calculational work slows down the decision procedure in the following sense: each automata-theoretic operation must be followed by a product and a minimization. (Under certain assumptions on restrictions, the need for minimization may disappear, see Section 4.) For example, complementation, which is normally very fast since it consists of flipping acceptance statuses of states, now would involve a product and a minimization operation.

In practice, the Mona implementation prior to the one implemented with the results of the present article used the *ad hoc strategy*: the restriction for variable p is conjoined only to atomic formulas where p occur and to the formula in the existential quantification introducing p . Note that this technique does not eliminate the problem of spurious behavior for intermediate formulas. For example, the conjunctively normalized atomic formula $p=0$ is not equivalent to the negation of the conjunctively normalized atomic formula $p\sim=0$, where $\sim =$ means “not equal.”

2.4 Emulation of string semantics in WS1S

We turn to the translation problem of how to use restrictions to efficiently translate the string-theoretic version of monadic second-order logics into a number-theoretic version. A simple choice of syntax for M2L(Str) is to make it identical to nutshell WS1S syntax. The satisfaction relation is now denoted \models_{string} ; it is the same as for WS1S except that quantification is changed to:

$$w \models_{string} \text{ex2 } P^i : \phi' \quad \text{iff} \quad \exists M \subseteq \{0, \dots, |w| - 1\} : w[P^i \mapsto M] \models \phi'$$

where the notation $w[P^i \mapsto M]$ now has a different meaning: it denotes the string w altered so that the P^i track describes M . Thus, the witness string $w[P^i \mapsto M]$ for the existential quantification has the same length as w . The interpretation of ϕ_0 on a string of w still does not depend on the individual tracks of w , but it *does* depend on the length of w . Thus we write $i \models_{string} \phi_0$ if ϕ_0 holds for a string w of length i . For example, a closed formula can be written that under these semantics holds if and only if w is of even length.

To emulate \models_{string} in \models , we must relativize all second-order terms to sets of numbers less than the last position in the string. Thus, we introduce a first-order variable $\$$ that simulates the entity $|w|$. (Of course, $\$$ really stands for some P_i variable that is otherwise unused and that is relativized to act like a first-order variable.) A $\$$ -restriction for a variable expresses that the variable is a subset of $\{0, \dots, \$ - 1\}$. Then, under the normalization strategy we conjoin $\$$ -constraints for all free variables of each subformula. The result is a WS1S formula ϕ' with one free variable $\$$ such that $m \models_{string} \phi \Leftrightarrow w \models \phi'$, where the $\$$ -track of w interprets $\$$ as m . For example, the formula $\text{ex1 } p : \text{ex1 } q : p = q$ becomes in WS1S:

$$\begin{aligned} & [\text{singleton}(\$) \& \\ & \text{ex2 } P : \text{ex2 } Q : \\ & \quad [\text{singleton}(P) \& \text{singleton}(Q) \& \text{singleton}(\$) \\ & \quad \& P < \$ \& Q < \$ \& P \text{sub } Q \& Q \text{sub } P]] \end{aligned}$$

as expressed in nutshell syntax, where each normalized subformula is enclosed in brackets. The M2L(Str) formulation is

$$\begin{aligned} & [\text{ex2 } P : \text{ex2 } Q : \\ & \quad [\text{singleton}(P) \& \text{singleton}(Q) \& \\ & \quad P \text{sub } Q \& Q \text{sub } P]] \end{aligned}$$

Proposition 3 Under the translation outlined above, the minimized, canonical automata arising during the M2L(Str) decision procedure are essentially the same as the ones arising during the WS1S procedure except for at most two additional states.

Proof (Sketch) First, we must establish the relationship between the meaning of a formula ϕ in M2L(Str) and the meaning of $\bar{\phi}$ in WS1S, where $\bar{\phi}$ is obtained by conjoining the $\$$ -restriction for each variable occurring free in ϕ . If there are n variables in ϕ , then the $\$$ -variable receives index $n + 1$. Let \bar{w} be a string interpreting these $n + 1$

variables. An inductive argument shows that $\bar{w} \models \bar{\phi}$ if and only if (1) the P^{n+1} -track is interpreted as a singleton $\{\ell\}$ and (2) for each free variable P^i in ϕ the P^i -track is interpreted as a set of numbers less than ℓ .

Second, we can use this knowledge to construct the WS1S automaton for $\bar{\phi}$ from the M2L(Str) automaton ϕ by adding states s_{accept} (an accepting state) and s_{reject} (a rejecting state). The transition relation of the new automaton is the same as for the old one as long as the additional P^{n+1} -component is 0. All old states are turned into rejecting states. When the $\$$ -component is 1, corresponding to the end of the string under the M2L(Str) representation, a transition is made to s_{accept} or s_{reject} according to the accept status of the state that would have been reached in the old automaton, provided that the P^i -component of all free variables is 0 (if the latter is not true, then a transition is made to s_{reject}). From s_{accept} , a transition is made to s_{reject} if any 1 occurs in the $\$$ -component (since $\$$ must be interpreted as a singleton set) or in the track corresponding to a free variable in ϕ . The s_{reject} state is connected to itself on all letters.

Finally, it can be shown that during minimization of the new automaton, every pair of any old states are still not equivalent with respect to the canonical equivalence relation: the transitions to the two new states induce the same partition of the old states, regarded as part of the new automaton, as the one defined by accepting or rejecting states of the old automaton. \square

Our practical experiments with running string-based examples translated into WS1S were based on the ad hoc strategy, where the restriction of variable is conjoined only to atomic formulas involving the variable and to the place where the variable is introduced by quantifier.

The problem that bloated automata may occur thanks to the non-robustness of the ad hoc strategy is not just a theoretical one. We discovered the following problem that was serious enough to prevent benign formulas from being decided.

Parity example Consider the formula $\phi_{\text{ODD}} = \text{ex1 } p : (p \text{ in } P^1 \oplus \dots \oplus p \text{ in } P^n)$ under the string-theoretic semantics, where \oplus denotes addition modulo 2 (properly formulated in nutshell syntax). The formula holds if and only if there is a position contained in an odd number of the sets P^i . Translated into nutshell WS1S under the ad hoc strategy, the formula becomes:

$$\begin{aligned}
& \text{singleton}(\$) \& \\
& P^1 < \$ \& \\
& \dots \\
& P^n < \$ \& \\
\text{ex2 } P : & (\text{singleton}(\$) \& \\
& ((P \text{ sub } P^1 \& \text{singleton}(P) \& \text{singleton}(\$) \& P^1 < \$) \\
& \oplus((\dots) \\
& \oplus (P \text{ sub } P^n \& \text{singleton}(P) \& \text{singleton}(\$) \& P^n < \$) \dots).
\end{aligned} \tag{2}$$

where \oplus is a binary operator defined in terms of $\&$ and \sim .

Proposition 4 The parity formula ϕ_{ODD} expressed as (2) produces intermediate automata whose size is doubly exponential in n when constructed according to Section 2.2.2. But if the $n + 2$ restrictions are conjoined to all subformula (that is, also to each intermediate \oplus formula), then all intermediate automata have at most 21 states.

Proof (Sketch) Initially, let us discuss the size of the minimal automaton for ϕ_{ODD} . The alphabet is \mathbb{B}^{n+2} , since the formula contains n variables P_i and the variables $\$$ and p , whose real names in nutshell logic are, say, P_{n+1} and P_{n+2} . But since p is quantified away, the p -track does not influence the way the minimal automaton works. So, we will regard the automaton as reading letters that are vectors of size $n + 1$. The automaton must check for each letter whether the number of ones among the first n tracks is odd. The automaton for $n = 2$ is shown in Figure 2 as generated by Mona and the Graphviz drawing program (the state labeled 0 is an artifact of Mona’s representation of Boolean variables—it can be ignored). Intuitively, the automata can be explained as follows, where states are named as in Figure 2. The initial state is 1, where the automaton can stay until it finds a letter with an odd number of ones among tracks P_i with $i \leq n$. On such a letter it proceeds to state 3 recording this fact, and it stays there. So far, we have assumed that the end of the emulated string has not been reached, that is, a 1 in the P_{n+1} -track has not yet occurred. When this 1 occurs, the automaton proceeds to either a rejecting state 2 from which it cannot escape or an accepting state 4—according to which of the two previously mentioned states it was in. The automaton leaves accepting state 4 if another occurrence of a 1 happens in track P_{n+1} , since that violates the constraint on $\$$ as a first-order variable.




Figure 2: Automaton emulating string semantics for $n = 2$ of parity example.

Now consider the subformula inside the quantifier $\text{ex}2 P$ in (2). It has P as a free variable. The translation of this formula results in a minimal automaton whose size is exponential in n . Intuitively, this explosion stems from the need to record the status of whether position m is in P^i , for each $1 \leq i \leq n$, where m is the first position in P_{n+1} —if this information is not recorded, then it is impossible for the automaton to figure out the eventual truth value of the i th summand. In fact, when $m \in P^i$, the truth value of the i th summand still becomes false if the P^i -track contains any 1s further out than the number designated by $\$$, because of the last conjoint in each summand.

Moreover, it can be seen that the subset construction applied in connection with eliminating the P -variable will yield a further exponential blow-up. Intuitively, what happens is that the subset construction results in automaton that records the set of all vectors seen, since it must guess the value of m .

While we do not present a formal proof, the experimental behavior supports the intuition just given: for $n = 1, 2$ and 3 , the size of the automaton corresponding to the subformula $\text{ex}2 P$ is 21, 265, and 65553, approximately 2^{2^2} , 2^{2^3} , and 2^{2^4} . It is not possible to calculate the automaton for $n = 5$.

Finally, we have experimentally found that for $n = 2, \dots, 10$ the maximum number of states occurring in any intermediate automaton is 21 or less under a normalizing semantics (with no increase from $n = 4$ to 10). These experiments are done based on a ternary logic; for the binary semantics, the maximum number of states will possibly be less (according to results later this in article.) The automata still grow linearly in size of n since the transitions, represented by BDDs in the Mona tool, become more complicated. We leave it to the reader to reflect on why the number of states is limited by a constant. \square

3 WS1S with relativizations

To give a deeper understanding of relativizations, we introduce *nutshell WS1S-R*, a variation on WS1S where relativizations are explicitly marked. Let ρ be a formula that is the *restriction* of variable P^i . Existential quantification will now take the form

$$\text{ex}2 P^i \text{ where } \rho: \phi', \tag{3}$$

where ϕ' is the *traditional part* of the quantified formula. The restriction ρ can be an arbitrary formula (as long as the main formula ϕ_0 remains closed). In general, we denote by $\rho(P^i)$ the formula ρ introduced by the existential quantification of P^i . For uniformity, we assume that each P^i is relativized in this way, possibly to the formula $P^i = P^i$, which is another way of saying true . Our goal in this section is to show how a ternary semantics allow restrictions to bubble up when needed—eliminating the need for normalization at every intermediate step.

To carry out inductive arguments, we define the partial ordering \triangleleft among subformulas as follows: $\phi \triangleleft \phi'$ if (1) ϕ is a proper subformula of ϕ' or (2) if there is a formula $\psi = \text{ex}2 P^i \text{ where } \rho(P^i): \phi''$ such that ϕ is a subformula of $\rho(P^i)$ and ϕ' is a subformula of ϕ'' . This definition is illustrated in Figure 3, where the root of the subtree for Case (2) is the node for the formula ψ , which has two children, the left one for the restriction and the right one for the traditional part. The partial ordering \triangleleft is well-founded: a post-order labeling of nodes with numbers $0, 1, \dots$ produces an ordering, where all children of a node are assigned a number less than that of the parent (making the labeling consistent with case (1)) and where any node that is a left descendant of some node is assigned a number less than a right descendant (making the labeling consistent with case (2)).




Figure 3: The two ways for $\phi \triangleleft \phi'$ to hold.

When we formulate semantics, variable occurrences in the traditional part of formulas will be subjected to restrictions. More precisely, a *restrained variable occurrence* of P^i is one inside the traditional part ϕ' of the formula (3) introducing P^i . For a formula ϕ , the set of *restrained direct variables* $\text{DRV}(\phi)$ is the set of variables who have restrained occurrences in ϕ .

Proposition 5 For each ϕ and each $P \in \text{DRV}(\phi)$, $\rho(P) \triangleleft \phi$.

Proof Let P^i be a restrained variable occurring in ϕ . Then, ϕ is a subformula of ϕ' , where $\text{ex2 } P^i \text{ where } \rho : \phi'$ is the formula introducing P^i with $\rho = \rho(P^i)$. Thus, $\rho = \rho(P^i) \triangleleft \phi$ holds according to (2) of the definition of \triangleleft . \square

The treatment of restrained variables require additional attention. Consider a variable Q that is relativized to the restriction $P = Q$. When a subformula ϕ mentioning Q is given meaning, we must include the requirement $P = Q$ in order to avoid later normalizations. But if P itself is relativized to $P = \emptyset$ in some outer existential quantification ψ , then the restriction on P itself must be included at some point. The situation is detailed here:

$$\psi = \text{ex2 } P \text{ where } P = \emptyset : \text{ex2 } Q \text{ where } P = Q : \underbrace{Q = Q}_{\phi} \tag{4}$$

In this case, we wish to enforce also implied restrictions such as $Q = \emptyset$.

To do so, we define for a formula ϕ the set $\text{RV}(\phi)$ of *restrained variables* to be the direct restrained variables $\text{DRV}(\phi)$ together with variables $\text{RV}(\rho(P))$ for $P \in \text{DRV}(\phi)$. This definition makes sense, because by virtue of Proposition 5, $\text{RV}(\phi)$ can be defined inductively on formulas ordered according to \triangleleft . For example, in (4) the restrained variables of ϕ are $\text{RV}(\phi) = \{P, Q\}$.

For any formula ϕ , we define the induced restriction $\rho^{\text{RV}(\phi)}$ to be the set of restrictions of restrained variables, that is, the set $\{\rho(P^i) \mid P^i \in \text{RV}(\phi)\}$. In the example, $\rho^{\text{RV}(\phi)}$ is $\{P = \emptyset, P = Q\}$ and the conjunction of these formulas imply $Q = \emptyset$ as

desired. By a union-wise extension of ρ^{RV} to an operator taking a set of formulas as an argument, the expression $\rho^{\text{RV}}(\rho^{\text{RV}}(\phi))$ also makes sense.

Proposition 6 (a) $\rho(P) \triangleleft \phi$ holds for $P \in \text{RV}(\phi)$.

(b) $\rho^{\text{RV}}(\phi) \supseteq \rho^{\text{RV}}(\rho^{\text{RV}}(\phi))$

Proof

(a) By Proposition 5, $\rho(P) \triangleleft \phi$ holds for $P \in \text{DRV}(\phi)$. Consider $P' \in \text{RV}(\rho(P))$ for some $P \in \text{DRV}(\phi)$. We may by induction according to \triangleleft assume that $\rho(P') \triangleleft \rho(P)$. Thus, by transitivity, $\rho(P') \triangleleft \phi$ holds.

(b) We need only to prove that $\text{RV}(\phi) \supseteq \text{RV}(\rho^{\text{RV}}(\phi))$ holds for $P^i \in \text{RV}(\phi)$. But this is a direct consequence of the definition of the restrained variables of a formula.

□

3.1 From binary to ternary semantics

Classic semantics There is an obvious way to define the semantics of Nutshell WS1S-R. We will state them using a meaning function $\llbracket \cdot \rrbracket^C$, which given a formula ϕ defines a value $\llbracket \phi \rrbracket^C \in \{0, 1\}$. This notation anticipates multi-valued semantics:

$$\begin{aligned} \llbracket \sim \phi' \rrbracket^C w &= \neg \llbracket \phi' \rrbracket^C w \\ \llbracket \phi' \ \& \ \phi'' \rrbracket^C w &= \llbracket \phi' \rrbracket^C w \wedge \llbracket \phi'' \rrbracket^C w \\ \llbracket \text{ex2 } P^i \text{ where } \rho: \phi' \rrbracket^C w &= \begin{cases} 1 & \text{if } \exists M : \llbracket \phi' \rrbracket^C w [P^i \mapsto M] = 1 \\ & \text{and } \llbracket \rho \rrbracket^C w [P^i \mapsto M] = 1 \\ 0 & \text{if } \forall M : \llbracket \phi' \rrbracket^C w [P^i \mapsto M] = 0 \\ & \text{or } \llbracket \rho \rrbracket^C w [P^i \mapsto M] = 0 \end{cases} \\ \llbracket P^i \text{ sub } P^j \rrbracket^C w &= \begin{cases} 1 & \text{if } w \models P^i \text{ sub } P^j \\ 0 & \text{if } w \not\models P^i \text{ sub } P^j \end{cases} \end{aligned}$$

(Again, we have included only one kind of atomic formula; the others also follow their logic semantics.) Of course, we have the following correspondence between WS1S and WS1S-R.

Proposition 7 Assume that for a formula ϕ in Nutshell WS1S-R, we denote by $\hat{\phi}$ the formula obtained by converting $\text{ex2 } P^i \text{ where } \rho: \phi'$ into $\text{ex2 } P^i : \rho \ \& \ \phi'$. Then, $w \models \hat{\phi}$ holds if and only if $\llbracket \phi \rrbracket^C w = 1$ holds.

Since the classic semantics behave like expected, we also write $w \models \phi$ for WS1S-R formulas that satisfy $\llbracket \phi \rrbracket^C w = 1$.

Conjunctively normalized semantics Recall from Section 2.3 that a conjunctive semantics solve the problem of spurious behavior of automata-based semantics. We define these semantics formally below, where we use the notation $\llbracket \rho^{\text{RV}}(P^i \text{ sub } P^j) \rrbracket^N w = 1$ to denote that for each $\rho \in \rho^{\text{RV}}(P^i \text{ sub } P^j)$ it holds that $\llbracket \rho \rrbracket^N w = 1$.

$$\begin{aligned}
\llbracket \sim \phi' \rrbracket^N w &= \neg(\llbracket \phi' \rrbracket^N w) \wedge (\llbracket \rho^{\text{RV}}(\phi') \rrbracket^N w = 1) \\
\llbracket \phi' \ \& \ \phi'' \rrbracket^N w &= \llbracket \phi' \rrbracket^N w \wedge \llbracket \phi'' \rrbracket^N w \\
\llbracket \text{ex2 } P^i \text{ where } \rho: \phi' \rrbracket^N w &= \begin{cases} 1 & \text{if } \exists M : \llbracket \phi' \rrbracket^N w[P^i \mapsto M] = 1 \\ 0 & \text{if } \forall M : \llbracket \phi' \rrbracket^N w[P^i \mapsto M] = 0 \end{cases} \\
\llbracket P^i \text{ sub } P^j \rrbracket^N w &= \begin{cases} 1 & \text{if } w \models P^i \text{ sub } P^j \\ & \text{and } \llbracket \rho^{\text{RV}}(P^i \text{ sub } P^j) \rrbracket^N w = 1 \\ 0 & \text{if } w \not\models P^i \text{ sub } P^j \\ & \text{or } \llbracket \rho^{\text{RV}}(P^i \text{ sub } P^j) \rrbracket^N w = 0 \end{cases}
\end{aligned}$$

Following our goal of using as few normalization operations as possible, we apply the restrictions of free variables to atomic formulas and negations only. For conjunctions, the free variables are the union of the free variables in the contents; so, the explicit normalization can be omitted. For existential quantification, any variable free in the quantified formula $\text{ex2 } P^i \text{ where } \rho: \phi'$ is also free in ρ or in ϕ' . One problem exists for the variable P^i : if the restriction ρ is not satisfiable and the variable P^i does not occur in ϕ' , then the quantified formula unexpectedly evaluates to 1. Anticipating a similar problem for the ternary semantics, we settle on the following assumption:

For all formulas of the form $\phi = \text{ex2 } P^i \text{ where } \rho: \phi'$ and for all w : it holds that

$$w \models \left(\bigwedge_{P \in \text{RV}(\phi)} \rho(P) \right) \Rightarrow \text{ex2 } P^i : \rho \tag{5}$$

where “ \Rightarrow ” stands for implication (by an obvious formula transformation). (Alternatively, we might demand that there is an occurrence of P^i in ρ for all existentially quantified formulas.) If we let $w \models \rho^{\text{RV}}(\phi)$ stand for the property that for all P^i in $\text{RV}(\phi)$ it holds that $w \models \rho(P^i)$, then we may state the correctness of the conjunctively normalized semantics as follows.

Proposition 8 Given assumption (5), $w \models \phi$ and $w \models \rho^{\text{RV}}(\phi)$ both hold if and only if $\llbracket \phi \rrbracket^N w = 1$.

Proof The proof is by induction.

Case $\phi = \sim \phi'$. We observe that $\text{RV}(\phi) = \text{RV}(\phi')$. Then, $\llbracket \sim \phi' \rrbracket^N w = 1$ if and only if $\llbracket \phi' \rrbracket^N w = 0$ and $\llbracket \rho^{\text{RV}}(\phi) \rrbracket^N w = 0$ if and only if (by applying induction hypothesis twice) $(w \not\models \phi' \text{ or } w \not\models \rho^{\text{RV}}(\phi))$ and $w \models \rho^{\text{RV}}(\phi)$ if and only if $w \not\models \phi'$ and $w \models \rho^{\text{RV}}(\phi)$ if and only if $w \models \phi$ and $w \models \rho^{\text{RV}}(\phi)$.

Case $\phi = \phi' \ \& \ \phi''$. We observe that $\text{RV}(\phi) = \text{RV}(\phi') \cup \text{RV}(\phi'')$ holds and that $\rho^{\text{RV}}(\phi) = 1$ holds if and only if both $\rho^{\text{RV}}(\phi') = 1$ and $\rho^{\text{RV}}(\phi'') = 1$ hold. Then,

$\llbracket \phi \rrbracket^N w = 1$ if and only if $\llbracket \phi' \rrbracket^N w = 1$ and $\llbracket \phi'' \rrbracket^N w = 1$ if and only if (by induction) $w \models \phi'$ and $w \models \rho^{\text{RV}}(\phi')$ and $w \models \phi''$ and $w \models \rho^{\text{RV}}(\phi'')$ if and only if $w \models \phi' \ \& \ \phi''$ and $w \models \rho^{\text{RV}}(\phi' \ \& \ \phi'')$.

Case $\phi = \text{ex2 } P^i$ where $\rho : \phi'$.

We need to consider the two directions separately.

(\Rightarrow) Assume that $w \models \phi$ and $w \models \rho^{\text{RV}}(\phi)$. Then, there is a set M such that $w[P^i \mapsto M] \models \phi'$ and $w[P^i \mapsto M] \models \rho$. But, since $\text{RV}(\phi)$ does not mention P^i , we also have that $w[P^i \mapsto M] \models \rho^{\text{RV}}(\phi)$. Taken together with $w[P^i \mapsto M] \models \rho$, we infer that $w[P^i \mapsto M] \models \rho^{\text{RV}}(\phi')$. Now, together with the fact $w[P^i \mapsto M] \models \phi'$, we use the induction hypothesis to establish that $\llbracket \phi' \rrbracket^N w[P^i \mapsto M] = 1$. Consequently, $\llbracket \phi \rrbracket^N w = 1$.

(\Leftarrow) Intuitively, this direction is valid because an occurrence of P^i in ϕ' guarantees that $\rho = \rho(P^i)$ holds thanks to the inclusion of restrictions on variable occurrences in atomic formulas; in the absence of such a variable occurrence, the assumption (5) can be used to find an alternative M that makes ρ hold without affecting the truth value of $\llbracket \phi' \rrbracket^N$.

We proceed more formally. Assume that $\llbracket \phi \rrbracket^N w = 1$. Then, there is a M such that $\llbracket \phi' \rrbracket^N w[P^i \mapsto M] = 1$. By inductive hypothesis, we infer that $w[P^i \mapsto M] \models \phi'$ and $w[P^i \mapsto M] \models \rho^{\text{RV}}(\phi')$. From $w[P^i \mapsto M] \models \rho^{\text{RV}}(\phi')$, we infer that $w \models \rho^{\text{RV}}(\phi)$, since P^i is not a free variable in $\rho^{\text{RV}}(\phi)$. That leaves us with the need to prove that $w \models \phi$ holds.

Now, if $P^i \in \text{RV}(\phi')$ holds, then $w[P^i \mapsto M] \models \rho$ holds—due to $w[P^i \mapsto M] \models \rho^{\text{RV}}(\phi')$ —and $w \models \phi$ holds.

On the other hand, if P^i is not a free variable in ϕ' , then we can use assumption (5) to find an M that makes $w[P^i \mapsto M] \models \rho$ hold without impacting the truth of $w[P^i \mapsto M] \models \phi'$. By induction hypothesis, it holds that $w[P^i \mapsto M] \models \phi'$ and $w[P^i \mapsto M] \models \rho$. Thus, we conclude that $w \models \phi$.

Case $\llbracket P^i \text{ sub } P^j \rrbracket^N$. $\llbracket P^i \text{ sub } P^j \rrbracket^N w = 1$ if and only if (by definition) $w \models P^i \text{ sub } P^j$ and $\llbracket \rho^{\text{RV}}(P^i \text{ sub } P^j) \rrbracket^N w = 1$ if and only if (by induction) $w \models P^i \text{ sub } P^j$ and $w \models \rho^{\text{RV}}(P^i \text{ sub } P^j)$ and $w \models \rho^{\text{RV}}(\rho^{\text{RV}}(P^i \text{ sub } P^j))$ if and only if (by Proposition 6(b)) $w \models P^i \text{ sub } P^j$ and $w \models \rho^{\text{RV}}(P^i \text{ sub } P^j)$. \square

The ternary semantics Let $\mathbb{B}^\perp = \mathbb{B} \cup \{\perp\}$ be the *extended Boolean domain*. We use \perp to denote a “don’t care” situation, one where not all the relativizations hold. Boolean operators \wedge^3 and \neg^3 are defined on this domain as for the usual case with the added rule that if any argument is \perp , then the result is \perp .

$$\begin{aligned}
\llbracket \sim \phi' \rrbracket^3 w &= \neg^3 \llbracket \phi' \rrbracket^3 w \\
\llbracket \phi' \ \& \ \phi'' \rrbracket^3 w &= \llbracket \phi' \rrbracket^3 w \wedge^3 \llbracket \phi'' \rrbracket^3 w \\
\llbracket \text{ex2 } P^i \text{ where } \rho: \phi' \rrbracket^3 w &= \begin{cases} 1 & \text{if } \exists M : \llbracket \phi' \rrbracket^3 w [P^i \mapsto M] = 1 \\ 0 & \text{if } \forall M : \llbracket \phi' \rrbracket^3 w [P^i \mapsto M] \neq 1 \text{ and} \\ & \exists M : \llbracket \phi' \rrbracket^3 w [P^i \mapsto M] = 0 \\ \perp & \text{if } \forall M : \llbracket \phi' \rrbracket^3 w [P^i \mapsto M] = \perp \end{cases} \\
\llbracket P^i \text{ sub } P^j \rrbracket^3 w &= \begin{cases} 1 & \text{if } w \models P^i \text{ sub } P^j \\ & \text{and } \llbracket \rho^{\text{RV}}(P^i \text{ sub } P^j) \rrbracket^3 w = 1 \\ 0 & \text{if } w \not\models P^i \text{ sub } P^j \\ & \text{and } \llbracket \rho^{\text{RV}}(P^i \text{ sub } P^j) \rrbracket^3 w = 1 \\ \perp & \text{if } \llbracket \rho^{\text{RV}}(P^i \text{ sub } P^j) \rrbracket^3 w \neq 1 \end{cases}
\end{aligned}$$

The case of existential quantification above reflects the intuition that if a witness M exists that makes ϕ' true, then ϕ is true; moreover, if there is an M that makes ϕ' false, and no M makes ϕ' true, then ϕ is false; and, finally, if all witnesses M make ϕ not satisfy the restrictions $\rho^{\text{RV}}(\phi')$, then w does not satisfy $\rho^{\text{RV}}(\phi)$.

Something seems to be missing in these semantics: the enforcement of a relativization. The proposition below shows that the relativization bubbles up automatically if needed:

Proposition 9 Given assumption (5), the following holds.

- (a) $w \models \rho^{\text{RV}}(\phi) \Leftrightarrow \llbracket \phi \rrbracket^3 w \neq \perp$.
- (b) $w \models \phi \ \& \ \rho^{\text{RV}}(\phi) \Leftrightarrow \llbracket \phi \rrbracket^3 w = 1$
- (c) $w \models \sim \phi \ \& \ \rho^{\text{RV}}(\phi) \Leftrightarrow \llbracket \phi \rrbracket^3 w = 0$

Proof (Idea) The proof is similar to that of Proposition 8 and is omitted. \square

Part (a) states that the truth of all syntactic restrictions $\rho^{\text{RV}}(\phi)$ is equivalent to $\llbracket \phi \rrbracket^3 w$ not being \perp .

3.2 Automata-theoretic realization of the ternary semantics

The decision procedure of Section 2.2 can be modified to reflect the ternary semantics.

First, we classify states as either 1, 0, or \perp states; that is, we replace Q^F of an automaton A , as defined in Section 2.2, with a valuation or labeling function $\lambda : Q \rightarrow \mathbb{B}^\perp$. For a word u , the valuation $\lambda(\text{LAST } A(u))$ of the last state of the run over u is the truth value *calculated* by A over u . We denote this value $A(u)$, and we call A a *ternary automaton*. Of course, the logic-automata connection can now be expressed: $A_\phi(u) = \llbracket \phi \rrbracket^3(u)$. Ternary automata can be minimized in the same way conventional deterministic automata are determinized with the difference that the initial state partition are the at most three maximal sets on which λ is constant.

Second, we modify the automata-theoretic constructions. For the case of conjunction, we simply change the way product states are labeled; such a state is 1 if both component states are one, it is 0 if both components are 0, and it is \perp if one of the components is \perp . Negation is handled by simply switching 1 states in 0 states and vice versa.

Third, we must generalize the projection operators and futurization to the ternary domain. We define $\text{PROJ}^{3,i}$, $\text{BPROJ}^{3,i}$, and $\text{FUT}^{3,i}$:

$$\text{PROJ}^{3,i}A(u) = \begin{cases} 1 & \text{if } \exists M : A(u[P^i \mapsto M]) = 1 \\ 0 & \text{if } \exists M : A(u[P^i \mapsto M]) = 0 \\ & \text{and } \forall M : A(u[P^i \mapsto M]) \neq 1 \\ \perp & \text{if } \forall M : A(u[P^i \mapsto M]) = \perp \end{cases}$$

$$\text{BPROJ}^{3,i}A(u) = \begin{cases} 1 & \text{if } \exists M \text{ with } \max M < |u| : A(u[P^i := M]) = 1 \\ 0 & \text{if } \exists M \text{ with } \max M < |u| : A(u[P^i := M]) = 0 \\ & \text{and } \forall M \text{ with } \max M < |u| : A(u[P^i := M]) \neq 1 \\ \perp & \text{if } \forall M \text{ with } \max M < |u| : A(u[P^i := M]) = \perp \end{cases}$$

$$\text{FUT}^{3,i}A(u) = \begin{cases} 1 & \text{if } \exists v \in \Sigma_0^{i*} : A(u \cdot v) = 1 \\ 0 & \text{if } \exists v \in \Sigma_0^{i*} : A(u \cdot v) = 0 \\ & \text{and } \forall v \in \Sigma_0^{i*} : A(u \cdot v) \neq 1 \\ \perp & \text{if } \forall v \in \Sigma_0^{i*} : A(u \cdot v) = \perp \end{cases}$$

Proposition 10

$$\text{PROJ}^{3,i}A = \text{BPROJ}^{3,i}(\text{FUT}^{3,i}A)$$

Proof Omitted. □

Fourth, we explain how these operators are implemented as automata algorithms. The construction of $\text{FUT}^{3,i}A$ is quite obvious: we label any state 1 for which a path along a null-but- i extension to a state labeled 1 exists; among the remaining states, we label those 0 for which a null-but- i extension exists that leads to a 0-labeled state in A ; and finally, those not yet labeled retain their \perp -label. In a subset implementation of $\text{BPROJ}^{3,i}$, a subset is labeled 1 if it contains a state labeled 1; when it contains no such state, it is labeled 0 if it contains a state labeled 0; and, when all states in the subset are labeled \perp , the subset state is also labeled \perp .

We call the resulting algorithm the *ternary decision procedure*.

4 Ternary valuations for restricted languages

All languages considered will be regular and over the alphabet $\Sigma = \mathbb{B}^k$. For a language L , the *canonical right-congruence* \sim_L is defined as $u \sim_L v$ if and only if $\forall w : u \cdot w \in L$.

$L \Leftrightarrow v \cdot w \in L$, where $u, v, w \in \Sigma^*$. The set of congruence classes is denoted Σ^*/\sim_L . This set can be regarded as the state set of a canonical, finite-state automaton recognizing L .

Consider languages L , sometimes called the *property*, and R , assumed non-empty, called a *restriction*. Thus, L_ϕ and $L_{\rho^{\text{RV}(\phi)}}$, for $\rho^{\text{RV}(\phi)}$ defined in Section 3.1, constitutes such a pair for any subformula ϕ of ϕ_0 . The *conjunctively normalized representation* is $L' = L \cap R$, and the *conjunctively normalized congruence* is $\sim_{L \cap R}$.

Note that Proposition 8 tells us that $\llbracket \phi \rrbracket^N$ is just the characteristic function for the set $L \cap R$ with $L = L_\phi$ and $R = L_{\rho^{\text{RV}(\phi)}}$.

For a string u , an *accepting extension* v is a string such that $u \cdot v \in L \cap R$. The *ternary valuation* determined by L and R is a function $\chi_{L,R}(u)$, defined to be 1 if $u \in L \cap R$, 0 if $u \in \overline{L} \cap R$, and \perp if $u \notin R$.

Note that Proposition 9 tells us that that $\llbracket \phi \rrbracket^3 = \chi_{L,R}$ with L and R identified as above.

The *ternary congruence* $\sim_{L,R}$ is then defined by $u \sim_{L,R} v$ if for all w it holds that $\chi_{L,R}(u \cdot w) = \chi_{L,R}(v \cdot w)$. The equivalence classes of a ternary congruence can be viewed as the states of a ternary automaton A . More precisely, A is $(\Sigma, Q, q_0, \rightarrow, \chi^F)$, where $Q = \Sigma^*/\sim_{L,R}$ along with $q^0 = \{u \mid u \sim_{L,R} \epsilon\}$ and \rightarrow are defined as usual; moreover, λ is defined so that $A(u) = \chi_{L,R}(u)$ holds—this make sense: the valuation function λ is constant on all strings reaching the same state. We say that A *recognizes* χ . Often, we will call A the $\chi_{L,R}$ -automaton. Note that this automaton respects \sim_R : for any state, there is a uniquely determined equivalence class of \sim_R that all strings reaching the state belongs to. This is just another way of saying that $\sim_{L,R}$ refines \sim_R . Obviously, $\sim_{L,R}$ also refines $\sim_{L \cap R}$.

4.1 Relating the conjunctive and ternary semantics

A *thin language* R is a non-empty set of strings such that

$$\forall u, v : u \not\sim_R v \Rightarrow \forall w : u \cdot w \notin R \vee v \cdot w \notin R \quad (6)$$

In particular, it can be seen that the canonical automaton for R has exactly one accepting state: just make u be a string that reaches one accepting state, v a string that reaches a different accepting state, and w the empty string.

Proposition 11

(a) The *first-order restriction* $R_{\text{singleton}(i)} = \{u \in \mathbb{B}^k \mid P^i\text{-track } i \text{ contains exactly one occurrence of a } 1\}$ is thin.

(b) The $\$$ -restriction

$$R_{\text{\$-restrict}(i)} = \{u \in \mathbb{B}^k \mid \text{the occurrences of } 1 \text{ in } P^i\text{-track are all in positions} \\ \text{no greater than that of the first occurrence of a } 1 \\ \text{in track } \$\} \cap R_{\text{singleton}(\$)}$$

is thin.

- (c) If R and R' are thin and $R \cap R' \neq \emptyset$, then $R \cap R'$ is thin.
- (d) Let R be thin, and let L be any language. If $u \sim_{L \cap R} v$ and u has an accepting extension, then $u \sim_{L,R} v$.
- (e) If u and v both have no accepting extensions, then $u \sim_R v \Leftrightarrow u \sim_{L,R} v$.
- (f) Assume that R is thin. Then, $\Sigma^*/\sim_{L,R}$ is a union of equivalence classes of $\Sigma^*/\sim_{L \cap R}$ and Σ^*/\sim_R . More precisely, either $\Sigma^*/\sim_{L,R} = \Sigma^*/\sim_{L \cap R}$ holds or

$$\Sigma^*/\sim_{L,R} = (\Sigma^*/\sim_{L \cap R} \setminus \{S_{L \cap R}\}) \cup \mathcal{S}$$

holds, where $S_{L \cap R} = \{u \mid \forall v : u \cdot v \notin L \cap R\}$ and $\mathcal{S} = \{S \in \Sigma^*/\sim_R \mid \exists u \in S : \forall v : u \cdot v \notin L \cap R\}$.

- (g) $|\Sigma^*/\sim_{L \cap R}| \leq |\Sigma^*/\sim_{L,R}| \leq |\Sigma^*/\sim_{L \cap R}| + |\Sigma^*/\sim_R| - 1$.

Proof

- (a) We note that there are three equivalence classes corresponding to this language: strings containing zero occurrences, exactly one occurrence, or two or more occurrences of a 1 in the P^i -track. Consider u and v such that $u \not\sim_R v$. Let w be any string, and assume that $u \cdot w \in R$. Then u contains at most one occurrence of a 1 in the P^i -track. We must prove that $v \cdot w \notin R$. *Case (1):* u contains no 1s in the P^i -track. Then by assumption that $u \cdot w \in R$, w contains exactly one occurrence of 1 in the P^i -track. But, by assumption that $u \not\sim_R v$, v contains either one 1 or two or more 1s. In all cases, $v \cdot w \notin R$. *Case (2):* u contains exactly one occurrence of a 1. Then w contains no occurrences and v contains zero or two or more occurrences. Thus, again it will hold that $v \cdot w \notin R$.
- (b) The proof is similar to the previous case. Also, it can be seen that the intersection with the singleton language is necessary for thinness.
- (c) Consider u and v such that $u \not\sim_{R \cap R'} v$. Then, either $u \not\sim_R v$ or $u \not\sim_{R'} v$. Assume the former and that $u \cdot w \in R \cap R'$. Then, $v \cdot w \notin R$ thanks to R being thin, and thus $u \cdot w \notin R \cap R'$.
- (d) Let w be the accepting extension of u . From $u \sim_{L \cap R} v$, we infer that $v \cdot w \in L \cap R$. Assume for a contradiction that $u \not\sim_R v$. Combined with the thinness of R , and the fact that $u \cdot w \in R$, we would conclude that $v \cdot w \notin R$. This contradicts that $v \cdot w \in L \cap R$. Thus, it must hold that $u \sim_R v$. It follows from elementary considerations that $u \sim_{L \cap R} v$ and $u \sim_R v$ implies that $u \sim_{L,R} v$.
- (e) Assume that u and v both have no accepting extensions and that $u \sim_R v$. Let w be any string. Then, $u \cdot w$ and $v \cdot w$ are both in R or none is. In the second case, $\chi_{L,R}(u \cdot w) = \chi_{L,R}(v \cdot w) = \perp$. In the first case, it follows from the assumption that u and v both have no accepting extensions that $\chi_{L,R}(u \cdot w) = \chi_{L,R}(v \cdot w) = 0$.

(f) If $S_{L \cap R}$ is empty, then all strings u have an accepting extension. Then it follows by 4. that $\Sigma^*/\sim_{L \cap R}$ and $\Sigma^*/\sim_{L,R}$ are identical. Otherwise, when $S_{L \cap R}$ is nonempty, it follows again by 4. that equivalence classes of $\sim_{L \cap R}$ and $\sim_{L,R}$ for strings that have accepting extensions coincide. Strings that do not have such extensions fall into equivalence classes of \sim_R according to 5.

(g) This follows from (f). □

From this proposition, it follows easily that for any ϕ the regular language representing the conjunction of the restrictions in $\rho^{\text{RV}}(\phi)$ is thin if variables are subjected to only first-order relativizations or to $\$$ -relativizations (or to both). The proposition also tells us that $\Sigma^*/\sim_{L,R}$ is pieced together from $\Sigma^*/\sim_{L \cap R}$ plus a subset of Σ^*/\sim_R .

4.2 The ternary decision procedure compared

Theorem 1 Assume that all relativizations are thin languages. If the automata of relativizations have at most N states, then the size of each intermediate, minimized automaton (representing some subformula) in the ternary decision procedure is the same, to within an additive constant of $N - 1$, as the size of corresponding automaton under the conjunctive semantics.

Proof This follows from (g) of the previous proposition. □

This result is the justification for the practical use of the ternary semantics since usually the number of first-order variables in simultaneous use is quite small. But note that the size of the additive constant is exponential in the number of free first-order variables. Also, thanks to Proposition 11(f), we see that the automata of the ternary approach are, apart from the Σ^*/\sim_R parts, the same as those that occur when the automaton of every subformula is normalized conjunctively.

5 The sexpartite approach

We show next how to get rid of the boundedness assumption in Theorem 1. We do so by re-introducing a normalization step at the usual place of the relativization where the variable is introduced by a quantifier. This will allow us to prune away states from the automaton representing the ternary semantics. We want to get rid of all states following a state if all these states have the same membership status with respect to L whenever they are in R .

5.1 Interesting strings, approximations, and sexpartification

To proceed more rigorously, we define a string u to be *interesting* (for L and R) if it has (a) some accepting extension and (b) some extension v , called a *rejecting extension*, such that $u \cdot v$ in $\overline{L} \cap R$. Also, a “*don't care*” extension is one that makes a string fall

outside R . Note that all prefixes of an interesting string are also interesting. In other words, an uninteresting string cannot be extended so as to become interesting. Note also that if $u \sim_{L \cap R} v$ and u is interesting then v is also interesting. Define the Boolean valuation $\iota_{L,R}(u)$ so that it denotes whether a string u is interesting. For an uninteresting string, let $\text{CUT}(u)$ be the shortest uninteresting prefix of u . The *approximation* $\alpha_{L,R}(u)$ of an uninteresting u is defined by

$$\alpha(u) = \begin{cases} 1 & \text{if } \text{CUT}(u) \text{ has an accepting extension} \\ 0 & \text{if } \text{CUT}(u) \text{ has a rejecting extension} \\ \perp & \text{if all extensions of } \text{CUT}(u) \text{ are "don't care"} \end{cases} \quad (7)$$

(These three cases are clearly mutually exclusive.)

For $u \neq \epsilon$, we define u^- to be the prefix v such that $v \cdot a = u$ for some $a \in \Sigma$.

We note that if $\text{CUT}(u) \neq \epsilon$, then $\text{CUT}(u)^-$ is interesting. Moreover, if $\text{CUT}(u) = \epsilon$, then $\alpha(u) \in \{0, 1\}$, since we assume that $R \neq \emptyset$.

When u is interesting, we define $\alpha_{L,R}(u)$ to be $\chi_{L,R}(u)$. We make the *separtite valuation* $\chi_{L,R}^6$ be the interest status together with the approximation: $\chi_{L,R}^6(u) = (\iota(u), \alpha(u))$. The valuation $\chi_{L,R}^6$ is also called the *separtification* of $\chi_{L,R}$ since it can be defined from $\chi_{L,R}$ alone. The *canonical separtite congruence* $\sim_{L,R}^6$ is defined from the valuation as for the ternary case.

5.1.1 The exactness property

When L and R are clear from the context, we often omit them as subscripts. The uninteresting equivalence classes of Σ^*/χ^6 are just that: if u is uninteresting, with $(\iota(u), \alpha(u)) = (0, X)$, then for any extension v , $(\iota(u), \alpha(v)) = (0, X)$. Thus, the only transition in Σ^*/\sim^6 from an uninteresting equivalence class is to itself. Moreover, there are at most three such classes.

The approximation $\alpha(u)$ satisfies the following properties: if u is interesting or $\chi(u)$ is not \perp , then $\alpha(u) = \chi(u)$ holds; otherwise, if u is uninteresting and $\chi(u)$ is \perp , then $\alpha(u)$ may take on any value in $\{0, 1, \perp\}$. The fact that $\alpha(u) = \chi(u)$ usually holds—with the only possible exception being that $\chi(u)$ is \perp for an uninteresting u —is called the *exactness property*.

We note that if u is uninteresting and $\alpha(u)$ is not \perp , then it is possible, and somewhat counterintuitive, that for all v the value $\chi(u \cdot v)$ is \perp . When $\alpha(u) \neq \perp$ holds, it is only for $u = \text{CUT}(u)$ that an extension v is guaranteed to exist such that $\chi(u \cdot v) \neq \perp$ holds.

5.1.2 Recovering χ^6 and χ from α

In the following, we will be concerned only with the approximation valuation α because of the following property:

Proposition 12 Let $\chi^6 = (\iota, \alpha)$ be a sextartite valuation. Then the following property obtains:

$$\iota(u) = 1 \iff \exists v^0, v^1 : \alpha(u \cdot v^0) = 0 \wedge \alpha(u \cdot v^1) = 1$$

Thus, the approximation α alone carries all the information of the sextartite valuation χ^6 .

We may even recuperate χ from α by using the operator \times_{\perp} on valuations that for a ternary χ and a binary ρ behaves according to

$$(\chi \times_{\perp} \rho)(u) = \begin{cases} \chi(u) & \text{if } \rho(u) = 1 \\ \perp & \text{if } \rho(u) = 0 \end{cases} \quad (8)$$

We write $\chi \times_{\perp} R$ to denote $\chi \times_{\perp} \rho$, where ρ is the characteristic function for R . Then, it is easy to see that χ can be recovered from α :

$$\chi_{L,R} = \alpha_{L,R} \times_{\perp} R$$

5.2 Sextartification algorithm

Let us consider some properties of $\sim_{L,R}$ through a study of the automaton A recognizing it. First, we note that the notions of “interesting,” “accepting extension,” “rejecting extension,” and “don’t care” extension apply to states of A as well. So, we can partition the states of A into at most four sets, namely the set of interesting states I and three sets of non-interesting states. The non-interesting sets are: N^1 , which consists of states allowing some accepting extension (but no rejecting extension); N^0 , which consists of states allowing some rejecting extension (but no accepting extension), and N^{\perp} , which consists of states allowing only don’t care extensions. Some of these sets may be empty; for simplicity, we assume that none is.

Second, we infer various properties from the definition of $\chi_{L,R}$. (1) N^{\perp} is a singleton set (assuming that A is minimal). (2) The only transitions among these sets of states are as follows: $I \rightarrow N^0$, $I \rightarrow N^1$, $I \rightarrow N^{\perp}$, $N^0 \rightarrow N^{\perp}$, and $N^1 \rightarrow N^{\perp}$, where $M \rightarrow N$ means that there is some transition from M to N . Thus, Figure 4 illustrates the structure of $\Sigma^* / \sim_{L,R}$.

With these observations, we can present an algorithm called *sexpartification* that from A recognizing $\chi_{L,R}$ calculates an automaton $\text{SEXP}(A)$ recognizing $\alpha_{L,R}$.

First, we construct from A a *collapsed* automaton A' by collapsing states in the set N^0 to one state, which for simplicity we also call N^0 . This state is defined to have a transition to itself on all letters; in particular, the outgoing transitions, all of which are to N^{\perp} , are removed. (Of course, the removal of transitions is at odds with the traditional ways of shrinking automata.) We do the same for N^1 . If there are no incoming transitions from I to N^{\perp} , we remove N^{\perp} . Each of the at most three states of the form N^Z is labeled Z . Each interesting state, i.e. each state in I , keeps its label. This completes the construction of A' .

Second, we minimize A' to obtain $\text{SEXP}(A)$. The shape of this automaton is depicted in Figure 5.




Figure 4: The transition structure of $\Sigma^*/\sim_{L,R}$.




Figure 5: The transition structure of $\Sigma^*/\sim_{L,R}^6$.

Proposition 13 $\text{SEXP}(A)$ calculated through sexpartification from an automaton A recognizing $\chi_{L,R}$ is the minimal automaton recognizing $\alpha_{L,R}$.

Proof We consider the collapsed automaton A' since the second step does not affect what is recognized. We consider only the case where the empty string is interesting. (Other cases are simpler.) It follows that the initial state is interesting. From the definition of $\chi_{L,R}^6$, it can be seen that the state reached in A' over a word u correctly identifies $\alpha_{L,R}(u)$ as long as the run stays within I . When a letter a is such that the run of A' over $u \cdot a$ leaves I on the transition on a , $u \cdot a$ is uninteresting, and it can be seen from the definition of $\chi_{L,R}^6$ that $\text{CUT}(u \cdot a) = u \cdot a$. Moreover, the N^X state

reached is such that X is the value of $\alpha_{L,R}(u \cdot a)$. Finally, we have already argued that the equivalence classes N^X allow no outgoing transitions except to themselves. Thus, A' so constructed recognizes $\alpha_{L,R}$. \square

We state a set of necessary and sufficient conditions for establishing that an automaton A is a sufficiently good approximation to $\chi_{L,R}$ for establishing that its sexpartification is that of $\alpha_{L,R}$:

Proposition 14 For establishing $\text{SEXP}(A) = \alpha_{L,R}$, it is necessary and sufficient that (i) for any interesting u , $\chi_{L,R}(u) = A(u)$ holds and (ii) for any minimal uninteresting u , the following three conditions obtain

- (N1) If $\text{LAST}\alpha_{L,R}(u) = N^1$, then there is an accepting extension for A of u and no don't care extensions.
- (N0) If $\text{LAST}\alpha_{L,R}(u) = N^0$, then there is an rejecting for A and no accepting extensions.
- (N \perp) If $\text{LAST}\alpha_{L,R}(u) = N^\perp$, then there are but don't care extensions A .

Proof “ \Leftarrow ” Condition (i) makes A behave just like $\chi_{L,R}$ for interesting states and conditions (N1), (N0), and (N \perp) ensure that uninteresting states are correctly identified.

“ \Rightarrow ” (i) if u is interesting, then $\text{SEXP}(A)(u) = \alpha_{L,R}(u) = \chi_{L,R}(u)$, but u will also be interesting relative to the state $\text{LAST}A(u)$, and sexpartification will ensure that $A(u) = \text{SEXP}(A)$. (ii) (N1) Assume that $\text{LAST}\alpha_{L,R}(u) = N^1$ and that u is minimal. Then there is a v such that $\chi_{L,R}(u \cdot v) = 1$. But $\text{SEXP}(A)(u \cdot v) = \alpha_{L,R}(u \cdot v) = 1$, so $A(u \cdot v)$ is also 1. A rejecting extension $u \cdot v$, i.e. one such that $A(u \cdot v) = 0$ would also be impossible since $\text{LAST}(\text{SEXP}(A))(u)$ is also the N^1 state. (N0) is similar to (N1). (N \perp) Since $\text{LAST}A(u)$ is N^\perp , there can be no v such that $A(u \cdot v) \neq \perp$. \square

For an automaton A satisfying the conditions of the proposition is called a *sexpartite automaton for L and R* .

Proposition 15 (a) $|\Sigma^* / \sim_{L,R}^6| \leq |\Sigma^* / \sim_{L,R}|$.

(b) $\sim_{L,R}$ may not refine $\sim_{L,R}^6$.

(c) For interesting strings, $\sim_{L,R}$ locally refines $\sim_{L,R}^6$, that is, if u is interesting, then $u \sim_{L,R} v$ implies that $u \sim_{L,R}^6 v$.

Proof Items (a) and (c) follow from the sexpartification construction. The fact that $\sim_{L,R}^6$ is not coarser than $\sim_{L,R}$ can be seen by construction of a χ corresponding to some L and R such that there is a string $u \cdot v$ with u bringing the χ -automaton into a N^0 -state and with v from there driving the automaton into a N^\perp -state (which will be unique). Also, there must be a string w bringing the automaton into the N^\perp -state without passing through any N^0 -state. Then, $u \cdot v \sim_{L,R} w$ holds, but $u \cdot v \not\sim_{L,R}^6 w$

does not hold, since after sexpartification $u \cdot v$ leads to the N^0 -state, but w leads to the N^\perp -state. \square

We shall later (in Section 8.2) introduce a property of R that strengthens (c) such that $\sim_{L,R}$ and $\sim_{L,R}^6$ locally coincides.

5.3 An example

Consider the alphabet $\Sigma = \{a, b\}$, $R_c = \{w \mid |w| = 0 \pmod 3\}$, and $L_c = a \cdot a \cdot \Sigma^* \cup b \cdot a \cdot a \cdot \Sigma^*$. Since R_c is the set of strings whose length is 0 modulo 3, it follows that R_c is thin. The automaton recognizing L_c is shown in Figure 6, along with automata for χ_{L_c, R_c} and χ_{L_c, R_c}^6 . In figure, we have used diamond shapes for the states that are labeled \perp . All states in the depiction of χ_{L_c, R_c}^6 are interesting, except for the states N^0 and N^1 .

This example is constructed to show that after sexpartification, interesting states may no longer respect \sim_R , even when the restriction is a thin language. (We have already noted that the states of χ respect \sim_R , i.e. that $\sim_{L,R}$ refines R .) To see this, consider the interesting χ -automaton states q' and q'' , which correspond to different \sim_{R_c} classes. They are fused as a result of sexpartification as can be seen in the depiction of the χ_{L_c, R_c}^6 -automaton. Incidentally, this automaton is isomorphic to the original one, except for the labeling of states. The example thus shows how sexpartification may, in lucky cases, completely undo the complication of a restriction that is conjoined to language.




Figure 6: Automata for L_c , χ_{L_c, R_c} , and χ_{L_c, R_c}^6 .

6 Conjunctions of sexpartite representations

In order to normalize restrictions to of the $R' \cap R''$, it is necessary to make further assumptions. We say that restrictions R' and R'' are *compatible* if the following holds:

$$\forall u : (\exists v' : u \cdot v' \in R') \wedge (\exists v'' : u \cdot v'' \in R'') \Rightarrow (\exists v : u \cdot v \in R' \cap R'')$$

Proposition 16 Any two restrictions R' and R'' that are intersections of sets of the form $R_{\text{singleton}(i)}$ and $R_{\text{\$-restrict}(i)}$ of Proposition 11 are compatible.

Proof Left to the reader. \square

Because χ can be recovered from α (see 5.1.2), we may easily establish that $\alpha_{L,R \cap R'} = \text{SEXP}(\alpha_{L,R} \times_{\perp} R \times_{\perp} R')$. But it is obvious to ask whether the product $\times_{\perp} R$ is really necessary; intuitively, it appears to be the case that the information is already built into $\alpha_{L,R}$. Although we will not need the result, we show that the product with the characteristic function of R in this identity is indeed unnecessary when R and R' are compatible:

Proposition 17 If R and R' are compatible, then

$$\alpha_{L,R \cap R'} = \text{SEXP}(\alpha_{L,R} \times_{\perp} R').$$

More generally, if we replace $\alpha_{L,R}$ with any sexpartite automaton B for L and R , then $\alpha_{L,R \cap R'} = \text{SEXP}(B \times_{\perp} R')$.

Proof We use the method of Proposition 14, where R of the proposition is $R \cap R'$ and A is $\alpha_{L,R} \times_{\perp} R'$. (For the more general formulation, the argument is similar, except that $\alpha_{L,R}$ is replaced by B . The exactness property and the notion of interesting remain the same.)

Case (i): The string u is interesting for L and $R \cap R'$. We must show that $\chi_{L,R \cap R'}(u) = A(u)$. We note that u is also interesting for L and R ; so, by the exactness property, $\chi_{L,R}(u) = \alpha_{L,R}(u)$ holds. By definition of $\chi_{L,R \cap R'}$, we have $\chi_{L,R \cap R'}(u) = (\chi_{L,R} \times_{\perp} R')(u) = (\alpha_{L,R} \times_{\perp} R')(u)$. Thus, $\chi_{L,R \cap R'}(u) = A(u)$ holds.

Case (ii) We now assume that u is uninteresting and that u^- is interesting.

Subcase (N1) Assume that $\text{LAST}\alpha_{L,R \cap R'}(u) = N^1$.

First, we prove that there is an accepting extension in A from $\text{LAST}A(u)$. Let v be such that $\chi_{L,R \cap R'}(u \cdot v) = 1$ (such a v exists because $\text{LAST}\alpha(u) = N^1$ and u is a minimal uninteresting string). Thus, $u \cdot v \in R \cap R'$ holds. In particular, $u \cdot v \in R$ holds, and by the exactness property, $\chi_{L,R}(u \cdot v) = \alpha_{L,R}(u \cdot v) = 1$ holds, whence we infer that $(\alpha_{L,R} \times_{\perp} R')(u \cdot v) = 1$. That is, there is an accepting extension in A .

Second, we prove that there is no rejecting extension in A from $\text{LAST}(u)$. To do so, we assume for contradiction that there is a v such that $(\alpha_{L,R} \times_{\perp} R')(u \cdot v) = 0$.

Now, if $u \cdot v$ is interesting for L and R , then $\chi_{L,R}(u \cdot v) = \alpha_{L,R}(u \cdot v) = 0$ (because of the exactness property and because we just assumed that $(\alpha_{L,R} \times_{\perp} R')(u \cdot v) = 0$ holds). Moreover, we know that $u \cdot v \in R'$. Thus, we also have $\chi_{L,R \cap R'}(u \cdot v) = 0$, but that contradicts the assumption that $\text{LAST}\alpha_{L,R \cap R'}(u) = N^1$.

So, we may assume that $u \cdot v$ is uninteresting for L and R . But, since u is interesting for L and R and $u \cdot v$ is not, there is a \hat{u} such that $u \cdot \hat{u}$ is uninteresting, $u \cdot \hat{u}$ is a prefix of $u \cdot v$, and $u \cdot \hat{u}^-$ is interesting. Moreover, there is a \hat{v} such that $u \cdot \hat{u} \cdot \hat{v} \in \overline{L} \cap R$, because of the minimality of $u \cdot \hat{u}$ as an uninteresting string and because $\alpha_{L,R}(u \cdot \hat{u})$ is also 0 by definition of the approximation function. (For the more general formulation, we note that $A(u \cdot \hat{u})$ may be either \perp or 0.)

By assumption that $(\alpha_{L,R} \times_{\perp} R')(u \cdot v) = 0$, we have that $u \cdot v \in R'$, whence we may find v' such that $u \cdot \hat{u} \cdot v' \in R'$. By the condition of compatibility, we find a \tilde{v} such that $u \cdot \hat{u} \cdot \tilde{v} \in R \cap R'$. By the exactness property applied twice, we find that $\chi_{L,R}(u \cdot \hat{u} \cdot \tilde{v}) = 1$ and $\chi_{L,R \cap R'}(u \cdot \hat{u} \cdot \tilde{v}) = 0$, that is $u \cdot \hat{u} \cdot \tilde{v}$ is in L and is not in L . This is a contradiction. Thus, there is no rejecting extension.

Subcase (N0) This case is similar to that of (N1).

Subcase (N \perp) Assume that $\text{LAST}_{\alpha_{L,R \cap R'}}(u) = N^{\perp}$. We must prove that every extension in A from $\text{LAST}(u)$ is rejecting. So, let v be a string so that $A(u \cdot v) \neq \perp$, say $A(u \cdot v) = 1$. Then, $u \cdot v$ is then R' , but if $u \cdot v \in R$ also held, then we could not have $\text{LAST}_{\alpha_{L,R \cap R'}}(u) = N^{\perp}$ by the exactness property; thus, $u \cdot v \notin R$ holds.

By our assumption $A(u \cdot v) = 1$, we infer that $\alpha_{L,R}(u \cdot v) = 1$ holds, and that $u \cdot v$ is uninteresting for L and R —if $u \cdot v$ was interesting, then $\alpha_{L,R}(u \cdot v)$ should be \perp by the exactness property, since $u \cdot v \notin R$ holds. But, since u is interesting for L and R , we may find a \hat{u} such that $u \cdot \hat{u}$ is uninteresting, $u \cdot \hat{u}^-$ is interesting, and $u \cdot \hat{u}$ is a prefix of $u \cdot v$. Hence, $\alpha_{L,R}(u \cdot \hat{u})$ is also 1 (or, in the general case, possibly \perp) and there is a \hat{v} such that $u \cdot \hat{u} \cdot \hat{v}$ is in $L \cap R$. Also, since $u \cdot v \in R'$, there is v' such that $u \cdot \hat{u} \cdot v' \in R'$. By assumption of compatibility, we can then find \tilde{v} such that $u \cdot \hat{u} \cdot \tilde{v}$ is in $R \cap R'$. In particular, since $u \cdot \hat{u} \cdot \tilde{v}$ is in R , we have that $\chi_{L,R}(u \cdot \hat{u} \cdot \tilde{v}) = 1$ by the exactness property. Also, by the exactness property and the assumption that $\text{LAST}_{\alpha_{L,R \cap R'}}(u) = N^{\perp}$ we have that $\chi_{L,R \cap R'}(u \cdot \hat{u} \cdot \tilde{v}) = \perp$, which is a contradiction. \square

The condition of compatibility in Proposition 17 is necessary. To see this, consider the languages L , R , and R' defined in Figure 7, where we have also shown the transition structure of an automaton for L . Each state is labeled with 1 or 0 according to whether it accepts or not; additionally, each state is marked with the membership status of a string reaching the state with respect to R and R' . The idea is to construct $\chi_{L,R \cap R'}$ so that the $\alpha_{L,R \cap R'}$ -automaton already on a enters the N^1 state, see Figure 8. It will do so because after a all extensions, except the empty string, lead to states not in $R \cap R'$ and because u itself is in L , R , and R' . Intuitively speaking, the future has been restricted to 1 except for most places, which are outside the restriction. In contrast, $\alpha_{L,R}$ “sees” more don’t-care extensions of u , including some that are 0. Therefore, $\alpha_{L,R}$ only goes into an uninteresting state after reading aa or ab . Moreover, we have arranged it so that aaa is in R' , but not in R (otherwise, $\alpha_{L,R \cap R'}$ would not enter the N^1 state on a); this is the source of the failure of $\text{SEXP}(\alpha_{L,R} \times_{\perp} R')$ to become identical to $\alpha_{L,R \cap R'}$ as can be seen in Figure 9. Technically, the compatibility requirement fails for the string aa : the empty extension brings it into R and the extension a brings it into R' , but there is no extension of aa that brings it into both.

Proposition 18 $\alpha_{L' \cap L'',R} = \text{SEXP}(\alpha_{L',R} \wedge^3 \alpha_{L'',R})$. More generally, if we replace $\alpha_{L',R}$ with any sexpartite automaton B' for L' and R and $\alpha_{L'',R}$ with any sexpartite automaton B'' for L'' and R , then $\alpha_{L' \cap L'',R} = \text{SEXP}(B' \wedge_{\perp} B'')$.

Proof Again, we model our proof on Proposition 14. For notational simplicity, let $\alpha = \alpha_{L' \cap L'',R}$, $\alpha' = \alpha_{L',R}$, and $\alpha'' = \alpha_{L'',R}$. We use of the conventions for the ternary valuations χ , χ' , and χ'' . (For the more general formulation, the argument is




Figure 7: An automaton for a language L where restrictions R and R' are not compatible.




Figure 8: The approximation automata for L .

similar, except that α' is replaced by B' and α'' is replaced by B'' . The exactness property and the notion of interesting remain the same.)

Case (i) Let u be interesting. Therefore, the α' -automaton on u is not in N^\perp and the α'' -automaton is not in N^\perp . Moreover, u is interesting for L' and R or for L'' and R ; assume the former, without loss of generality. Then, $\chi'(u) = \alpha'(u)$ holds. Since α'' on u is not in N^\perp , the following holds: if $\alpha''(u)$ is \perp , then u is not in R and if $\alpha''(u)$ is not \perp , then $\alpha''(u) = \chi''(u)$; in either case, we have that $\chi(u) = \alpha'(u) \wedge^3 \alpha''(u)$.

Case (ii). Let u be uninteresting with u^- interesting.

Subcase (N1) Assume that $\text{LAST}_{\alpha_{L, R \cap R'}}(u) = N^1$. Then, there is a v such that $\chi(u \cdot v)$ is 1, whence $\chi'(u \cdot v)$ and $\chi''(u \cdot v)$ are also 1. By the exactness property, we infer that $\alpha'(u \cdot v)$ and $\alpha''(u \cdot v)$ also are 1.

Assume now for a contradiction that there is a v such that $(\alpha_{L', R}(u \cdot v) \wedge^3 \alpha_{L'', R}(u \cdot v)) = 0$ holds. Thus, $\alpha_{L', R}(u \cdot v)$ and $\alpha_{L'', R}(u \cdot v)$ are both 0. If the α' -automaton has entered the N^0 state after u (in the general case, read “a N^0 state”) or after a longer prefix of $u \cdot v$, then there is a \hat{u} and a \hat{v} such that $u \cdot \hat{v}$ is a prefix of $u \cdot v$ and $u \cdot \hat{v} \cdot \hat{v}$




Figure 9: The approximation automaton $\alpha_{L,R} \times_{\perp} R'$ and its sexpartification.

is in $\overline{L'} \cap R$. But then, $\chi(u \cdot \hat{v} \cdot \hat{v})$ must be equal to 0, which is inconsistent with the subcase assumption that $\text{LAST}_{\alpha_{L,R} \cap R'}(u) = N^1$. So, $u \cdot \hat{v}$ is interesting for L' and R and, of course by similar reasoning, for L'' and R . Hence, we have that $\chi'(u \cdot v)$ and $\chi''(u \cdot v)$ are both 0, but that contradicts again the subcase assumption.

Subcase (N0) This case is similar to subcase (N1).

Subcase (N \perp) Assume that $\text{LAST}_{\alpha_{L,R} \cap R'}(u) = N^{\perp}$. Then, for all v , $u \cdot v$ is not in R , whence u is uninteresting for L' and R . In other words, both the α' -automaton and the α'' -automaton are in an uninteresting state after reading u .

Now fix v . We must prove that $(\alpha_{L',R}(u \cdot v) \wedge^3 \alpha_{L'',R}(u \cdot v)) = \perp$. So, the only interesting cases are that both automata entered interesting states that are not N^{\perp} . For example, we may assume that both entered the N^1 state (in the general case, read “an N^1 state”). (The other three cases are similar.) Then, there must exist a prefix u' of u such that u' is uninteresting for L' and R and minimal for this property. A similar uninteresting prefix u'' exists for L'' and R . Without loss of generality, we may assume that u' is a prefix of u'' . We know that for u''^- there are extensions in both $L' \cap L'' \cap R$ and $L' \cap \overline{L''} \cap R$. Moreover, for u''^- all extensions are in \overline{R} or in $L' \cap L'' \cap R$ and there is an extension in R . Therefore, the α -automaton must enter the N^1 state on u . That contradicts the subcase assumption. \square

7 Sexpartite semantics for WS1S

We define $\llbracket \phi \rrbracket^6$ to be the approximation function α for the ternary valuation $\chi = \llbracket \phi \rrbracket^3$. The challenge is to calculate the approximation function of a composite formula from those of its constituents. For basic formulas, the approach is evident: we define $\llbracket \phi \rrbracket^6$ to be the automaton (or function) $\text{SEXP}(\llbracket \phi \rrbracket^3)$, where we adopt the convention that sexpartification calculates the approximation function, not the sexpartite valuation, as

explained in the remarks after Proposition 13. The other cases are treated below.

7.1 Sexpartite negation

For negation $\phi = \neg\phi'$, we define $\llbracket \neg\phi' \rrbracket^6 = \neg^3 \llbracket \phi' \rrbracket^6$. There is an obvious automata-theoretic algorithm for achieving this operation. It holds for reasons of symmetry that if $\llbracket \phi' \rrbracket^6 = \text{SEXP}(\llbracket \phi' \rrbracket^3)$, then $\llbracket \neg\phi' \rrbracket^6 = \text{SEXP}(\llbracket \neg\phi' \rrbracket^3) = \neg^3 \llbracket \phi' \rrbracket^6$.

7.2 Sexpartite conjunction

For conjunction $\phi = \phi' \wedge \phi''$, the case is more complicated, since $\llbracket \phi' \rrbracket^3$ and $\llbracket \phi'' \rrbracket^3$ may be based on different restrictions. Assume that $R' = L\rho^{\text{RV}(\phi')}$, $R'' = L\rho^{\text{RV}(\phi')}$, and that $R = R' \cap R'' = L\rho^{\text{RV}(\phi)}$. Moreover, define $L' = L_{\phi'}$, $L'' = L_{\phi''}$, and $L = L_{\phi} = L_{\phi'} \cap L_{\phi''}$.

We can express the semantics of conjunction according to:

$$\llbracket \phi' \wedge \phi'' \rrbracket^6 = \text{SEXP}((\llbracket \phi' \rrbracket^6 \wedge^3 \llbracket \phi'' \rrbracket^6) \times_{\perp} \rho^{\text{RV}(\phi)}) \quad (9)$$

This is correct: if $\llbracket \phi' \rrbracket^6 = \alpha_{L',R'}$ and $\llbracket \phi'' \rrbracket^6 = \alpha_{L'',R''}$ hold, then we have the following identities:

$$\begin{aligned} & \text{SEXP}((\llbracket \phi' \rrbracket^6 \wedge^3 \llbracket \phi'' \rrbracket^6) \times_{\perp} \rho^{\text{RV}(\phi)}) \\ &= \text{SEXP}((\llbracket \phi' \rrbracket^6 \times_{\perp} \rho^{\text{RV}(\phi')}) \wedge^3 (\llbracket \phi'' \rrbracket^6 \times_{\perp} \rho^{\text{RV}(\phi'')})) \\ &= \text{SEXP}((\alpha_{L',R'} \times_{\perp} R') \wedge^3 (\alpha_{L'',R''} \times_{\perp} R'')) \\ &= \text{SEXP}(\chi_{L',R'} \wedge^3 \chi_{L'',R''}) \\ &= \text{SEXP}(\llbracket \phi' \rrbracket^3 \wedge^3 \llbracket \phi'' \rrbracket^3) \\ &= \text{SEXP}(\llbracket \phi \rrbracket^3) \\ &= \llbracket \phi' \wedge \phi'' \rrbracket^6 \end{aligned}$$

7.3 Sexpartite projection

We fix an index i as found in the existential quantification $\phi = \text{ex}2 P^i$ where $\rho: \phi'$. To simplify notation, we write u^M for the string $u[i \mapsto M]$.

Proposition 19 For a given ternary valuation $\chi' = \chi_{L',R'}$, we define $\chi = \text{PROJ}^{3,i}\chi'$ and $R = \{u \mid \chi(u) \neq \perp\}$. Assume that there is some R^i with $R' = R \cap R^i$. Moreover, assume that for all u

$$u \in R \Rightarrow \forall M : u^M \in R \quad (10)$$

holds. Then these equalities between functions holds:

$$\chi = \text{PROJ}^{3,i}\chi' = \text{PROJ}^{3,i}(\text{SEXP}\chi' \times_{\perp} R^i) \times_{\perp} R$$

Proof We use the following notation: $\alpha' = \text{SEXP}\chi'$, $R' = R \cap R^i$, $B' = \alpha' \times_{\perp} R^i$, and $B = \text{PROJ}^{3,i}B'$. We must prove that $\chi = B \times_{\perp} R$.

Claim 1 The following three properties hold:

- i. $\chi'(u) \neq \perp \Rightarrow B'(u) = \chi'(u)$
- ii. $u \in R \Rightarrow B'(u) = \chi'(u)$
- iii. $u \in R \Rightarrow B(u) = \chi(u)$

Proof

- i. Assume that $\chi'(u) \neq \perp$ holds. Then, we know by the exactness property that $\alpha'(u) = \chi'(u)$ holds. Moreover, from assumptions $R' = R \cap R^i$ and $\chi'(u) \neq \perp$, we know that u is in R^i . Thus, we conclude $B'(u) = (\alpha' \times_{\perp} R^i)(u) = \alpha'(u) = \chi'(u)$.
- ii. Assume that u is in R . If u is also in R^i , then $\chi'(u) \neq \perp$, and we use (i). If u is not in R^i , then $B'(u) = (\alpha' \times_{\perp} R^i)(u) = \perp = \chi'(u)$.
- iii. Assume that $u \in R$. Consider any M . Then u^M is in R by (10), so by (ii), $B'(u^M) = \chi'(u^M)$. Thus, according to the definition of projection, we obtain that $B(u) = \chi(u)$.

□

From (iii) of the Claim, we infer that $\chi(u) = (B \times_{\perp} R)(u)$, since the \times_{\perp} -product with R takes care of the situations when u is not in R . □

Then, the following identity holds:

$$\llbracket \text{ex2 } P^i \text{ where } \rho : \phi' \rrbracket^6 = \text{SEXP}(\text{PROJ}^{3,i}(\llbracket \phi' \rrbracket^6 \times_{\perp} \rho) \times_{\perp} \rho^{\text{RV}}(\phi)) \quad (11)$$

To see the correctness of this identity, note that $R = \{u \mid \chi(u) \neq \perp\}$ is the set of w such that $w \vDash \rho^{\text{RV}}(\phi)$ (by Proposition 9(a)); similarly, $R' = \{u \mid \chi'(u) \neq \perp\}$ is described by $\rho^{\text{RV}}(\phi')$. Additionally, R^i can be chosen to be the set represented by ρ . Thus, $\rho^{\text{RV}}(\phi')$ is equivalent to $\rho^{\text{RV}}(\phi) \wedge \rho$ (where we for convenience we assume that P^i occurs in ϕ'). Thus, the identity $R' = R \cap R^i$ holds. Also, the condition (10) holds since $\rho^{\text{RV}}(\phi)$ does not depend on the P^i -track. Thus, from the Proposition the following equalities obtain:

$$\begin{aligned} & \text{SEXP}(\text{PROJ}^{3,i}(\llbracket \phi' \rrbracket^6 \times_{\perp} \rho) \times_{\perp} \rho^{\text{RV}}(\phi)) \\ &= \text{SEXP}(\text{PROJ}^{3,i}(\alpha' \times_{\perp} R^i) \times_{\perp} R) \\ &= \text{SEXP}(\text{PROJ}^{3,i} \chi') \\ &= \text{SEXP}(\llbracket \text{ex2 } P^i \text{ where } \rho : \phi' \rrbracket^3) \\ &= \llbracket \text{ex2 } P^i \text{ where } \rho : \phi' \rrbracket^6 \end{aligned}$$

8 The sexpartite decision procedure

Propositions 18 and 19 are not quite as appealing as we would like: both involve extensive normalization. Fortunately, for first-order and $\$$ -restrictions, we can show that the normalizations are superfluous or can be made computationally inexpensive.

8.1 Orthogonality of conjunctions of thin languages

In the sequel, we will often see that the state of the R -automaton is implicitly determined by the state of a sexpartite automaton. Sometimes, knowing the state r is important, since a product construction with the R -automaton is to be carried out from r . Since it is the goal to avoid the construction of the full R -automaton, we need an effective way of representing its states. We will take advantage of the following property.

Proposition 20 Let $R = R_1 \cap \dots \cap R_n$, where each R_i is a thin language. If $u \sim_R v$ and there is a w such that $u \cdot w \in R$, then for all i , $1 \leq i \leq n$, it holds that $u \sim_{R_i} v$.

Proof Assume that $u \sim_R v$ and $u \cdot w \in R$ hold. Moreover, for a contradiction, assume that for some i , $u \not\sim_{R_i} v$ holds. By assumption that R_i is a thin language, either $u \cdot w \notin R_i$ or $v \cdot w \notin R_i$ holds, whence from the fact $u \cdot w \in R_i$, we infer that $v \cdot w \notin R_i$ holds. But then, we would have that $v \cdot w \notin R$, which contradicts the assumptions that $u \cdot w \in R$ and $u \sim_R v$. \square

The proposition entails that if we know that a state q in an automaton A determines a state r of the automaton for R (in the sense that $\forall u : \text{LAST}A(u) = q \Rightarrow \text{LAST}R(u) = r$) and r is not the rejecting sink state, then r is determined by the tuple of states of the R^i -automaton.

8.2 Crispness

The identification of different interesting states (after collapse and minimization) that we saw in the example of Section 5.3 makes it difficult in general to recover the state of the R -automaton from the sexpartification. Fortunately, this phenomenon does not occur for the restrictions that we interested in.

Let us consider languages L and R as expressed through $\chi = \chi_{L,R}$. Let us say that string u is *readily distinguished* from v if there is a letter $a \in \Sigma$ such that for all w it holds that $u \cdot a \cdot w \notin R$ but there is a \hat{w} such that $u \cdot a \cdot \hat{w} \in R$. Intuitively speaking, the automaton for R is in a non-accepting sink (N^0 or N^\perp) after reading $u \cdot a$, but it goes into a state from where it can still accept after reading $v \cdot a$. We say that R is *crisp* if for all u and v with $u \not\sim_R v$, either u is readily distinguished from v or v is readily distinguished from u .

Crispness of R ensures that sexpartification does not merge interesting states of χ that are not equivalent with respect to \sim_R .

Proposition 21 Assume that R is crisp.

- (a) Then, $\forall u, v : \iota_{L,R}(u) \wedge u \sim_{L,R}^6 v \Rightarrow u \sim_R v$ holds.
- (b) Moreover, if R is also thin and u is interesting, then $u \sim_{L,R} v \Leftrightarrow u \sim_{L,R}^6 v \Leftrightarrow u \sim_{L \cap R} v$.
- (c) Consider the automaton A for $L \cap R$ and the α -automaton representing the approximation valuation for L and R . Then there is a transition-respecting mapping f from the states of A , except the rejecting sink state, to the states of the α -automaton except for the N^0 and N^\perp states. On this domain and codomain, f is surjective.

Proof

- (a) The proof is by contraposition. We assume that $\iota(u)$ and $u \not\sim_R v$. If $\iota(v)$ does not hold, then $u \not\sim_{L,R}^6 v$ holds, because $\iota(u)$ holds by assumption. Thus, we may further assume that $\iota(v)$ holds. Let q' be the state reached in the χ -automaton on u , and let q'' be defined similarly for v . We may assume that it is u that is not readily distinguished from v according to some letter a . Then, it will be the case that (q', a, N^\perp) and (q'', a, q) are transitions of the χ -automaton, where q is different from N^\perp . Consequently, the states q' and q'' are not united during the minimization phase of sexpartification.
- (b) The first bi-implication follows from (a) and Proposition 15(c). For the second bi-implication, the relation \Rightarrow holds because $u \cdot w \in L \cap R$ if and only if ($u \cdot w \in R$ implies $u \cdot w \in L$) if and only if (by the exactness property and by the antecedent $u \sim_{L,R}^6 v$) if and only if ($v \cdot w \in R$ implies $v \cdot w \in L$) if and only if $v \cdot w \in L \cap R$ for; the direction \Leftarrow holds by virtue Proposition 11(d), which relies on thinness of R .
- (c) It can be seen that the subautomaton induced by the interesting states of the $\alpha_{L,R}$ -automaton is identical, except for state labels, to a subautomaton of A . (A subautomaton maybe incomplete: for some states and labels, no outgoing transitions may be defined.) Consequently, the notion of interesting state make sense for A . A transition out of this set off interesting states corresponds in $\chi_{L,R}^6$ to a transition to the N^1 , N^\perp , or N^0 state. In the latter two cases, the A automaton proceeds to a rejecting sink state (this is why we cannot map the rejecting sink state to a state of the $\alpha_{L,R}$ -automaton). In the first case, the A automaton state reached is mapped to N^1 ; so is every other state, except the rejecting sink state, that is further reachable from this state.

□

Part (c) of this Proposition is a key property: it tells us that the automaton of the conjunctive representation is the same as the automaton of the sexpartite representation except that the latter represents by at most three states the uninteresting states.

Proposition 22 Any language $R \subseteq \Sigma^k$ that is the conjunction of $R_{\mathcal{S}\text{-restrict}(i)}$ for $i < k$ and $R_{\text{singleton}(i)}$ for $i \in \mathcal{S}$, where $\mathcal{S} \in \{0, \dots, k-1\}$ and where the P^k -track is used to encode the variable \mathcal{S} , is crisp.

Proof We proceed somewhat informally by directly studying the automaton A for R . In particular, we fix $k = 3$ and we assume that $\mathcal{S} = \{1, 2\}$. The automaton A is shown in Figure 10. We note that it has an accepting state, reached only after the first 1 in the \mathcal{S} -track has occurred and after each of the first-order variable tracks has seen exactly one occurrence of a 1. It also has a rejecting sink. The remaining 2^2 states keep track of which 1s have occurred for first-order tracks. All states except the sink state has a transition to the sink state in addition to a transition to a non-sink state. \square




Figure 10: The automaton for first-order and \mathcal{S} -restrictions on P^1 and P^2 .

The two preceding propositions tell us that the sextuplet representation of $M2L(\text{Str})$ subformulas in an WS1S encoding is identical to that of the ternary representation, except for the collapsing of N^Z states in the ternary automaton. In particular, each interesting state specifies the state of each restriction automaton. This fact will help us formulate algorithms for conjunction and projection that avoid most explicit normalization.

8.3 Algorithm for conjunction

Given automata for $\alpha' = \alpha_{L', R'}$ and $\alpha'' = \alpha_{L'', R''}$, which are approximations based on crisp, thin, and compatible R' and R'' , a product automaton $B = (\Sigma, Q, q_0, \rightarrow, \lambda)$

recognizing $\alpha_{L' \cap L'', R' \cap R''}$ can be constructed as follows. We assume that the α' -automaton is $(\Sigma, Q', q'_0, \rightarrow', \lambda')$ and the α'' -automaton is $(\Sigma, Q'', q''_0, \rightarrow'', \lambda'')$. The language R' is represented by an automaton $(\Sigma, Q_{R'}, q_{R'}^0, \rightarrow_{R'}, Q_{R'}^F)$ and R'' is represented by an automaton $(\Sigma, Q_{R''}, q_{R''}^0, \rightarrow_{R''}, Q_{R''}^F)$. The state of the R' -automaton that is a rejecting sink is denoted r'_{rej} (it exists by the assumption that R' is crisp); the state r''_{rej} of the R'' -automaton is defined similarly.

By the above assumptions and Proposition 22(b), any interesting state q' of the α' -automaton determines a state r' of the R' -automaton; a similar observation holds for the α'' -automaton. We now let

$$Q = Q' \times Q'' \cup Q' \times Q_{R'} \cup Q_{R'} \times Q'' \cup \{N^1, N^0, N^\perp\}.$$

Assuming again that the initial states are interesting (to avoid various special cases), we let $q_0 = (q'_0, q''_0)$. To define the transition relation \rightarrow of B , we introduce the relation $\overset{a}{\rightsquigarrow}'$ on $Q' \cup (Q_{R'} \setminus \{r'_{\text{rej}}\}) \cup \{N^0, N^\perp\}$ defined according to:

$$\begin{aligned} q' \overset{a}{\rightsquigarrow}' \hat{q}' & \text{ if } (q', a, \hat{q}') \in \rightarrow' \text{ and } \hat{q}' \text{ is interesting} \\ q' \overset{a}{\rightsquigarrow}' \hat{r}' & \text{ if } (q', a, \hat{q}') \in \rightarrow', \hat{q}' \text{ is the } N^1\text{-state, and } \hat{r}' \text{ is determined by } \hat{q}' \\ q' \overset{a}{\rightsquigarrow}' N^\perp & \text{ if } (q', a, \hat{q}') \in \rightarrow' \text{ and } \hat{q}' \text{ is the } N^\perp\text{-state} \\ q' \overset{a}{\rightsquigarrow}' N^0 & \text{ if } (q', a, \hat{q}') \in \rightarrow' \text{ and } \hat{q}' \text{ is the } N^0\text{-state} \\ r' \overset{a}{\rightsquigarrow}' \hat{r}' & \text{ if } (r', a, \hat{r}') \in \rightarrow_{R'} \text{ and } \hat{r}' \neq r'_{\text{rej}} \\ r' \overset{a}{\rightsquigarrow}' N^\perp & \text{ if } (r', a, \hat{r}') \in \rightarrow_{R'} \text{ and } \hat{r}' = r'_{\text{rej}} \end{aligned}$$

Note that in the second line above \hat{r}' is not r'_{rej} because \hat{q}' is the N^1 -state (and \hat{q}' thus determines a state of the R' -automaton that is not r'_{rej}). The transition relation $\overset{a}{\rightsquigarrow}''$ on $Q'' \cup (Q_{R''} \setminus \{r''_{\text{rej}}\}) \cup \{N^0, N^\perp\}$ is defined similarly.

Now, define \rightarrow of the automaton B to consist of the following transitions for a letter a , where $s' \overset{a}{\rightsquigarrow}' \hat{s}'$ and $s'' \overset{a}{\rightsquigarrow}'' \hat{s}''$:

$$\begin{aligned} ((s', s''), a, (\hat{s}', \hat{s}'')) & \text{ if } s', s'' \notin \{N^0, N^\perp\} \text{ and } (s' \in Q' \text{ or } s'' \in Q'') \\ ((s', s''), a, N^\perp) & \text{ if } \hat{s}' = N^\perp \text{ or } \hat{s}'' = N^\perp \\ ((s', s''), a, N^1) & \text{ if } s' \in Q_{R'} \text{ and } s'' \in Q_{R''} \\ ((s', s''), a, N^0) & \text{ if } (\hat{s}' = N^0 \text{ and } \hat{s}'' \neq N^\perp) \text{ or } (\hat{s}'' = N^0 \text{ and } \hat{s}' \neq N^\perp) \end{aligned}$$

Here, N^0 , N^1 , and N^\perp are sink states, for which λ takes on the value 0, 1, and \perp , respectively. For states of the form (q', q'') , we define $\lambda((q', q'')) = \lambda'(q') \wedge^3 \lambda''(q'')$. For states of the form (q', r'') , we define $\lambda((q', r'')) = \lambda'(q')$, if $r'' \in Q_{R''}^F$, and $\lambda((q', r'')) = \perp$, if $r'' \notin Q_{R''}^F$. As in the usual algorithmic treatment of product automata, we consider in the following only the states of B that are reachable from the initial state.

Proposition 23 The automaton B constructed above recognizes $\alpha_{L' \cap L'', R' \cap R''}$. Moreover, there is a surjective mapping from the reachable states of the non-minimized product automaton for $(L \cap R')$ and $(L'' \cap R'')$, except for the rejecting sink state, to the reachable states of B (except for one state).

Proof (Idea) Using the exactness property, we see that $\chi_{L' \cap L'', R' \cap R''} = (\alpha' \wedge^3 \alpha'') \times_\perp (R' \cap R'') = (\alpha' \times_\perp R'') \wedge^3 (\alpha'' \times_\perp R')$. Thus, it holds that $\alpha_{L' \cap L'', R' \cap R''} =$

$\text{SEXP}((\alpha' \times_{\perp} R'') \wedge^3 (\alpha'' \times_{\perp} R'))$. Now using Proposition 17, we may evaluate $\text{SEXP}((\alpha' \times_{\perp} R'') \wedge^3 (\alpha'' \times_{\perp} R'))$ as

$$\text{SEXP}(\text{SEXP}(\alpha' \times_{\perp} R'') \wedge^3 \text{SEXP}(\alpha'' \times_{\perp} R')) \quad (12)$$

To reflect the expression $\text{SEXP}(\alpha' \times_{\perp} R'') \wedge^3 \text{SEXP}(\alpha'' \times_{\perp} R')$ in (12), we make an automaton C based on tuples of the form $((q', r''), (q'', r'))$ when states q' and q'' are both interesting. The states r' and r'' are determined by q' and q'' , respectively. The assumption of compatibility ensures that if both q' and q'' are interesting then (q', r'') of $\alpha' \times_{\perp} R''$ is also interesting. The pair (q', r'') can be viewed as a state of $\text{SEXP}(\alpha' \times_{\perp} R'')$ because the interesting part of $\alpha' \times_{\perp} R''$ is locally isomorphic to $\text{SEXP}(\alpha' \times_{\perp} R'')$ (due to crispness and thinness and Proposition 21(b)).

But since r' and r'' are determined already by q' and q'' , respectively, we may generate an isomorphic subautomaton for states of the form $((q', r''), (q'', r'))$ by just considering tuples of the form (q', q'') . This is how the automaton B acts for interesting states: it simulates C . We note that during the construction of the $Q' \times Q''$ -part of the state space of B it is possible to simultaneously keep track of the determined r' and r'' states—information that is needed in some cases when q' or q'' turn to N^1 states. We omit a detailed discussion, but we remark that this auxiliary information is needed only for the frontier or queue of not-yet-fully explored product states; the information does not need to be stored along with the states.

We are not done describing the simulation. Say that on some transition the pair (q', q'') becomes (\hat{q}', \hat{q}'') with \hat{q}' being the N^1 -state and \hat{q}'' still interesting. The sexpartification of $\alpha' \times_{\perp} R''$ then yields an N^1 -state as well. Therefore, the automaton C enters a state $(N^1, (\hat{q}'', \hat{r}'))$. Thus, to continue simulating the C automaton using pairs, we let the B automaton enter the state (\hat{r}', \hat{q}'') ; this is explained in detail through the rules that define its transitions. If instead a transition turns (q', q'') into (\hat{q}', \hat{q}'') with \hat{q}' being the N^0 -state, the the simulation of the C automaton may be stopped, since the outer sexpartification in (12) will yield either an N^0 or an N^{\perp} state. Other cases are explained in a similar vein.

Finally, the surjective mapping from the product automaton of the conjunctive semantics is constructed from the mappings f' and f'' that exist according to Proposition 21(c). \square

8.4 Algorithm for projection

To formulate an algorithm for projection that largely does away with normalizations, we substitute the identity $\text{PROJ}^{3,i}A = \text{BPROJ}^{3,i}(\text{FUT}^{3,i}A)$ of Proposition 10 in the reformulation of projection in Proposition 19. Thus, our starting point is the identities

$$\begin{aligned} \alpha_{L,R} &= \text{SEXP}(\text{PROJ}^{3,i}\chi_{L',R'}) \\ &= \text{SEXP}(\text{PROJ}^{3,i}(\alpha_{L',R'} \times_{\perp} R^i) \times_{\perp} R) \\ &= \text{SEXP}(\text{BPROJ}^{3,i}(\text{FUT}^{3,i}(\alpha_{L',R'} \times_{\perp} R^i)) \times_{\perp} R) \end{aligned}$$

where we use the notations of Proposition 19. The challenge is to construct a subset automaton E recognizing $\text{BPROJ}^{3,i}(\text{FUT}^{3,i}(\alpha_{L',R'} \times_{\perp} R^i))$ and to avoid the subse-

quent product $\times_{\perp} R$. To do so we focus on representing $\alpha_{L',R'} \times_{\perp} R^i$ as an equivalent automaton D . We note that each interesting state q of the α' -automaton (where $\alpha' = \alpha_{L',R'}$) determines a state r^i of R^i . So, in analogy with the construction of the automaton for conjunction, we may omit in D the explicit construction of pairs corresponding to $\alpha_{L',R'} \times_{\perp} R^i$. And, when the $\alpha_{L',R'}$ -automaton exits to an N^1 -state, the automaton D simulates R^i whose rejecting states are then relabeled \perp . A state from this copy of R^i is denoted r^i_+ . Similarly, when the $\alpha_{L',R'}$ -automaton exits to an N^0 -state, D also simulates R^i , but the accepting states are now labeled \emptyset and the rejecting states are labeled \perp . A state from this copy of R^i is denoted r^i_- . When $\alpha_{L',R'}$ -automaton exits to N^{\perp} , D also enters a N^{\perp} -state. To avoid the $\times_{\perp} R$ product, we note that any subset of D containing an interesting state q of α' -automaton determines the state r of R ; also, any subset containing no interesting states can be replaced by some N^Z -state, and the $\times_{\perp} R$ -product is again unnecessary, because of the outer sexpartification. Thus, the effect of the $\times_{\perp} R$ -product can be effectuated by a simultaneous traversal of the subset automaton and the automaton R , where subset states not satisfying R are labeled \perp . We also note that if a subset contains say both N^0 and N^{\perp} then N^{\perp} can be removed without changing the accepted language; more generally, we may prune all subsets states so that they contain at most one non-interesting state.

Proposition 24 Given the assumptions of Proposition 19 and the further assumption that restrictions are crisp and thin and that the automaton E described above recognizes $\text{SEXP}(\text{PROJ}^{3,i}\chi_{L',R'})$.

Moreover, the number of states of E is at most $N \cdot 2^{|R_i|+1}$, where N is the number of subset states encountered when determinizing $\text{BPROJ}^i(\text{FUT}^{3,i}F)$, where F recognizes $L' \cap R'$.

Proof (Idea) We have already outlined the reasons why the above construction is correct. As regards the size of E , we study the partial mapping f from states of the automaton F recognizing $L' \cap R'$ to states of the α' -automaton that exists according to Proposition 21(c). We note that this mapping is undefined only for the rejecting sink state q_{rej} of F , since there are two corresponding states, N^0 and N^{\perp} , in the α' -automaton.

We may define a mapping g that maps states q of F to subsets of C defined as follows: if q is interesting then let $g(q) = \{q'\}$; if q is not interesting, but allows an accepting extension, then let r^i be the state of the R^i -automaton determined (by Proposition 11(d) and Proposition 20) and let $g(q) = \{r^i_+\}$; and if q is q_{rej} , then $g(q) = \{N^{\perp}, N^0\} \cup \{r^i_-\} \mid r^i \neq r_{\text{rej}}\}$. Then, for any string u , the subset N of states of E reachable on u is related to the subset M of reachable states of $\text{FUT}^{3,i}F$ on u :

$$N = \cup_{q \in M} \sigma(q)$$

for some function σ that selects a nonempty subset of $g(q)$, where $q \in M$. The numeric bound of the proposition follows from combinatorial properties of g : only one q is mapped to a non-singleton set and that set has $|R_i| + 1$ members. \square

8.5 The sextartite decision procedure compared

Theorem 2 For first-order and $\$$ -restrictions, WS1S can be decided in a way such that

- (a) The sizes of the intermediate automata occurring during the sextartite decision procedure are at most those of the conjunctively normalized semantics, except for the subset construction (and before minimization), where the automata of the sextartite decision may be up to 32 times bigger. These bounds ignore an additive constant of 1.
- (b) The conjunctively normalized automata may be exponentially bigger than the sextartite automata.

Proof

- (a) That the property holds for minimized automata follows from Proposition 21(b). The bounds on the number of states come from Proposition 23 and Proposition 24, where we have used the fact that the biggest R^i -automata stem from a simultaneous $\$$ -restriction and a first-order restriction, which yield an automaton with 4 states. Thus, the factor of Proposition 24 is $2^{4+1} = 32$.
- (b) It suffices to consider a subformula of the form $\phi_{1st-vars} = ((\dots(p^1 = p^1 \ \& \ \dots) \ \& \ p^j = p^j)$, where each p^i is a first-order variable. The sextartite automaton has one state, namely an N^1 state, whereas the automaton for the conjunctively normalized semantics has $2^j + 2$ states. To see this, for $j = 2$, contemplate the automaton of Figure 10, which is identical to the automaton obtained under the conjunctively normalized semantics and which has $2^2 + 2 = 6$ states. The sextartification for this formula operates on the same automaton except that the 5 rejecting states are turned into \perp -states. Consequently, a single N^1 state arises from the collapse.

□

9 In practice

We showed experimental evidence in [8] that WS1S could be as fast a way to decide string-theoretic problems as M2L(Str) but only after sometimes solving by hand state explosion problems like the one discussed in Section 2.4.

Since June 1998, the Mona tool has been based on the ternary semantics for WS1S, and our state explosion problems stemming from running M2L(Str) formulas through WS1S have disappeared. Moreover, with a default relativization mechanism that we have added to Mona, M2L(Str) formulas can be directly embedded in WS1S. The running times under these semantics are in all non-contrived cases the same (to within 5% or so) as for the ad hoc semantics we used before. (In practice, we used first-order relativizations that are not thin languages, but which enjoy similar properties.) Thus, we can state that also from a practical point of view the techniques presented here solve

both the translation problem and the first-order semantics problem. We have not yet implemented the sexpartite semantics.

Future work should look at factorization techniques that would be more generally applicable. Consider for example the formula $P = Q \ \& \ \phi_{1st-vars}$, where $\phi_{1st-vars}$ simply introduces a number j of first-order variables (as described in the proof of Theorem 2). This formula produces automata with a number of states exponential in j , even under sexpartification.

Also the sexpartite semantics should be investigated in practice; the factor 32 blow-up of subset automata before minimization may turn out to be a theoretical limit that is never encountered in practice—perhaps, a tighter analysis will show that it is too pessimistic.

Acknowledgements Anders Møller implemented the ideas presented here and contributed many useful insights. Jacob Elgaard found exploding Mona code from which the parity example was derived. Ken McMillan kindly discussed relativization issues with me. I thank the referees for their many positive and detailed encouragements to complete the job. The patience of the editor, Olivier Danvy, is also kindly acknowledged.

References

- [1] Abdelwaheb Ayari and David A. Basin. Bounded model construction for monadic second-order logics. In *Computer Aided Verification*, pages 99–112, 2000.
- [2] David Basin and Nils Klarlund. Automata based symbolic reasoning in hardware verification. *Formal Methods in System Design*, pages 255–288, 1998. Extended version of “Hardware verification using monadic second-order logic,” *Computer aided verification : 7th International Conference, CAV '95*, LNCS 939, 1995.
- [3] John Bell and Moshe Machover. *A course in Mathematical Logic*. North-Holland, 1977.
- [4] R. E. Bryant. Symbolic Boolean manipulation with ordered binary-decision diagrams. *ACM Computing surveys*, 24(3):293–318, September 1992.
- [5] J.R. Büchi. Weak second-order arithmetic and finite automata. *Z. Math. Logik Grundl. Math.*, 6:66–92, 1960.
- [6] C.C. Elgot. Decision problems of finite automata design and related arithmetics. *Trans. Amer. Math. Soc.*, 98:21–52, 1961.
- [7] J.G. Henriksen, J. Jensen, M. Jørgensen, N. Klarlund, B. Paige, T. Rauhe, and A. Sandholm. Mona: Monadic second-order logic in practice. In *Tools and Algorithms for the Construction and Analysis of Systems, First International Workshop, TACAS '95, LNCS 1019*, 1996.
- [8] Anders Møller Jacob Elgaard, Nils Klarlund. Mona 1.x: new techniques for ws1s and ws2s. In *Computer Aided Verification, CAV '98, Proceedings*, volume 1427 of LNCS. Springer Verlag, 1998.
- [9] P. Kelb, T. Margaria, M. Mendler, and C. Gsottberger. Mosel: a flexible toolset for Monadic Second-order Logic. In *Computer Aided Verification, CAV '97, Proceedings*, LNCS 1217, 1997.
- [10] N. Klarlund. Mona & Fido: the logic-automaton connection in practice. In *CSL '97 Proceedings*. LNCS 1414, Springer-Verlag, 1998.
- [11] Nils Klarlund and Anders Møller. *MONA Version 1.3 User Manual*. BRICS, 1998. URL: <http://www.brics.dk/mona>.
- [12] Nils Klarlund, Anders Møller, and Michael I. Schwartzbach. MONA implementation secrets. *International Journal of Foundations of Computer Science*, 13(4):571–586, 2002.
- [13] Anders Møller and Michael Schwartzbach. The Pointer Assertion Logic Engine. In *Proceedings of ACM SIGPLAN Conference of Programming Language Design and Implementation*, 2001.

- [14] Mark Smith and Nils Klarlund. Verification of a sliding window protocol using IOA and MONA. In *FORTE/PSTV 2000: IFIP TC6 WG6.1 Joint International Conference on Formal Description Techniques for Distributed Systems and Communication Protocols (FORTE XIII), and Protocol Specification, Testing, and Verification (PSTV XX)*, pages 19–34. Kluwer Academic Publishers, 2000.
- [15] Howard Straubing. *Finite Automata, Formal Logic, and Circuit Complexity*. Birkhäuser, 1994.
- [16] Wolfgang Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, chapter Languages, automata, and logic. Springer Verlag, 1997.
- [17] B.A. Trakhtenbrot. Finite automata and the logic of one-place predicates. *Sib. Math. J.*, 3:103–131, 1962. In Russian. English translation: *AMS Transl.*, 59 (1966), pp. 23-55.
- [18] Nathan Vaillette. Logical specification of finite-state transductions for NLP. *Natural Language Engineering*, ??? to appear.