

Convergence Testing in Term-Level Bounded Model Checking

Randal E. Bryant Shuvendu K. Lahiri Sanjit A. Seshia
June 2003
CMU-CS-03-156

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

A shorter version of this paper will appear at CHARME '03.

Abstract

We consider the problem of bounded model checking of systems expressed in a decidable fragment of first-order logic. While model checking is not guaranteed to terminate for an arbitrary system, it converges for many practical examples, including pipelined processors. We give a new formal definition of convergence that generalizes previously stated criteria. We also give a sound semi-decision procedure to check this criterion based on a translation to quantified separation logic. Preliminary results on simple pipeline processor models are presented.

This research was supported in part by the Semiconductor Research Corporation, Contract RID 1029 and by ARO grant DAAD 19-01-1-0485.

The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained in this document are those of the authors, and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Department of Defense or the U.S. Government.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE JUN 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE Convergence Testing in Term-Level Bounded Model Checking				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University, School of Computer Science, Pittsburgh, PA, 15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 21	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Keywords: Term-level verification — Convergence in Model Checking — Symbolic Simulation — Uninterpreted functions — Second-order Logic — Decision procedures — Quantified Separation Logic — Processor verification

1 Introduction

Systems with parameters of finite but arbitrary or large size are often modeled as infinite-state systems. Such systems include superscalar processors, communication protocols with unbounded channels, and networks of an arbitrary number of identical processes. While state elements can still be of Boolean type, richer data types such as unbounded integers or unbounded arrays of integers are also used. Employing this richer expressive power is one approach to tackling the state explosion problem.

In the area of hardware verification, the logic of Equality with Uninterpreted Functions and Memories (EUFM) has been successfully used for the automated verification of pipelined processor designs [7, 3]. The more general logic of Counter Arithmetic with Lambda Expressions and Uninterpreted Functions [4] (CLU) has been used for bounded model checking and inductive invariant checking of out-of-order microprocessors with unbounded resources [14]. Bounded model checking proceeds by symbolically simulating the system for a finite number of steps starting from an initial state, checking on each step that a state property holds. As the state elements can be terms in a first-order logic, we will refer to this technique as *term-level bounded model checking*. Since term-level models can express Turing machines [12], the symbolic simulation might never reach a fixpoint in general. However, in many practical cases, the simulation does converge. It is therefore necessary to check, after each simulation step, whether the simulation has converged. Term-level bounded model checking is also useful in combination with other techniques such as Burch-Dill style verification [7], since it provides a way to compute the most general reachable state in which to initialize the system when using those techniques.

In this paper, we make two main contributions. First, we give a new formal definition of convergence for term-level bounded model checking, where CLU logic is used as the modeling formalism. The convergence criterion is formulated as a quantified second-order formula with one quantifier alternation, and is undecidable in general. Second, we give two semi-decision procedures for this class of second-order formulas, the first being sound and the second being complete. Our procedures are based on a translation to a decidable fragment of first-order logic called *quantified separation logic* (QSL). QSL formulas are quantified Boolean combinations of Boolean variables and predicates of the form $x_i < x_j + c$ or $x_i = x_j + c$, where x_i and x_j are real or integer variables, and c is a constant. The QSL formulas are then decided by a translation to quantified Boolean logic [16]. Although we use the semi-decision procedures for convergence checking, our results are also more generally applicable to automated theorem proving of second-order formulas.

Previous term-level model checkers vary in expressiveness of the underlying logic, and either use syntactic convergence criteria or approximation techniques that guarantee convergence at the cost of completeness. Hojati et al. [12] presented a modeling formalism called ICS which is similar in expressiveness to EUFM. They showed that ICS models do not converge in general, except under highly restrictive assumptions that are not of practical interest. Isles et al. [13] built on this work, giving a conservative, syntactic definition of convergence of ICS models, and using it to verify versions of the DLX pipeline. Our logic is more expressive than ICS. Also, as we show in Section 5.2, their convergence criterion is a special case of the one we present in this paper. Corella et al. [8] have used Multiway Decision Graphs (MDGs) for term-level model checking. MDGs are BDD-like data structures used for representing formulas in quantifier-free logics such as EUFM and CLU; the exact logic represented depends on the set of interpreted function symbols used in the model. Thus, Corella et al. use MDGs to represent the characteristic function of the set of states of a term-level model. Unlike our work, their models cannot have variables of function type, and hence

cannot verify systems with embedded memories. However, they address a more general class of properties expressible in a first order temporal logic. With respect to convergence checking, Corella et al. use syntactic rewriting techniques similar to those employed for ICS [13]. Bultan et al. [5] have used Presburger arithmetic for verifying concurrent algorithms. Checking convergence for systems expressed in Presburger arithmetic is decidable; however, since the model checking might not converge in general, they conservatively approximate the fixpoint, allowing the possibility of spurious counterexamples. In comparison, our use of CLU logic allows us to use uninterpreted functions and also lets us model richer systems with memories. This expressive power, however, results in convergence checking becoming undecidable.

The rest of the paper is organized as follows. Section 2 presents CLU logic and our system modeling formalism. Section 3 defines the term-level bounded model checking problem. In Section 4, we formally define the convergence criterion. Section 5 describes how we check this criterion. Finally, we conclude in Section 6 with some preliminary results with pipelined processor models. Detailed proofs of the theorems can be found in the appendix.

2 Preliminaries

2.1 CLU Logic

Syntax. The syntax includes four classes of *expressions*, representing computations of truth values or integers, as well as functions over integers yielding truth values or integers. We use *symbols* to

$$\begin{aligned}
\textit{bool-expr} & ::= \mathbf{true} \mid \mathbf{false} \mid \textit{bool-symbol} \mid \neg \textit{bool-expr} \mid (\textit{bool-expr} \wedge \textit{bool-expr}) \\
& \quad \mid (\textit{int-expr} = \textit{int-expr}) \mid (\textit{int-expr} < \textit{int-expr}) \\
& \quad \mid \textit{predicate-expr}(\textit{int-expr}, \dots, \textit{int-expr}) \\
\textit{int-expr} & ::= \textit{lambda-var} \mid \textit{int-symbol} \mid \textit{ITE}(\textit{bool-expr}, \textit{int-expr}, \textit{int-expr}) \\
& \quad \mid \textit{int-expr} + \textit{int-constant} \mid \textit{function-expr}(\textit{int-expr}, \dots, \textit{int-expr}) \\
\textit{predicate-expr} & ::= \textit{predicate-symbol} \mid \lambda \textit{lambda-var}, \dots, \textit{lambda-var} . \textit{bool-expr} \\
\textit{function-expr} & ::= \textit{function-symbol} \mid \lambda \textit{lambda-var}, \dots, \textit{lambda-var} . \textit{int-expr}
\end{aligned}$$

Figure 1: **Expression Syntax.** Expressions can denote computations of Boolean values, integers, or functions yielding Boolean values or integers.

represent abstract values and functions. Symbols are written with a typewriter font, such as **a** or **f**. Associated with each symbol is a *type* indicating what kind of value it represents (truth, integer, function, or predicate). For function and predicate symbols, the type includes its *arity* indicating the number of arguments it takes. For function symbol **f**, we write its arity as *arity*(**f**). For a set of symbols \mathcal{A} , we let $E(\mathcal{A})$ denote the set of all expressions that can be formed using these symbols, obeying the usual rules on type matching.

The syntax includes integer *lambda variables*. These only serve to represent the arguments to lambda expressions. Note also that the lambda expression syntax is constrained so that they cannot have functions as arguments, and they cannot express any form of looping or recursion.

Sets of Expressions. We use two ways to refer to sets of expressions in which we must identify the different elements. The first is a *vector notation*, in which we index the elements with integer

subscripts. We use the notation \overline{e}_n to denote a vector with elements e_1, \dots, e_n . The second is a *named-element notation*, in which we have a set of symbolic names \mathcal{A} and write a set of expressions e as having an element e_a for each $a \in \mathcal{A}$.

With both notations, we can indicate the syntactic substitution of elements for symbols or variables in an expression. That is, the expression $s[\overline{e}_n/\overline{x}_n]$ denotes the expression where each instance of x_i in s is replaced by the expression e_i for $1 \leq i \leq n$. These substitutions are performed in parallel, so there is no ambiguity of some expression e_i contains the symbol x_j . Similarly, $s[\overline{e}/\mathcal{A}]$ indicates the result of replacing each instance of a symbol $a \in \mathcal{A}$ with the expression e_a .

Semantics. For a set of symbols \mathcal{A} , we let $\sigma_{\mathcal{A}}$ indicate an *interpretation* of each of these symbols. That is, $\sigma_{\mathcal{A}}$ maps each symbol to an integer, a truth value, or a function according to the symbol type. For any expression $e \in E(\mathcal{A})$, we define its *evaluation under interpretation* $\sigma_{\mathcal{A}}$, denoted $\langle e \rangle_{\sigma_{\mathcal{A}}}$ as the value obtained by evaluating e when each symbol a is replaced by its interpretation $\sigma_{\mathcal{A}}(a)$. We omit the detailed definition.

A truth expression $e \in E(\mathcal{A})$ is said to be *universally valid* when it evaluates to **true** for all interpretations of its symbols, i.e., when $\langle e \rangle_{\sigma_{\mathcal{A}}} = \mathbf{true}$ for all $\sigma_{\mathcal{A}}$.

As a final notation, for disjoint symbol sets \mathcal{A} and \mathcal{B} , each having interpretations $\sigma_{\mathcal{A}}$ and $\sigma_{\mathcal{B}}$, we let $\sigma_{\mathcal{A}} \cdot \sigma_{\mathcal{B}}$ denote the interpretation over the symbols in $\mathcal{A} \cup \mathcal{B}$ obtained by applying the respective interpretations to the symbols in \mathcal{A} and \mathcal{B} .

As noted earlier, our syntax for function applications requires all arguments to be integer expressions. We can therefore transform any integer or truth expression containing lambda expressions into an equivalent lambda-free one by performing *Beta reduction*, in which the actual parameter expressions are syntactically substituted in parallel with the actual parameter expressions.

2.2 System Model

We model the system as having a number of *state elements*, where each state element may be a truth or integer value, or a function or predicate. This latter class of state elements allows us to describe various forms of memories. For example, a conventional random-access memory can be modeled as a function that yields an integer data value given an integer address as argument. We use symbolic names to represent the different state elements giving the set of *state symbols* \mathcal{S} . We also introduce a set of *input symbols* \mathcal{T} , representing a set of input signals that can be set to different values on each step of operation. That is, on each step i , we introduce a symbol a_i for each input symbol a . We refer to such signals as the *indexed input symbols*. We introduce two more sets of symbols \mathcal{K} and \mathcal{I} to allow one run by the verifier to compute the behavior of systems with different functionality operating with different initial state and input values. The symbols in \mathcal{K} parameterize system functionality. This could include, for example, function symbols for the ALU, and the contents of the instruction memory. The symbols in \mathcal{I} parameterize the initial state and system input sequence. These could include a function symbol to encode the initial state of a memory. They also include the indexed input symbols.

The overall system operation is characterized by an *initial state* s^0 and a *transition behavior* δ . The initial state contains an expression for each state element. The initial value of state element a is given by an expression $s_a^0 \in E(\mathcal{I})$. The transition behavior consists of an expression for each state element. The behavior for state element a is given by an expression $\delta_a \in E(\mathcal{K} \cup \mathcal{S} \cup \mathcal{T})$. In this expression, we use the state element symbols to represent the current system state, and the input symbols to represent the current values of the inputs. The expression then gives the new state for

that state element.

From these expressions, we define the *state sequence* for the system s^0, \dots, s^i, \dots , where the state at step i consists of an expression for each state element $s_{\mathbf{a}}^i \in E(\mathcal{K} \cup \mathcal{I})$. This expression is given by performing the double substitution

$$s_{\mathbf{a}}^i = \delta_{\mathbf{a}} [s^{i-1}/\mathcal{S}, t^i/\mathcal{T}], \quad (1)$$

where the input expression t^i has $t_{\mathbf{a}}^i = \mathbf{a}_i$ for each $\mathbf{a} \in \mathcal{I}$. As mentioned earlier, we always perform Beta reduction following a substitution such as this. We use the shorthand $s^i = \delta(s^{i-1}, t^i)$ to indicate this process of generating the expressions for the state at step i .

3 Property Checking

A *system property* P is represented as a Boolean expression over the state elements $P \in E(\mathcal{S})$. Typically we want to determine whether P holds at some particular step k , or whether P holds at every step. We can determine whether P holds at some particular step k by applying a decision procedure for CLU logic. However, our interest here is to prove that P holds for every step $i \geq 0$. In general, this task is undecidable. The problem remains undecidable even if we restrict the class of systems to ones with only integer state elements, and where the system behavior is described using a logic of equality with uninterpreted functions [12].

Instead, we focus on a more restricted class of systems that satisfy a property we call *k-convergence*. With these systems, every reachable state can be reached within k steps for some combination of initial state and inputs, for some fixed bound k . If we can prove that a system is *k-convergent*, then we can guarantee property P holds on every step by verifying that it holds on every step up through s^k .

Formally, we say that a system with initial state s^0 and transition behavior δ *converges in k steps*, when for every interpretation $\sigma_{\mathcal{I}}$ of the initial state and inputs and for every interpretation $\sigma_{\mathcal{K}}$ of the system parameters, there exists a step $i \leq k$ and an alternate interpretation $\theta_{\mathcal{I}}$ of the initial state and inputs, such that for every state symbol $\mathbf{a} \in \mathcal{S}$

$$\langle s_{\mathbf{a}}^i \rangle_{\theta_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \langle s_{\mathbf{a}}^{k+1} \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}}. \quad (2)$$

We use the shorthand $\langle s^i \rangle_{\theta_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \langle s^{k+1} \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}}$ to indicate this equality for every state element. Property (2) states that by step $k+1$, the system will not reach any new states. That is, for every possible interpretation of the system parameters $\theta_{\mathcal{K}}$, and for every possible operation of the system for $k+1$ steps, as determined by the interpretation $\sigma_{\mathcal{I}}$ of the initial state and indexed input symbols \mathcal{I} , there is some alternate initial state and input sequence, given by interpretation $\theta_{\mathcal{I}}$ that would have led to the exact state in i steps for some $0 \leq i \leq k$.

We show that this property guarantees that the system will not reach new states beyond step k .

Theorem 1 *If a system converges in k steps, then for any $j \geq 0$ and any interpretation $\sigma_{\mathcal{K}}$ of the system parameters, there exists a step $i \leq k$ and an alternate interpretation $\theta_{\mathcal{I}}$ of the initial state and inputs, such that*

$$\langle s^i \rangle_{\theta_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \langle s^j \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}}. \quad (3)$$

Before we prove Theorem 1, we highlight a key property of our system model.

Proposition 1 *For any interpretations $\sigma_{\mathcal{I}}$ and $\sigma_{\mathcal{K}}$ and any step i*

$$\langle s^{i+1} \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \left\langle \delta \left[\langle s^i \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} / \mathcal{S}, \langle t^{i+1} \rangle_{\sigma_{\mathcal{I}}} / \mathcal{T} \right] \right\rangle_{\sigma_{\mathcal{K}}} \quad (4)$$

By way of explanation, (4) combines a basic property of symbolic simulation with some specific characteristics of our model. On the right hand side, we evaluate state s^i under an interpretation of symbols in $\mathcal{K} \cup \mathcal{I}$, yielding an integer or Boolean value, or an integer or Boolean function for each state element. Similarly, we evaluate the indexed inputs at step $i + 1$, but these depend only on the interpretation of symbols in \mathcal{I} . Now we substitute these values for the state element symbols and input symbols in the expressions for the transition behavior δ . Finally, we apply an interpretation to each system parameter symbol in \mathcal{K} and evaluate the results, giving a new value for each state element. The left hand side gives a value for each state element by applying the same interpretations to the expressions reached after $i + 1$ steps of symbolic simulation. Our claim is that either route leads to the same values.

The proposition follows from the definition of s^{i+1} , the property that the transition behavior is independent of the values assigned to the symbols \mathcal{I} , since these only encode the initial state and the input values, and the values of inputs t^{i+1} are independent of the values of the system parameterization symbols.

We now prove Theorem 1.

Proof: The proof proceeds by induction on j . For $j \leq k$, the condition holds trivially by letting $i = j$. Let us assume it holds for j . That is, there is some $i' \leq k$ such that $\langle s^{i'} \rangle_{\theta_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \langle s^j \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}}$.

We first show that state s^{j+1} must be equivalent to the state at step $i' + 1$ under an alternate interpretation of the initial state and indexed input symbols. First, we apply (4) and (3) to expand state s^{j+1} and apply the induction hypothesis, giving

$$\begin{aligned} \langle s^{j+1} \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} &= \langle \delta [s^j / \mathcal{S}, t^{j+1} / \mathcal{T}] \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} \\ &= \left\langle \delta \left[\langle s^j \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} / \mathcal{S}, \langle t^{j+1} \rangle_{\sigma_{\mathcal{I}}} / \mathcal{T} \right] \right\rangle_{\sigma_{\mathcal{K}}} \\ &= \left\langle \delta \left[\langle s^{i'} \rangle_{\theta_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} / \mathcal{S}, \langle t^{j+1} \rangle_{\theta_{\mathcal{I}}} / \mathcal{T} \right] \right\rangle_{\sigma_{\mathcal{K}}} \end{aligned}$$

Now let us define an interpretation $\theta'_{\mathcal{I}}$ that is identical to $\theta_{\mathcal{I}}$, except that for each symbol $\mathbf{a}_{i'+1}$ representing the value of input \mathbf{a} at step $i' + 1$, let $\theta'_{\mathcal{I}}(\mathbf{a}_{i'+1}) = \theta_{\mathcal{I}}(\mathbf{a}_{j+1})$. State expression $s^{i'}$ does not have any indexed input symbols with step index $i' + 1$, and hence it will evaluate to the same set of values under interpretations $\theta_{\mathcal{I}}$ and $\theta'_{\mathcal{I}}$. We can therefore continue the derivation as follows:

$$\begin{aligned} \langle s^{j+1} \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} &= \left\langle \delta \left[\langle s^{i'} \rangle_{\theta_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} / \mathcal{S}, \langle t^{j+1} \rangle_{\theta_{\mathcal{I}}} / \mathcal{T} \right] \right\rangle_{\sigma_{\mathcal{K}}} \\ &= \left\langle \delta \left[\langle s^{i'} \rangle_{\theta'_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} / \mathcal{S}, \langle t^{i'+1} \rangle_{\theta'_{\mathcal{I}}} / \mathcal{T} \right] \right\rangle_{\sigma_{\mathcal{K}}} \\ &= \left\langle \delta \left[s^{i'} / \mathcal{S}, t^{i'+1} / \mathcal{T} \right] \right\rangle_{\theta'_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} \\ &= \langle s^{i'+1} \rangle_{\theta'_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} \end{aligned}$$

For $i' < k$, we can let $i = i' + 1 \leq k$ be the earlier step and $\theta'_{\mathcal{I}}$ be the alternate interpretation to prove the induction hypothesis.

For $i' = k$, we have shown that $\langle s^{j+1} \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \langle s^{k+1} \rangle_{\theta'_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}}$. Applying the convergence criterion, there must be some step $i \leq k$ and some alternate interpretation $\eta_{\mathcal{I}}$ such that $\langle s^{j+1} \rangle_{\sigma_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \langle s^{k+1} \rangle_{\theta'_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}} = \langle s^i \rangle_{\eta_{\mathcal{I}} \cdot \sigma_{\mathcal{K}}}$, to show that the state at step $j + 1$ is identical to the state at step i under alternate interpretation $\eta_{\mathcal{I}}$.

Note how this proof relied on the structure of our model. We encode variations in the system behavior and operation symbolically. On each step, the inputs can change arbitrarily (since we introduce a new set of symbols on each step), but the system behavior remains fixed (since it is parameterized by the fixed set of symbols \mathcal{K}).

4 Formulation of the Convergence Criterion

We now reach the main topic of this paper: determining whether a system is k -convergent for some value of k . We can express this as a problem in second-order logic as follows. Introduce a symbol set \mathcal{J} consisting of a symbol \mathbf{a}' for each initial state symbol $\mathbf{a} \in \mathcal{I}$, and a symbol $\mathbf{a}'_i \in \mathcal{I}$ for each indexed input signal \mathbf{a}_i , for $1 \leq i \leq k$. Rewrite each state expression s^i , for $0 \leq i \leq k$ to an expression r^i , by replacing each symbol in \mathcal{I} with its counterpart in \mathcal{J} .

Using the notation of predicate calculus, we consider the symbols in \mathcal{I} , \mathcal{J} , and \mathcal{K} to be quantified *variables*, either first-order (for integer or Boolean symbols) or second-order (for function or predicate symbols). We can then write the convergence criterion as:

$$\forall \mathcal{K} \forall \mathcal{I} \exists \mathcal{J} \left[\bigvee_{0 \leq i \leq k} \bigwedge_{\mathbf{a} \in \mathcal{S}} r^i_{\mathbf{a}} = s^{k+1}_{\mathbf{a}} \right] \quad (5)$$

With these quantifiers, we are really quantifying over the possible interpretations of the symbols. Note that this formula cannot be expressed in first-order logic, because we have existentially quantified function symbols.

Example 1: Consider a system with the integer state variables \mathbf{x} , \mathbf{y} and Boolean state variable \mathbf{b} . The operations are defined by:

$$\begin{array}{lll} \text{init}[\mathbf{x}] & = & \mathbf{c}_0 & \text{init}[\mathbf{y}] & = & \mathbf{c}_0 & \text{init}[\mathbf{b}] & = & \mathbf{true} \\ \text{next}[\mathbf{x}] & = & \mathbf{f}(\mathbf{x}) & \text{next}[\mathbf{y}] & = & \mathbf{f}(\mathbf{y}) & \text{next}[\mathbf{b}] & = & (\mathbf{x} = \mathbf{y}) \end{array}$$

where \mathbf{c}_0 is an integer symbol and \mathbf{f} is an uninterpreted function symbol. Using our notation, the sets of symbols are defined as follows — $\mathcal{S} = \{\mathbf{x}, \mathbf{y}, \mathbf{b}\}$, $\mathcal{K} = \{\mathbf{f}\}$, $\mathcal{I} = \{\mathbf{c}_0\}$ and $\mathcal{J} = \{\mathbf{c}'_0\}$.

After simulating the system for one step, the convergence condition (given by equation 5, where $k = 0$) becomes:

$$\forall \mathbf{f} \forall \mathbf{c}_0 \exists \mathbf{c}'_0 [\mathbf{c}'_0 = \mathbf{f}(\mathbf{c}_0) \wedge \mathbf{c}'_0 = \mathbf{f}(\mathbf{c}_0) \wedge \mathbf{true} = (\mathbf{f}(\mathbf{c}_0) = \mathbf{f}(\mathbf{c}_0))]$$

which simplifies to $\forall \mathbf{f} \forall \mathbf{c}_0 \exists \mathbf{c}'_0 [\mathbf{c}'_0 = \mathbf{f}(\mathbf{c}_0)]$, which is clearly valid, with \mathbf{c}'_0 taking the value $\mathbf{f}(\mathbf{c}_0)$.

Therefore the system converges after one step of simulation. As expected, the state variable \mathbf{b} is always **true** in the reachable set of states.

For a function or predicate state element F , the expression $r_F^i = s_F^{k+1}$ is a *second-order equation*—it states that two functions or predicates are identical for all possible arguments.

For systems without function or predicate state elements, our convergence criterion yields a formula with the quantification structure shown in (5), with only first-order equations. Even for the simple case of a system with one integer symbol in \mathcal{I} , one function symbol of arity 2 in \mathcal{K} , deciding the truth of a formula with this structure is undecidable [2].

Again we find ourselves facing an undecidable property. We deal with this by 1) using syntactic transformations to eliminate the second-order equations for function and predicate state elements, and 2) using a sound, but incomplete decision procedure for second-order formulas of the form shown in (5). Our procedure is quite simple, but it seems to work well for the formulas arising in our convergence testing.

5 Checking Convergence

5.1 Function and Predicate State Elements

We can convert our convergence formula (5) to one containing only first-order equations by introducing a set of *argument symbols* $\mathcal{Z} = z_1, \dots, z_n$, where n is the maximum arity of any predicate or function state element. Suppose state element F has arity $\text{arity}(F) = m$. Then define $\tilde{r}_F^i \doteq r_F^i(z_1, \dots, z_m)$, and similarly define $\tilde{s}_F^i \doteq s_F^i(z_1, \dots, z_m)$. Then we can rewrite the convergence criterion as:

$$\forall \mathcal{K} \forall \mathcal{I} \exists \mathcal{J} \forall \mathcal{Z} \left[\bigvee_{0 \leq i \leq k} \bigwedge_{\mathbf{a} \in \mathcal{S}} \tilde{r}_{\mathbf{a}}^i = \tilde{s}_{\mathbf{a}}^k \right] \quad (6)$$

Unfortunately, we have no general approach to handle formulas with this quantifier structure. Instead, we use rewriting techniques to handle limited forms of function and predicate state elements. Our technique is sufficient to handle random-access memories, including the data memory and register file of a microprocessor.

A random-access memory is modeled as a function state element Mem where the argument is an address, and the function returns the value stored at that address. Consider a memory with address input Adr , data input Dat and write-enable signal Wrt . We describe the memory operation in our term-level modeling language as:

$$\begin{aligned} \text{init}[\text{Mem}] &= \mathbf{m}_0 \\ \text{next}[\text{Mem}] &= \lambda x . \text{ITE}(\text{Wrt} \wedge x = \text{Adr}, \text{Dat}, \text{Mem}(x)) \end{aligned}$$

where \mathbf{m}_0 is an uninterpreted function giving the initial memory contents. Note the restricted class of expressions that will result when modeling the operation of this memory over time to generate the expression \tilde{r}_{Mem}^i . At the base is an uninterpreted function, which can be assigned an interpretation that matches any desired functionality. There will then be a bounded number of updates due to write operations, but these will each be to a single (symbolic) address.

Suppose we wish to determine whether the system has converged for some fixed time point i , so that Equation 6 reduces to

$$\forall \mathcal{K} \forall \mathcal{I} \exists \mathcal{J} \forall \mathcal{Z} \left[\bigwedge_{\mathbf{a} \in \mathcal{S}} \tilde{r}_{\mathbf{a}}^i = \tilde{s}_{\mathbf{a}}^k \right] \quad (7)$$

Then the convergence criterion for state element **Mem** will have the general form:

$$\forall \mathcal{A} \exists \mathcal{B} \forall \mathbf{z} F'(\mathbf{z}) = F(\mathbf{z}) \quad (8)$$

where expression F has only symbols in \mathcal{A} , while expression F' has symbols from both \mathcal{B} and \mathcal{A} .

We apply a set of rewrites to the symbols in \mathcal{B} and generate a set of verification conditions that guarantees (8) holds, based on the structure of expression F' . In general, our rules apply to equations of the form $P(\mathbf{z}) \implies F'(\mathbf{z}) = F(\mathbf{z})$, where P is a predicate expression with symbols from both \mathcal{B} and \mathcal{A} . At the top level, we start with P being an expression that always yields **true**.

1. For equations of the form $P(\mathbf{z}) \implies \mathbf{f}'(\mathbf{z}) = F(\mathbf{z})$, where \mathbf{f}' is a function symbol in \mathcal{B} , rewrite all occurrences of \mathbf{f}' in \tilde{r}^i to be $\lambda x . \text{ITE}(P(x), F(x), \mathbf{f}'(x))$.
2. For equations of the form $P(\mathbf{z}) \wedge \mathbf{z} = E \implies F'(\mathbf{z}) = F(\mathbf{z})$, where E is an expression with symbols from both \mathcal{B} and \mathcal{A} , reduce the equation to $P(E) \implies F'(E) = F(E)$. This eliminates any reference to \mathbf{z} in the equation.
3. For equations of the form $P(\mathbf{z}) \implies [\lambda x . \text{ITE}(Q(x), G'(x), H'(x))] (\mathbf{z}) = F(\mathbf{z})$, where Q , G' , and H' are predicate and function expressions containing symbols in both \mathcal{A} and \mathcal{B} , we generate two verification conditions: $P(\mathbf{z}) \wedge Q(\mathbf{z}) \implies G'(\mathbf{z}) = F(\mathbf{z})$, and $P(\mathbf{z}) \wedge \neg Q(\mathbf{z}) \implies H'(\mathbf{z}) = F(\mathbf{z})$, and solve these recursively.
4. For equations of the form $P(\mathbf{z}) \implies \mathbf{f}(\mathbf{z}) = F(\mathbf{z})$, where \mathbf{f} is a function symbol in \mathcal{A} , we recursively analyze the structure of F .
 - If F is of the form $\text{ITE}(Q(x), G(x), H(x))$, where Q , G , and H are predicate and function expressions containing symbols in \mathcal{A} , we generate two verification conditions: $P(\mathbf{z}) \wedge Q(\mathbf{z}) \implies \mathbf{f}(\mathbf{z}) = G(\mathbf{z})$, and $P(\mathbf{z}) \wedge \neg Q(\mathbf{z}) \implies \mathbf{f}(\mathbf{z}) = H(\mathbf{z})$, and solve these recursively.
 - If F is of the form $\mathbf{g}(\mathbf{z})$, then the symbols \mathbf{f} and \mathbf{g} need to be the same. If the two symbols are different, we return **false** which implies that no rewrite exists.
5. For equations of the form $P(\mathbf{z}) \implies F'(\mathbf{z} + c) = F(\mathbf{z})$ with integer constant c , transform the equation to be $P(\mathbf{z} - c) \implies F'(\mathbf{z}) = F(\mathbf{z} - c)$, and solve it recursively.

Similar rules hold for equations of the form $P \implies F'(\mathbf{z}) = F(\mathbf{z})$, i.e., P is a Boolean expression independent of \mathbf{z} .

Given the special form of the expressions describing the updating of a random-access memory, we can see that by repeated application of these rules, we can eliminate all occurrences of symbol \mathbf{z} in (7). The first rule handles the uninterpreted function representing the initial memory state. The second rule handles updates to individual memory addresses. The third rule lets us split based on the case structure of the expression. The last two rules would be required for more complex memory structures.

Note that CLU logic can be used to model memories in which multiple entries can be updated in parallel [14]. The rewriting techniques proposed in this section do not work for such memories.

5.2 Convergence with First-Order Equations

Assume we have applied transformation rules to eliminate all second-order equations, and hence the convergence criterion is expressed by an equation of the form shown in (5) with only first-order equations. We would therefore like to decide the validity of a formula ψ of the form

$$\psi \doteq \forall \mathcal{A} \exists \mathcal{B} \phi \tag{9}$$

where ϕ does not contain any quantifiers. In fact, ϕ is a CLU formula, and we can assume that transformations have been applied to eliminate all *ITE* operations¹ and lambda applications.

Our system model is sufficiently general that we can generate any second-order formula having the structure shown in (9) as part of a convergence test. To see this, let the variables in ϕ be $\mathcal{A} = \overline{a_n}$ and $\mathcal{B} = \overline{b_m}$. Introduce a set of $m + 1$ state elements, consisting of an element \mathbf{q}_i for each existentially quantified variable $b_i \in \mathcal{B}$, and a final truth-valued state element \mathbf{q}_{m+1} . For each universally quantified variable $a_i \in \mathcal{A}$, introduce a system parameter \mathbf{a}_i . Let the system have transition behavior δ such that $\delta_{\mathbf{q}_{n+1}} \doteq \phi[\overline{\mathbf{q}_m}/\overline{b_m}, \overline{\mathbf{a}_n}/\overline{a_n}]$, and $\delta_{\mathbf{q}_i} = \mathbf{q}_i$ for $1 \leq i \leq m$. Finally, let the initial state $s_{\mathbf{q}_i}^0$ of each state element \mathbf{q}_i for $1 \leq i \leq m$ be \mathbf{a}_i , and the initial state of \mathbf{q}_{m+1} be **true**. Then the system is 0-convergent if and only if the formula $\forall \mathcal{A} \exists \mathcal{B} \phi$ is valid.

This construction shows that we cannot assume any particular restrictions on the formulas we must decide to prove convergence, other than the quantifier structure shown in (9).

5.2.1 Syntactic Approach.

Previous approaches to convergence have been based on finding syntactic similarities between the earlier state r^i and the current state s^{k+1} . The convergence criterion given by Isles et al. [13] is a more conservative check than the criterion we give in Equation 6, and hence is less general. We can see that their syntactic substitution-based technique is simply a strategy for proving the validity of a formula with the structure shown in (9) as follows.

Proposition 2 *Let b denote a set containing an expression $b_a \in E(\mathcal{A})$ for each $a \in \mathcal{B}$. If $\forall \mathcal{A} \phi[b/\mathcal{B}]$ is valid, then so is $\forall \mathcal{A} \exists \mathcal{B} \phi$.*

The proof of this proposition follows by instantiating any symbol $\mathbf{a} \in \mathcal{B}$ with the value $\langle b_{\mathbf{a}} \rangle_{\sigma_{\mathcal{A}}}$.

With this approach, we can prove convergence by using a decision procedure for CLU logic to prove the universal validity of $\phi[b/\mathcal{B}]$. The challenge, of course, is to find an appropriate set of substitutions to the symbols in \mathcal{B} .

5.2.2 Semantic Approach.

We describe two ways to transform formulas of the structure $\psi \doteq \forall \mathcal{A} \exists \mathcal{B} \phi$ into a formula in the logic we call *Quantified Separation Logic* (QSL). QSL consists of quantified Boolean and integer variables, Boolean connectives, and predicates of the form $\mathbf{x} = \mathbf{y} + c$ and $\mathbf{x} < \mathbf{y} + c$, where \mathbf{x} and \mathbf{y} are integer variables, and c is an integer constant. Our first translation $T_s(\psi)$ (for “sound”) yields a formula that is valid only if ψ is valid. Our second translation $T_c(\psi)$ (for “complete”) yields a

¹These can be eliminated by the “push to the leaves” transformation [17].

formula that is valid if ψ is valid. The two formulas are very similar to each other. They differ in the ordering of quantifiers and an additional set of clauses in the antecedent of the second formula. By deciding the validity of the first translation we can test for definite convergence, while we can test for possible convergence by deciding the validity of the second translation.

$$\begin{aligned}
\text{bool-atom} & ::= \text{bool-symbol} \\
& \quad | \text{predicate-symbol}(\text{int-atom} + \text{int-constant}, \dots, \text{int-atom} + \text{int-constant}) \\
\text{int-atom} & ::= \text{int-symbol} \\
& \quad | \text{function-symbol}(\text{int-atom} + \text{int-constant}, \dots, \text{int-atom} + \text{int-constant}) \\
\text{bool-expr} & ::= \text{bool-atom} \mid \mathbf{true} \mid \mathbf{false} \\
& \quad | \neg \text{bool-expr} \mid (\text{bool-expr} \wedge \text{bool-expr}) \\
& \quad | (\text{int-atom} = \text{int-atom} + \text{int-constant}) \\
& \quad | (\text{int-atom} < \text{int-atom} + \text{int-constant})
\end{aligned}$$

Figure 2: **Normal Form Syntax.** Any integer or Boolean expression in CLU can be rewritten into this form.

1. Preserving Soundness. As shown in Figure 2, we can rewrite any Boolean or integer expression in CLU into a *normal form*, in which all *ITE* operations have been eliminated, and the additions of integer constants are grouped together. Define an *atomic expression* as either an integer expression following the rules for syntactic type *int-atom* shown in the figure, or a Boolean expression following the rules for syntactic type *bool-atom*. We can see that an arbitrary Boolean expression consists of Boolean atoms, equality and ordering predicates applied to integer atoms (possibly with a constant offset), and Boolean connectives.

Without loss of generality, let us assume ϕ is in normal form. We start by enumerating all of the atomic expressions occurring in ϕ as a sequence g_1, \dots, g_n . Let $\text{top}(g_i)$ denote the top-level symbol in subexpression g_i . We can see that each atomic expression g_i must be of one of the following forms:

1. Boolean symbol. $g_i \doteq \mathbf{b}$, giving $\text{top}(g_i) = \mathbf{b}$.
2. Predicate application. $g_i \doteq \mathbf{p}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$, giving $\text{top}(g_i) = \mathbf{p}$.
3. Integer symbol. $g_i \doteq \mathbf{x}$, giving $\text{top}(g_i) = \mathbf{x}$.
4. Function application. $g_i \doteq \mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$, giving $\text{top}(g_i) = \mathbf{f}$.

We require the sequence to be ordered according to subexpression containment. That is, for the function and predicate application forms listed above, we require $i_l < i$ for $1 \leq l \leq k$. The soundness property of translation T_s holds for any such ordering, but we get a tighter bound by listing the subexpressions having top-level symbols in \mathcal{A} as early as possible. That is, if $\text{top}(g_i) \in \mathcal{A}$ and $\text{top}(g_j) \in \mathcal{B}$, then $i < j$, unless g_j is a subexpression of g_i .

Now introduce a sequence of symbols $\overline{\mathbf{v}_n} \doteq \mathbf{v}_1, \dots, \mathbf{v}_n$, where \mathbf{v}_i is an integer (respectively, Boolean) symbol when $\text{top}(g_i)$ is an integer or function symbol (respectively, Boolean or predicate symbol). We generate two formulas $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$, each of which is a conjunction of consistency constraints by

considering each pair of subexpressions g_i and g_j , with $i < j$ and $\text{top}(g_i) = \text{top}(g_j)$. These are the same constraints used by Ackermann for removing function applications from a formula [1]. For subexpression g_i of the form $\mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$, and g_j of the form $\mathbf{f}(g_{j_1} + c_{j,1}, \dots, g_{j_k} + c_{j,k})$, we include the constraint

$$\mathbf{v}_{i_1} = \mathbf{v}_{j_1} + (c_{j,1} - c_{i,1}) \wedge \dots \wedge \mathbf{v}_{i_k} = \mathbf{v}_{j_k} + (c_{j,k} - c_{i,k}) \implies \mathbf{v}_i = \mathbf{v}_j \quad (10)$$

This constraint is included in either $C_{\mathcal{A}}$ or $C_{\mathcal{B}}$ according to whether $\mathbf{f} \in \mathcal{A}$ or $\mathbf{f} \in \mathcal{B}$. Similar constraints are generated when the top-level symbol in g_i and g_j is a predicate symbol \mathbf{p} .

Let $\hat{\phi}$ be the formula generated by replacing each atomic expression g_i in ϕ with the symbol \mathbf{v}_i . We always replace maximal subexpressions, so that the resulting formula no longer contains any symbols from ϕ .

Let quantifier Q_i be \forall when $\text{top}(g_i) \in \mathcal{A}$, and \exists when $\text{top}(g_i) \in \mathcal{B}$.

The soundness-preserving translation of ψ is given by

$$T_s(\psi) \doteq Q_1 \mathbf{v}_1 Q_2 \mathbf{v}_2 \dots Q_n \mathbf{v}_n \left[C_{\mathcal{A}} \implies (C_{\mathcal{B}} \wedge \hat{\phi}) \right] \quad (11)$$

Theorem 2 *For any formula ψ having the structure $\psi \doteq \forall \mathcal{A} \exists \mathcal{B} \phi$, if $T_s(\psi)$, as given by (11), is valid, then so is ψ .*

Proof: First, we use Skolemization to transform $T_s(\psi)$ into a formula where the existential quantifiers all come before the universal ones [10]. For $0 \leq i \leq n$, define $m(i)$ to be the number of universal quantifiers in the sequence Q_1, \dots, Q_i . Letting u be the number of symbols in \mathcal{V}_a , we have $m(n) = u$. Let $m^{-1}(i)$ be the position of the i th universal quantifier. (By convention, $m^{-1}(0) = 0$). For any i such that $\mathbf{v}_i \in \mathcal{V}_a$, we have $m^{-1}(m(i)) = i$. For any i such that $\mathbf{v}_i \in \mathcal{V}_e$, we have $m^{-1}(m(i)) < i$.

Let $\mathbf{y}_1, \dots, \mathbf{y}_u$ be a set of integer and Boolean symbols, where symbol \mathbf{y}_i has the same type as $\mathbf{v}_{m^{-1}(i)}$. For each i such that $\mathbf{v}_i \in \mathcal{V}_e$, introduce Skolem function symbol (when \mathbf{v}_i is an integer symbol) or predicate symbol (when \mathbf{v}_i is a Boolean symbol) \mathbf{f}_i having arity $m(i)$.

Generate formulas $C_{\mathcal{A}}^*$, $C_{\mathcal{B}}^*$, and $\hat{\phi}^*$ from $C_{\mathcal{A}}$, $C_{\mathcal{B}}$, and $\hat{\phi}$ by replacing each symbol \mathbf{v}_i by $\mathbf{y}_{m(i)}$ when $\mathbf{v}_i \in \mathcal{V}_a$ and by $\mathbf{f}_i(\mathbf{y}_1, \dots, \mathbf{y}_{m(i)})$ when $\mathbf{v}_i \in \mathcal{V}_e$. Then the Skolemized form of ψ , which we call $T_{sk}(\psi)$, is defined as

$$T_{sk}(\psi) \doteq \exists \mathcal{F} \forall \mathcal{Y} \left[C_{\mathcal{A}}^* \implies (C_{\mathcal{B}}^* \wedge \hat{\phi}^*) \right], \quad (12)$$

where \mathcal{F} is the set of all Skolem function and predicate symbols, and \mathcal{Y} is the set of symbols $\{\mathbf{y}_1, \dots, \mathbf{y}_u\}$. Formula $T_{sk}(\psi)$ is valid iff $T_s(\psi)$ is valid.

With this transformation, we shift the problem to one of showing that if $T_{sk}(\psi)$, given by (12), is valid, then so is formula $\psi \doteq \forall \mathcal{A} \exists \mathcal{B} \phi$. Assume (12) is valid, and that we are given some interpretation $\sigma_{\mathcal{A}}$ of the symbols in \mathcal{A} . We need to generate an interpretation $\sigma_{\mathcal{B}}$ of the symbols in \mathcal{B} , such that $\langle \phi \rangle_{\sigma_{\mathcal{A}} \cdot \sigma_{\mathcal{B}}} = \mathbf{true}$. Let $\sigma_{\mathcal{F}}$ be an interpretation of the Skolem function and predicate symbols in \mathcal{F} that satisfies (12). We construct a sequence of integer and Boolean values $\overline{a_n} \doteq a_1, \dots, a_n$ as follows:

1. For $\mathbf{v}_i \in \mathcal{V}_a$, when subexpression g_i is of the form \mathbf{x} (either an integer or Boolean symbol), we must have $\mathbf{x} \in \mathcal{A}$. Let $a_i = \sigma_{\mathcal{A}}(\mathbf{x})$. When g_i is of the form $\mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$, we have $\mathbf{f} \in \mathcal{A}$ (either a predicate or a function symbol). Let $a_i = \sigma_{\mathcal{A}}(\mathbf{f})(a_{i_1} + c_{i,1}, \dots, a_{i_k} + c_{i,k})$.

2. For $\mathbf{v}_i \in \mathcal{V}_e$, let $a_i = \sigma_{\mathcal{F}}(\mathbf{f}_i)(a_{m^{-1}(1)}, \dots, a_{m^{-1}(m(i))})$.

Let $\sigma_{\mathcal{Y}}$ be the interpretation of the symbols in \mathcal{Y} where $\sigma_{\mathcal{Y}}(\mathbf{y}_i) = a_{m^{-1}(i)}$. We can see that the sequence a_1, \dots, a_n consists of the values for the symbols in \mathcal{Y} and the result of applying the Skolem functions to these values. By (12), we are guaranteed that $\langle C_{\mathcal{A}}^* \implies (C_{\mathcal{B}}^* \wedge \hat{\phi}^*) \rangle_{\sigma_{\mathcal{F}} \cdot \sigma_{\mathcal{Y}}} = \mathbf{true}$.

Given the close relation between formulas $C_{\mathcal{A}} \implies (C_{\mathcal{B}} \wedge \hat{\phi})$ and $C_{\mathcal{A}}^* \implies (C_{\mathcal{B}}^* \wedge \hat{\phi}^*)$, and the way we generated the sequence \overline{a}_n , we can see that using the \overline{a}_n as the values for the symbols $\overline{\mathbf{v}}_n$ will satisfy our constraint formula. That is, if we perform the substitution

$$\left(C_{\mathcal{A}} \implies [C_{\mathcal{B}} \wedge \hat{\phi}] \right) [\overline{a}_n / \overline{\mathbf{v}}_n]$$

and then evaluate this formula, the result will equal **true**.

We can also see that when we perform the substitution $C_{\mathcal{A}} [\overline{a}_n / \overline{\mathbf{v}}_n]$, the resulting expression will evaluate to **true**, since we generated the sequence a_1, \dots, a_n based on a consistent interpretation of the function and predicate symbols in \mathcal{A} . From this, we can infer that the expressions $C_{\mathcal{B}} [\overline{a}_n / \overline{\mathbf{v}}_n]$ and $\hat{\phi} [\overline{a}_n / \overline{\mathbf{v}}_n]$ will evaluate to **true** as well.

Define interpretation $\sigma_{\mathcal{B}}$ such that for any g_i of the form \mathbf{x} , where \mathbf{x} is an integer or Boolean symbol in \mathcal{B} , we let $\sigma_{\mathcal{B}}(\mathbf{x}) = a_i$. For any g_i of the form $\mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$, where \mathbf{f} is a function or predicate symbol in \mathcal{B} , let $\sigma_{\mathcal{B}}(\mathbf{f})(a_{i_1} + c_{i,1}, \dots, a_{i_k} + c_{i,k}) = a_i$. No conflicts can arise in defining this interpretation, since $C_{\mathcal{B}}$ holds when the symbols $\overline{\mathbf{v}}_n$ are assigned the values \overline{a}_n . Complete the interpretation of \mathbf{f} by defining for any argument values x_1, \dots, x_k not covered already, the value of $\sigma_{\mathcal{B}}(\mathbf{f})(x_1, \dots, x_k)$ to be either 0 (when \mathbf{f} is a function) or **false** (when \mathbf{f} is a predicate.)

We can readily see that under the interpretation we have constructed, we will have $\langle g_i \rangle_{\sigma_{\mathcal{A}} \cdot \sigma_{\mathcal{B}}} = a_i$, for $1 \leq i \leq n$. From this, we can infer that $\langle \phi \rangle_{\sigma_{\mathcal{A}} \cdot \sigma_{\mathcal{B}}} = \mathbf{true}$, showing that $\forall \mathcal{A} \exists \mathcal{B} \phi$ is valid.

2. Preserving Completeness. To generate the completeness preserving transformation, let π be the permutation of $1, \dots, n$, that moves all of the universal quantifiers in the sequence Q_1, \dots, Q_n to the left, while otherwise preserving the relative orderings of symbols. That is, when we write the sequence $Q_{\pi(1)}, \dots, Q_{\pi(n)}$, we will have a sequence of the form $\forall^u \exists^{n-u}$, where u is the number of universal quantifiers. In addition, for i and j with $i < j$ and $Q_i = Q_j$, we have $\pi(i) < \pi(j)$.

Divide the symbols $\overline{\mathbf{v}}_n$ into two sets: those that are universally quantified $\mathcal{V}_a \doteq \{\mathbf{v}_{\pi(1)}, \dots, \mathbf{v}_{\pi(u)}\}$, and those that are existentially quantified $\mathcal{V}_e \doteq \{\mathbf{v}_{\pi(u+1)}, \dots, \mathbf{v}_{\pi(n)}\}$.

We generate an additional set of quantified antecedent clauses C_t to ensure completeness in the presence of some argument consistency constraints. Suppose for $i < j$ that subexpressions g_i and g_j are of the form $g_i \doteq \mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$, and $g_j \doteq \mathbf{f}(g_{j_1} + c_{j,1}, \dots, g_{j_k} + c_{j,k})$, where $\mathbf{f} \in \mathcal{A}$. Then, for this pair of subexpressions we add the constraint

$$\begin{aligned} \mathbf{v}_i \neq \mathbf{v}_j \wedge \bigwedge_{\substack{1 \leq l \leq k \\ \mathbf{v}_{i_l}, \mathbf{v}_{j_l} \in \mathcal{V}_a}} \mathbf{v}_{i_l} = \mathbf{v}_{j_l} + (c_{j_l} - c_{i_l}) \\ \implies \exists \mathbf{v}_{\pi(u+1)} \cdots \exists \mathbf{v}_{\pi(n)} \bigwedge_{1 \leq l \leq k} \mathbf{v}_{i_l} = \mathbf{v}_{j_l} + (c_{j_l} - c_{i_l}) \end{aligned} \quad (13)$$

to the set of clauses C_t . Note that the quantifiers in the consequent of this constraint take precedent over the quantifiers that are global to the overall formula.

We can now write the completeness preserving translation of ψ as

$$T_c(\psi) \doteq \forall \mathbf{v}_{\pi(1)} \cdots \forall \mathbf{v}_{\pi(u)} \exists \mathbf{v}_{\pi(u+1)} \cdots \exists \mathbf{v}_{\pi(n)} \left[(C_{\mathcal{A}} \wedge C_t) \implies (C_{\mathcal{B}} \wedge \hat{\phi}) \right] \quad (14)$$

Theorem 3 *For any formula ψ having the structure $\psi \doteq \forall \mathcal{A} \exists \mathcal{B} \phi$, if ψ is valid, then so is $T_c(\psi)$, as given by (14).*

Proof: Suppose we are given values $a'_{\pi(1)}, \dots, a'_{\pi(u)}$ for the universally quantified symbols $\mathbf{v}_{\pi(1)}, \dots, \mathbf{v}_{\pi(u)}$. Let A denote the set of all assignments $\overline{a_n}$ to the symbols $\overline{\mathbf{v}_n}$ such that $a_{\pi(i)} = a'_{\pi(i)}$, for $1 \leq i \leq u$. Then we must find a vector $\overline{a_n} \in A$ such that when we perform the substitution

$$\left([C_{\mathcal{A}} \wedge C_t] \implies [C_{\mathcal{B}} \wedge \hat{\phi}] \right) [\overline{a_n}/\overline{\mathbf{v}_n}] \quad (15)$$

the resulting formula will evaluate to **true**.

Our first strategy is to try to find a vector that violates a consistency constraint in C_t or in $C_{\mathcal{A}}$. This requires having two subexpressions of the form $g_i \doteq \mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$ and $g_j \doteq \mathbf{f}(g_{j_1} + c_{j,1}, \dots, g_{j_k} + c_{j,k})$, where $a'_i \neq a'_j$, and \mathbf{f} is either an integer or Boolean function in \mathcal{A} . It also requires that $a_{i_l} = a_{j_l} + (c_{j,l} - c_{i,l})$ for all $1 \leq l \leq k$ such that $\mathbf{v}_{i_l}, \mathbf{v}_{j_l} \in \mathcal{V}_a$.

Given that these conditions hold, then we can show that one of the two types of antecedent constraints will be violated. If there is some $\overline{a_n} \in A$ such that $a_{i_l} = a_{j_l} + (c_{j,l} - c_{i,l})$ for all $1 \leq l \leq k$, then we can use this as an assignment to the symbols $\overline{\mathbf{v}_n}$ that violates the consistency constraint (10) in $C_{\mathcal{A}}$. If no such $\overline{a_n}$ exists, then argument constraint (13) in C_t will be violated. In either case, the antecedent will be false, and hence (15) will evaluate to **true**.

Otherwise, we can assume that for every pair of subexpressions of the form $g_i \doteq \mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$ and $g_j \doteq \mathbf{f}(g_{j_1} + c_{j,1}, \dots, g_{j_k} + c_{j,k})$, where \mathbf{f} is a Boolean or integer function in \mathcal{A} , we have either $a'_i = a'_j$ or there is some argument position l , with $\mathbf{v}_{i_l}, \mathbf{v}_{j_l} \in \mathcal{V}_a$ and $a_{i_l} \neq a_{j_l} + (c_{j,l} - c_{i,l})$. We can therefore generate an interpretation $\sigma_{\mathcal{A}}$ of all of the symbols in \mathcal{A} such that for every subexpression $g_i \doteq \mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$, where $\mathbf{f} \in \mathcal{A}$, we have $\langle \mathbf{f} \rangle_{\sigma_{\mathcal{A}}}(a_{i_1} + c_{i,1}, \dots, a_{i_k} + c_{i,k}) = a_i$ for all $\overline{a_n} \in A$.

More precisely, we define $\langle \mathbf{f} \rangle_{\sigma_{\mathcal{A}}}(x_1, \dots, x_k)$ for arbitrary values of x_1, \dots, x_k by considering every subexpression of the form $g_i \doteq \mathbf{f}(g_{i_1} + c_{i,1}, \dots, g_{i_k} + c_{i,k})$. If for some such subexpression, we have $x_l = a_{i_l} + c_{i,l}$ for every argument position l such that $\mathbf{v}_{i_l} \in \mathcal{V}_a$, then we define $\langle \mathbf{f} \rangle_{\sigma_{\mathcal{A}}}(x_1, \dots, x_k) \doteq a_i$. If there is no such subexpression, then we define $\langle \mathbf{f} \rangle_{\sigma_{\mathcal{A}}}(x_1, \dots, x_k)$ to either equal **false**, when \mathbf{f} is a Boolean function, or 0, when \mathbf{f} is an integer function.

To complete the proof of Theorem 3, if we assume $\psi \doteq \forall \mathcal{A} \exists \mathcal{B} \phi$ is valid, then we can use our interpretation $\sigma_{\mathcal{A}}$ as an assignment of values to the symbols in \mathcal{A} . We are then guaranteed that there is some assignment of values to the symbols in \mathcal{B} such that ϕ holds. Use this assignment to define an interpretation $\sigma_{\mathcal{B}}$. Then we define a_i for $1 \leq i \leq n$ as $a_i \doteq \langle g_i \rangle_{\sigma_{\mathcal{A}}, \sigma_{\mathcal{B}}}$. We can see that $\overline{a_n} \in A$, since we will have $a_i = a'_i$ for each i such that $\mathbf{v}_i \in \mathcal{A}$. Since this assignment was derived from a consistent interpretation of the symbols in ϕ , all of the constraints in $C_{\mathcal{B}}$ will be satisfied for this assignment. Formula $\hat{\phi}$ will also evaluate to **true** under this assignment, since it is derived from an interpretation of the symbols in ϕ that makes it evaluate to **true**. From this we can infer that (15) will evaluate to **true**.

We therefore conclude that translation T_c preserves completeness.

We now give some examples to demonstrate the capabilities and limitations of our two translation methods.

Example 1: Our first example is a case where we successfully prove soundness.

$$\forall \mathbf{f}, \mathbf{y} [\forall \mathbf{x} \mathbf{x} = \mathbf{f}(\mathbf{x})] \implies \mathbf{y} = \mathbf{f}(\mathbf{f}(\mathbf{y})) \quad (16)$$

To get this into the required form, we rewrite it as

$$\forall \mathbf{f}, \mathbf{y} \exists \mathbf{x} [\neg(\mathbf{x} = \mathbf{f}(\mathbf{x})) \vee \mathbf{y} = \mathbf{f}(\mathbf{f}(\mathbf{y}))]$$

We write the subexpressions as follows. To make the resulting formulas more readable, we introduce symbols with names based on the subexpressions, rather than the more generic $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$:

Subexpression	g_1	g_2	g_3	g_4	g_5
	\mathbf{y}	$\mathbf{f}(\mathbf{y})$	$\mathbf{f}(\mathbf{f}(\mathbf{y}))$	\mathbf{x}	$\mathbf{f}(\mathbf{x})$
Symbol	\mathbf{y}	\mathbf{fy}	\mathbf{ffy}	\mathbf{x}	\mathbf{fx}

For $C_{\mathcal{A}}$ we then get

$$(\mathbf{x} = \mathbf{y} \implies \mathbf{fx} = \mathbf{fy}) \wedge (\mathbf{x} = \mathbf{fy} \implies \mathbf{fx} = \mathbf{ffy}) \wedge (\mathbf{y} = \mathbf{fy} \implies \mathbf{fy} = \mathbf{ffy})$$

For formula $C_{\mathcal{B}}$ we get **true**, while for $\hat{\phi}$ we get

$$\neg(\mathbf{x} = \mathbf{fx}) \vee \mathbf{y} = \mathbf{ffy}$$

and the overall quantifier structure is:

$$\forall \mathbf{y} \forall \mathbf{fy} \forall \mathbf{ffy} \exists \mathbf{x} \forall \mathbf{fx}$$

To see that the QSL formula is valid, consider a game played between opponents Bob and Alice. Bob has control over the universally quantified symbols and is attempting to make the formula to evaluate to **false**, while Alice has control over the existentially quantified symbols and is attempting to make the formula evaluate to **true**. They take turns instantiating symbols according to the quantifier structure. If Alice always has a winning strategy, then the formula is valid.

In this example, Bob must give values for \mathbf{y} , \mathbf{fy} , and \mathbf{ffy} . He must choose values such that $\mathbf{y} \neq \mathbf{ffy}$ to avoid satisfying $\hat{\phi}$, and must have either $\mathbf{y} \neq \mathbf{fy}$ or $\mathbf{fy} = \mathbf{ffy}$ to avoid falsifying the third consistency constraint. In the latter case, we also have $\mathbf{y} \neq \mathbf{fy}$.

Alice now sets $\mathbf{x} = \mathbf{y}$. This forces Bob to set $\mathbf{fx} = \mathbf{fy}$ to avoid falsifying the first consistency constraint. Combining these we get $\mathbf{x} = \mathbf{y} \neq \mathbf{fy} = \mathbf{fx}$, implying that $\hat{\phi}$ is satisfied. Alice has a winning strategy, showing that the quantified formula is valid.

Example 2: Our second example illustrates a case where the formula is valid, but the soundness-preserving transformation fails to show this.

$$\forall \mathbf{f} [\forall \mathbf{x} \mathbf{f}(\mathbf{x}) < \mathbf{f}(\mathbf{x} + 1)] \implies [\forall \mathbf{y} \mathbf{f}(\mathbf{y}) < \mathbf{f}(\mathbf{y} + 2)] \quad (17)$$

To get this into the required form, we rewrite it as

$$\forall \mathbf{f} \forall \mathbf{y} \exists \mathbf{x} \neg(\mathbf{f}(\mathbf{x}) < \mathbf{f}(\mathbf{x} + 1)) \vee \mathbf{f}(\mathbf{y}) < \mathbf{f}(\mathbf{y} + 2)$$

We write the subexpressions as follows.

Subexpression	g_1	g_2	g_3	g_4	g_5	g_6
	y	$f(y)$	$f(y+2)$	x	$f(x)$	$f(x+1)$
Symbol	y	fy	$fy2$	x	fx	$fx1$

For C_A we then get

$$(x=y \implies fx=fy) \wedge (x=y-1 \implies fx1=fy) \wedge (x=y+2 \implies fx=fy2) \wedge (x=y+1 \implies fx1=fy2)$$

For formula C_B we get **true**, while for $\hat{\phi}$ we get

$$\neg(fx < fx1) \vee fy < fy2$$

and the overall quantifier structure is:

$$\forall y \forall fy \forall fy2 \exists x \forall fx \forall fx1$$

This formula is not valid.

This example shows the limited capability of our translation T_s . It does not do the multiple instantiations of x required to replace the quantified antecedent in (17) with $f(y) < f(y+1) \wedge f(y+1) < f(y+2)$.

The completeness-preserving translation of this formula is identical, except that it yields a quantifier structure

$$\forall y \forall fy \forall fy2 \forall fx \forall fx1 \exists x$$

This formula can be shown to be valid.

In this case, Bob must choose values for all of his symbols, and then Alice gets to pick a value for x . She will be able to satisfy the antecedent of any of the four consistency constraints, so Bob must attempt to satisfy all of the consequents, giving $fx = fy = fx1 = fy2$, but this would imply that $fx \not< fx1$, satisfying $\hat{\phi}$. We conclude that Alice can always win.

Example 3: Our third example illustrates a case where the completeness-preserving transformation is overly optimistic.

$$\forall f \forall x \exists y f(x, y) = f(y, x+1) \tag{18}$$

This formula is clearly not valid.

We write the subexpressions as follows.

Subexpression	g_1	g_2	g_3	g_4
	x	y	$f(x, y)$	$f(y, x+1)$
Symbol	x	y	$f1$	$f2$

For C_A we then get

$$x=y \wedge y=x+1 \implies f1=f2$$

The above antecedent is unsatisfiable, and hence C_A reduces to **true**. Similarly, C_B is **true**. For the argument constraints we get

$$\neg(f1=f2) \implies \exists y (x=y \wedge y=x+1)$$

Since the consequent in this formula is unsatisfiable, this constraint reduces to $f1=f2$. Formula $\hat{\phi}$ is also $f1=f2$, and hence the translation T_c simply yields

$$\forall f1 \forall f2 [f1=f2 \implies f1=f2]$$

which reduces to **true**.

This example shows how much the set of argument constraints weakens the precision of translation T_c when the arguments have a structure where any possible instantiation of the existentially quantified symbols would yield conflicts.

To date, we have been unable to devise an example that illustrates the need for the argument consistency constraints C_t . This requires a formula that is valid, but T_c would be false without C_t in the antecedent.

6 Results & Discussion

We have implemented a prototype of the convergence testing framework within the UCLID [4] verification tool. Currently, we have only implemented the soundness-preserving translation to QSL. For deciding the resulting QSL formula, we used Difference Decision Diagrams [15] and a BDD-based implementation of a QSL solver that translates a QSL formula to a quantified Boolean formula (QBF) [16]. All the experiments are performed on a 2GHz Pentium-4 running Linux, with 1 GB of memory.

In this section, we describe our experience with the convergence testing framework for a three-stage arithmetic pipeline given in figure 3. This example originated with the first work on symbolic model checking [6], and has subsequently become a standard for verification research [9, 13]. In our version, we make use of both stalling and forwarding to resolve read-after-write hazards in the pipeline. Previous versions used only forwarding, with the result that a new result is written to the register file on each step of operation.




Figure 3: **Pipelined Version of ALU Circuit.** The three stages of the pipeline: fetch, execute and write-back. Read-after-write hazards are resolved for the first operand by stalling and for the second by forwarding. The dashed lines indicate Boolean control and the solid lines represent the flow of integer values.

The state elements of the pipeline include a function state variable, an unbounded register file pRF . The integer state elements include the different register identifiers, namely $eSRC2$, $eDEST$ and $wDEST$, the data values $eARG1$, $eARG2$ and $wVAL$, and the program counter pPC . The Boolean

state elements consist of the write enable registers $eWRT$ and $wWRT$. The system functionality is parameterized by uninterpreted function symbols for decoding an instruction, updating the program counter and the ALU. The Boolean state elements are initialized to **false** and the rest of the state elements take on arbitrary initial values.

The pipeline was symbolically simulated starting from the initial state. The QSL formula produced by the soundness preserving translation was **false** after $k = 1$ and $k = 2$ steps of simulation. A look at the Boolean state elements indicated that the system indeed does not converge within two steps. However, after $k = 3$ steps of simulation, the QSL formula produced was too large to be solved with a BDD-based implementation of our QSL solver [16] or with Difference Decision Diagrams [15]. The formula had 53 quantified integer variables, with 6 levels of quantifier alternations, 836 nodes in a Directed Acyclic Graph (DAG) representation of the formula, and the BDD representing the QBF formula exceeds 1 GB of memory. However, we have been able to prove the convergence of two simplified versions of the pipeline processor.

1. For the first case, we removed the data-path components of the processor including the register file, operand values and the write-back value. The remaining pipeline still contains the entire control complexity of the original pipeline including the stalling and the forwarding mechanisms. This model converges after $k = 3$ steps of simulation and our decision procedure detects so within 2 seconds with less than 11 MB of memory. The QSL formula contains 27 quantified integer variables, with 4 levels of quantifier alternations and 249 nodes in the DAG form. Notice that this example contains uninterpreted function symbols but does not contain any function state elements.
2. For the second case, we combined the execute and the write-back stages of the pipeline into a single stage (making the pipeline 2-stage), but retained the register file pRF and the data-path. The pipeline was modified to accommodate both stalling and forwarding of data. This example converges after $k = 2$ steps of simulation and our decision procedure takes 8 seconds to prove it valid. The memory consumption was about 80 MB. The QSL formula contains 29 quantified integer variables, with 4 levels of quantifier alternations and 203 nodes in the DAG form.

We are currently working on an alternate SAT-based implementation of our QSL solver and hope to test the convergence of the pipeline with a few optimizations. We are also experimenting with enumeration based QBF solvers including Quaffle [18]. The BDD-based implementation might also benefit from early quantification heuristics.

Discussion. The notion of k -convergence is not useful for systems with unbounded buffers, since many such systems do not converge. Moreover, our preliminary results indicate that the convergence criterion we present is precise, but computationally difficult to check. Abstraction techniques, such as predicate abstraction [11], allow for greater efficiency at the expense of using an approximate notion of convergence, and are a promising area for future work.

References

- [1] W. Ackermann. *Solvable Cases of the Decision Problem*. North-Holland, Amsterdam, 1954.
- [2] Egon Börger, Erich Grädel, and Yuri Gurevich. *The Classical Decision Problem*. Springer-Verlag, 1997.

- [3] R. E. Bryant, S. German, and M. N. Velev. Processor verification using efficient reductions of the logic of uninterpreted functions to propositional logic. *ACM Transactions on Computational Logic*, 2(1):1–41, January 2001.
- [4] R. E. Bryant, S. K. Lahiri, and S. A. Seshia. Modeling and verifying systems using a logic of counter arithmetic with lambda expressions and uninterpreted functions. In E. Brinksma and K. G. Larsen, editors, *Computer-Aided Verification (CAV 2002)*, LNCS 2404, pages 78–92. Springer-Verlag, 2002.
- [5] T. Bultan, R. Gerber, and W. Pugh. Symbolic model checking of infinite state systems using Presburger arithmetic. In Orna Grumberg, editor, *Computer-Aided Verification (CAV '97)*, LNCS 1254, pages 400–411. Springer-Verlag, 1997.
- [6] J. R. Burch, E. M. Clarke, K. L. McMillan, and D. L. Dill. Sequential circuit verification using symbolic model checking. In *27th Design Automation Conference (DAC '90)*, pages 46–51, 1990.
- [7] J. R. Burch and D. L. Dill. Automated verification of pipelined microprocessor control. In D. L. Dill, editor, *Computer-Aided Verification (CAV '94)*, LNCS 818, pages 68–80. Springer-Verlag, 1994.
- [8] Francisco Corella, Z. Zhou, Xiaoyu Song, Michel Langevin, and Eduard Cerny. Multiway decision graphs for automated hardware verification. *Formal Methods in System Design*, 10(1):7–46, 1997.
- [9] D. Cyrluk and P. Narendran. Ground temporal logic: a logic for hardware verification. In D. Dill, editor, *Computer-Aided Verification (CAV '94)*, LNCS 818, pages 247–259. Springer-Verlag, 1994.
- [10] Herbert B. Enderton. *A Mathematical Introduction to Logic*. Academic Press, 1972.
- [11] S. Graf and H. Saïdi. Construction of abstract state graphs with PVS. In O. Grumberg, editor, *Computer-Aided Verification (CAV '97)*, LNCS 1254, pages 72–83. Springer-Verlag, 1997.
- [12] R. Hojati, A. Isles, D. Kirkpatrick, and R. K. Brayton. Verification using finite instantiations and uninterpreted functions. In M. Srivas and A. Camilleri, editors, *Formal Methods in Computer-Aided Design (FMCAD '96)*, LNCS 1166, pages 218–232. Springer-Verlag, 1996.
- [13] A. J. Isles, R. Hojati, and R. K. Brayton. Computing reachable control states of systems modeled with uninterpreted functions and infinite memory. In A. J. Hu and M. Y. Vardi, editors, *Computer-Aided Verification (CAV '98)*, LNCS 1427, pages 256–267. Springer-Verlag, 1998.
- [14] Shuvendu K. Lahiri and Randal E. Bryant. Deductive verification of advanced out-of-order microprocessors. In *Computer-Aided Verification (CAV 2003)*, LNCS 2725, pages 341–354. Springer-Verlag, 2003.
- [15] Jesper B. Møller. DDDLIB: A library for solving quantified difference inequalities. In Andrei Voronkov, editor, *Conference on Automated Deduction (CADE 2002)*, LNCS 2392, pages 129–133. Springer-Verlag, 2002.

- [16] Sanjit A. Seshia and Randal E. Bryant. Unbounded, fully symbolic model checking of timed automata using Boolean methods. In *Computer-Aided Verification (CAV 2003)*, LNCS 2725, pages 154–166. Springer-Verlag, 2003.
- [17] M. N. Velev and R. E. Bryant. Effective use of Boolean satisfiability procedures in the formal verification of superscalar and VLIW microprocessors. In *38th Design Automation Conference (DAC '01)*, pages 226–231, 2001.
- [18] Lintao Zhang and Sharad Malik. Towards a symmetric treatment of satisfaction and conflicts in quantified Boolean formula evaluation. In Pascal Van Hentenryck, editor, *Principles and Practice of Constraint Programming (CP '02)*, LNCS 2470, pages 200–215. Springer-Verlag, 2002.