

Xs are for Trajectory Evaluation, Booleans are for Theorem Proving (Extended Version)

MARK D. AAGAARD,¹ THOMAS F. MELHAM,² JOHN W. O'LEARY¹

¹Strategic CAD Labs
Intel Corporation, JFT-102
5200 NE Elam Young Parkway
Hillsboro, OR 97124, USA
{maagaard,joleary}@ichips.intel.com

²FATA Research Group
Department of Computing Science
University of Glasgow
Glasgow, Scotland, G12 8QQ
tfm@dcs.gla.ac.uk

Abstract

This paper describes a semantic connection between the symbolic trajectory evaluation model-checking algorithm and relational verification in higher-order logic. We prove a theorem that translates correctness results from trajectory evaluation over a four-valued lattice into a shallow embedding of temporal operators over Boolean streams. This translation connects the specialized world of trajectory evaluation to a general-purpose logic and provides the semantic basis for connecting additional decision procedures and model checkers.

M. D. Aagaard, T. F. Melham and J. W. O'Leary, 'Xs are for Trajectory Evaluation, Booleans are for Theorem Proving (Extended Version)', Technical Report TR-2000-52, Department of Computing Science, University of Glasgow (January 2000).

This report is an extended version of M. D. Aagaard, T. F. Melham, and J. W. O'Leary, 'Xs are for Trajectory Evaluation, Booleans are for Theorem Proving', in *Correct Hardware Design and Verification Methods: 10th IFIP WG10.5 Advanced Research Working Conference: Bad Herrenalb, September 1999: Proceedings*, edited by L. Pierre and T. Kropf, Lecture Notes in Computer Science (Springer-Verlag, 1999), pp. 202–218.

Contents

1	Introduction	5
1.1	Organization of the Paper	7
1.2	Related Work	8
2	Symbolic Trajectory Evaluation	8
2.1	The Four Valued Lattice	9
2.2	Circuit Models	9
2.3	Trajectory Evaluation Logic	10
3	Trajectory Logic over Boolean Streams	12
3.1	Definitions and Axioms	12
3.2	Correctness of Boolean Valued Trajectory Evaluation	13
4	Trajectory Logic as Relations Over Streams	15
4.1	The Shallow Embedding	16
4.2	The Deep Embedding	17
5	Simple Example Revisited	18
6	Voss Implementation of Trajectory Evaluation	19
7	Conclusion	20
	References	21

1 Introduction

The well-known limits to BDD-based model-checking techniques have motivated a great deal of interest in combining model-checking with theorem proving [3, 12, 10, 6]. The foundation of any such hybrid verification approach is a semantic connection between the logic of properties in the model checker and the logic of the theorem prover. Symbolic trajectory evaluation [17] is a highly effective model checker for datapath verification. It has been combined with theorem proving and the combination has been used effectively on complex industrial circuits [1, 13]. But two of the features that make trajectory evaluation so effective as a model checker create difficulties or limitations in the theorem proving domain. Trajectory evaluation’s temporal logic has limited expressability, and it operates over a lattice of values containing notions of contradiction (\top) and unknown (\mathbf{X}).

In this paper, we establish a semantic link from symbolic trajectory evaluation (STE) to higher-order logic (HOL). This link allows trajectory evaluation to be used as a decision procedure without encumbering the theorem proving world with the complications and limitations of trajectory evaluation. The trajectory evaluation temporal operators are defined in a shallow embedding of predicates over streams and the lattice domain is converted to simple Booleans. This translates trajectory evaluation results into the conventional ‘relations-over-Boolean-streams’ approach to hardware modeling in higher-order logic [7].¹

We believe that the relational world is the right domain for connecting model checking engines. Each model checking algorithm typically has its own temporal logic. By translating results into higher-order logic, the vestiges of the individual model checkers are removed, allowing the full power of general-purpose theorem proving to be brought to bear.

To give some intuition about the two worlds we reason about and the connection between them, consider the NAND-DELAY circuit in Figure 1.




Figure 1: Simple example circuit NAND-DELAY

We will use this circuit as a running example throughout this paper. In this example, for simplicity, we consider the NAND gate to be zero delay. Trajectory evaluation typically uses a more detailed timing model of circuits.

Figure 2 presents simple correctness statements for both the trajectory evaluation and relational styles. Trajectory evaluation is based on symbolic simulation. Correctness statements are of the form $\models_{ckt} [ant \implies cons]$, where \implies is similar, but not identical to, implication (details are given in Section 2). The *antecedent* ant gives an initial state and input stimuli to the circuit ckt , while the *consequent* $cons$ specifies the desired response of

¹In the rest of the paper we will take the phrase ‘relational’ style to mean ‘relations over Boolean streams’ with a shallow embedding of temporal operators as predicates.

$$\begin{array}{l}
i \stackrel{\text{def}}{=} "i" \\
o \stackrel{\text{def}}{=} "o" \\
\vdash_{\text{ckt}} \left[\begin{array}{l} (i \text{ is } b_1) \text{ and} \\ (o \text{ is } b_2) \end{array} \implies (\mathbf{N} (o \text{ is } \neg(b_1 \wedge b_2))) \right]
\end{array}
\qquad
\begin{array}{l}
\text{ckt}_{\text{HOL}} i o \stackrel{\text{def}}{=} \\
\exists m. \\
(\forall t. m t = \text{NAND}(i t)(o t)) \wedge \\
(\forall t. o (t + 1) = m t) \\
\text{spec}_{\text{HOL}} i o \stackrel{\text{def}}{=} \\
\forall t. o (t + 1) = \text{NAND}(i t)(o t) \\
\forall i, o. \text{ckt}_{\text{HOL}} i o \implies \text{spec}_{\text{HOL}} i o
\end{array}$$

Trajectory-evaluation verification
Relational style verification

Figure 2: Example stream and trajectory evaluation verifications

the circuit. Circuits are black boxes—their implementations are not user-visible. Circuit *nodes* are named by strings. In the example, the antecedent drives the nodes “*i*” and “*o*” with the Boolean variables b_1 and b_2 at the initial step of the verification. The consequent says that at the next time step the node “*o*” has the value $\neg(b_1 \wedge b_2)$.

In relational verification, correctness statements are of the form $\text{ckt } i o \implies \text{spec } i o$, where \implies is true implication. Signals (e.g. i , m , and o) are modelled by *streams*, which are functions from time to values. Both the circuit and the specification are relations over these infinite streams. A stream satisfies a circuit if it is a series of values that could be observed on the corresponding signals in the circuit. The correctness criterion for the example says that if the streams i and o satisfy the circuit model, then they must conform to the specification.

At a very cursory level, a mapping from the trajectory evaluation result in Figure 2 to the relational world would result in:

$$\begin{array}{l}
\forall i, o. \forall t. \\
((i t) = b_1) \wedge ((o t) = b_2) \wedge (\text{ckt } i o) \\
\implies \\
(o (t + 1)) = \neg(b_1 \wedge b_2)
\end{array}$$

Substituting for b_1 and b_2 throughout the expression gives:

$$\forall i, o. \forall t. \text{ckt } i o \implies (o (t + 1)) = \neg((i t) \wedge (o t))$$

Technically, this description is not quite correct, but it does give the intuition behind our result that correctness statements in trajectory evaluation imply relational correctness statements. Section 5 shows the actual results for the NAND-DELAY circuit. The difficulties arise in details such as translation between different semantic domains (e.g. Boolean and lattice valued streams) and the treatment of free variables in trajectory formulas.

Our main result is a formal translation from trajectory evaluation’s temporal operators over lattices to a shallow embedding of the temporal operators over Boolean streams. We prove that any result verified by the trajectory evaluation algorithm will hold in the relational world. This allows trajectory evaluation to be used as a decision procedure in a theorem prover without changing the relational style of verification used in the theorem prover.

It is interesting to note that our result is an implication, not an ‘if-and-only-if’; that is, we do not guarantee that every statement provable in the relational world will also hold in trajectory evaluation. The problem stems from the differences in how the relational world and the trajectory evaluation world handle contradictions in the circuit and antecedent. Joyce and Seger gave an extra constraint on trajectory evaluation that can be used to prove an if-and-only-if relationship [11]. The constraint requires reasoning about contradictions and top values. Because we use trajectory evaluation as a decision procedure, we are able to avoid the burden of reasoning about contradictions.

1.1 Organization of the Paper

Figure 3 is a roadmap of the paper. Section 2 begins with a presentation of trajectory assertions, the specifications for symbolic trajectory evaluation, over the standard four-valued lattice. Our translation then follows two major steps: from the four-valued lattice to Booleans, and then from meta-logical Booleans to a shallow embedding of trajectory formulas in higher order logic.

In Section 3 we introduce our definition of trajectory assertions over Booleans and prove that a trajectory assertion over lattice-valued streams implies the same result over Boolean-valued streams. The proof that relates trajectory assertions over lattices to trajectory assertions over Booleans is a meta-logical proof about the semantics of the two languages and links the free (Boolean) variables that appear in these assertions.

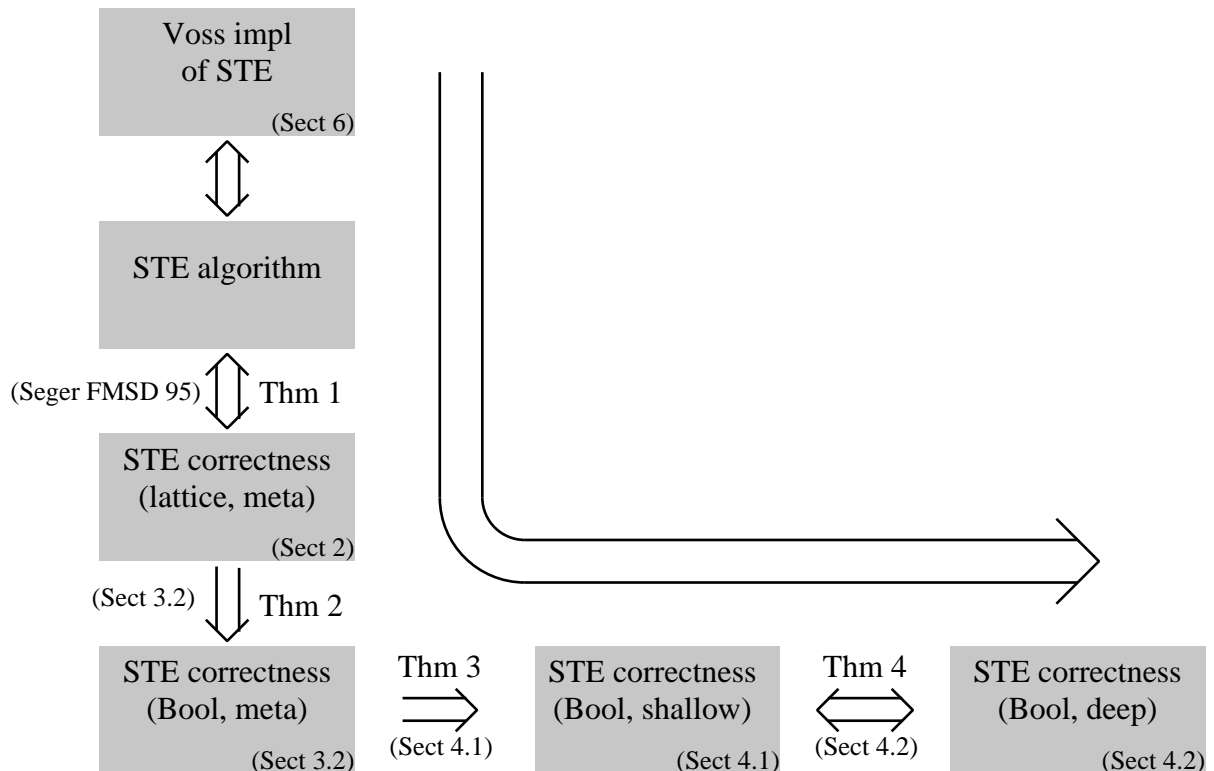


Figure 3: Roadmap of the paper

In Section 4.1, we then prove that a shallow embedding of trajectory formulas as relations over streams follows in higher order logic from an STE result in lattice-valued trajectory logic. Finally, in Section 4.2, we transfer our results to a deep embedding.

1.2 Related Work

An early experiment in combining model checking and theorem proving was the HOL-Voss System [11]. Within HOL [7], Joyce and Seger proved the correspondence between a deep embedding of the Voss [8] ‘5-tuple’ implementation of trajectory assertions (see Section 6) and a simple specification language that was deeply embedded in HOL. This allowed them to use Voss as an external decision procedure for HOL. Our work focuses on the logical content of trajectory assertions, independent of any particular implementation (e.g. Voss), and connects this to a shallow embedding of temporal operators in higher-order logic.

The modal mu-calculus has been embedded in HOL [2] and in PVS [16]. In the PVS work, Rajan *et al* implemented a shallow embedding upon which they defined the $\forall\text{CTL}^*$ temporal logic. They connected a mu-calculus model checker and verified a number of abstraction operations. It should be possible to translate a subset of the $\forall\text{CTL}^*$ formulas into predicates over Boolean streams, but given the complexity of $\forall\text{CTL}^*$, it is difficult to estimate the feasibility of this approach.

Chou has given a set-theoretic semantics of trajectory evaluation, focusing on extensions to the basic algorithm [5]. Our work translates trajectory evaluation results into a form that can be seamlessly integrated with current practice in higher-order-logic theorem proving. It would be interesting to explore connections between the work presented here and Chou’s, to find a seamless connection from extended trajectory evaluation to theorem-proving in higher-order logic.

2 Symbolic Trajectory Evaluation

This section presents the logic of trajectory assertions. Our results are meta-logical, but for convenience we use a **bold** face logical-style notation to state our results. Our presentation of trajectory evaluation has three parts. After a few preliminary definitions, we proceed to describe the four-valued lattice in Section 2.1. Section 2.2 then overviews the circuit model used in trajectory evaluation. Finally, Section 2.3 describes the specification logic of trajectory evaluation.

First, some preliminaries. We suppose there is a set of *nodes*, naming observable points in circuits. A *stream* is a function from natural numbers representing time to data values in the stream. A *sequence* takes a node and returns the stream for that node. A *state* is a mapping from *nodes* to values:


$$\begin{array}{ll}
 \alpha \text{ stream} & \stackrel{\text{type}}{=} \mathcal{N} \rightarrow \alpha \\
 \alpha \text{ sequence} & \stackrel{\text{type}}{=} \text{node} \rightarrow \alpha \text{ stream} \\
 \alpha \text{ state} & \stackrel{\text{type}}{=} \text{node} \rightarrow \alpha
 \end{array}$$

We typically use σ for sequences and s for states. Two convenient sequence operations are taking the suffix and transposing a sequence so that it is a stream of states:

$$\begin{aligned} \text{suffix: } \quad \sigma_i &\stackrel{\text{def}}{=} \lambda n. \lambda t. (\sigma \ n \ (t + i)) \\ \text{transpose: } \sigma^\top &\stackrel{\text{def}}{=} \lambda t. \lambda n. (\sigma \ n \ t) \\ \sigma^\top &:: (\alpha \ \text{state}) \ \text{stream} \end{aligned}$$

2.1 The Four Valued Lattice

In this paper, the only lattice we use is the four-valued lattice shown below. The theory of trajectory evaluation works equally well over all complete lattices. But this lattice simplifies the presentation and is the lattice used by the Voss implementation of trajectory evaluation [8]. Our mathematical development also is based on this lattice—generalizing our results would be of theoretical interest, but not of immediate practical benefit.



The ordering over lattice values shown above defines the ordering relation \sqsubseteq , which we lift pointwise and overload over streams and states.

We inject the set of Boolean values to lattice values with the postfix operator \downarrow (read ‘drop’, Definition 1), which maps the Boolean values **T** and **F** to their counterparts in the lattice. Drop is lifted pointwise to states, to sequences, and to state streams.

Definition 1 *Dropping from Boolean to lattice values.*

$$\begin{aligned} \mathbf{F}\downarrow &\stackrel{\text{def}}{=} 0 \\ \mathbf{T}\downarrow &\stackrel{\text{def}}{=} 1 \end{aligned}$$

2.2 Circuit Models

In our description of circuit models we refer to Table 1, which gives the lattice transition function for the example NAND-DELAY circuit. In trajectory evaluation, the circuit model is given by a next state function \mathbf{Y} that takes a circuit and maps states to states:

$$\mathbf{Y} :: \text{ckt} \rightarrow \text{lattice state} \rightarrow \text{lattice state}$$

A lattice state is an assignment to circuit nodes of values drawn from the four valued lattice. The first argument to \mathbf{Y} identifies the particular circuit of interest, and for the present purposes may be regarded as an uninterpreted constant. Intuitively, the next state function expresses a constraint on the set of possible states into which the circuit may go for any given state. Suppose the circuit is in state s , then $\mathbf{Y}(s)$ will give the least specified state the system can transition to. Here, ‘least specified’ means that if a node can take on both 1 and 0 values in the next state, then $\mathbf{Y}(s)$ will assign the value **X** to that node.

$\langle i \ o \ \rangle$	$\langle i' \ o' \ \rangle$	
0 0	X 1	
0 1	X 1	
0 X	X 1	← o can initially be X, and o' is still defined
1 0	X 1	
1 1	X 0	
1 X	X X	← o' is unknown, because o is unknown
X 0	X 1	← i can initially be X, and o' is still defined
X 1	X X	← o' is unknown, because i is unknown
X X	X X	

Table 1: Lattice transition function for NAND-DELAY circuit

A critical requirement for trajectory evaluation is that the next-state function be monotonic, which is captured in Axiom 1.

Axiom 1 *Monotonicity of \mathbf{Y} .*

For all s, s' . ($s \sqsubseteq s'$) implies ($\mathbf{Y}s \sqsubseteq \mathbf{Y}s'$)

Monotonicity can be seen in the NAND-DELAY circuit by comparing a transition in which one of the current state variables (e.g. o) is X with a transition in which o is either 0 or 1. A bit of the algorithmic efficiency of trajectory evaluation is illustrated here. The initial value for some of the circuit nodes can be X and a meaningful result can still be verified. In this way, the lattice often allows trajectory evaluation to prove results with fewer BDD variables than would otherwise be needed.

A sequence σ is said to be *in the language* of a circuit (Definition 2) if the set of behaviors that the sequence encodes is a subset of the behaviors that the circuit can exhibit. This means that the result of applying \mathbf{Y} to any element of the state stream σ^\top is no more specified (with respect to the \sqsubseteq ordering) than the succeeding element of σ^\top .

Definition 2 *Sequence is in the language of a circuit.*

$$\sigma \in \mathcal{L} \text{ ckt} \stackrel{\text{def}}{=} \text{For all } t \geq 0. (\mathbf{Y} \text{ ckt } (\sigma^\top t)) \sqsubseteq (\sigma^\top (t+1))$$

2.3 Trajectory Evaluation Logic

Trajectory evaluation correctness statements (known as *trajectory assertions*) are written

$$\models_{\text{ckt}} [ant \implies\!\!\!\implies cons]$$

where *ant* and *cons* are *trajectory formulas*. The intuition is that the antecedent *ant* provides stimuli to nodes in the circuit and the consequent *cons* specifies the values expected on nodes in the circuit. Before further describing trajectory assertions, we define trajectory formulas (Definition 3) and what it means for a sequence to satisfy a trajectory formula (Definition 5).

Definition 3 *Trajectory formulas.*

f	$\stackrel{def}{=} n \text{ is } 0$	$// n \text{ has value } 0$
	$ n \text{ is } 1$	$// n \text{ has value } 1$
	$ f_1 \text{ and } f_2$	$// \text{ conjunction of formulas}$
	$ f \text{ when } g$	$// f \text{ is asserted only when } g \text{ is true}$
	$ \mathbf{N} f$	$// f \text{ holds in the next time step}$

where f, f_1, f_2 range over formulas; n ranges over the node names of the circuit; and g is a Boolean expression, commonly called a *guard*.

Trajectory formulas are guarded expressions defining values on nodes in the sequence. Guards may contain free variables. In fact, guards are the only place that free variables are allowed in the primitive definition of trajectory formulas. But the following Syntactic sugar for *is* allows variables to appear in the value field as well:

Definition 4 *Boolean values.*

$$n \text{ is } b \stackrel{def}{=} ((n \text{ is } 1) \text{ when } b) \text{ and } ((n \text{ is } 0) \text{ when } \neg b)$$

where n ranges over node names and b is a Boolean variable.

The use of this abbreviation is illustrated in Figure 2 and Section 5.

Definition 5 describes when a sequence σ satisfies a trajectory formula f . Satisfaction is defined with respect to an assignment ϕ of Boolean values to the variables that appear in the guards of the formula.

Definition 5 *Sequence satisfies a trajectory formula.*

$(\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} (n \text{ is } 0)$	$\stackrel{def}{=} \sigma n 0 \sqsupseteq 0$
$(\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} (n \text{ is } 1)$	$\stackrel{def}{=} \sigma n 0 \sqsupseteq 1$
$(\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} (f_1 \text{ and } f_2)$	$\stackrel{def}{=} ((\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} f_1) \text{ and } ((\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} f_2)$
$(\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} (f \text{ when } g)$	$\stackrel{def}{=} (\phi \models g) \text{ implies } ((\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} f)$
$(\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} (\mathbf{N}f)$	$\stackrel{def}{=} (\phi, \sigma_1) \stackrel{def}{\models}_{\text{STE}} f$

where $\phi \models g$ means that the assignment that ϕ makes to the free variables in g renders g true.

We now have sufficient notation to define a *trajectory assertion* (Definition 6). In trajectory evaluation, correctness criteria are formulated as trajectory assertions.

Definition 6 *Trajectory assertion.*

$$\phi \models_{ckt} [ant \implies\!\!\!\implies cons] \stackrel{def}{=} \text{For all } \sigma. (\sigma \in \mathcal{L} \text{ ckt}) \text{ implies } ((\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} ant) \text{ implies } ((\phi, \sigma) \stackrel{def}{\models}_{\text{STE}} cons)$$

The fundamental theorem of trajectory evaluation [17] says the trajectory evaluation algorithm (*STE ckt ant cons*) computes the Boolean condition e on the free variables

in *ant* and *cons* if and only if any assignment ϕ satisfying e also proves the trajectory assertion $\phi \models_{ckt} [ant \implies\!\!\implies cons]$ (Theorem 1).

Theorem 1 *Correctness of STE algorithm.*

For all circuits ckt , antecedents ant , and consequences $cons$, the implementation of the STE algorithm returns the value e (that is, $e = \text{STE } ckt \text{ ant } cons$) if and only if:

For all ϕ . $\phi \models e$ implies $\phi \models_{ckt} [ant \implies\!\!\implies cons]$

3 Trajectory Logic over Boolean Streams

In this section we give a definition of trajectory logic over Boolean streams (as opposed to the standard lattice-valued streams in Section 2) and prove that trajectory evaluation results that hold over the four valued lattice also hold over Boolean streams. Boolean identifiers (e.g. next-state relations, sequences, and languages) will be distinguished from their lattice valued counterparts by marking them with a \circ , as in \mathring{Y} , $\mathring{\sigma}$, and $\mathring{\mathcal{L}}$.

3.1 Definitions and Axioms

In the Boolean world, circuit behavior is modeled as a *relation* between current and next states. In contrast, circuit behavior in trajectory evaluation is modeled as a next state *function*. The Boolean next state relation, denoted \mathring{Y} , has the type:

$\mathring{Y} :: \text{ckt} \rightarrow \text{bool state} \rightarrow \text{bool state} \rightarrow \text{bool}$

We write $\mathring{Y} \text{ ckt}$ as an infix operator, as in $s (\mathring{Y} \text{ ckt}) s'$.

As a concrete example, the next state relation for the NAND-DELAY circuit of Figure 1 is defined by Table 2, where the vectors $\langle i, o \rangle$ and $\langle i', o' \rangle$ denote the current and next states of the input and output. Note that the non-determinism that was represented by Xs in Y (Table 1) appears as multiple next states with the same current state in Table 2.

$\langle i, o \rangle$	$\langle i', o' \rangle$
F F	F T
F F	T T
F T	F T
F T	T T
T F	F T
T F	T T
T T	F F
T T	F T

Table 2: Boolean next-state relation for NAND-DELAY

Given a circuit's next state relation \mathring{Y} , we say that a Boolean sequence $\mathring{\sigma}$ is *in the language* of the circuit when consecutive states in the state stream $\mathring{\sigma}^T$ are included in the next-state relation (Definition 7).

Definition 7 *Boolean sequence is in the language of a circuit.*

$$\mathring{\sigma} \in \mathring{\mathcal{L}} \text{ ckt} \stackrel{\text{def}}{=} \text{For all } t \geq 0. (\mathring{\sigma}^\top t) (\mathring{\mathbf{Y}} \text{ ckt}) (\mathring{\sigma}^\top(t+1))$$

We now define when a Boolean sequence $\mathring{\sigma}$ satisfies a trajectory formula f (Definition 8). The only distinction between \models° and satisfaction over lattice sequences ($\models_{\mathbb{S}\mathbb{T}\mathbb{E}}$) is that for the formulas $(n \text{ is } 0)$ and $(n \text{ is } 1)$, satisfaction is defined in terms of values in the Boolean domain rather than the lattice domain.

Definition 8 *Boolean sequence satisfies a trajectory formula.*

$$\begin{aligned} (\phi, \mathring{\sigma}) \models^\circ (n \text{ is } 0) &\stackrel{\text{def}}{=} (\mathring{\sigma} \ n \ 0) = \mathbb{F} \\ (\phi, \mathring{\sigma}) \models^\circ (n \text{ is } 1) &\stackrel{\text{def}}{=} (\mathring{\sigma} \ n \ 0) = \mathbb{T} \\ (\phi, \mathring{\sigma}) \models^\circ (f_1 \text{ and } f_2) &\stackrel{\text{def}}{=} ((\phi, \mathring{\sigma}) \models^\circ f_1) \text{ and } ((\phi, \mathring{\sigma}) \models^\circ f_2) \\ (\phi, \mathring{\sigma}) \models^\circ (f \text{ when } g) &\stackrel{\text{def}}{=} (\phi \models g) \text{ implies } ((\phi, \mathring{\sigma}) \models^\circ f) \\ (\phi, \mathring{\sigma}) \models^\circ (\mathbf{N} \ f) &\stackrel{\text{def}}{=} (\phi, \mathring{\sigma}_1) \models^\circ f \end{aligned}$$

3.2 Correctness of Boolean Valued Trajectory Evaluation

To link the worlds of lattice and Boolean based trajectory evaluation, we require that the two models of the circuit behavior (\mathbf{Y} and $\mathring{\mathbf{Y}}$) describe the same behavior. Axiom 2 says that if two Boolean states s and s' satisfy the next-state relation $\mathring{\mathbf{Y}}$, then the result of applying the next-state function \mathbf{Y} to the dropped versions of s results in a state that is no higher in the lattice than $s' \downarrow$ (\mathbf{Y} is a ternary extension of $\mathring{\mathbf{Y}}$).

Axiom 2 *Relating the next-state relation and the next-state function.*

$$\text{For all ckt, } s, s'. \\ s (\mathring{\mathbf{Y}} \text{ ckt}) s' \text{ implies } (\mathbf{Y} \text{ ckt } (s \downarrow)) \sqsubseteq (s' \downarrow)$$




Figure 4: Illustration of Axiom 2

Axiom 2, which is illustrated in Figure 4, says that any Boolean next state possible in the relational model is consistent with the next state in the lattice-valued model. It also constrains $\mathring{\mathbf{Y}} \text{ ckt}$ to yield false whenever it is applied to two states s, s' that are inconsistent with $\mathbf{Y} \text{ ckt}$ (e.g. unreachable states). Inconsistency is manifested by $\mathbf{Y} \text{ ckt } s$ giving the lattice value \mathbb{T} (top).

Theorem 2 makes the connection between trajectory assertions over lattice ‘values’ and Boolean trajectory assertions. If a trajectory assertion holds over lattice-valued streams, then the same antecedent leads to the same consequent over Boolean-valued streams. This is the crux of connection from the lattice world to the Boolean world.

Theorem 2 *Translate trajectory logic from lattice to Boolean sequences.*

For all $ckt, ant, cons$.

For all ϕ, σ .

$\sigma \in \mathcal{L} \text{ ckt}$ **implies**

For all $t \geq 0$. $((\phi, \sigma_t) \models_{\text{STE}} ant)$ **implies** $((\phi, \sigma_t) \models_{\text{STE}} cons)$

implies

For all $\phi, \overset{\circ}{\sigma}$.

$\overset{\circ}{\sigma} \in \overset{\circ}{\mathcal{L}} \text{ ckt}$ **implies**

For all $t \geq 0$. $((\phi, \overset{\circ}{\sigma}_t) \models^{\circ} ant)$ **implies** $((\phi, \overset{\circ}{\sigma}_t) \models^{\circ} cons)$

The proof of Theorem 2 relies on the two simple lemmas given below. The first, Lemma 1, says that if a Boolean sequence $\overset{\circ}{\sigma}$ is in the Boolean language of a circuit then the dropped version $\overset{\circ}{\sigma} \downarrow$ is in the lattice-valued language of the circuit. The proof of Lemma 1 is done by unfolding the definition of $\in \overset{\circ}{\mathcal{L}}$ and using Axiom 2.

Lemma 1 *Relationship between Boolean and ternary sequences.*

For all $\overset{\circ}{\sigma}$. $\overset{\circ}{\sigma} \in \overset{\circ}{\mathcal{L}} \text{ ckt}$ **iff** $\overset{\circ}{\sigma} \downarrow \in \mathcal{L} \text{ ckt}$

Proof:

$$\begin{aligned} & \overset{\circ}{\sigma} \in \overset{\circ}{\mathcal{L}} \text{ ckt} \\ \text{iff} & \{ \text{Definition 7 } (\in \overset{\circ}{\mathcal{L}}) \} \\ & \forall t. (\overset{\circ}{\sigma}^{\top}(t)) (\overset{\circ}{\mathbf{Y}} \text{ ckt}) (\overset{\circ}{\sigma}^{\top}(t+1)) \\ \text{iff} & \{ \text{Axiom 2 (relating } \mathbf{Y}, \overset{\circ}{\mathbf{Y}}) \} \\ & \forall t. (\mathbf{Y} \text{ ckt } ((\overset{\circ}{\sigma}^{\top}(t)) \downarrow)) \sqsubseteq (\overset{\circ}{\sigma}^{\top}(t+1)) \downarrow \\ \text{iff} & \{ \text{Definition 1 (drop)} \} \\ & \forall t. (\mathbf{Y} \text{ ckt } (\overset{\circ}{\sigma}^{\top} \downarrow (t))) \sqsubseteq (\overset{\circ}{\sigma}^{\top} \downarrow (t+1)) \\ \text{iff} & \{ \text{Definition 2 } (\in \mathcal{L}) \} \\ & \overset{\circ}{\sigma} \downarrow \in \mathcal{L} \text{ ckt} \end{aligned}$$

End of proof (Lemma 1).

Lemma 2 relates satisfaction over Boolean sequences to satisfaction over lattice-valued sequences. Its proof is by induction over the structure of trajectory formulas, unfolding the definitions of \models° and \models_{STE} , and employing properties of the drop operator.

Lemma 2 *Satisfaction over Boolean and ternary sequences.*

For all $\phi, f, \overset{\circ}{\sigma}$. $(\phi, \overset{\circ}{\sigma}) \models^{\circ} f$ **iff** $(\phi, \overset{\circ}{\sigma} \downarrow) \models_{\text{STE}} f$

Proof: The proof proceeds by induction over the structure of the trajectory formula f .

In the base case where f is ‘ n is 0’ we have, for arbitrary ϕ, σ :

$$\begin{aligned}
& (\phi, \sigma) \models^{\circ} (n \text{ is } 0) \\
\text{iff } & \{\text{Definition 8 } (\models^{\circ})\} \\
& (\sigma \ n \ 0) = \mathbf{F} \\
\text{iff } & \{\text{Definition 1 (drop)}\} \\
& (\sigma \ n \ 0)_{\downarrow} = 0 \\
\text{iff } & \{\text{Definition 1 (drop)}\} \\
& ((\sigma_{\downarrow}) \ n \ 0) = 0 \\
\text{iff } & \{\text{Definition 5 } (\models_{\text{STE}})\} \\
& (\phi, \sigma_{\downarrow}) \models_{\text{STE}} (n \text{ is } 0)
\end{aligned}$$

The proof for the case when f is ‘ n is 1’ is almost identical.

There are three step cases, one for each of **and**, **when**, and **N**. The cases are similar, so we show only the case when f is ‘ f' when g ’, which is the most interesting of the three. The induction hypothesis is

$$(\phi, \sigma) \models^{\circ} f' \quad \text{iff} \quad (\phi, \sigma_{\downarrow}) \models_{\text{STE}} f'$$

and the proof proceeds as follows:

$$\begin{aligned}
& (\phi, \sigma) \models^{\circ} (f' \text{ when } g) \\
\text{iff } & \{\text{Definition 8 } (\models^{\circ})\} \\
& (\phi \models g) \text{ implies } ((\phi, \sigma) \models^{\circ} f') \\
\text{iff } & \{\text{Induction hypothesis}\} \\
& (\phi \models g) \text{ implies } ((\phi, \sigma_{\downarrow}) \models_{\text{STE}} f') \\
\text{iff } & \{\text{Definition 5 } (\models_{\text{STE}})\} \\
& (\phi, \sigma_{\downarrow}) \models_{\text{STE}} (f' \text{ when } g)
\end{aligned}$$

End of proof (Lemma 2).

Theorem 2 follows by some simple first-order reasoning from these lemmas. We note that Theorem 2 is an implication and not an if-and-only-if result. The reason is that Lemma 2 is universally quantified over Boolean sequences only, and does not relate satisfaction by arbitrary lattice-valued sequences to satisfaction by Boolean ones. Considering the case where the trajectory formula f contains both n is 1 and n is 0 illustrates why. No Boolean sequence satisfies both n is 1 and n is 0, but a lattice-valued sequence in which $\sigma \ n \ 0 = \top$ would satisfy them both.

4 Trajectory Logic as Relations Over Streams

This section begins with a description of a shallow embedding of trajectory assertions in a higher-order logic² version of Boolean sequences. In the shallow embedding, trajectory formulas are predicates over Boolean sequences; later, in Section 4.2, we link the shallow embedding with a deep embedding of trajectory formulas.

²This is not a mechanized implementation, but rather a paper description that could be implemented in a higher-order logic proof system. In what follows, we therefore use the term ‘HOL’ to refer to higher order logic generally.

Although it is simple to do, the decision to develop both shallow and deep embeddings was made for important reasons. While we much prefer the shallow embedding for reasoning about properties of circuits stated as trajectory assertions, the deep embedding enables reasoning about trajectory logic itself (in particular, it allows quantification over trajectory formulas). We consider both activities important.

4.1 The Shallow Embedding

The focus of this paper is on the connection from trajectory evaluation over lattice values to relations over Boolean streams. Formalizing this connection forces us to reason about three different worlds: trajectory assertions with both lattice and Boolean values, and higher-order logic. A completely formal representation of these worlds (in particular, the semantics of higher-order logic [7]) and mappings between them would obfuscate the focus of our work. To maintain focus, we gloss over some of the semantic mappings between these different worlds. In particular, we use the same representation of trajectory formulas for both the metalogical results and a shallow embedding of trajectory formulas in logic.

Circuits are ordinarily represented as relations over streams in the relational style of modelling; checking that a sequence is in the language of a circuit is done by simply applying the circuit relation to the sequence. In this work, however, we treat circuits as uninterpreted and axiomatize the relationship between circuits in our shallow embedding and the language of circuits over Boolean sequences.

Axiom 3 *Relating languages in deep and shallow embeddings.*

A sequence $\mathring{\sigma}$ is in the language of an embedded, uninterpreted circuit ckt ($\mathring{\sigma}$ in_lang ckt) if and only if $\mathring{\sigma} \in \mathring{\mathcal{L}}(ckt)$.

This is an axiom, rather than a lemma, because we do not give interpretations of circuits. Indeed, for much practical work we only require the ability to distinguish verification results obtained on different circuits, and for this purpose it is sufficient to leave circuits uninterpreted. A complete implementation of the work described here would need to prove that the implementation of circuits satisfies Axiom 3.

Definition 9 presents the trajectory formula type and the trajectory formula constructs `is`, `and`, `when`, and `N` for the shallow embedding in a higher-order logic. In the shallow embedding, a sequence satisfies a formula if applying it to the sequence yields true.

Definition 9 *Shallow embedding in HOL of trajectory formulas in Boolean streams.*

$$\begin{aligned} \text{traj_form} & \stackrel{\text{type}}{=} (\text{bool sequence}) \rightarrow \text{bool} \\ n \text{ is } 1 & \stackrel{\text{def}}{=} \lambda \mathring{\sigma}. \mathring{\sigma} \ n \ 0 = \text{T} \\ n \text{ is } 0 & \stackrel{\text{def}}{=} \lambda \mathring{\sigma}. \mathring{\sigma} \ n \ 0 = \text{F} \\ f_1 \text{ and } f_2 & \stackrel{\text{def}}{=} \lambda \mathring{\sigma}. (f_1 \ \mathring{\sigma}) \wedge (f_2 \ \mathring{\sigma}) \\ f \text{ when } g & \stackrel{\text{def}}{=} \lambda \mathring{\sigma}. g \implies (f \ \mathring{\sigma}) \\ \text{N } f & \stackrel{\text{def}}{=} \lambda \mathring{\sigma}. f \ \mathring{\sigma}_1 \end{aligned}$$

We are now in a position to use the results of the previous section to derive the key relationship between the standard trajectory logic and our shallow embedding of trajectory logic as predicates over Boolean streams in HOL (Theorem 3).

Theorem 3 *Translate STE result to shallow embedding of Boolean streams.*

For all circuits ckt , antecedents ant , and consequences $cons$, if an implementation of the STE algorithm returns e :

$$e = \text{STE } ckt \text{ ant } cons$$

then the following is a true formula in HOL:

$$\models_{\text{HOL}} e \implies \forall \sigma. \sigma \text{ in_lang } ckt \implies \forall t \geq 0. ant(\sigma_t) \implies cons(\sigma_t)$$

The occurrences of ant and $cons$ in the HOL formula refer, of course, to the shallow embedding of the corresponding STE formulas.

Theorem 3 follows directly from the results of the previous sections (Theorems 1 and 2), together with the following lemma concerning the ability to ‘time shift’ trajectory evaluation results. Let $\mathbf{N}^n f$ denotes n applications of the next-time temporal operator \mathbf{N} to the formula f . Then we have the following lemma.

Lemma 3 *Trajectory evaluation time shifting.*

$$\phi \models_{ckt} [ant \implies\!\!\implies cons] \quad \mathbf{iff} \quad (\text{For all } t \geq 0. \phi \models_{ckt} [\mathbf{N}^t cons \implies\!\!\implies \mathbf{N}^t cons])$$

Proof: For the ‘if’ direction, note that $\mathbf{N}^0 f = f$. The ‘only if’ direction was proved by Hazelhurst and Seger in previous work (Theorem III.2 of [9]).

End of proof (Lemma 3).

4.2 The Deep Embedding

We now overlay onto our previous results a simple deep embedding of trajectory formulas in HOL. Definitions for the deep embedding of trajectory formulas and satisfaction (is, and, when, N and sat) are not shown because they are direct implementations of the metalogical presentation in Section 3. We simply identify the deeply embedded operators by underlining them.

The following defines a translation $\llbracket \cdot \rrbracket$ from deeply-embedded to shallowly-embedded trajectory formulas.

Definition 10 *Translation from deep to shallow embedding.*

$$\begin{aligned} \llbracket (n \text{ is } 1) \rrbracket & \stackrel{def}{=} n \text{ is } 1 \\ \llbracket (n \text{ is } 0) \rrbracket & \stackrel{def}{=} n \text{ is } 0 \\ \llbracket (f_1 \text{ and } f_2) \rrbracket & \stackrel{def}{=} \llbracket f_1 \rrbracket \text{ and } \llbracket f_2 \rrbracket \\ \llbracket (f \text{ when } g) \rrbracket & \stackrel{def}{=} \llbracket f \rrbracket \text{ when } g \\ \llbracket (\mathbf{N} \text{ } f) \rrbracket & \stackrel{def}{=} \mathbf{N} \llbracket f \rrbracket \end{aligned}$$

The core of the relationship between trajectory formulas in the deep and shallow embeddings is captured in Theorem 4. The theorem says that translating a deeply embedded formula f' to a shallow embedding via $\llbracket \cdot \rrbracket$ (Definition 10) results in a trajectory formula that is satisfied by exactly the same set of sequences as f' . The proof is done by induction over the structure of trajectory formulas in the deep embedding.

Theorem 4 *Translate trajectory logic over Booleans in logic from deep to shallow embedding.*

$$\forall f, \dot{\sigma}. (\llbracket f \rrbracket \dot{\sigma}) \iff (\dot{\sigma} \text{ \underline{sat} } f)$$

We can use this correspondence to derive our final result, Theorem 5, which says that a trajectory evaluation result over Boolean streams holds if and only if it holds in logic using a deep embedding of trajectory formulas.

Theorem 5 *Translate trajectory logic over Booleans to deep embedding in higher-order logic.*

For all circuits ckt , antecedents ant , and consequences $cons$, if

For all ϕ .

$\phi \models e$ **implies**

For all $\dot{\sigma}$.

$\phi \models (\dot{\sigma} \in \mathcal{L} \text{ ckt})$ **implies**

For all $t \geq 0$. $((\phi, \dot{\sigma}_t) \models^{\circ} ant)$ **implies $((\phi, \dot{\sigma}_t) \models^{\circ} cons)$**

then the following is a true formula in HOL:

$$\begin{aligned} \models_{\text{HOL}} e \implies \\ \forall \dot{\sigma}. \dot{\sigma} \text{ in_lang } ckt \implies \forall t \geq 0. (\dot{\sigma}_t \text{ \underline{sat} } ant) \implies (\dot{\sigma}_t \text{ \underline{sat} } cons) \end{aligned}$$

The theorem follows directly from Theorem 3 using Theorem 4. The occurrences of ant and $cons$ in the HOL formula refer to the deep embedding of the corresponding STE formulas.

5 Simple Example Revisited

We now revisit the simple NAND-DELAY circuit first introduced in Section 1. When we introduced the circuit, we showed an intuitive, but not quite technically correct, translation from a standard trajectory assertion to a relations over Boolean streams result. We now demonstrate how the semantic link provided by our principal theorem (Theorem 3) induces a formally correct and intuitively satisfying connection between user-level verifications in trajectory evaluation and higher-order logic.

We begin with the trajectory formula that first appeared in Figure 2.

$$\begin{aligned} i &\stackrel{\text{def}}{=} "i" \\ o &\stackrel{\text{def}}{=} "o" \end{aligned}$$

$$\text{NAND-DELAY} \models \left[\begin{array}{l} (i \text{ is } b_1) \text{ and} \\ (o \text{ is } b_2) \end{array} \implies\!\!\!\implies (\text{N } (o \text{ is } \neg(b_1 \wedge b_2))) \right]$$

Expanding with the definition of ‘is’ (Definition 4) results in

$$\text{NAND-DELAY} \models \left[\begin{array}{c} \left(\begin{array}{l} \text{“}i\text{” is 1 when } b_1 \text{ and “}i\text{” is 0 when } \neg b_1 \text{ and} \\ \text{“}o\text{” is 1 when } b_2 \text{ and “}o\text{” is 0 when } \neg b_2 \end{array} \right) \\ \implies \\ \left(\mathbb{N} \left(\begin{array}{l} \text{“}o\text{” is 1 when } \neg(b_1 \wedge b_2) \text{ and} \\ \text{“}o\text{” is 0 when } \neg\neg(b_1 \wedge b_2) \end{array} \right) \right) \end{array} \right]$$

Running the STE algorithm establishes that the NAND-DELAY circuit does satisfy the specification, and indeed satisfies it for all valuations of the Boolean variables b_1 and b_2 . By instantiating e with \top in Theorem 3, we can justify postulating the HOL axiom shown below:

$$\begin{aligned} \vdash \forall \overset{\circ}{\sigma}. \overset{\circ}{\sigma} \text{ in_lang NAND-DELAY} &\implies \\ \forall t \geq 0. & \\ \overset{\circ}{\sigma}_t &\models \left(\begin{array}{l} \text{“}i\text{” is 1 when } b_1 \text{ and “}i\text{” is 0 when } \neg b_1 \text{ and} \\ \text{“}o\text{” is 1 when } b_2 \text{ and “}o\text{” is 0 when } \neg b_2 \end{array} \right) \implies \\ \overset{\circ}{\sigma}_t &\models \left(\begin{array}{l} \text{“}o\text{” is 1 when } \neg(b_1 \wedge b_2) \text{ and} \\ \text{“}o\text{” is 1 when } \neg\neg(b_1 \wedge b_2) \end{array} \right) \end{aligned}$$

Unfolding the definition of \models (Definition 8) and performing some routine simplification results in

$$\begin{aligned} \vdash \forall \overset{\circ}{\sigma}. \overset{\circ}{\sigma} \text{ in_lang NAND-DELAY} &\implies \\ \forall t \geq 0. & \\ \left(\begin{array}{l} (\overset{\circ}{\sigma} \ t \ \text{“}i\text{”}) = b_1 \ \wedge \\ (\overset{\circ}{\sigma} \ t \ \text{“}o\text{”}) = b_2 \end{array} \right) &\implies ((\overset{\circ}{\sigma} \ (t+1) \ \text{“}o\text{”}) = \neg(b_1 \wedge b_2)) \end{aligned}$$

Because b_1 and b_2 are free variables, we can universally quantify over them, and then use the equalities in the antecedent to substitute for b_1 and b_2 through the consequent. This finally leaves us with the result we had hoped for, namely the following intuitive user-level HOL theorem.

$$\begin{aligned} \vdash \forall \overset{\circ}{\sigma}. \overset{\circ}{\sigma} \text{ in_lang NAND-DELAY} &\implies \\ \forall t \geq 0. ((\overset{\circ}{\sigma} \ \text{“}o\text{”} \ (t+1)) = \neg((\overset{\circ}{\sigma} \ \text{“}i\text{”} \ t) \wedge (\overset{\circ}{\sigma} \ \text{“}o\text{”} \ t))) & \end{aligned}$$

6 Voss Implementation of Trajectory Evaluation

In this section we briefly touch on the Voss implementation of trajectory logic. Information on Voss has been published previously [8], we include this section to make the connection to an implementation of trajectory evaluation more concrete. In Voss, the four-valued lattice is represented by a pair of Booleans, called a dual-rail value (Definition 11). In Voss, Booleans are implemented as BDDs and may therefore be symbolic. A Boolean value v is translated to a dual-rail value by putting v on the high rail and its negation on the low rail ($v \downarrow \stackrel{def}{=} (v, \neg v)$).

Definition 11 *Implementation of dual-rail lattice in Voss.*

$$\begin{aligned} \top &\stackrel{def}{=} (\mathbf{F}, \mathbf{F}) \\ 1 &\stackrel{def}{=} (\mathbf{T}, \mathbf{F}) \\ 0 &\stackrel{def}{=} (\mathbf{F}, \mathbf{T}) \\ \times &\stackrel{def}{=} (\mathbf{T}, \mathbf{T}) \end{aligned}$$

Trajectory formulas are represented in Voss by a very simple, but slightly indirect, deep embedding. Rather than representing trajectory formulas by a data type that mirrors the abstract syntax tree of the formulas, a trajectory formula is a list of ‘5-tuples’ (Figure 5). The five elements of each tuple are the guard, a node name, a Boolean value (which may be a variable), a natural-number start-time, and a natural-number end-time. The meaning is ‘if the guard is true, then the node takes on the given value from the start-time to the end-time.’

```
//      guard      node      value      start      end
traj_form  $\stackrel{type}{=}$  ( bool × string × bool × nat × nat ) list

n is v       $\stackrel{def}{=}$  [(T, n, v, 0, 1)]
f1 and f2  $\stackrel{def}{=}$  f1 append f2
f when g     $\stackrel{def}{=}$  map (λ(g', n, v, t0, t1). (g' ∧ g, n, v, t0, t1)) f
N f          $\stackrel{def}{=}$  map (λ(g, n, v, t0, t1). (g, n, v, t0 + 1, t1 + 1)) f

f from t0  $\stackrel{def}{=}$  map (λ(g, n, v, z, t1). (g, n, v, t0, t1)) f
f to t1    $\stackrel{def}{=}$  map (λ(g, n, v, t0, z). (g, n, v, t0, t1)) f
```

Figure 5: Implementation of trajectory formulas in Voss

As the definitions in Figure 5 show, every formula in the trajectory logic can be represented quite simply by a 5-tuple list. Moreover, it is clear that the list representation does not add anything new to the expressive power of our formulas: any 5-tuple whose duration spans more than one-time unit can be expressed by an appropriate conjunction of applications of the next time operator.

7 Conclusion

The motivation of our work is to combine the automatic deductive power of finite-state model checking with the expressive power and flexibility of general-purpose theorem proving. Key elements of such an architecture are semantic links between the general logic of the theorem prover and the more specialised model-checker logic. To this end, we have established a semantic link between correctness results in the temporal logic of symbolic trajectory evaluation and relational hardware specifications in higher order logic.

Ultimately, our aim is to combine results from several different model checkers in a single system. The expressive power of typed higher-order logic, which is a foundational formalism, and the generality of the relational approach are an excellent ‘glue logic’ for this. For many other linear-time temporal logics, at least, it seems quite reasonable to expect smooth embeddings into the relational world. Indeed, there have been several such embeddings in the literature [19, 4]. Branching time logics may be more challenging, but the PVS results cited earlier are encouraging. More general temporal logics, such as Extended Temporal Logic [18] will provide a further challenge.

Semantically, our work does not prejudge the architectural question of ‘black box’ versus ‘white box’ integration of the verification tools. However, experience with a variety of techniques for combining trajectory evaluation with theorem proving [11, 1] has shown that a purely black box integration is not an effective solution for gate-level hardware verification. Large projects, such as PROSPER [15] and Ω MEGA [14], that are integrating multiple formal tools are pursuing a variety of approaches.

We also do not prescribe that the model checker is used solely as a decision procedure to dispose of subgoals generated by the theorem prover in the course of a verification. The general idea is that verification problems can be decomposed into finite state proof obligations, which are then vulnerable to attack by efficient model checking engines. Deductive reasoning in the theorem prover justifies the decomposition and stitches model-checking results together. But the strategy can be viewed as way to extend the reach of model checkers (by combining results to yield properties not provable monolithically) or to improve model checking efficiency (by transforming specifications so they are simpler to verify automatically), as well as simply automating proof of subgoals.

Indeed, the theorem-proving component can be quite lightweight and still add enormous value. One approach being pursued by the PROSPER group is to view the theorem prover and its logic as a kind of ‘scripting language’ for embedded proof engines built from reasoning tools components. Higher-order logic provides a mathematically rigorous interface between the specialised logics involved. And in an LCF-style system like HOL, the functional metalanguage (ML) provides very powerful and flexible control language for orchestrating verifications and tools.

Acknowledgments

This paper benefited from the helpful comments of Carl Seger, Ching-Tsun Chou, and the anonymous referees of the conference version.

References

- [1] M. D. Aagaard, R. B. Jones, and C.-J. H. Seger. Formal verification using parametric representations of Boolean constraints. In *ACM/IEEE Design Automation Conference*, July 1999.

- [2] S. Agerholm and H. Skjødt. Automating a model checker for recursive modal assertions in HOL. Technical Report DAIMI IR-92, Computer Science Department, Aarhus University, 1990.
- [3] A. Cheng and K. Larsen, editors. *Program and Abstracts of the BRICS Autumn School on the Verification*, Aug. 1996. BRICS Notes Series NS-96-2.
- [4] C.-T. Chou. Predicates, temporal logic, and simulations. In C.-J. H. Seger and J. J. Joyce, editors, *HOL User's Group Workshop*, pages 310–323. Springer Verlag; New York, Aug. 1994.
- [5] C.-T. Chou. The mathematical foundation of symbolic trajectory evaluation. In *Workshop on Computer-Aided Verification*. Springer Verlag; New York, 1999.
- [6] G. C. Gopalakrishnan and P. J. Windley, editors. *Formal Methods in Computer-Aided Design*. Springer Verlag; New York, Nov. 1998.
- [7] M. J. C. Gordon and T. F. Melham, editors. *Introduction to HOL: a theorem proving environment for higher order logic*. Cambridge University Press, New York, 1993.
- [8] S. Hazelhurst and C.-J. H. Seger. Symbolic trajectory evaluation. In T. Kropf, editor, *Formal Hardware Verification*, chapter 1, pages 3–78. Springer Verlag; New York, 1997.
- [9] S. Hazelhurst and C.-J. H. Seger, A Simple Theorem Prover Based on Symbolic Trajectory Evaluation and BDDs', *IEEE Transactions on the Computer Aided Design*, April 1995.
- [10] A. J. Hu and M. Y. Vardi, editors. *Computer Aided Verification*. Springer Verlag; New York, July 1998.
- [11] J. Joyce and C.-J. Seger. Linking BDD based symbolic evaluation to interactive theorem proving. In *ACM/IEEE Design Automation Conference*, June 1993.
- [12] M. Newey and J. Grundy, editors. *Theorem Proving in Higher Order Logics*. Springer Verlag; New York, Sept. 1998.
- [13] J. O'Leary, X. Zhao, R. Gerth, and C.-J. H. Seger. Formally verifying IEEE compliance of floating-point hardware. *Intel Technical Journal*, First Quarter 1999. Online at <http://developer.intel.com/technology/itj/>.
- [14] C. Benzmüller, L. Cheikhrouhou, D. Fehrer, A. Fiedler, X. Huang, M. Kerber, M. Kohlhase, A. Meirer, E. Melis, W. Schaarschmidt, J. Siekmann, and V. Sorge, Ω MEGA, Towards a mathematical assistant. in *14th Conference on Automated Deduction*, edited by W. McCune, Lecture Notes in Artificial Intelligence, Vol. 1249, (Springer-Verlag, 1997), pp. 252–255.
- [15] Proof and specification assisted design environments, ESPRIT LTR project 26241, 1999. <http://www.dcs.gla.ac.uk/prosper/>.

- [16] S. Rajan, N. Shankar, and M. Srivas. An integration of model checking automated proof checking. In *Workshop on Computer-Aided Verification*. Springer Verlag; New York, 1996.
- [17] C.-J. H. Seger and R. E. Bryant. Formal verification by symbolic evaluation of partially-ordered trajectories. *Formal Methods in System Design*, 6(2):147–189, Mar. 1995.
- [18] M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115:1–37, 1994.
- [19] J. von Wright. Mechanising the temporal logic of actions in HOL. In M. Archer, J. J. Joyce, K. N. Levit, and P. J. Windley, editors, *International Workshop on the HOL Theorem Proving System and its Applications*, pages 155–159. IEEE Computer Society Press, Washington D.C., Aug. 1991.