

STeP: A Tool for the Development of Provably
Correct Reactive and Real-Time Systems
— Final Technical Report —

P.I.: Prof. Zohar Manna
Computer Science Department
Stanford University
Stanford, CA. 94305-9045

June, 1999

U.S. Army Research Office
Grant Number: DAAH04-95-1-0317

Approved for Public Release;

Distribution Unlimited.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Blank reporting burden for this 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 and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 06/30/99	3. REPORT TYPE AND DATES COVERED FINAL REPORT: 6/1/95-3/31/99	
4. TITLE AND SUBTITLE STeP: A tool for the Development of Provably Correct Reactive and Real-Time Systems		5. FUNDING NUMBERS DAAH04-95-1-0317	
6. AUTHOR(S) P.I.: Professor Zohar Manna		8. PERFORMING ORGANIZATION REPORT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Computer Science Department Stanford University Stanford CA 94305-9045		10. SPONSORING/MONITORING AGENCY REPORT NUMBER ARO 34456.8-MA	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P. O. Box 12211 Research Triangle Park, NC 27709-2211		11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.	
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This research is directed towards the implementation of a comprehensive toolkit for the development and verification of high assurance reactive systems, especially concurrent, real-time, and hybrid systems. For this, we have designed and implemented the STeP (Stanford Temporal Prover) verification system. STeP is a tool for the computer-aided formal verification of reactive systems, including real-time and hybrid systems, based on their temporal specification. STeP integrates model checking and deductive methods to allow the verification of a broad class of systems, including parameterized (N-component) circuit designs, parameterized (N-process) programs, and programs with infinite data domains.			
14. SUBJECT TERMS		15. NUMBER OF PAGES	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED		16. PRICE CODE	
18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	
20. LIMITATION OF ABSTRACT UL			

Contents

1	Introduction	4
2	Describing Reactive Systems	4
2.1	Transition Systems	5
2.2	SPL Programs	5
2.3	Real-Time Systems	6
2.4	Hybrid Systems	7
2.5	Modularity	7
2.6	Hardware Description	8
3	Verification Methodology	8
4	Deductive Verification and Model Checking	9
4.1	Property Specification: Linear-Time Temporal Logic	10
4.2	Deductive Verification	11
4.3	Model Checking	12
4.4	Modular Verification	12
5	Combining Deductive and Algorithmic Methods	13
5.1	Automatic Invariant Generation	13
5.2	Decision Procedures	14
5.3	Generalized Verification Diagrams	16
5.4	Constructing Finite-State Abstractions	18
6	Applications	19
6.1	Real-time Systems	19
6.2	Fault-tolerant Systems	19
6.3	Hybrid Systems	19
6.4	Case Study: Steam boiler	19
6.5	Educational Use	20
7	Implementation	20
7.1	Stand-alone components	21
7.2	External systems	21
7.3	Obtaining STeP	21
8	Publications	22
9	Participating Scientific Personnel	24

10 Report of Inventions

25

11 Bibliography

26

1 Introduction

Reactive systems have an ongoing interaction with their environment, and their computations are infinite sequences of states. A large number of systems can be seen as reactive systems, including hardware, concurrent programs, network protocols, and embedded systems. Temporal logic provides a convenient language for expressing properties of reactive systems. A temporal verification methodology provides procedures for proving that a given system satisfies a given temporal property.

This research is directed towards the implementation of a comprehensive toolkit for the development and verification of high assurance reactive systems, especially concurrent, real-time, and hybrid systems. For this, we have designed and implemented the STeP (Stanford Temporal Prover) verification system.

STeP is a tool for the computer-aided formal verification of reactive systems, including real-time and hybrid systems, based on their temporal specification. STeP integrates model checking and deductive methods to allow the verification of a broad class of systems, including parameterized (N -component) circuit designs, parameterized (N -process) programs, and programs with infinite data domains.

Figure 1 presents an outline of the STeP system. The main inputs are a reactive system and a property to be proven for it, expressed as a temporal logic formula. The system can be a hardware or software description, and include real-time and hybrid components (Section 2). Verification is performed by model checking or deductive means (Section 4), or a combination of the two (Section 5).

2 Describing Reactive Systems

The various systems STeP can verify differ in their time model—discrete, real-time, or hybrid—as well as in the domain of their state variables, which can be finite or infinite. Furthermore, systems can be *parameterized* in the number of processes that compose them (N -process systems). All of these systems can be modeled, however, using the same underlying computational model: (fair) transition systems [MP95]. This basic model is extended in appropriate ways to allow for modular structures, hardware-specific components, clocks, or continuous variables. Figure 2 describes the scope of STeP, classified along these three main dimensions.

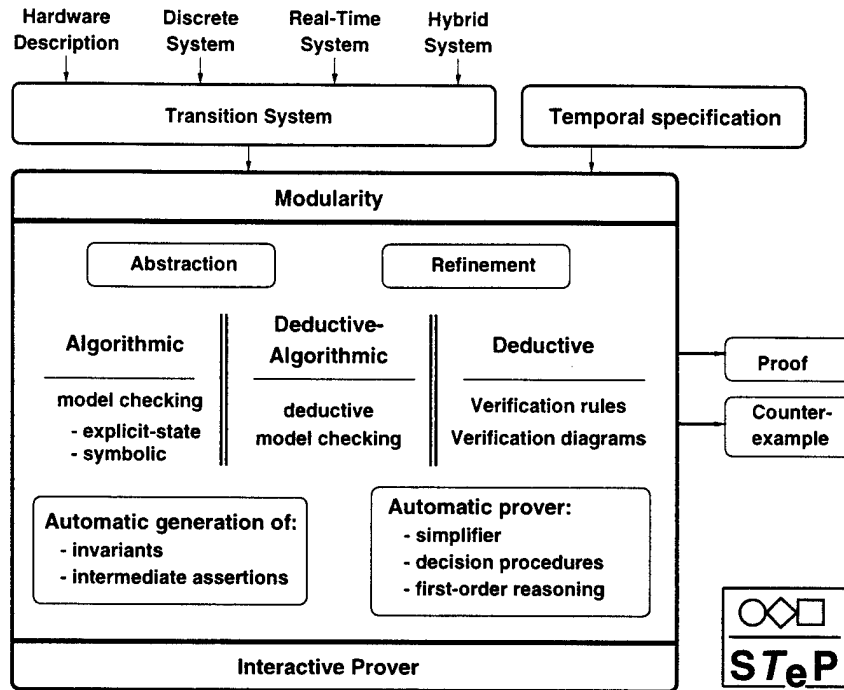


Figure 1: An outline of the STeP system

2.1 Transition Systems

The basic system representation in STeP uses a set of *transitions*. Each transition is a relation over unprimed and primed system variables, expressing the values of the system variables at the current and next state. Transitions can thus be represented as general first-order formulas, though more specialized notations for guarded commands and assignments is also available. In the discrete case, transitions can be labeled as *just* or *compassionate*; such fairness constraints are relevant to the proof of progress properties (see [MP95]).

2.2 SPL Programs

For convenience, discrete systems can be described in the Simple Programming Language (SPL) of [MP95]. SPL programs are automatically translated into the corresponding fair transition systems, which are then used as the basis for verification.

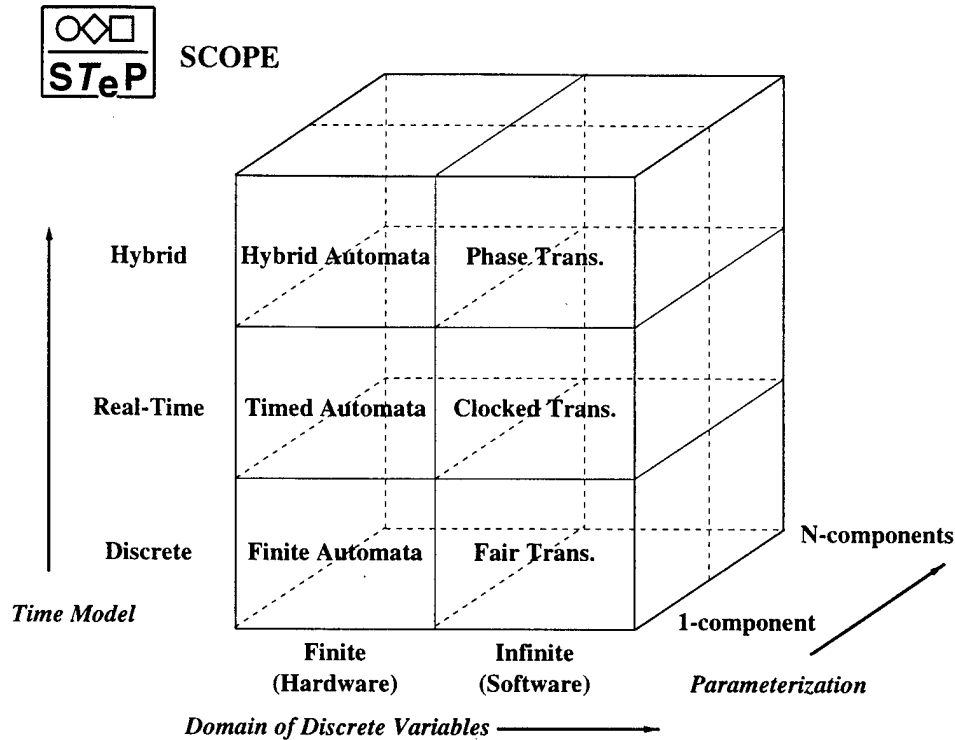


Figure 2: Scope of STeP

2.3 Real-Time Systems

STeP can verify properties of real-time systems, using the computational model of *clocked transition systems* [MP96]. Clocked transition systems consist of standard instantaneous transitions that can reset auxiliary clocks, and a *progress condition* that limits the time that the system can stay in a particular discrete state. Clocked transition systems are converted into discrete transition systems by including a *tick* transition that advances time, constrained by the progress condition. The tick transition is parameterized by a positive real-valued duration of the time step.

Temporal logic properties can refer to the global clock, and the auxiliary ones, to specify real-time properties; the underlying temporal logic remains the same. Since the transition system framework is also retained by representing *tick* as a discrete parameterized transition of the time-step, this

representation allows STeP to reuse existing verification rules for untimed temporal logic can be reused for [KMP96].

2.4 Hybrid Systems

Hybrid transition systems generalize clocked transition systems, by allowing real-valued variables other than clocks to vary continuously over time. The evolution of continuous variables is described by a set of constraints, which can be sets of differential equations or differential inclusions. Similar to clocked transition systems, hybrid transition systems are converted into discrete transition systems by including a tick transition, parameterized by the duration of the time step. However, for hybrid systems the tick transition must not only update the values of the clocks, which is straightforward, but must also determine the value of the continuous variables at the end of the time step. The updated value of the continuous variables is represented symbolically; axioms and invariants, generated based on the constraints, are used to determine the actual value or the range of values at the time they are needed.

Other formalisms such as timed transition systems, timed automata and hybrid automata can be easily translated into hybrid and clocked transition systems [MP96]. Furthermore, the transition system representation allows discrete variables that range over infinite-domains, as opposed to automata-based formalisms where the domain of discrete variables must be finite-state.

2.5 Modularity

Complex systems are built from smaller components. Most modern programming languages and hardware description languages therefore provide the concept of *modularity*. STeP includes facilities for modular specification and verification [FMS98], based on *modular transition systems*, which can concisely describe complex transition systems. Each module has an *interface* that determines the observability of module variables and transitions. The interface may also include an *environment assumption*, a relation over primed and unprimed interface variables that limits the possible environments the module can be placed in. The module can only be composed with other modules that satisfy the environment assumption. Communication between a module and its environment can be asynchronous, through shared variables, and synchronous, through synchronization of labeled transitions.

More complex modules can be constructed from simpler ones by possibly recursive module expressions, allowing the description of hierarchical

systems of unbounded depth. Module expressions can refer to modules defined earlier, or instances of parameterized modules, enabling the reuse of code and of properties proven about these modules. Besides the usual hiding and renaming operations, the language provides a construct to augment the interface with new variables that provide a summary value of multiple variables within the module. Symmetrically, a restriction operation allows the module environment to combine or rearrange the variables it presents to the module.

Real-time and hybrid systems: Real-time and hybrid systems can also be described as modular systems; discrete, real-time and hybrid modules may be combined into one system. The evolution constraints of hybrid modules may refer to continuous variables of other modules, thus enabling the decomposition of systems into smaller modules. To enable proofs of nontrivial properties over such modules, we allow arbitrary constraints on these external continuous variables in the environment assumption.

2.6 Hardware Description

A Verilog hardware description language front-end has recently been added to STeP. Its main component is a compiler that takes Verilog input and produces a fair transition system, which can then be analyzed using the deductive and algorithmic tools of STeP.

The goal of this compiler is to produce a faithful representation of the input program, taking into account the delays and events that are part of the Verilog semantics. The compiler extends the Verilog language by allowing parameters to be left unspecified. These parameters can be used to declare bit vectors of arbitrary size, or to compose an array of lower-level modules. These features cater to the deductive component of STeP, which can verify properties of general infinite-state systems.

3 Verification Methodology

STeP is best viewed as providing a toolkit of verification methods based on a common system description language and specification language. A given system can be analyzed in a number of ways. Depending on the system and property to be proved, different tools will be applicable or more appropriate:

Model Checking: If the system is finite-state, arbitrary temporal properties can be automatically established or refuted, using explicit-state or

symbolic model checking (Section 4.3). Symbolic model checking is applicable only if all variables have finite-domain. Some classes of infinite-state systems can be checked with the explicit-state model checker, which is not guaranteed to terminate in this case. For parameterized and infinite-state systems, finite-state instances of the system can be model checked to quickly search for bugs.

Invariant Generation: For most deductive verification proofs, system invariants of increasing strength must be collected, where previous invariants are used to establish subsequent ones. Automatic invariant generation is used to establish a basic initial set of invariants (Section 5.1).

Verification Rules: To prove simple safety properties, deductive rules can be used, with the user providing adequate strengthenings (intermediate assertions) where necessary (Section 4.2). Previously established invariants are used to prove the required verification conditions, which are automatically generated by the system.

Verification Diagrams: A *verification diagram* can be provided by the user, as a system abstraction that proves a particular property in question. Verification conditions justify the correctness of the diagram, while an algorithmic procedure checks that the diagram, indeed, proves the property at hand (Sect. 5.3).

Abstraction: A *finite-state abstraction* of the system can be generated, such that properties model checked for the abstract system will hold of the original system as well (Sect. 5.4). Since the abstraction is finite-state, it can be model checked. Available invariants improve the quality of the abstraction, allowing more properties to be proved.

4 Deductive Verification and Model Checking

STeP provides a comprehensive, integrated environment to prove temporal properties over reactive systems. The STeP Session Editor, presented in Figure 3, keeps track of the main properties of interest throughout the verification session, including axioms, assumptions, previously proven properties, and automatically generated invariants, as well as the module to which each applies. Thus, it can handle multiple systems and proofs simultaneously. Properties can be activated or deactivated to control the extent of their use in automatic theorem-proving.

File System Diagram Property Trace

☐ ☐ ☐

Systems			
Name	Type	File	
Boiler	FTS	/manet/u2/step/examples/steamboiler/SEP/Boiler.fts	
Controller	FTS	/manet/u2/step/examples/steamboiler/SEP/Controller.fts	
BoilerSystem	FTS	/manet/u2/step/examples/steamboiler/SEP/BoilerSystem.fts	

Properties						
Name	Type	Syst...	Module	Text	Proven	Acti...
readBoilerSensors	Goal	Cont...	Controller	pc = readBoilerSensors ∧ programstatus = programrunning ...	<input type="checkbox"/>	<input type="checkbox"/>
readValvePos	Goal	Cont...	Controller	pc = readValvePos ∧ programstatus = programrunning ==> ..	<input type="checkbox"/>	<input type="checkbox"/>
readPumpState	Goal	Cont...	Controller	pc = readPumpState ∧ programstatus = programrunning ∧ ..	<input type="checkbox"/>	<input type="checkbox"/>
readPumpState	Goal	Cont...	Controller	pc = readPumpState ∧ programstatus = programrunning ∧ ..	<input type="checkbox"/>	<input type="checkbox"/>
consistency	Goal	maint	maint	[maint] = ((¬)(eqstate = broken ∧ M.maintstate = ok) ==> ..	<input type="checkbox"/>	<input type="checkbox"/>
returntoService	Goal	maint	maint	[maint] = ((¬)(eqstate = broken ∧ M.maintstate = ok) ==> ..	<input type="checkbox"/>	<input type="checkbox"/>
operations	Goal	oper...	operations	□(O.opstate = preparing ∧ plantstate = idle ∨ O.opstate = ..	<input type="checkbox"/>	<input type="checkbox"/>
respondtoEmergency	Goal	oper...	operations	O.emergency ∨ gotoEmstop ==> <>(O.opstate = preparing / ..	<input type="checkbox"/>	<input type="checkbox"/>
mingrad	Axiom	Boile...	BoilerSystem	□mingrad * delta < 0	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
maxgrad	Axiom	Boile...	BoilerSystem	□0 < maxgrad * delta	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
steamflow	Axiom	Boile...	BoilerSystem	□0 <= B.sf	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 3: STeP Session Editor

4.1 Property Specification: Linear-Time Temporal Logic

We use *linear-time temporal logic* (LTL) to represent properties of reactive systems [MP95]. A model of LTL is an infinite sequence of states. We use the usual temporal operators, such as $\Box p$ (p is always true), $\Diamond p$ (p is eventually true), $p \mathcal{U} q$ (p is true until q is true, which eventually happens), and $p \mathcal{W} q$ (p awaits q — p is true at least until q is true, but q need not eventually happen).

We distinguish between *safety* and *progress* properties. Informally, safety properties say that certain “bad states” will never be reached, e.g. as in an invariance $\Box p$ for an assertion p .

Progress properties, on the other hand, can say that “good” states will eventually be reached (perhaps recurrently). Safety properties do not depend on the fairness constraints of the system, whereas progress properties require the justice or compassion of particular transitions in order to be proved.

To specify properties of real-time and hybrid systems, temporal-logic properties can refer to the global and auxiliary clocks, and to the continuous

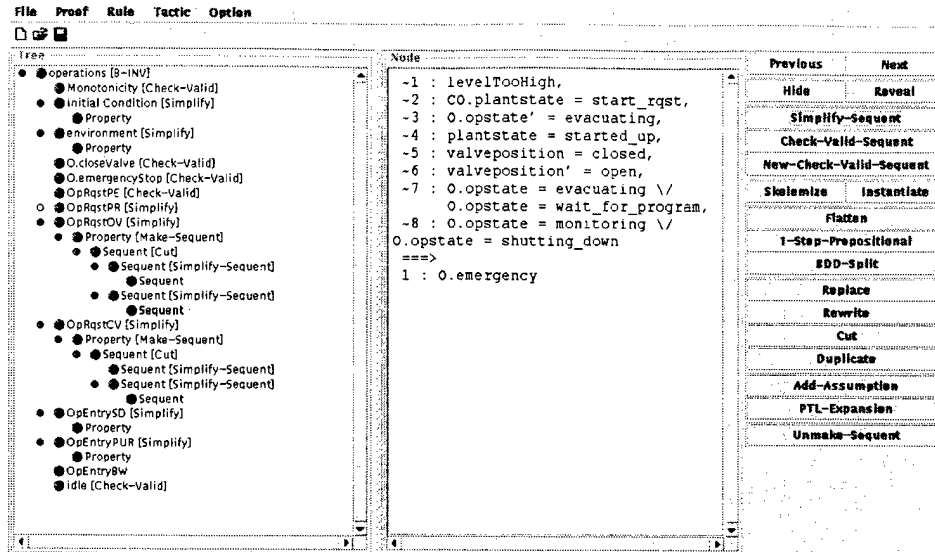


Figure 4: STeP Proof Editor

variables; the underlying temporal logic remains the same.

4.2 Deductive Verification

The deductive methods of STeP verify temporal properties of systems by means of verification rules and verification diagrams. *Verification rules* reduce temporal properties of systems to first-order verification conditions [MP95]. *Verification diagrams* [MP94] provide a visual language for guiding, organizing, and displaying proofs, and automatically generating the appropriate verification conditions as well (see Section 5.3).

As clocked and hybrid transition systems are converted into fair transition systems, verification rules and diagrams are uniformly applicable to discrete, real-time and hybrid systems. However, due to the parameterization of the tick transition, the resulting verification conditions for real-time and hybrid systems are usually more complex than those for (unparameterized) discrete systems.

Figure 4 shows the STeP Proof Editor, which is used to apply the basic deductive temporal verification rules as well as the Gentzen-style interactive theorem proving rules. In a typical deductive verification effort, the top-level

goal is a temporal formula to be proven valid for a given system. Verification rules or diagrams are used to generate verification conditions, as subgoals, which together imply the system validity of the original temporal property. These subgoals are then established automatically using decision procedures (Section 5.2) or interactively using the Gentzen-style rules. Model checking is also initiated by the Proof Editor.

4.3 Model Checking

STeP features automatic explicit-state and symbolic model checking for linear-time temporal logic. The *explicit-state model checker* performs an incremental (depth-first) search of the state-space, directed by the temporal tableau (automaton) for the negated specification. Thus, only those states that can potentially violate the specification are visited. This enables the use of the explicit-state model checker on some infinite-state systems, though it is not guaranteed to terminate for these systems. The *symbolic model checker* uses a breadth-first search through sets of states represented by ordered binary decision diagrams (OBDDs). Thus, it is limited to finite-state systems, whose variables range over a fixed, finite number of values.

When transitions can be expressed as guarded commands (i.e., the system is a set of deterministic actions), symbolic model checking is optimized using techniques for computing predecessor states without computing the entire transition relation. A specialized backwards search for proving invariants is also available. The set of states visited in the backwards search is constrained by auxiliary invariants, which may have been formulated and verified before, or generated automatically (see Section 5.1).

The symbolic and explicit-state model checkers complement each other. Although limited to finite-state systems, the symbolic model checker can be considerably more efficient, particularly when the state-space is large and the transition relation and fixed points are amenable to representation by OBDD's [McM93]. On the other hand, the explicit-state model checker is often faster on systems with relatively few reachable states.

4.4 Modular Verification

Different components of a large system may require the application of different verification methodologies, depending on their specific type (real-time or discrete, finite- or infinite-state). Using the notion of modular validity, modular properties can be established by the same set of methods as global properties, accounting for environment transitions. Automatic property in-

heritance then ensures that such properties can be used as lemmas in proofs over composite modules. In the case of recursively defined systems, properties can be established by structural induction.

Many properties are not directly guaranteed by a module, but hold only under certain assumptions. STeP's proof management allows assumptions to be used before their proof is available, checking the resulting dependency diagram to avoid unsound circular reasoning. Assumptions about the environment can be made when proving a modular property, and subsequently discharged when the module is composed with another. The search for appropriate assumptions can be guided by constructing verification diagrams for each module and attempting to prove the associated verification conditions [FMS98, MCF⁺98].

5 Combining Deductive and Algorithmic Methods

STeP includes formalisms that combine deductive and algorithmic verification in a number of different ways, which differ in the degree and type of intervention that is required from the user.

Besides model checking, described in Section 4.3, STeP provides two basic automatic tools that support deductive and deductive-algorithmic verification: automatic invariant generation and decision procedures. Both are used extensively in the combinations of deductive and algorithmic verification presented in Sections 5.3 and 5.4.

5.1 Automatic Invariant Generation

Deductive verification is usually an incremental process: simple properties of the system being verified are proved first and then used to help establish more complex ones. STeP implements techniques for the *automatic generation of invariants*, as described in [BBM97]. Invariant generation is based on approximate propagation, starting from the set of initial states, through the state-space of the system until a fixpoint is reached, based on the framework of abstract interpretation [CC77]. Depending on the approximation method used, different types of invariants can be generated:

- *Local invariants* result from analyzing the possible values of individual variables, as well as the relation between control locations and data values.
- *Linear invariants* express linear relationships between system variables.

- *Polyhedral invariants* generalize linear invariants, expressing polyhedral constraints over sets of system variables.

These invariant generation methods are now being specialized to the case of real-time and hybrid systems [MS98]. Invariant generation can also be used for the modular verification of real-time systems, as shown in [BMSU98].

For real-time and hybrid systems STeP provides an alternative technique of invariant generation, also based on forward propagation of system behavior through the state space, but now starting from the entire state space [BMSU98]. In this case every propagation step leads to an invariant; no fixpoint needs to be computed. For hybrid systems these techniques have been further optimized to take advantage of the structure of the constraints, resulting in stronger invariants. In [MS98] we show an example where the invariants thus generated are sufficiently strong to prove the main property of interest.

5.2 Decision Procedures

The verification conditions generated in deductive verification refer to the domain of computation of the system being verified. To establish verification conditions in the most automatic and efficient manner, STeP includes *decision procedures* for a number of theories frequently used in computation domains, and thus common in formal verification [Bjø98].

The basic integration of decision procedures is a variant of Shostak's congruence closure-based algorithm [Sho84, CLS96, Bjø98]. At the top-level, an algorithm based on congruence closure propagates equality constraints through function symbols. It invokes the other decision procedures as auxiliary simplifiers and solvers. The theories supported in this way include:

- *Partial orders*. Beyond basic equality, partial orders are a more expressive constraint language to specify relations between variables.
- *Linear and non-linear arithmetic*. STeP provides Fourier's quantifier elimination procedure to deal with formulas involving linear arithmetic; this procedure also extracts implied equalities.

Some systems, including most hybrid systems, require reasoning about formulas featuring multiplication and division. Formulas involving some use of multiplication and division however arise naturally from

even simple hybrid systems. For instance

$$\left(\begin{array}{l} a/v_1 + a \cdot 2/v_r < x_2 + i/v_1 + i \cdot 2/v_r \\ \wedge \quad v_1 > 0 \wedge v_r > 0 \wedge p \geq i \wedge p \leq a \end{array} \right) \rightarrow p/v_r \leq x_2 + i/v_r$$

was encountered during verification of a symbolic temperature controller [MS98]. STeP therefore includes basic techniques for eliminating first- and second-degree variables as described in [Wei97]. Aggressive use of abstract interpretation techniques based on sign attributes keeps the search space manageable. For instance in the example above, the information $v_r > 0$ is used to simplify $p/v_r \leq x_2 + i/v_r$ to $p \leq x_2 \cdot v_r + i$. Asserted and implied equalities are used to eliminate variables by rewriting equalities into the form $x = t$, where x does not occur in t , and applying the substitution $[x \mapsto t]$ to the constraints. The equality $x = t$ is also passed to the congruence closure which uses it for simplifying complex (non-arithmetical) terms containing x .

- *Bit-vectors.* Reasoning about bit-vectors is essential for hardware verification. STeP includes decision procedures for fixed-size bit-vectors with boolean bitwise operations and concatenation, and for non-fixed size bit-vectors with concatenation. These procedures are described in [BP98].
- *Lists, queues, and word decision procedures.* Lists and queues are common data structures, especially in systems using abstract datatypes or asynchronous channels. Both lists and queues can be viewed as special cases of words, with concatenation being the basic operation. Although the known decision procedures for word equalities have prohibitive complexity, the special cases of lists and queues can be solved efficiently.
- *Recursive data-types.* STeP supports equality reasoning for general recursive datatypes, which allow the specification of S-expressions and other tree-like structures. Enumeration types and records are treated as special cases of recursive datatypes.

Co-inductive data-types, such as lazy lists, are also supported. Both equality constraints and subterm relationships are supported in the integration of decision procedures.

Recursive data-types are given initial term algebra semantics. On the other hand, co-recursive data-types contain infinite and cyclic terms.

- *Set theory.* STeP provides basic support for Multi-level Syllogistic Set-theory (MLSS) [CFO89, CZ98]. MLSS terms range over sets, and operations include union, intersection, set-difference, and finite set-enumeration. Atomic relations include set equality, inclusion and membership.

STeP uses decision procedures not only to check validity, but to *simplify* formulas as well, rewriting them to smaller, logically equivalent ones. Efficient formula simplification can make verification conditions more readable and manageable, and improves the efficiency of subsequent validity checking.

The above decision procedures check validity of *ground* formulas, where no first-order quantification is present. STeP extends this combination of ground decision procedures to include theory-specific unification algorithms, which find quantifier instantiations needed for first-order validity checking [BSU97].

As mentioned in Section 4.2, an interactive Gentzen-style theorem prover is available as part of the Proof Editor (Figure 4) to establish verification conditions that are not proved automatically.

5.3 Generalized Verification Diagrams

Generalized verification diagrams [BMS95, MBSU98] are an extension of verification diagrams that allow the verification of arbitrary temporal properties. Like temporal formulas and ω -automata, generalized verification diagrams have an associated set of computations, constrained by an acceptance condition. Diagrams can be seen as intermediaries between the system and the property to be proven. A set of verification conditions is proved, deductively, to show that the diagram faithfully represents computations of the system: initiality and consecution conditions, associated with individual nodes of the diagram, ensure that all runs of the system have corresponding paths of the diagram. Acceptance conditions, based on well-founded orders, prove that system computations fulfill the acceptance requirements associated with the diagram. An algorithmic check then establishes that the diagram corresponds to the formula being proved. Together, these two stages show that all computations of the system are models of the temporal property.

The STeP Diagram Editor, shown in Figure 5, allows the user to draw a diagram and then prove, using the Proof Editor, the associated verification conditions. In STeP 2.0, the Diagram Editor and the Proof Editor are more tightly coupled, to facilitate the incremental development of diagrams. The

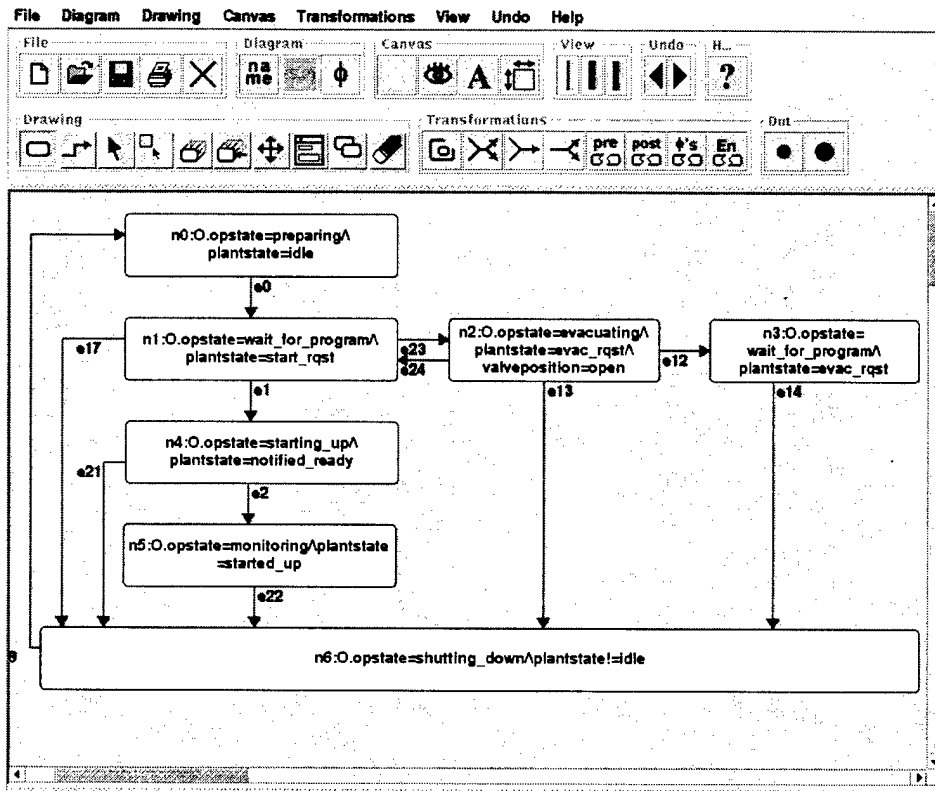


Figure 5: STeP Diagram Editor

user can draw an initial version and try to prove the associated verification conditions. If they fail, the user can make local corrections to the diagram (or discover something wrong with the system) and attempt the proof again.

The verification conditions are *local* to the diagram; failed verification conditions point to missing edges or nodes, weak assertions, or possible bugs in the system. Since local changes to a diagram do not affect the verification conditions elsewhere, much of the work from the previous iteration can be saved. Using feedback from the Proof Editor, the Diagram Editor can highlight proved and unproved edges and nodes in the diagram, helping the user correct the diagram. A change to the diagram automatically invalidates the verification conditions in the Proof Editor that are affected by the change.

Deductive model checking [SUM99] uses diagrams to explore and refine the state-space of possibly infinite-state systems, searching for a counterexample computation by transforming the diagram. The STeP Diagram Editor supports some of these diagram transformations for interactive state-space exploration. We will include a more comprehensive implementation in upcoming releases.

5.4 Constructing Finite-State Abstractions

Temporal properties can be proved for a complex system by finding a simpler *abstract system* such that if the abstract system satisfies a related property, then the original *concrete system* satisfies the original one as well. If the abstract system is finite-state, its temporal properties can be established automatically using a model checker. We have developed methods for automatically generating finite-state abstractions of possibly infinite-state systems [CU98, Uri98], using the decision procedures in STeP (Section 5.2).

The abstraction algorithm compositionally abstracts the transitions of the system, expressed as first-order relations, relative to a given, fixed set of assertions which define the abstract state-space. The number of validity checks is proportional to the size of the system description, rather than the size of the abstract state-space.

Once the finite-state abstraction is generated, it can be model checked, explicitly or symbolically (see Section 4.3). The generated abstractions are *weakly preserving* for universal (\forall CTL*) temporal properties, including LTL [DGG97]. This means that validity at the abstract level implies the validity of the original property over the concrete system; however, if the abstract property fails, the original property might still hold. In this case, we say that the abstraction was not fine enough. An abstract counterexample can be used, manually, to determine if a corresponding concrete counterexample

exists, or else to build a finer abstraction.

6 Applications

6.1 Real-time Systems

In [BMSU98], we present a modular framework for proving temporal properties of real-time systems, based on clocked transition systems and linear-time temporal logic. We show how deductive verification rules, verification diagrams, and automatic invariant generation can be used to establish properties of real-time systems in this framework. We also discuss global and modular proofs of the branching-time property of non-Zenoness. As an example, we present the mechanical verification of the *generalized railroad crossing* case study using STeP.

6.2 Fault-tolerant Systems

In [BLM97], a parameterized fault-tolerant leader-election algorithm recently proposed in [GM96] is modeled and verified using STeP. Our methods settle the general N -process correctness for the algorithm, which had been previously verified only for $N = 3$. We formulate the notion of *Uniform Compassion* to model progress in faulty systems more faithfully, and combine it with the more standard notions of fairness. We also show how the correctness proofs generalize to different channel models by a reduction to a simple channel model.

6.3 Hybrid Systems

In [MS98] we present invariant generation methods for hybrid systems, and verify a simple hybrid water level controller using STeP. In [MCF⁺98], we show how deductive verification tools, and the combination of finite-state model checking and abstraction, allow the verification of infinite-state systems featuring data types commonly used in software specifications, including real-time and hybrid systems.

6.4 Case Study: Steam boiler

The incorporation of modularity and abstraction in STeP has enabled us to analyze much larger systems than was previously possible. An example is the *steam boiler* case study [ABL96], a benchmark for specification and verification methods for hybrid controlled systems. At the time of its

appearance we developed a comprehensive model of this system, including both the plant and the controller. The model consisted of some 1000 lines of SPL code and contained eight parallel processes. However, verification proved impractical and further analysis was suspended. Recently the case study was revived. The system was rewritten as a modular transition system consisting of ten modules with a total of 80 transitions, 18 real-valued variables and 28 finite-domain variables.

Modularity allowed us to prove properties over selected subsystems and inherit them for the full system, thus reducing the number of verification conditions to be proven. In some cases, involving discrete finite-state modules only, the model checker could be applied, making the verification fully automatic. In our previous implementation finite-state components could not be separated from the infinite-state ones, and thus use of the model checker was not possible.

Assertion-based abstraction (see Section 5.4) enabled us to indirectly apply the model checker to infinite-state modules as well, by eliminating the real-valued variables. The relationships between the relevant real-valued variables captured by a small set of assertions were sufficient to let us prove the properties.

A more detailed description of the case study is given in [MCF⁺98].

6.5 Educational Use

STeP has been used on several occasions for teaching a graduate-level introductory course on temporal verification of reactive systems, at Stanford University and at the Technion (Israel Institute of Technology, Haifa).

STeP is freely available for research and educational use. The STeP manual is available as [BBC⁺95], and system descriptions are provided in [BBC⁺96, MBB⁺99]. For information on obtaining STeP, see the STeP web pages at:

<http://www-step.stanford.edu/>

7 Implementation

The parsing, theorem-proving and invariant generation components of STeP are implemented in Standard ML of New Jersey. The graphical user interface for STeP 2.0 was developed in Java. The Java graphical packages and the inheritance features of the language are well-suited for the implementation

of a variety of visual formalisms with common features, as in the STeP Diagram Editor.

The explicit-state model checker is implemented in C, while the symbolic model checker uses ML linked together with external OBDD libraries, written in C. Similarly, the polyhedral invariant generation uses external polyhedra manipulation routines, implemented in C [HP95].

7.1 Stand-alone components

Many of the components of STeP described above can be used in batch mode, as stand-alone components, including:

- the parser/translator from SPL into fair transition systems (Section 2)
- the explicit-state and symbolic model checkers (Section 4.3)
- the linear, local and polyhedral invariant generators (Section 5.1)
- the validity checker and formula simplifier (Section 5.2)

These options are available as command-line options to a separate executable binary file.

7.2 External systems

Verification conditions in STeP can be output to external theorem provers or decision procedures. In particular, the MONA package for monadic second order logic can be used; output to the OTTER and Gandalf resolution theorem provers is also provided. STeP interacts with these provers by invoking them in the background and digesting their output to check if the verification conditions have been discharged.

7.3 Obtaining STeP

STeP is freely available for research and educational use. It has been downloaded by more than 180 registered users in 28 countries around the world. For more information, the STeP web pages at

<http://www-step.stanford.edu>

8 Publications

The following publications were supported (partially or in whole) by this grant:

1. N.S. Bjørner, A. Browne, E.S. Chang, M. Colón, A. Kapur, Z. Manna, H.B. Sipma, and T.E. Uribe. STeP: The Stanford Temporal Prover, User's Manual. Technical Report STAN-CS-TR-95-1562, Computer Science Department, Stanford University, November 1995.
2. N.S. Bjørner, A. Browne, E.S. Chang, M. Colón, A. Kapur, Z. Manna, H.B. Sipma, and T.E. Uribe. STeP: Deductive-algorithmic verification of reactive and real-time systems. In R. Alur and T.A. Henzinger, editors, *Proc. 8th Intl. Conference on Computer Aided Verification*, vol. 1102 of *LNCS*, pages 415–418. Springer-Verlag, July 1996.
3. N.S. Bjørner, A. Browne, and Z. Manna. Automatic generation of invariants and intermediate assertions. *Theoretical Computer Science*, 173(1):49–87, February 1997. Preliminary version appeared in 1st *Intl. Conf. on Principles and Practice of Constraint Programming*, vol. 976 of *LNCS*, pp. 589–623, Springer-Verlag, 1995.
4. N.S. Bjørner. *Integrating Decision Procedures for Temporal Verification*. PhD thesis, Computer Science Department, Stanford University, November 1998.
5. N.S. Bjørner, U. Lerner, and Z. Manna. Deductive verification of parameterized fault-tolerant systems: A case study. In *Intl. Conf. on Temporal Logic*. Kluwer, 1997. To appear.
6. A. Browne, Z. Manna, and H.B. Sipma. Generalized temporal verification diagrams. In *15th Conference on the Foundations of Software Technology and Theoretical Computer Science*, vol. 1026 of *LNCS*, pages 484–498. Springer-Verlag, 1995.
7. N.S. Bjørner, Z. Manna, H.B. Sipma, and T.E. Uribe. Deductive verification of real-time systems using STeP. Technical Report STAN-CS-TR-98-1616, Computer Science Department, Stanford University, January 1998. To appear in *Theoretical Computer Science*. Preliminary version appeared in *4th Intl. AMAST Workshop on Real-Time Systems*, vol. 1231 of *LNCS*, pages 22–43. Springer-Verlag, May 1997.

8. N.S. Bjørner and M.C. Pichora. Deciding fixed and non-fixed size bit-vectors. In *4th Intl. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, vol. 1384 of *LNCS*, pages 376–392. Springer-Verlag, 1998.
9. N.S. Bjørner, M.E. Stickel, and T.E. Uribe. A practical integration of first-order reasoning and decision procedures. In *Proc. of the 14th Intl. Conference on Automated Deduction*, vol. 1249 of *LNCS*, pages 101–115. Springer-Verlag, July 1997.
10. M.A. Colón and T.E. Uribe. Generating finite-state abstractions of reactive systems using decision procedures. In A.J. Hu and M.Y. Vardi, editors, *Proc. 10th Intl. Conference on Computer Aided Verification*, vol. 1427 of *LNCS*, pages 293–304. Springer-Verlag, July 1998.
11. B. Finkbeiner, Z. Manna, and H.B. Sipma. Deductive verification of modular systems. In W.P. de Roever, H. Langmaack, and A. Pnueli, editors, *Compositionality: The Significant Difference, COMPOS'97*, vol. 1536 of *LNCS*, pages 239–275. Springer-Verlag, December 1998.
12. Y. Kesten, Z. Manna, and A. Pnueli. Verifying clocked transition systems. In R. Alur, T.A. Henzinger, and E.D. Sontag, editors, *Hybrid Systems III*, vol. 1066 of *LNCS*, pages 13–40. Springer-Verlag, 1996.
13. Z. Manna, N.S. Bjørner, A. Browne, M. Colón, B. Finkbeiner, M. Pichora, H.B. Sipma, and T.E. Uribe. An update on STeP: Deductive-algorithmic verification of reactive systems. In R. Berghammer and Y. Lakhnech, editors, *Tool Support for System Specification, Development and Verification*, Advances in Computing Science, pages 174–188. Springer-Verlag, 1999.
14. Z. Manna, A. Browne, H.B. Sipma, and T.E. Uribe. Visual abstractions for temporal verification. In A. Haeberer, editor, *Algebraic Methodology and Software Technology (AMAST'98)*, vol. 1548 of *LNCS*, pages 28–41. Springer-Verlag, December 1998.
15. Z. Manna, M.A. Colón, B. Finkbeiner, H.B. Sipma, and T.E. Uribe. Abstraction and modular verification of infinite-state reactive systems. In M. Broy, editor, *Requirements Targeting Software and Systems Engineering (RTSE)*, LNCS. Springer-Verlag, 1998. To appear.

16. Z. Manna and A. Pnueli. Clocked transition systems. Technical Report STAN-CS-TR-96-1566, Computer Science Department, Stanford University, April 1996.
17. Z. Manna and H.B. Sipma. Deductive verification of hybrid systems using STeP. In T. Henzinger and S. Sastry, editors, *Hybrid Systems: Computation and Control*, vol. 1386 of *LNCS*, pages 305–318. Springer-Verlag, April 1998.
18. H.B. Sipma, T.E. Uribe, and Z. Manna. Deductive model checking. To appear in *Formal Methods in System Design*, 1999. Preliminary version appeared in *Proc. 8th Intl. Conference on Computer Aided Verification*, vol. 1102 of *LNCS*, Springer-Verlag, pp. 208–219, 1996.
19. T.E. Uribe. *Abstraction-based Deductive-Algorithmic Verification of Reactive Systems*. PhD thesis, Computer Science Department, Stanford University, December 1998. Technical Report STAN-CS-TR-99-1618.

9 Participating Scientific Personnel

(Partially supported by this grant.)

Graduated Master's Students:

1. Mark Pichora (1998)

Graduated Ph.D. Students:

1. Arjun Kapur (1997). Thesis: *Interval and Point-based Approaches to Hybrid System Verification*.
2. Luca de Alfaro (1997). Thesis: *Formal Verification of Probabilistic Systems*.
3. Nikolaj Bjørner (1998). Thesis: *Integrating Decision Procedures for Temporal Verification*.
4. Tomás E. Uribe (1998). Thesis: *Abstraction-Based Deductive-Algorithmic Verification of Reactive Systems*.
5. Henny B. Sipma (1999). Thesis: *Diagram-based Verification of Discrete, Real-time and Hybrid Systems*.

Other Personnel:

1. P.I.: Prof. Zohar Manna
2. Scientific Programmer: Anca Browne
3. Ph.D. Student: Michael Colón
4. Ph.D. Student: Bernd Finkbeiner

10 Report of Inventions

A negative Final Report of Inventions and Subcontracts (DD Form 882) is enclosed.

11 Bibliography

References

- [ABL96] J.R. Abrial, E. Boerger, and H. Langmaack, editors. *Formal Methods for Industrial Applications: Specifying and Programming the Steam Boiler Control*, vol. 1165 of LNCS. Springer-Verlag, 1996.
- [BBC⁺95] N.S. Bjørner, A. Browne, E.S. Chang, M. Colón, A. Kapur, Z. Manna, H.B. Sipma, and T.E. Uribe. STeP: The Stanford Temporal Prover, User's Manual. Technical Report STAN-CS-TR-95-1562, Computer Science Department, Stanford University, November 1995.
- [BBC⁺96] N.S. Bjørner, A. Browne, E.S. Chang, M. Colón, A. Kapur, Z. Manna, H.B. Sipma, and T.E. Uribe. STeP: Deductive-algorithmic verification of reactive and real-time systems. In R. Alur and T.A. Henzinger, editors, *Proc. 8th Intl. Conference on Computer Aided Verification*, vol. 1102 of LNCS, pages 415–418. Springer-Verlag, July 1996.
- [BBM97] N.S. Bjørner, A. Browne, and Z. Manna. Automatic generation of invariants and intermediate assertions. *Theoretical Computer Science*, 173(1):49–87, February 1997. Preliminary version appeared in 1st Intl. Conf. on Principles and Practice of Constraint Programming, vol. 976 of LNCS, pp. 589–623, Springer-Verlag, 1995.
- [Bjø98] N.S. Bjørner. *Integrating Decision Procedures for Temporal Verification*. PhD thesis, Computer Science Department, Stanford University, November 1998.
- [BLM97] N.S. Bjørner, U. Lerner, and Z. Manna. Deductive verification of parameterized fault-tolerant systems: A case study. In *Intl. Conf. on Temporal Logic*. Kluwer, 1997. To appear.
- [BMS95] A. Browne, Z. Manna, and H.B. Sipma. Generalized temporal verification diagrams. In *15th Conference on the Foundations of Software Technology and Theoretical Computer Science*, vol. 1026 of LNCS, pages 484–498. Springer-Verlag, 1995.

- [BMSU98] N.S. Bjørner, Z. Manna, H.B. Sipma, and T.E. Uribe. Deductive verification of real-time systems using STeP. Technical Report STAN-CS-TR-98-1616, Computer Science Department, Stanford University, January 1998. To appear in *Theoretical Computer Science*. Preliminary version appeared in *4th Intl. AMAST Workshop on Real-Time Systems*, vol. 1231 of *LNCS*, pages 22–43. Springer-Verlag, May 1997.
- [BP98] N.S. Bjørner and M.C. Pichora. Deciding fixed and non-fixed size bit-vectors. In *4th Intl. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, vol. 1384 of *LNCS*, pages 376–392. Springer-Verlag, 1998.
- [BSU97] N.S. Bjørner, M.E. Stickel, and T.E. Uribe. A practical integration of first-order reasoning and decision procedures. In *Proc. of the 14th Intl. Conference on Automated Deduction*, vol. 1249 of *LNCS*, pages 101–115. Springer-Verlag, July 1997.
- [CC77] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *4th ACM Symp. Princ. of Prog. Lang.*, pages 238–252. ACM Press, 1977.
- [CFO89] D. Cantone, A. Ferro, and E. Omodeo. *Computable Set Theory*. Oxford Science Publications, 1989.
- [CLS96] D. Cyrluk, P. Lincoln, and N. Shankar. On Shostak's decision procedure for combinations of theories. In *Proc. of the 13th Intl. Conference on Automated Deduction*, vol. 1104 of *LNCS*, pages 463–477. Springer-Verlag, 1996.
- [CU98] M.A. Colón and T.E. Uribe. Generating finite-state abstractions of reactive systems using decision procedures. In A.J. Hu and M.Y. Vardi, editors, *Proc. 10th Intl. Conference on Computer Aided Verification*, vol. 1427 of *LNCS*, pages 293–304. Springer-Verlag, July 1998.
- [CZ98] D. Cantone and C.G. Zarba. A new fast tableau-based decision procedure for an unquantified fragment of set theory. In *Int. Workshop on First-Order Theorem Proving (FTP'98)*, 1998.

- [DGG97] D.R. Dams, R. Gerth, and O. Grumberg. Abstract interpretation of reactive systems. *ACM Transactions on Programming Languages and Systems*, 19(2):253–291, 1997.
- [FMS98] B. Finkbeiner, Z. Manna, and H.B. Sipma. Deductive verification of modular systems. In W.P. de Roever, H. Langmaack, and A. Pnueli, editors, *Compositionality: The Significant Difference, COMPOS'97*, vol. 1536 of *LNCS*, pages 239–275. Springer-Verlag, December 1998.
- [GM96] H. Garavel and L. Mounier. Specification and verification of various distributed leader election algorithms for unidirectional ring networks. Rapport de recherche 2986, INRIA, Rhone-Alpes, France, September 1996.
- [HP95] N. Halbwachs and Y.E. Proy. *POLyhedra desK cAlculator (POLKA)*. VERIMAG, Montbonnot, France, September 1995.
- [KMP96] Y. Kesten, Z. Manna, and A. Pnueli. Verifying clocked transition systems. In R. Alur, T.A. Henzinger, and E.D. Sontag, editors, *Hybrid Systems III*, vol. 1066 of *LNCS*, pages 13–40. Springer-Verlag, 1996.
- [MBB⁺99] Z. Manna, N.S. Bjørner, A. Browne, M. Colón, B. Finkbeiner, M. Pichora, H.B. Sipma, and T.E. Uribe. An update on STeP: Deductive-algorithmic verification of reactive systems. In R. Berghammer and Y. Lakhnech, editors, *Tool Support for System Specification, Development and Verification*, Advances in Computing Science, pages 174–188. Springer-Verlag, 1999.
- [MBSU98] Z. Manna, A. Browne, H.B. Sipma, and T.E. Uribe. Visual abstractions for temporal verification. In A. Haeberer, editor, *Algebraic Methodology and Software Technology (AMAST'98)*, vol. 1548 of *LNCS*, pages 28–41. Springer-Verlag, December 1998.
- [MCF⁺98] Z. Manna, M.A. Colón, B. Finkbeiner, H.B. Sipma, and T.E. Uribe. Abstraction and modular verification of infinite-state reactive systems. In M. Broy, editor, *Requirements Targeting Software and Systems Engineering (RTSE)*, LNCS. Springer-Verlag, 1998. To appear.
- [McM93] K.L. McMillan. *Symbolic Model Checking*. Kluwer Academic Pub., 1993.

- [MP94] Z. Manna and A. Pnueli. Temporal verification diagrams. In M. Hagiya and J.C. Mitchell, editors, *Proc. International Symposium on Theoretical Aspects of Computer Software*, vol. 789 of *LNCS*, pages 726–765. Springer-Verlag, 1994.
- [MP95] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer-Verlag, New York, 1995.
- [MP96] Z. Manna and A. Pnueli. Clocked transition systems. Technical Report STAN-CS-TR-96-1566, Computer Science Department, Stanford University, April 1996.
- [MS98] Z. Manna and H.B. Sipma. Deductive verification of hybrid systems using STeP. In T. Henzinger and S. Sastry, editors, *Hybrid Systems: Computation and Control*, vol. 1386 of *LNCS*, pages 305–318. Springer-Verlag, April 1998.
- [Sho84] R.E. Shostak. Deciding combinations of theories. *J. ACM*, 31(1):1–12, January 1984.
- [SUM99] H.B. Sipma, T.E. Uribe, and Z. Manna. Deductive model checking. To appear in *Formal Methods in System Design*, 1999. Preliminary version appeared in *Proc. 8th Intl. Conference on Computer Aided Verification*, vol. 1102 of *LNCS*, Springer-Verlag, pp. 208–219, 1996.
- [Uri98] T.E. Uribe. *Abstraction-based Deductive-Algorithmic Verification of Reactive Systems*. PhD thesis, Computer Science Department, Stanford University, December 1998. Technical Report STAN-CS-TR-99-1618.
- [Wei97] V. Weispfenning. Quantifier elimination for real algebra—the quadratic case and beyond. In *Applied Algebra and Error-Correcting Codes (AAECC) 8*, pages 85–101, 1997.