InVeSt : A Tool for the Verification of Invariants*

S. Bensalem† Y. Lakhnech‡ S. Owre§

1 Methodology

There are basically two approaches to the verification of reactive systems, the *algorithmic* approach on one hand and the *deductive* approach on the other hand. The algorithmic approach is based on the computation of fix-points, on effective representations of sets of states, and on decision procedures to solve the inclusion problem of sets of states. For example the *backward procedure* is an instance of this approach. To prove that a set of states P is an invariant of a system S, the backward procedure computes the largest set Q of states satisfying $Q \subseteq P$ and $Q \subseteq \text{WP}(\tau, Q)$, for every transition $\tau \in \mathcal{T}$ of S. Here $\text{WP}(\tau, Q)$ is the weakest pre-condition of τ with respect to Q. Then, P is an invariant of S if and only if every initial state of S satisfies Q. In general, the algorithmic approach is based on an effective representation $\mathcal{R}$ for sets of states, effective boolean operations, a procedure for deciding inclusion in $\mathcal{R}$, effective predicate transformers to guarantee recursiveness of the method, and convergence of fix-points to guarantee completeness.

In general, in case of infinite state systems, first-order logic with Peano arithmetic is considered as representation $\mathcal{R}$. In fact, it can be proved that any weaker logic is not expressive enough (e.g. [7]), when the considered system contains variables that range over infinite domains. Thus, one has effective boolean operations and can define predicate transformers, but inclusion is undecidable. Moreover, convergence of fix-points is not guaranteed. Consequently, the algorithmic approach cannot be applied in general to infinite state systems. On the other hand, the deductive approach is very powerful and gives a complete method even for infinite state systems. It relies upon finding *auxiliary invariants* and proving validity of first-order formulas, called *verification conditions*. The deductive approach is, however, in contrast to the algorithmic approach, difficult to apply. Indeed, it is in general a hard task to find suitable auxiliary invariants and time consuming to discharge all generated verification conditions. Therefore, there is a strong need for tools that support both tasks. InVeSt is such a tool as it supports the verification of invariance properties of infinite state systems.

The salient feature of InVeSt is that it combines the algorithmic with the deductive approaches to program verification in two different ways:

1. It integrates the principles underlying the algorithmic (e.g. [3, 18]) and the deductive methods (e.g. [14]) in the sense that it uses fix-point calculation as in the algorithmic approach but also the reduction of the invariance problem to a set of first-order formulas as in the deductive approach.

2. It integrates the theorem prover PVS [17] with the model-checker SMV [15] through the automatic computation of finite abstractions. That is, it provides the ability to automatically compute finite abstractions of infinite state systems which are then analyzed by SMV or, alternatively, by the model-checker of PVS.

InVeSt supports the proof of invariance properties using the method based on induction and auxiliary invariants (e.g. [14]) as well as the method based on abstraction techniques [4, 12, 6, 10, 11, 5].

*This work has been partly performed while the first two authors were visiting the Computer Science Laboratory, SRI International. Their visits were funded by NSF Grants No. CCR-9712383 and CCR-9509931.

†VERIMAG, Centre Equation – 2, avenue de Vignate, F-38610 Giéres, France. Bensalem@imag.fr

‡Institut für Informatik und Praktische Mathematik, Christian-Albrechts-Universität zu Kiel, Preußerstr. 1-9, D-24105 Kiel, Germany. yl@informatik.uni-kiel.de

§Computer Science Laboratory, SRI International, Menlo Park, CA 94025, USA. owre@csl.sri.com

InVeSt's approach to finding auxiliary invariants. We use calculation of pre-fix-points by applying the body of the backward procedure a finite number of times and use techniques for the automatic generation of invariants (cf. [14, 13, 2]) to support the search for auxiliary invariants. The tool provides strategies which allow to derive *local invariants*, that is, predicates attached to control locations and which are satisfied whenever the computation reaches the corresponding control point. InVeSt includes strategies for deriving local invariants for sequential systems as well as a composition principle that allows to combine invariants generated for sequential systems to obtain invariants of a composed system. Consider a composed system $S_1 \parallel S_2$ and control locations l_1 and l_2 of S_1 and S_2, respectively. Suppose that we generated the local invariants P_1 and P_2 at l_1 and l_2, respectively. Let us call P_i *interference independent*, if P_i does not contain a free variable that is written by S_j with $j \neq i$. Then, depending on whether P_i is interference independent we compose the local invariants P_1 and P_2 to obtain a local invariant at (l_1, l_2) as follows: if P_i is interference independent, then we can affirm that P_i is an invariant at (l_1, l_2) and if both P_1 and P_2 are interference dependent, then $P_1 \vee P_2$ is an invariant at (l_1, l_2). This composition principle proved to be useful in the examples we considered. However, examples showed that predicates obtained by this composition principle can become very large. Therefore, we also consider the alternative option where local invariants are not composed until they are needed in a verification condition. Thus, we assign to each component of the system two lists of local invariants. The first corresponds to interference independent local invariants and the second to interference dependent ones. Then, when a verification condition is considered, we use heuristics to determine which local invariants are useful when discharging the verification condition. A useful heuristic concerns the case when the verification condition is of the form $(pc(1) = l_1 \wedge pc(2) = l_2) \Rightarrow \phi$, where $pc(1) = l_1 \wedge pc(2) = l_2$ asserts that computation is at the local control locations l_1 and l_2. In this case, we combine the local invariants associated to l_1 and l_2 and add the result to the left-hand side of the implication.

InVeSt's approach to computing abstractions. InVeSt provides also a module that allows to compute an abstract system from a given concrete system and an abstraction function. The method underlying this module is presented in [1]. The main features of this method is that it is automatic and compositional. Moreover, it generates an abstract system which has the same structure as the concrete one. This gives the ability to apply further abstractions and techniques to reduce the state explosion problem and facilitates the debugging of the concrete system. The computed abstract system is optionally represented in the specification language of PVS or in that of SMV. A graphical interface allows to interact with InVeSt and SMV in a uniform way.

Finally, it is important to understand that our use of the theorem-prover PVS is limited to discharging the verification conditions. This shows a difference of our approach to the approach followed in most of the work using theorem-proving for verifying invariance properties (e.g. [8, 9]). That is we do not encode the invariance problem in the specification language of the considered theorem-prover and then use the theorem-prover to solve it, but we use the deductive approach to reduce the problem to a set of first-order formulas whose validity is proved using the theorem-prover. Moreover, the construction of the verification conditions as well as the generation of an auxiliary invariant are performed outside the theorem-prover.

2 Design Principles

The structure of InVeSt is motivated by a number of design decisions. These decisions are :

- **Minimization of user's intervention.** This decision is motivated by our belief that the success of the algorithmic approach is partly due to the fact that it does not require user's intervention. To support this choice we have
  - developed techniques for the generation of auxiliary invariants,
  - implemented the strengthening method and its refined version,
  - investigated strategies for proving first-order predicates, and
  - developed a method for computing abstractions of infinite state systems.

- **Use of an existing theorem-prover.** We build on an existing theorem-prover for the following reasons. The first is that we want to rely on a widely used tool; this increases our trust and confidence in the prover. Our particular choice is to use PVS, since PVS gives us the possibility to combine decision procedures and interactive proofs (see [16]). The examples we have considered show that this feature is a prerequisite to reach a high level of automation.

- **Theorem-prover as a "Decision Procedure".** Most of the theorem-provers, including PVS, are general purpose provers. This means that they have general specification languages and if they include pre-defined strategies, then these are in general not tuned to a particular application. This is often a source of inefficiency and prevents a higher level of automation. There is always a trade off between generality on one side and efficiency and automation on the other side. We use the theorem-prover in a particular way and for a particular task, namely to discharge the verification conditions. This should be seen in contrast to the alternative approach where one encodes the verification problem within the specification language of PVS and tries to solve it completely within this theorem-prover, usually by expanding the definitions of the semantics of the programs and the definition of invariance and using an induction argument. Our approach is different, since we use deductive rules to reduce the invariance problem to a set of first-order formulas whose validity is proved using PVS. This design decision allows us to implement the components of our tool outside PVS.

- **Modularity with respect to the theorem-prover.** Our tool builds on PVS in two different ways. The first obvious point is that it uses PVS to discharge verification conditions as explained above. There is another, less transparent dependency, which lies in the fact that we use the internal representations of PVS of all objects constructed by the PVS type-checker including programs, actions, formulas, expressions, etc.. In order to be modular with respect to PVS, the interface between the components of the tool and PVS itself has to be defined precisely. As interface, we use a module which contains functions that allow to access internal PVS variables. This ensures that even when the data structure used in PVS is modified, the functioning of our tool is still guaranteed as long as the accessor functions maintain their semantics.

3 Tool structure

The main components of InVeSt are: a front-end that translates guarded command-like programs into a PVS-theory, a module of functions for generating invariants, a module of proof strategies, and a module for computing abstractions. In the sequel, we discuss each of these components.

- **Front-end:** As formalism for describing systems we consider a language similar to Unity. Variables are allowed to be of any type of the specification language of PVS. The role of the front-end consists of translating the extended transition system into a PVS-theory that can be type-checked. The components of the system are translated into PVS-constants of type "Program" consisting of a list of actions. An action can be a function or a relation between states. When the PVS-theory corresponding to the system is type-checked a list of LISP-objects corresponding to the declarations in the PVS-theory is constructed. By accessing to the elements of this list, we have at our disposal PVS-representations corresponding to the relevant syntactic objects of the system. Our tool works with these representations. Accessors functions have been implemented that allow to access these representations and their components. For instance, there is a function GET-GUARD that is used to extract the guard of an action, and a function GET-AFFECTED-VARIABLES to extract the list of variables to which a value is assigned by the action.

- **Generation of Invariants:** A central component in the tool is a module consisting of functions that implement strategies used to automatically generate invariants. Basically, we implemented the strategies presented in [2] which allow us to derive *local invariants*, that is, predicates attached to control locations such that these predicates are satisfied whenever the computation reaches the corresponding control point. We have strategies for deriving local invariants for sequential systems and a composition principle that allows to combine invariants generated for sequential systems to obtain an invariant of a composed system.

- **Proof strategies:** A proof strategy determines which verification conditions are constructed and how to handle failed proofs. The choice of a strategy is determined by:

  1 whether auxiliary invariants are generated automatically,

  2 whether in case the proof of a verification condition fails, the strengthening method [14] or its refined version [2] is applied.

  In case auxiliary invariants are generated, verification conditions are weakened by taking the generated predicate as assumption in the left-hand side of the implication. Thus, if φ is the generated invariant, then to prove that P is preserved by a transition τ it suffices to prove $(\varphi \wedge P) \Rightarrow \text{WP}(\tau, P)$. We implemented a function that takes as arguments two predicates φ and ψ, a transition τ, and a PVS-proof strategy *str* and which calls the PVS-prover on the formula $(\varphi \wedge P) \Rightarrow \text{WP}(\tau, \psi)$ with the strategy *str*. The result of this formula is the list of remaining subgoals after applying the strategy on the formula. In case this list is empty, i.e. the proof succeeded, the next transition is considered. Otherwise, the remaining subgoals are considered in the next strengthening step.

- **Computing Abstractions:** The abstraction module implements the method presented in [1] for computing abstractions of infinite state systems. For a given concrete system and a given abstraction function, it computes an abstraction of the concrete system compositionally and automatically. The process of generation of the abstract system does not depend on the assumed semantics of the parallel operator; it works for the synchronous as well as for the asynchronous computation model. The generated abstract system has the same structure as the concrete one and there is a clear correspondence between the transitions of both systems. This does not only allow to apply further abstractions and techniques to mitigate the state explosion problem but also facilitates the debugging of the concrete system. Indeed, traces of the abstract system can be transformed into concrete traces, which are then checked whether they are behaviors of the concrete system. The generated abstract system is represented optionally in the specification language of PVS or of SMV.

4 Concluding remarks

InVeSt is a tool that supports the verification of invariance properties of infinite state systems. This tool builds on the deductive method to generate verification conditions and uses the theorem-prover PVS to discharge them. It also includes a module of functions for generating invariants, implements the backward strengthening method (cf. [14]) and its refined version [2], and a new method that automatically and compositionally computes abstractions for infinite system. The tool has been used for the verification of various mutual exclusion algorithms and protocols. These include the BRP and the Alternating bit protocol, the Bakery and Peterson's mutual exclusion algorithms, the reader writer example, and a simplified version of the futurebus+ cache coherence protocol.

References

[1] S. Bensalem, Y. Lakhnech, and S. Owre. Computing abstractions of infinite state systems automatically and compositionally. Accepted in CAV'98, 1998.

[2] S. Bensalem, Y. Lakhnech, and H. Saidi. Powerful techniques for the automatic generation of invariants. In *CAV'96*, volume 1102 of *LNCS*. Springer-Verlag, 1996.

[3] E.M. Clarke, E.A. Emerson, and E. Sistla. Automatic verification of finite state concurrent systems using temporal logic specifications: A practical approach. In *10th ACM symp. of Prog. Lang.* ACM Press, 1983.

[4] E.M. Clarke, O. Grumberg, and D.E. Long. Model checking and abstraction. *ACM Transactions on Programming Languages and Systems*, 16(5), 1994.

[5] D. Dams. *Abstract interpretation and partition refinement for model checking.* PhD thesis, Technical University of Eindhoven, 1996.

[6] D. Dams, R. Gerth, and O. Grumberg. Abstract interpretation of reactive systems: Abstractions preserving ACTL*, ECTL* and CTL*. In *PROCOMET*. IFIP Transactions, North-Holland/Elsevier, 1994.

[7] J.W. de Bakker. *Mathematical Theory of Program Cortrectness.* Prentice-Hall, NJ., 1980.

[8] K. Havelund and N. Shankar. Experiments in theorem proving and model checking for protocol verification. In *FME'96*, volume 1051 of *LNCS*. Springer-verlag, 1996.

[9] J. Hooman. Verifying part of the access.bus protocol using PVS. In *Proc. 15th Conference on the Foundations of Software Technology and Theoretical Computer Science*, volume 1026 of *LNCS*. Springer-Verlag,, 1995.

[10] R.P. Kurshan. *Computer-Aided Verification of Coordinating Processes, the automata theoretic approach.* Princeton Series in Computer Science. 1994.

[11] C. Loiseaux, S. Graf, J. Sifakis, A. Bouajjani, and S. Bensalem. Property preserving abstractions for the verification of concurrent systems. *Formal Methods in System Design*, 6(1), 1995.

[12] D. E. Long. *Model Checking, Abstraction, and Compositional Reasoning.* PhD thesis, Carnegie Mellon, 1993.

[13] Z. Manna, A. Anuchitanukul, N. Bjøner, A. Browne, E. Chang, M. Colon, L. de Alfaro, H. Devarajan, H. Sipma, and T. Uribe. STeP : The Stanford Temporal Prover. Technical report, Stanford Univ., Stanford, CA, 1994.

[14] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety.* Springer-Verlag, 1995.

[15] K.L. McMillan. *Symbolic model checking.* Kluwer Academic Publishers, Boston, 1993.

[16] S. Owre, S. Rajan, J.M. Rushby, N. Shankar, and M. Srivas. PVS: Combining specification, proof checking, and model checking. volume 1102 of *LNCS*, pages 411–414. Springer-Verlag, 1996.

[17] S. Owre, J. Rushby, N. Shankar, and F. von Henke. Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering*, 21(2):107–125, Feb. 1995.

[18] J. P. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. In *Proc. 5th Int. Sym. on Programming*, volume 137 of *Lecture Notes in Computer Science*, pages 337–351. Springer-Verlag, 1982.