D-Finder: A Tool for Compositional Deadlock Detection and Verification

Saddek Bensalem, Marius Bozga, Thanh-Hung Nguyen, Joseph Sifakis

► To cite this version:

Saddek Bensalem, Marius Bozga, Thanh-Hung Nguyen, Joseph Sifakis. D-Finder: A Tool for Compositional Deadlock Detection and Verification. Computer Aided Verification, 21st International Conference, CAV 2009, Jun 2009, Grenoble, France. pp.614-619, 10.1007/978-3-642-02658-4_45 . hal-00722550

HAL Id: hal-00722550

<https://hal.science/hal-00722550>

Submitted on 2 Aug 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

D-Finder: A Tool for Compositional Deadlock Detection and Verification

Saddek Bensalem Marius Bozga Thanh-Hung Nguyen Joseph Sifakis

Verimag Laboratory, Université Joseph Fourier Grenoble, CNRS.

Abstract. D-Finder tool implements a compositional method for the verification of component-based systems described in BIP language encompassing multi-party interaction. For deadlock detection, D-Finder applies proof strategies to eliminate potential deadlocks by computing increasingly stronger invariants.

1 Methodology

Compositional verification techniques are used to cope with state explosion in concurrent systems. The idea is to apply divide-and-conquer approaches to infer global properties of complex systems from properties of their components. Separate verification of components limits state explosion. Nonetheless, components mutually interact in a system and their behavior and properties are inter-related. This is a major difficulty in designing compositional techniques. As explained in [1], compositional rules are in general of the form

$$\frac{B_1 < \Phi_1 >, B_2 < \Phi_2 >, C(\Phi_1, \Phi_2, \Phi)}{B_1 \parallel B_2 < \Phi >} \quad (1)$$

That is, if two components with behaviors B_1, B_2 meet individually properties Φ_1, Φ_2 respectively, and $C(\Phi_1, \Phi_2, \Phi)$ is some condition taking into account the semantics of parallel composition operation and relating the individual properties with the global property, then the system $B_1 \parallel B_2$ resulting from the composition of B_1 and B_2 will satisfy a global property Φ .

- **Compositional verification by *assume-guarantee*** . In this approach properties are decomposed into two parts. One is an assumption about the global behavior of the environment of the component; the other is a property guaranteed by the component when the assumption about its environment holds. This approach has been extensively studied (see for example [2–9]). Many issues make the application of assume-guarantee rules difficult. These are discussed in detail in a recent paper [10] which provides an evaluation of automated assume-guarantee techniques. The main difficulties are finding decompositions into sub-systems and choosing adequate assumptions for a particular decomposition.

- **D-Finder's approach to compositional verification** . We present a different approach for compositional verification of invariants based on the following rule:

$$\frac{\{B_i < \Phi_i >\}_i, \Psi \in II(\|\gamma\{B_i\}_i, \{\Phi_i\}_i), (\bigwedge_i \Phi_i) \wedge \Psi \Rightarrow \Phi}{\|\gamma\{B_i\}_i < \Phi >} \quad (2)$$

The rule allows to prove invariance of Φ for systems obtained by using a n-ary composition operation parameterized by a set of interactions γ . It uses global invariants which are the conjunction of individual invariants of components Φ_i and an interaction invariant Ψ . The latter expresses constraints on the global state space induced by interactions between components. It can be computed automatically from abstractions of the system to be verified. These are the composition of finite state abstractions B_i^α of the components B_i with respect to their invariants Φ_i . They can be represented as a Petri net whose transitions correspond to interactions between components. Interaction invariants correspond to traps [11] of the Petri net and are computed symbolically as solutions of a set of boolean equations.

Figure 1 illustrates the method for a system with two components, invariants Φ_1 and Φ_2 and interaction invariant Ψ . Our method differs from assumption-guarantee methods in that it avoids combinatorial explosion of the decomposition and is directly applicable to systems with multiparty (not only binary) interactions. Furthermore, it needs only guarantees for components. It replaces the search for adequate assumptions for each component by the use of interaction invariants. These can be computed automatically from given component invariants (guarantees). Interaction invariants correspond to a “*cooperation test*” in the terminology of [12] as they allow to eliminate product states which are not feasible by the semantics.

Fig. 1.

1.1 Checking Deadlock-freedom and Invariance Properties

D-Finder provides a method for automated verification of component-based systems described in BIP (Behavior-Interaction-Priority) language [13]. In BIP, a system is the composition of a set of atomic components which are automata extended with data and functions written in C. To prove a global invariant Φ for a system $\gamma(B_1, \dots, B_n)$, obtained by composing a set of atomic components $B_1, \dots, B_n$ by using a set of interactions γ , we use the rule (2) above, where $B_i < \Phi_i >$ means that Φ_i is an invariant of component B_i and Ψ is an interaction invariant of $\gamma(B_1, \dots, B_n)$ computed automatically from Φ_i and $\gamma(B_1, \dots, B_n)$. A key issue in the application of this rule is finding component invariants Φ_i . If the components B_i are finite state, then we can take $\Phi = Reach(B_i)$, the set of

reachable state of B_i , or any upper approximation of $Reach(B_i)$. If the components are infinite state, $Reach(B_i)$ is approximated using techniques presented in [14, 15].

- **Checking Invariance Properties.** We give a sketch of a semi-algorithm allowing to prove invariance of Φ by iterative application of the rule (2). The semi-algorithm takes a system $\langle \gamma(B_1, \dots, B_n), Init \rangle$ and a predicate Φ . It iteratively computes invariants of the form $\mathcal{X} = \Psi \wedge (\bigwedge_{i=1}^n \Phi_i)$ where Ψ is an interaction invariant and Φ_i an invariant of component B_i . If $\mathcal{X}$ is not strong enough for proving that Φ is an invariant ($\mathcal{X} \wedge \neg\Phi = false$) then either a new iteration with stronger Φ_i is started or we stop. In this case, we cannot conclude about invariance of Φ .
- **Checking Deadlock-Freedom.** Checking global deadlock-freedom of a system $\gamma(B_1, \dots, B_n)$ is a particular case of proving invariants - proving invariance of the predicate $\neg DIS$, where DIS is the set of the states of $\gamma(B_1, \dots, B_n)$ from which all interactions are disabled.

1.2 Generating Component Invariants and Interaction Invariants

D-Finder provides methods for computing component invariants, particularly useful for checking deadlock-freedom. It also provides a general method for computing interaction invariants for $\gamma(B_1, \dots, B_n)$ from a given set of component invariants Φ_i .

- **Computing Component Invariants.** Invariants for atomic components are generated by simple forward analysis of their behavior. A key issue is efficient computation of such invariants as the precise symbolic computation of reachable states requires quantifier elimination. An alternative to quantifier elimination is to compute over-approximations based on syntactic analysis of the predicates occurring in guards and actions. In this case, the obtained invariants may not be inductive. D-Finder uses different strategies which allow to derive local assertions, that is, predicates attached to control locations and which are satisfied whenever the computation reaches the corresponding control location. A more detailed presentation, as well as the techniques implemented in D-Finder for generating component invariants are given in [16, 17].
- **Computing Interaction Invariants.** Interaction invariants express global synchronization constraints between atomic components. Their computation consists of the following steps. 1) For given component invariants Φ_i of the atomic components B_i , we compute a finite-state abstractions $B_i^{\alpha_i}$ of B_i where α_i is the abstraction induced by the elementary predicates occurring in Φ_i . This step is necessary only for components B_i which are infinite state. 2) The system $\gamma(B_1^{\alpha_1}, \dots, B_n^{\alpha_n})$ which is an abstraction of $\gamma(B_1, \dots, B_n)$, can be considered as a safe Petri net. The set of the traps of the Petri net defines a global invariant which we compute symbolically. 3) The concretization of this invariant gives an interaction invariant of the initial system.

2 Tool Structure

D-Finder consists of a set of modules interconnected as shown in Figure 2.

It takes as input a BIP program and progressively find and eliminate potential deadlocks. It basically works as follows. First, it constructs the predicate characterizing the set of deadlock states (DIS generation module). Second, iteratively, it constructs increasingly stronger local invariants of components (Φ_i generation module) and using them, finer finite state abstractions and increasingly stronger global interaction invariants (Abstraction and Ψ generation module). Third, it checks deadlock freedom by checking satisfiability of $\bigwedge \Phi_i \wedge \Psi \wedge DIS$ (satisfiability module). If it succeeds, the system is proven deadlock-free, else it may continue or give up, according to the user choice. For doing all this, D-Finder is con-

Fig. 2. DeadlockFinder tool

nected with several external tools. It uses Omega [18] for quantifier elimination and Yices [19] for checking satisfiability of predicates. It is also connected to the state space exploration tool of the BIP platform, for finer analysis when the heuristic fails to prove deadlock-freedom. We provide non trivial examples showing the capabilities of D-Finder as well as the efficiency of the method.

3 Experimentation and Concluding remarks

We provide experimental results for four examples. The first example is Utopar, an industrial case study of the European Integrated project SPEEDS (<http://www.speeds.eu.com/>) about an automated transportation system. A succinct description of Utopar can be found at <http://www.combest.eu/home/?link=Application2>. The system is the composition of three types of components: autonomous vehicles, called U-cars, a centralized Automatic Control System and Calling Units. The latter two types have (almost exclusively) discrete behavior. U-cars are equipped with a local controller, responsible for handling the U-cars sensors and performing various

routing and driving computations depending on users' requests. We analyzed a simplified version of Utopar by abstracting from data exchanged between components as well as from continuous dynamics of the cars. In this version, each U-Car is modeled by a component having 7 control locations and 6 integer variables. The Automatic Control System has 3 control locations and 2 integer variables. The Calling Units have 2 control locations and no variables. In the second example, we consider Readers-Writer systems in order to evaluate how the method scales up for components without data. The third example is Gas Station in order to compare with other compositional method *assume-guarantee* [10]. Finally, as a last example, we consider Dinning Philosophers which is a well-known classical example.

The table below provides an overview of the experimental results obtained for these examples. In this table, n is the number of BIP components in the example, q is the total number of control locations, x_b (resp. x_i) is the total number of boolean (resp. integer) variables, D_{ci} is the number of deadlock configurations remaining in $DIS \wedge CI \wedge II$ and t is the total time for computing invariants and checking for satisfiability of $DIS \wedge CI \wedge II$. Detailed results are available at <http://www-verimag.imag.fr/~thnguyen/tool>.

<i>example</i>	n	q	x_b	x_i	D_{ci}	t
Utopar System (40 U-Cars, 256 Calling Units)	297	795	40	242	0	3m46s
Utopar System (60 U-Cars, 625 Calling Units)	686	1673	60	362	0	25m29s
Readers-Writer (7000 readers)	7002	14006	0	1	0	17m27s
Readers-Writer (10000 readers)	10002	20006	0	1	0	36m10s
Gas station (100 pumps - 2000 customers)	1101	4302	0	0	0	14m06s
Gas station (300 pumps - 3000 customers)	3301	12902	0	0	0	33m02s
Philosophers (2000 Philos)	4000	10000	0	0	3	32m14s
Philosophers (3001 Philos)	6001	15005	0	0	1	54m34s

We did some comparison with some well-known monolithic verification tools such as NewSMV (NuSMV) and Spin. All the experimentations are done on a Linux machine Intel Pentium 4 3.0 GHz and 1G Ram.

Fig. 3. Comparison with NuSmv and Spin on Dinning Philosopher

Fig. 4. Comparison with NuSmv on Gas Station

The first comparison between NuSmv, Spin and DFinder is on Dining Philosopher example. We increase the number of Philosophers and compare the verification time between these three tools (figure 3). In this figure, Spin runs out of memory at the size 17 (Philosophers); NuSmv runs out of memory at the size 150 while DFinder can go much further until the size 3000.

The second comparison between NuSmv and DFinder is on Gas Station example. We consider a system with 3 pumps and increase the number of customers. The comparison of verification time is in figure 4. In this figure, NuSmv runs out of memory at the size 180 (Customers) while DFinder can go much further until the size 3000.

References

1. Kupferman, O., Vardi, M.Y.: Modular model checking. LNCS **1536** (1998) 381–401
2. Alur, R., Henzinger, T.: Reactive modules. In: Proceedings of the 11th Annual Symposium on LICS, IEEE Computer Society Press (1996) 207–208
3. Abadi, M., Lamport, L.: Conjoining specifications. *Toplas* **17**(3) (1995) 507–534
4. Clarke, E., Long, D., McMillan, K.: Compositional model checking. In: Proceedings of the 4th Annual Symposium on LICS. (1989) 353–362
5. Chandy, K., J. Misra: Parallel program design: a foundation. Addison-Wesley Publishing Company (1988)
6. Grumberg, O., Long, D.E.: Model checking and modular verification. *ACM Transactions on Programming Languages and Systems* **16**(3) (1994) 843–871
7. McMillan, K.L.: A compositional rule for hardware design refinement. In: CAV ’97, Springer-Verlag (1997) 24–35
8. Pnueli, A.: In transition from global to modular temporal reasoning about programs. (1985) 123–144
9. Stark, E.W.: A proof technique for rely/guarantee properties. In: FSTTCS: proceedings of the 5th conference. Volume 206., Springer-Verlag (1985) 369–391
10. Cobleigh, J.M., Avrunin, G.S., Clarke, L.A.: Breaking up is hard to do: An evaluation of automated assume-guarantee reasoning. *ACM Transactions on Software Engineering and Methodology* **17**(2) (2008)
11. Peterson, J.: Petri Net theory and the modelling of systems. Englewood-Cliffs: Prentice Hall (1981)
12. Apt, K.R., Francez, N., de Roever, W.P.: A proof system for communicating sequential processes. *ACM Trans. Program. Lang. Syst.* **2**(3) (1980) 359–385
13. Basu, A., Bozga, M., Sifakis, J.: Modeling heterogeneous real-time components in bip. In: SEFM. (2006) 3–12
14. Lakhnech, Y., Bensalem, S., Berezin, S., Owre, S.: Incremental verification by abstraction. In: TACAS. (2001) 98–112
15. Bradley, A.R., Manna, Z.: Checking safety by inductive generalization of counterexamples to induction. In: FMCAD. (2007) 173–180
16. Bensalem, S., Lakhnech, Y.: Automatic generation of invariants. *FMSD* **15**(1) (July 1999) 75–92
17. Bensalem, S., Bozga, M., Sifakis, J., Nguyen, T.H.: Compositional verification for component-based systems and application. In: ATVA. (2008) 64–79
18. Team, O.: The omega library. Version 1.1.0 (November 1996)
19. Dutertre, B., de Moura, L.: A fast linear-arithmetic solver for DPLL(T). In: CAV’06. Volume 4144 of LNCS. (2006) 81–94