

CADP 2006: A Toolbox for the Construction and Analysis of Distributed Processes (Tool Paper)

Hubert Garavel, Radu Mateescu, Frédéric Lang, and Wendelin Serwe

INRIA

Centre de recherche Rhône-Alpes / VASY

655, avenue de l'Europe, Montbonnot, 38 334 Saint Ismier Cedex, France

{Hubert.Garavel,Radu.Mateescu,Frederic.Lang,Wendelin.Serwe}@inria.fr

1 Introduction

CADP (*Construction and Analysis of Distributed Processes*)¹ [2,3] is a toolbox for specification, rapid prototyping, verification, testing, and performance evaluation of asynchronous systems (concurrent processes with message-passing communication). The developments of CADP during the last five years led to a new release named CADP 2006 “*Edinburgh*” (as a tribute to the achievements in concurrency theory of the Laboratory for Foundations of Computer Science) that supersedes the previous version CADP 2001.

2 Modular Integration of Verification Techniques

CADP 2006 includes a complete range of functionalities for the design of asynchronous systems: code generation, rapid prototyping, random, interactive, or guided simulation, explicit-state verification, test case generation, and performance evaluation. CADP accepts as input either networks of communicating automata or higher-level specification languages, such as the ISO standard LOTOS.

Many CADP tools operate on Labeled Transition Systems (LTSS), which are represented either *explicitly*, as compact binary files encoded in the BCG (*Binary Coded Graphs*) format, or *implicitly*, as C programs implementing the transition relation according to the OPEN/CÆSAR API (*Application Programming Interface*). Three forms of verification are supported by CADP: *visual checking* (graphical inspection of an LTS), *model checking* (satisfaction of a modal μ -calculus formula by an LTS), and *equivalence checking* (comparison of two LTSS with respect to some equivalence/preorder relation).

To address the state space explosion problem, CADP 2006 provides the following verification techniques²:

- *Compositional verification* builds the LTS of a concurrent system incrementally by generating, minimizing, and recomposing the LTSS of individual

¹ <http://www.inrialpes.fr/vasy/cadp>

² Related verification tools are listed at <http://anna.fi.muni.cz/yahoda>

processes. Refined compositional verification uses interface constraints (specified manually or synthesized automatically) that restrict the behaviour of a process depending on its environment, thus limiting the size of intermediate LTSS.

- *On-the-fly verification* avoids the complete construction of an LTS by exploring only the portion relevant for verification. Model checking and equivalence checking are reformulated in terms of Boolean Equation Systems (BESs), which are solved on-the-fly using specialized linear-time algorithms.
- *Partial order reductions*, performed on-the-fly, reduce the LTS size by eliminating redundancies arising from the interleaving of independent transitions, still preserving various equivalence relations (branching bisimulation, weak trace equivalence, etc.).
- *Static analysis* aims at reducing the size of the LTSS generated from a system description, still preserving strong bisimulation.
- *Massively parallel verification* handles very large LTSS by using the computing resources of machine clusters and grids.

A key feature of CADP is to allow all these techniques to be combined in a highly modular way within the same software environment. Examples of such combinations that cumulate the reductions and scale up to larger systems are:

1. One can use a grid to generate the LTS corresponding to a LOTOS specification, while applying static analysis and on-the-fly τ -confluence reduction simultaneously.
2. One can model check on-the-fly whether a μ -calculus formula is satisfied by a network of communicating LTSS, while applying partial order reduction.
3. One can compare on-the-fly an LTS to a network of communicating LTSS that have been previously generated (taking into account refined interface constraints) and minimized for some bisimulation.

The tools of CADP can be invoked either interactively, by using the EUCALYPTUS graphical user interface, or in batch mode, by describing the desired verification scenario as a script in the SVL language.

3 New and Enhanced Tools in CADP 2006

CADP 2006 offers 42 tools and 20 generic software libraries dedicated to verification. Numerous tools existing in CADP 2001 were entirely rewritten or significantly improved, and 15 new tools and code libraries were added:

- BCG_GRAPH generates several useful kinds of LTSS in the BCG format, such as bags, FIFO queues, and chaos LTSS.
- BCG_MERGE [5] produces a single BCG file from a PBG file (see the DISTRIBUTOR tool below). It is equipped with a graphical tool to monitor the LTS generation interactively.

- BCG_STEADY and BCG_TRANSIENT [7] perform steady-state and transient analysis of an (extended) CTMC (*Continuous-Time Markov Chain*) encoded in the BCG format.
- BISIMULATOR [11,1] compares two LTSS on-the-fly modulo one out of 14 behavioural relations (strong, branching, observational, $\tau^*.a$, safety, trace, or weak trace equivalences — and their associated preorders). The equivalence checking problem is reformulated as a BES, which is solved using the CÆSAR_SOLVE library below. When both LTSS are not related, BISIMULATOR generates a counterexample, i.e., an acyclic LTS containing distinguishing transition sequences. Compared with former CADP tools, BISIMULATOR is more efficient and generates smaller counterexamples.
- CÆSAR 7.0 [6] is a compiler for the behaviour part of LOTOS. Among other improvements, CÆSAR 7.0 implements a static analysis based on live variable analysis, which assigns a canonical value to variables that are no longer used, thus avoiding to distinguish states that only differ by values of variables not used in the future. Compared to its previous version 6.2, CÆSAR 7.0 can reduce LTS size by several orders of magnitude (e.g., 10^4), thus allowing larger LOTOS specifications to be handled.
- CÆSAR.BDD uses Binary Decision Diagrams to perform various structural analyses on basic Petri nets, such as exploring reachable markings to determine the set of “dead” transitions and the pairs of “concurrent” units.
- CÆSAR_SOLVE [11] is a generic software library for on-the-fly resolution of BESS, represented as implicit boolean graphs similarly to LTSS in OPEN/CÆSAR. It offers 6 resolution algorithms, based on breadth-first search or depth-first search (with memory-efficient variants for acyclic or disjunctive/conjunctive BESS) strategies, which also generate diagnostics (boolean subgraphs) explaining the truth value of boolean variables.
- DETERMINATOR [7] takes as input an extended CTMC encoded in the BCG format and tries to extract on-the-fly a pure CTMC (i.e., containing only stochastic transitions). Doing so, the tool checks a sufficient condition ensuring that the resulting CTMC is unique, or returns an error otherwise.
- DISTRIBUTOR [5] performs distributed LTS generation and on-the-fly reduction by τ -compression and τ -confluence [10] using several machines connected by a network. It launches a remote process on each machine to generate a fragment of the entire LTS. The result of the distributed generation is a PBG (*Partitioned BCG Graph*), i.e., a set of LTS fragments located on remote machines. DISTRIBUTOR is equipped with a graphical tool to monitor the PBG generation in real-time.
- EVALUATOR 3.5 [12,11] evaluates on-the-fly, on an LTS, temporal properties expressed in regular alternation-free μ -calculus. The model checking problem is reformulated as a BES, which is solved using the CÆSAR_SOLVE library. The tool generates full diagnostics (examples and counterexamples) and optimizes memory consumption (i.e., does not store the transitions, but only the states of the LTS) for a large spectrum of properties. EVALUATOR 3.5 is 3 times faster and consumes 3 times less memory than EVALUATOR 3.0.

- EXP.OPEN 2.0 [8] maps communicating LTSS composed using synchronization vectors, parallel composition operators (from CCS, CSP, μ CRL, LOTOS, and E-LOTOS), and/or generalized hide, rename, and cut operators onto the OPEN/CÆSAR API. It implements partial order reductions preserving various equivalences (e.g., branching, stochastic, weak trace equivalence). It can also translate communicating LTSS into the PEP, TINA, and FC2 formats, and can synthesize interface constraints for refined compositional verification [9]. EXP.OPEN 2.0 uses 2 times less memory and is up to 45 times faster than EXP.OPEN 1.0.
- PROJECTOR 2.0 [13] reduces an LTS on-the-fly with respect to interface constraints represented by another LTS and a set of labels. Among its new features, PROJECTOR 2.0 allows to describe the set of labels more compactly by the way of regular expressions. Experiments indicate that PROJECTOR 2.0 is up to 4 times faster than PROJECTOR 1.0.
- REDUCTOR 5.0 reduces an LTS on-the-fly, either partially or totally, modulo one out of 8 relations (trace, weak trace, $\tau^*.a$, and safety equivalences, τ -confluence, τ -compression, τ -divergence [10], and strong bisimulation), possibly displaying the equivalence classes of the quotient graph. Some of these reductions perform a local resolution of a BES using CÆSAR_SOLVE. The algorithm used by REDUCTOR 5.0 to eliminate invisible transitions has a lower average complexity than in earlier versions.
- SEQ.OPEN [4] maps traces onto the OPEN/CÆSAR API. It implements an optimized representation of a trace as an LTS without storing the whole trace in memory, which is also significantly more efficient (up to 50 times faster) than converting the trace first into a BCG file.

Notice that the ALDÉBARAN tool (available since the origin of CADP) was replaced in CADP 2006 by a (upward-compatible) shell script that invokes BISIMULATOR, BCG_INFO, BCG_MIN, and REDUCTOR to provide the same functionalities as ALDÉBARAN.

4 Conclusion

CADP 2006 “*Edinburgh*” is the result of five years of intensive R&D in verification technology. Four computing platforms are currently supported: SPARC/SOLARIS, INTEL/LINUX, INTEL/WINDOWS, and POWERPC/MAC OS. As regards impact, 366 organizations have signed the CADP license agreement already, and CADP has been installed on 820 machines in the world during year 2006. CADP was used for 94 case-studies³ and 29 research tools⁴ are connected to CADP. Most notably, in the FORMALFAME project, CADP was successfully used to validate crucial parts of Bull’s NOVASCALE machines, which form the core of TERA10, Europe’s most powerful supercomputer.

CADP will continue to be enhanced in the next years. In particular, we plan to apply CADP in three main areas: software environments for critical embedded

³ listed at <http://www.inrialpes.fr/vasy/cadp/case-studies>

⁴ listed at <http://www.inrialpes.fr/vasy/cadp/software>

systems (French projects OPENEMBEDD and TOPCASED), high-performance multiprocessor architectures (French project MULTIVAL), and bio-informatics (European project EC-MOAN, where CADP will contribute to the understanding of a bacterial model system).

Acknowledgements. Damien Bergamini, David Champelovier, Adrian Curic, Nicolas Descoubes, Holger Hermanns, Christophe Joubert, Bruno Ondet, Gordon Pace, Frédéric Perret, Irina Smarandache-Sturm, Gilles Stragier, Frédéric Tronel, and Marie Vidal contributed to CADP 2006. The authors are also grateful to the 79 CADP users⁵ whose suggestions led to numerous CADP enhancements.

References

1. Bergamini, D., Descoubes, N., Joubert, C., Mateescu, R.: BISIMULATOR: A Modular Tool for On-the-Fly Equivalence Checking. In: Ziarko, W., Yao, Y. (eds.) RSTC 2000. LNCS (LNAI), vol. 2005, Springer, Heidelberg (2001)
2. Fernandez, J.-C., Garavel, H., Kerbrat, A., Mateescu, R., Mounier, L., Sighireanu, M.: CADP (CÆSAR/ALDEBARAN Development Package): A Protocol Validation and Verification Toolbox. In: Alur, R., Henzinger, T.A. (eds.) CAV 1996. LNCS, vol. 1102, Springer, Heidelberg (1996)
3. Garavel, H., Lang, F., Mateescu, R.: An Overview of CADP 2001. EASST Newsletter 4, 13–24 (2002) Also available as INRIA Technical Report RT-0254
4. Garavel, H., Mateescu, R.: SEQ.OPEN: A Tool for Efficient Trace-Based Verification. In: Gelbukh, A. (ed.) CICLing 2001. LNCS, vol. 2004, Springer, Heidelberg (2001)
5. Garavel, H., Mateescu, R., Bergamini, D., Curic, A., Descoubes, N., Joubert, C., Smarandache-Sturm, I., Stragier, G.: DISTRIBUTOR and BCG_MERGE: Tools for Distributed Explicit State Space Generation. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006 and ETAPS 2006. LNCS, vol. 3920, Springer, Heidelberg (2006)
6. Garavel, H., Serwe, W.: State Space Reduction for Process Algebra Specifications. Theoretical Computer Science 351(2), 131–145 (2006)
7. Hermanns, H., Joubert, C.: A Set of Performance and Dependability Analysis Components for CADP. In: Garavel, H., Hatcliff, J. (eds.) ETAPS 2003 and TACAS 2003. LNCS, vol. 2619, Springer, Heidelberg (2003)
8. Lang, F.: EXP.OPEN 2.0: A Flexible Tool Integrating Partial Order, Compositional, and On-the-fly Verification Methods. In: Romijn, J.M.T., Smith, G.P., van de Pol, J. (eds.) IFM 2005. LNCS, vol. 3771, Springer, Heidelberg (2005)
9. Lang, F.: Refined Interfaces for Compositional Verification. In: Najm, E., Pradat-Peyre, J.F., Donzeau-Gouge, V.V. (eds.) FORTE 2006. LNCS, vol. 4229, Springer, Heidelberg (2006)
10. Mateescu, R.: On-the-fly State Space Reductions for Weak Equivalences. In: Proc. of FMICS '05, ACM Computer Society Press, New York (2005)
11. Mateescu, R.: CAESAR_SOLVE: A Generic Library for On-the-Fly Resolution of Alternation-Free Boolean Equation Systems. Springer International Journal on Software Tools for Technology Transfer (STTT) 8(1), 37–56 (2006)

⁵ listed at <http://www.inrialpes.fr/vasy/cadp/news6.html>

12. Mateescu, R., Sighireanu, M.: Efficient On-the-Fly Model-Checking for Regular Alternation-Free Mu-Calculus. *Science of Computer Programming* 46(3), 255–281 (2003)
13. Pace, G., Lang, F., Mateescu, R.: Calculating τ -Confluence Compositionally. In: Hunt Jr., W.A., Somenzi, F. (eds.) *CAV 2003*. LNCS, vol. 2725, Springer, Heidelberg (2003)