

Fighting State Space Explosion: Review and Evaluation

Radek Pelánek*

Department of Information Technology, Faculty of Informatics
Masaryk University Brno, Czech Republic
xpelanek@fi.muni.cz

Abstract. In order to apply formal methods in practice, the practitioner has to comprehend a vast amount of research literature and realistically evaluate practical merits of different approaches. In this paper we focus on explicit finite state model checking and study this area from practitioner's point of view. We provide a systematic overview of techniques for fighting state space explosion and we analyse trends in the research. We also report on our own experience with practical performance of techniques. Our main conclusion and recommendation for practitioner is the following: be critical to claims of dramatic improvement brought by a single sophisticated technique, rather use many different simple techniques and combine them.

1 Introduction

If you are a practitioner who wants to apply formal methods in industrial critical systems, you have to address following problems: Which of the many approaches should I use? If I want to improve the performance of my tool, which of the techniques described by researchers should I use? Which techniques are worth the implementation effort? These are important questions which are not easy to answer, nevertheless they are seldom addressed in research papers.

Research papers rather propose a steady flow of novel techniques, improvements, and optimizations. However, the experimental work reported in research papers has often poor quality [59] and it is difficult to judge the practical merit of proposed techniques. The goal of this paper is to provide an overview of research and realistic assessment of practical merits of techniques. The paper should serve as a guide for a practitioner who is trying to answer the above given questions.

It is not feasible to realize this goal for the whole field of formal verification at once. Therefore, we focus on one particular area (explicit model checking) and give an overview of research in this area and report on practical experience. Even though our discussion is focused on one specific area, we believe that our main recommendation — that it is better to combine several simple techniques rather than to focus on one sophisticated one — is applicable to many other areas of formal methods.

* Partially supported by GA ČR grant no. 201/07/P035.

Explicit model checking is principally very simple — a brute force traversal of all possible model states. Despite the simplicity of the basic idea, explicit model checking is still the best approach for many practically important areas of application, e.g., verification of communication protocols and software. The popularity of the approach is illustrated by large number of available tools (e.g., Spin, CADP, mCRL2, Uppaal, Divine, Java PathFinder, Helena) and widespread availability of courses and textbooks on the topic (e.g., [10]).

The main obstacle in applying explicit model checking in practice is the state space explosion problem. Hence, the research focuses mainly on techniques for fighting state space explosions — during the last 15 years more than 100 papers have been published on the topic, proposing various techniques for fighting state space explosion. What are these techniques and how can we classify them? What is the real improvement brought by these techniques? Which techniques are practically useful? Which techniques should a practitioner study and use?

We try to answer these questions, particularly we provide the following:

- We overview techniques for fighting state space explosion in explicit model checking and divide them into four main areas (Section 2).
- We review and analyse research on fighting state space explosion, and discuss main trends in this research (Section 3).
- We report on our own practical experience with application and evaluation of techniques for fighting state space explosion (Section 4).
- Based on the review of literature and our experience, we provide specific recommendation for practitioner in industry (Section 5).

The main aim of this paper is to present and support the following message: Rather than optimizing the performance of a single sophisticated technique, we should use many different simple techniques, study how to combine them, and how to run them effectively in parallel.

2 Overview of Techniques for Fighting State Space Explosion

Fig. 1. gives an algorithm EXPLORE which explores the reachable part of the state space. This basic algorithm can be directly used for verification of simple safety properties; for more complex properties, we have to use more sophisticated algorithms (e.g., cycle detection [72]). Nevertheless, the basic ideas of techniques for fighting state space explosion are similar. For clarity, we discuss these techniques mainly with respect to the basic EXPLORE algorithm.

The main problem of explicit state space exploration is state space explosion problem and consequently memory and time requirements of the algorithm EXPLORE. Techniques for fighting state space explosion can be divided into four main groups:

```

proc EXPLORE( $M$ )
   $Wait = \{s_0\}; Visited = \emptyset$ 
  while  $Wait \neq \emptyset$  do
    get  $s$  from  $Wait$ 
    explore state  $s$ 
    foreach  $s' \in$  successors of  $s$  do
      if  $s' \notin Visited$  then
        add  $s'$  to  $Wait$ 
        add  $s$  to  $Visited$  fi od
    od
end

```

Fig. 1. The basic algorithm which explores all reachable states. Data structure *Wait* (also called open list) holds states to be visited, data structure *Visited* (also called visited list, closed list, transposition table, or just hash table) stores already explored states.

1. Reduce the number of states that need to be explored.
2. Reduce the memory requirements needed for storing explored states.
3. Use parallelism or distributed environment.
4. Give up the requirement on completeness and explore only part of the state space.

In the following we discuss these four types of approaches and for each of them we list examples of specific techniques.

2.1 State Space Reductions

When we inspect some simple models and their state spaces, we quickly notice significant redundancy in these state spaces. So the straightforward idea is to try to exploit this redundancy and reduce the number of states visited during the search. In order to exploit this idea in practice, we have to specify which states are omitted from the search and we have to show the correctness of the approach, i.e., prove that the visited part of the state space is equivalent to the whole state space with respect to some equivalence (typically bisimulation or stutter equivalence).

State based reductions. State based reductions exploit observation that if two states are bisimilar then it is sufficient to explore successors of only one of them. The reduction can be performed either on-the-fly during the exploration or by a static modification of the model before the exploration. Examples of such reductions are symmetry reduction [12,20,43,44,70,74], live variable reduction [21,69], cone of influence reductions, and slicing [19,35].

Path based reductions. Path based reductions exploit observation that sometimes it is sufficient to explore only one of two sequences of actions because they are just different linearizations of “independent” actions and therefore have the

same effect. These reductions try to reduce the number of equivalent interleavings. Examples of such reductions are transition merging [18,48], partial order reduction [27,33,40,63,64], τ -confluence [11], and simultaneous reachability analysis [55].

Compositional methods. Systems are often specified as a composition of several components. This structure can be exploited in two ways: compositional generation of the state space [46] and assume-guarantee approach [22,32,66].

2.2 Storage Size Reductions

The main bottleneck of model checking are usually memory requirements. Therefore, we can save some memory at the cost of using more time, i.e., by employing some kind of time-memory trade-off. The main source of memory requirements of the algorithm EXPLORE is the structure *Visited* which stores previously visited states. Hence, techniques, which try to lower memory requirements, focus mainly on this structure.

State compression. During the search, each state is represented as a byte vector which can be quite large (e.g., 100 bytes). In order to save space, this vector can be compressed [25,26,30,39,49,56,73] or common components can be shared [38]. Instead of compressing individual states, we can also represent the whole structure *Visited* implicitly as a minimized deterministic automaton [41].

Caching and selective storing. Instead of storing all states in the structure *Visited*, we can store only some of these states — this approach can lead to revisits of some states and hence can increase runtime, but it saves memory. Techniques of this type are for example:

- caching [24,28,65], which deletes some currently stored states when the memory is full,
- selective storing [9,49], which stores only some states according to given heuristics,
- sweep line method [15,54,68], which uses so called progress function; this function guarantees that some states will not be revisited in the future and hence these states can be deleted from the memory.

Use of magnetic disk. Simple use of magnetic disk leads to an extensive swapping and slows down the computation extremely. So the magnetic disk have to be used in a sophisticated way [7,8,71] in order to minimize disk operations.

2.3 Parallel and Distributed Computation

Another approach to manage a large number of states is to use even more brute force — more processors.

Networks of workstations. Distributed computation can be realized most easily by network of workstations connected by fast communication medium

(i.e., workstations communicate by message passing). In this setting the state space is partitioned among workstations (i.e., each workstation stores part of the data structure *Visited*) and workstations exchange messages about states to be visited (*Wait* structure), see e.g., [23,50,51]. The application of distributed environment for verification of liveness properties is more complicated, because classical algorithms are based on depth-first search, which cannot be easily adapted for distributed environment. Hence, for verification of liveness properties we have to use more sophisticated algorithms, see e.g., [1,2,3,5,13,14].

Multi-core processors. Recently, multi-core processors become widely available. Multi-core processors provide parallelism with shared memory, i.e., the possibility to reduce run-time of the verification by parallel exploration of several states, see e.g., [4,42].

2.4 Randomized Techniques and Heuristics

If the memory requirements of the search are too large even after the application of above given techniques, we can use randomized techniques and heuristics. These techniques explore only part of the state space. Therefore, they can help only in the detection of an error; they cannot assist us in proving correctness.

Heuristic search (also called directed or guided search). States are visited in an order given by some heuristics, i.e., *Wait* list is implemented as priority queue [31,47,67]. Different heuristic approach is to use genetic algorithm which tries to ‘evolve’ a path to a goal state [29].

Random walk and partial search. Random walk does not store any information and always visits just one successor of a current state [34,60]. This basic strategy can be extended in several ways, e.g., by visiting a subset of all successors (instead of just one state), storing some states in the *Visited* structure, or combining random walk with local breadth-first search, see e.g., [36,45,52,53,60].

Bitstate hashing. The algorithm does not store whole states but only one bit per state in a large hash table [37]. In a case of collision some states are omitted by the search. A more involved version of this technique is based on Bloom filters [16,17].

3 Research Analysis

What are the trends in the research literature about techniques for fighting state space explosion? Is the quality of experimental evidence improving? How significant is the improvement reported in research papers? How is this improvement changing over time?

3.1 Research Papers

In order to answer the above given questions, we have collected and analyzed large set of research papers. More specifically, we collected research papers that

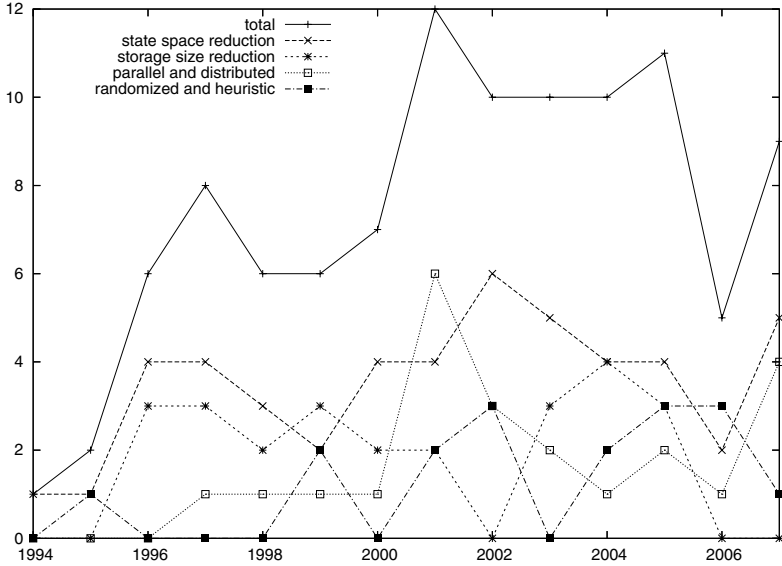


Fig. 2. Numbers of publications; note that some papers can be counted in two categories

describe techniques for fighting state space explosion in explicit model checking of finite state systems¹.

The collection contains more than 100 papers – these papers were obtained by systematically collecting papers from the most relevant conferences and by citation tracking. The full list of reviewed papers, which includes all papers referenced above, is freely available². The collection is certainly not complete, but we believe that it is a good sample of research in the area.

Fig. 2. shows the number of publications in each year during the last 13 years. Although there are rises and downfalls, the overall flow of publications on the topic is rather steady. The figure also shows that all four areas described in the previous section are pursued concurrently.

3.2 Quality of Experiments

Although some of the considered research is rather theoretically oriented (e.g., partial order reduction), all considered techniques are in fact heuristics which aim at improving performance of model checking tools. So what really matters is the practical improvement brought by each technique. To assess the improvement

¹ In few cases we also include techniques which are not purely explicit, but target the similar application domain (i.e., the experiments are done on same models as for other included papers).

² <http://www.fi.muni.cz/~xpelane/amase/reductions.bib>

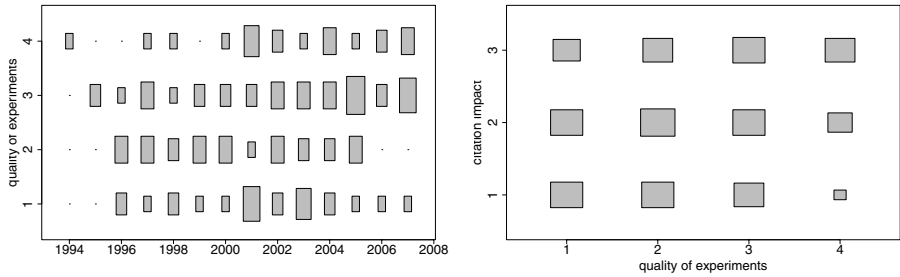


Fig. 3. The first graph shows the quality of experiments reported in model checking papers during time. The size of a box corresponds to a number of published papers in a given year and quality category. The second graph shows the relation between experiment quality and citation impact; citation impact is divided into three categories: less than 10 citations, 10–30 citations, more than 30 citations; only publication before 2004 are used.

it is necessary to perform experimental evaluation. Only good experiments can provide realistic evaluation of practical merits of proposed techniques.

In order to study the quality of experiments, we classify experiments in each paper into one of four classes, depending on the number and type of used models³:

1. Random inputs or few toy models.
2. Several toy models (possibly parametrized) or few simple models.
3. Several simple models (possibly parametrized) or one large case study.
4. Exhaustive study of parametrized simple models or several case studies.

Fig 3. presents the quality of experiments in papers from our sample. The figure shows that the quality on average is not very good and, what is even more disappointing, that there is slow progress in time, although many realistic case studies are available (see [59] for more detailed discussion of these issues).

Since there is a large number of techniques, it is important to compare performance of novel techniques with previously studied one. However, analysis of our research sample shows that only about 40% papers contain some comparison with similar techniques; this ratio is improving with time, but only slowly. Moreover, the comparison is usually only shallow.

Our analysis also shows one encouraging trend. Fig. 3. shows that there is a relation between quality of experiments and citation impact of a paper — research with better experiments is more cited.

3.3 Reported Improvement

Before the discussion of improvements reported in research papers, we clarify the terminology that we use to measure this improvement. We use the notion

³ The classification is clearly slightly subjective. Nevertheless, we believe that the main conclusions of our analysis do not depend on the subjective factor.




Fig. 4. Reported improvement with respect to time and quality of experiments

‘reduction ratio’ to denote the ratio between the memory consumption of the technique for fighting state space explosion and the memory consumption of the standard reachability (exploration of full reachable state space). Some authors report ‘reduced by’ factor, i.e., if we report ‘reduction ratio’ 80%, it means that the memory consumption was ‘reduced by’ 20%. Note that in this section we analyze reduction ratios as reported by authors, not what we consider realistic reduction ratios of techniques.

For clarity of presentation, we again divide the reported reduction ratios into four classes:

1. Reported reduction ratio is 50% or worse (or sometimes good but sometimes worse than 100%).
2. Reported reduction ratio is in most cases 10%-50%.
3. Reported reduction ratio is in most cases 1%-10%.
4. Reported reduction ratio is better than 1% (or exponential improvement is reported or only out-of memory for full search is reported, i.e., reduction ratio is impossible to assess).

Fig. 4. shows that in most cases the reported reduction is in the second category (reduction between 10% and 50%). The relation with the quality of experiments clearly demonstrates that this is also the most realistic evaluation — better results are often caused by poor experiments, not by special features of techniques.

There is no clear trend with respect to time, i.e., it seems that novel techniques do not significantly improve on performance of previous techniques. This does not automatically mean that the recent research is misguided. In some cases novel technique provides principally different way how to obtain the reduction and can be combined with previously proposed techniques in orthogonal way. Novel technique can also extend the application domain of previously studied techniques.

As we already mentioned, the research in this domain is purely practically motivated. However, the amount of research into certain topic is not really related to its practical merit. For example, our experience (described in next section) shows that the dead variable reduction brings similar improvement as partial order reduction. Nevertheless, there are significantly more research papers about

partial order reduction than about dead variable reduction. This is probably not due to the practical merits of partial order reduction, but because it can be extensively studied theoretically.

4 Practical Experience

In this section we report on our experience with techniques for fighting state space explosion. Our experience is based on large-scale research studies, which are described in stand-alone publications. Here we provide only a brief description of these studies and present their main conclusions. Technical details can be found in cited papers.

Our experience report obviously does not cover all techniques for fighting state space explosion. However, we cover all four areas described in Section 2 and the main conclusions are in all cases similar, so we believe that it is reasonable to generalize our experience.

4.1 On-the-fly State Space Reductions

We evaluated several techniques for on-the-fly state space reductions. Setting of this study (see [57] for details):

- Implementation: publicly available implementations of explicit model checkers (Spin, Murphy, DiVinE).
- Models: models included in tool distributions plus few more publicly available case studies.
- Techniques: dead variable reduction, partial order reduction, transition merging, symmetry reduction.

When we measured the performance of techniques over realistic models, we found that the reduction ratio is usually worse than what is reported in research papers — research papers often use simple models with artificially high values of model parameters. More specifically, the main results of our evaluation are the following: dead variable reduction works on nearly all models, reduction ratio is usually between 10% and 90%; partial order reduction works only in some cases, reduction ratio is between 5% and 90%; transition merging works in similar cases as partial order reduction, it is weaker but easier to realize, reduction ratio is usually between 50% and 95%; symmetry reduction works only for few models (symmetric ones), reduction ratio is usually between 8% and 50%.

Our main conclusion from this study are the following:

- Each technique is applicable only to some types of models. No technique works really universally; more specialized techniques yield better reduction.
- On real models, no single technique is able to achieve reduction ratio significantly under 5%. Claims about drastic reduction, which occur in some papers, are not really appropriate.
- Since there are many techniques and many of them are orthogonal, most models can be reduced quite significantly.

4.2 Caching and Compression

From the area of ‘storage size reduction’ techniques we evaluated two techniques for reducing memory consumption of the data structure *Visited*. Setting of this study (see [62] for details):

- Implementation: all techniques are implemented in uniform way using the DiVinE environment [6] (source codes are publicly available).
- Models: 120 models from BEEM (BENchmarks for EXPLICIT Model checkers) [59].
- Techniques: state caching with 7 different caching strategies, state compression with Huffman coding (two variants: static code and code computed by training runs).

In the study we also reviewed previous research on storage size reduction techniques. We found that using proper parameter values with our simple and easy-to-implement techniques, we were able to achieve very similar results to those reported in other works which use far more sophisticated approaches. Concrete results of the evaluation are the following:

- Caching strategies are to a certain degree complementary. Using an appropriate state caching strategy, the reduction ratio is in most cases 10% to 30%.
- Using state compression, the reduction ratio is usually around 60%.
- The two techniques combine well.

4.3 Distributed Exploration

From the area of parallel and distributed techniques we report on the basic distributed approach to explicit model checking: we have a network of workstations connected by fast Ethernet, workstations communicate via message passing (MPI library), state space is partitioned among workstations. In this setting the reduction ratio is clearly bounded by $1/n$, where n is the number of workstations. In practice the reduction ratio is worse because of communication overhead. Here we report on results of our evaluation, however the results are rather typical in this area.

Setting of this study (see [58] for details):

- Implementation: the DiVinE tool (public version).
- Models: 120 models from BEEM (BENchmarks for EXPLICIT Model checkers) [59].
- Techniques: distributed reachability on 20 workstations.

In this study the speedup varies from 2 to 12, typical value of the speedup is between 4 and 6 (i.e., reduction ratio around 20%). We also found that the speedup is negatively correlated with the speed of successor generation by the tool.

4.4 Error Detection Techniques

From the area of ‘randomized techniques and heuristics’ we have chosen 9 techniques and evaluated their performance. In this case we do not study the reduction ratio, because it is not known — the experiments are done on models for which the standard reachability is not feasible. Therefore, we focus on relative performance of techniques and on the issue of complementarity.

Setting of this study (see [61] for details):

- Implementation: all techniques are implemented in uniform way using the DiVinE environment [6] (source codes are publicly available),
- Models: 54 models (with very large state space) from BEEM [59].
- Techniques: breadth-first search, depth-first search, randomized DFS, two variants of random walk, bitstate hashing with repetition, two variants of directed search, and under-approximation refinement based on partial order reduction.

For the evaluation we used several performance measures: number of steps needed to find an error, length of reported counterexample, and coverage metrics. The main results of this study are the following:

- There is no single best technique. Results depend on used performance metrics, even for a given metric, the most successful technique is the best one only over 25% of models.
- It is important to focus on complementarity of techniques, not just on their overall (average) performance. For example, in our study the random walk technique had rather poor overall performance, but it was successful on models where other techniques fail, i.e., it is a useful technique which we should not discard.

5 Conclusions

This paper is concerned with techniques for fighting state space explosion problem in explicit model checking. We review the research in the area during the last 15 years (more than 100 research papers) and report on our practical experience. As a result of our review we identify four main groups of techniques: state space reductions, storage size reductions, parallel and distributed computation, randomization and heuristics. These four groups are rather orthogonal and can be combined; within each group techniques are often based on similar ideas and their combination can be difficult.

The review of research shows that despite a steady flow of publications on the topic, the progress is not very significant — in fact the reduction ratio reported in research papers stays practically the same over the last 15 years. This analysis stresses the need for good practical evaluation. However, realistic evaluation of research progress is complicated by rather poor experimental standards and by unjustified claims by researchers.

Results reported in research papers often make an impression of dramatic improvements. Our practical experience suggests that it is not realistic to get better reduction ratio than 5% with a single technique, in fact in most cases the obtained reduction ratio is between 20% and 80%. Nevertheless, this does not mean that techniques for fighting state space explosion are not useful. Techniques of different types can be combined, and together they might be able to bring a significant improvement.

Our experience also suggest that simple techniques are often sufficient. The performance obtained by sophisticated techniques is often similar to performance of basic techniques from each area. Complicated techniques often achieve better results only for specialized application domains. This observation can be also supported by analysis of techniques implemented in model checking tools. Tools usually implement basic versions of many techniques, sophisticated techniques are often implemented only in a tool used by authors of the technique.

To summarise, we propose following recommendations for those who want to apply model checking in practice:

- Use large number of simple techniques of different types.
- Do not try to find ‘the best’ technique of a specific type. Try to find a set of simple complementary techniques and run all of them (preferably in parallel).
- Be critical to claims in research papers, particularly if the experimental evidence is poor.
- Use sophisticated techniques only if they are specifically targeted at your domain of application.
- Focus on combination of orthogonal techniques.

Researchers, we believe, should focus not just on the development of novel techniques, but also on issues of techniques combination, selection, and efficient scheduling: How to select right technique for a given model? In what order we should try available techniques? Can information gathered by one technique be used by another techniques?

Acknowledgment

I thank Václav Rosecký, Pavel Moravec, and Jaroslav Šeděnka for cooperation on practical evaluation of techniques.

References

1. Barnat, J., Brim, L., Černá, I.: Property driven distribution of nested DFS. In: Proc. of Workshop on Verification and Computational Logic, number DSSE-TR-2002-5 in DSSE Technical Report, pp. 1–10. University of Southampton, UK (2002)
2. Barnat, J., Brim, L., Chaloupka, J.: Parallel breadth-first search LTL model-checking. In: Proc. of Automated Software Engineering (ASE 2003), pp. 106–115. IEEE Computer Society, Los Alamitos (2003)

3. Barnat, J., Brim, L., Chaloupka, J.: From distributed memory cycle detection to parallel LTL model checking. *ENTCS* 133(1), 21–39 (2005)
4. Barnat, J., Brim, L., Rockai, P.: Scalable multi-core LTL model-checking. In: Bošnački, D., Edelkamp, S. (eds.) *SPIN 2007*. LNCS, vol. 4595, pp. 187–203. Springer, Heidelberg (2007)
5. Barnat, J., Brim, L., Štříbrná, J.: Distributed LTL model-checking in SPIN. In: Dwyer, M.B. (ed.) *SPIN 2001*. LNCS, vol. 2057, pp. 200–216. Springer, Heidelberg (2001)
6. Barnat, J., Brim, L., Černá, I., Moravec, P., Rockai, P., Šimeček, P.: Di-VinE - a tool for distributed verification. In: Ball, T., Jones, R.B. (eds.) *CAV 2006*. LNCS, vol. 4144, pp. 278–281. Springer, Heidelberg (2006), <http://anna.fi.muni.cz/divine>
7. Barnat, J., Brim, L., Šimeček, P.: I/o efficient accepting cycle detection i/o efficient accepting cycle detection. In: Damm, W., Hermanns, H. (eds.) *CAV 2007*. LNCS, vol. 4590, pp. 281–293. Springer, Heidelberg (2007)
8. Barnat, J., Brim, L., Šimeček, P., Weber, M.: Revisiting resistance speeds up i/o-efficient ltl model checking. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS 2008*. LNCS, vol. 4963, pp. 48–62. Springer, Heidelberg (2008)
9. Behrmann, G., Larsen, K.G., Pelánek, R.: To store or not to store. In: Hunt Jr., W.A., Somenzi, F. (eds.) *CAV 2003*. LNCS, vol. 2725, pp. 433–445. Springer, Heidelberg (2003)
10. Ben-Ari, M.: *Principles of the SPIN Model Checker*. Springer, Heidelberg (2008)
11. Blom, S., van de Pol, J.: State space reduction by proving confluence. In: Brinksma, E., Larsen, K.G. (eds.) *CAV 2002*. LNCS, vol. 2404, pp. 596–609. Springer, Heidelberg (2002)
12. Bosnacki, D.: A light-weight algorithm for model checking with symmetry reduction and weak fairness. In: Ball, T., Rajamani, S.K. (eds.) *SPIN 2003*. LNCS, vol. 2648, pp. 89–103. Springer, Heidelberg (2003)
13. Brim, L., Černá, I., Krčál, P., Pelánek, R.: Distributed LTL model checking based on negative cycle detection. In: Hariharan, R., Mukund, M., Vinay, V. (eds.) *FSTTCS 2001*. LNCS, vol. 2245, pp. 96–107. Springer, Heidelberg (2001)
14. Černá, I., Pelánek, R.: Distributed explicit fair cycle detection. In: Ball, T., Rajamani, S.K. (eds.) *SPIN 2003*. LNCS, vol. 2648, pp. 49–73. Springer, Heidelberg (2003)
15. Christensen, S., Kristensen, L.M., Mailund, T.: A sweep-line method for state space exploration. In: Margaria, T., Yi, W. (eds.) *TACAS 2001*. LNCS, vol. 2031, pp. 450–464. Springer, Heidelberg (2001)
16. Dillinger, P.C., Manolios, P.: Bloom filters in probabilistic verification. In: Hu, A.J., Martin, A.K. (eds.) *FMCAD 2004*. LNCS, vol. 3312, pp. 367–381. Springer, Heidelberg (2004)
17. Dillinger, P.C., Manolios, P.: Fast and accurate bitstate verification for SPIN. In: Graf, S., Mounier, L. (eds.) *SPIN 2004*. LNCS, vol. 2989, pp. 57–75. Springer, Heidelberg (2004)
18. Dong, Y., Ramakrishnan, C.R.: An optimizing compiler for efficient model checking. In: *Proc. of Formal Description Techniques for Distributed Systems and Communication Protocols (FORTE XII) and Protocol Specification, Testing and Verification (PSTV XIX)*, pp. 241–256. Kluwer, Dordrecht (1999)
19. Dwyer, M.B., Hatcliff, J., Hoosier, M., Ranganath, V.P.: Evaluating the effectiveness of slicing for model reduction of concurrent object-oriented programs. In: Hermanns, H., Palsberg, J. (eds.) *TACAS 2006*. LNCS, vol. 3920, pp. 73–89. Springer, Heidelberg (2006)

20. Emerson, E.A., Wahl, T.: Dynamic symmetry reduction. In: Halbwachs, N., Zuck, L.D. (eds.) TACAS 2005. LNCS, vol. 3440, pp. 382–396. Springer, Heidelberg (2005)
21. Fernandez, J.C., Bozga, M., Ghirvu, L.: State space reduction based on live variables analysis. *Journal of Science of Computer Programming (SCP)* 47(2-3), 203–220 (2003)
22. Flanagan, C., Qadeer, S.: Thread-modular model checking. In: Ball, T., Rajamani, S.K. (eds.) SPIN 2003. LNCS, vol. 2648, pp. 213–224. Springer, Heidelberg (2003)
23. Garavel, H., Mateescu, R., Smarandache, I.: Parallel state space construction for model-checking. In: Dwyer, M.B. (ed.) SPIN 2001. LNCS, vol. 2057, pp. 217–234. Springer, Heidelberg (2001)
24. Geldenhuys, J.: State caching reconsidered. In: Graf, S., Mounier, L. (eds.) SPIN 2004. LNCS, vol. 2989, pp. 23–38. Springer, Heidelberg (2004)
25. Geldenhuys, J., de Villiers, P.J.A.: Runtime efficient state compaction in SPIN. In: Dams, D.R., Gerth, R., Leue, S., Massink, M. (eds.) SPIN 1999. LNCS, vol. 1680, pp. 12–21. Springer, Heidelberg (1999)
26. Geldenhuys, J., Valmari, A.: A nearly memory-optimal data structure for sets and mappings. In: Ball, T., Rajamani, S.K. (eds.) SPIN 2003. LNCS, vol. 2648, pp. 136–150. Springer, Heidelberg (2003)
27. Godefroid, P.: Partial-order methods for the verification of concurrent systems: an approach to the state-explosion problem. In: Godefroid, P. (ed.) *Partial-Order Methods for the Verification of Concurrent Systems*. LNCS, vol. 1032, p. 142. Springer, Heidelberg (1996)
28. Godefroid, P., Holzmann, G.J., Pirottin, D.: State space caching revisited. In: Probst, D.K., von Bochmann, G. (eds.) CAV 1992. LNCS, vol. 663, pp. 178–191. Springer, Heidelberg (1993)
29. Godefroid, P., Khurshid, S.: Exploring very large state spaces using genetic algorithms. In: Katoen, J.-P., Stevens, P. (eds.) TACAS 2002. LNCS, vol. 2280, pp. 266–280. Springer, Heidelberg (2002)
30. Gregoire, J.: State space compression in spin with GETSs. In: *Proc. Second SPIN Workshop*, Rutgers University, New Brunswick, New Jersey (1996)
31. Groce, A., Visser, W.: Heuristics for model checking java programs. *Software Tools for Technology Transfer (STTT)* 6(4), 260–276 (2004)
32. Grumberg, O., Long, D.E.: Model checking and modular verification. *ACM Transactions on Programming Languages and Systems* 16(3), 843–871 (1994)
33. Gueta, G., Flanagan, C., Yahav, E., Sagiv, M.: Cartesian partial-order reduction. In: Bošnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 95–112. Springer, Heidelberg (2007)
34. Haslum, P.: Model checking by random walk. In: *Proc. of ECSEL Workshop* (1999)
35. Hatcliff, J., Dwyer, M.B., Zheng, H.: Slicing software for model construction. *Higher Order Symbol. Comput.* 13(4), 315–353 (2000)
36. Holzmann, G.J.: Algorithms for automated protocol verification. *AT&T Technical Journal* 69(2), 32–44 (1990)
37. Holzmann, G.J.: An analysis of bitstate hashing. In: *Proc. of Protocol Specification, Testing, and Verification*, pp. 301–314. Chapman & Hall, Boca Raton (1995)
38. Holzmann, G.J.: State compression in SPIN: Recursive indexing and compression training runs. In: *Proc. of SPIN Workshop* (1997)
39. Holzmann, G.J., Godefroid, P., Pirottin, D.: Coverage preserving reduction strategies for reachability analysis. In: *Proc. of Protocol Specification, Testing, and Verification* (1992)

40. Holzmann, G.J., Peled, D.: An improvement in formal verification. In: Proc. of Formal Description Techniques VII, pp. 197–211. Chapman & Hall, Ltd., Boca Raton (1995)
41. Holzmann, G.J., Puri, A.: A minimized automaton representation of reachable states. *Software Tools for Technology Transfer (STTT)* 3(1), 270–278 (1998)
42. Holzmann, G.J., Bosnacki, D.: The design of a multicore extension of the spin model checker. *IEEE Transactions on Software Engineering* 33(10), 659–674 (2007)
43. Iosif, R.: Symmetry reduction criteria for software model checking. In: Bošnački, D., Leue, S. (eds.) *SPIN 2002*. LNCS, vol. 2318, pp. 22–41. Springer, Heidelberg (2002)
44. Ip, C.N., Dill, D.L.: Better verification through symmetry. *Formal Methods in System Design* 9(1–2), 41–75 (1996)
45. Jones, M.D., Sorber, J.: Parallel search for LTL violations. *Software Tools for Technology Transfer (STTT)* 7(1), 31–42 (2005)
46. Krimm, J.P., Mounier, L.: Compositional state space generation from Lotos programs. In: Brinksma, E. (ed.) *TACAS 1997*. LNCS, vol. 1217, pp. 239–258. Springer, Heidelberg (1997)
47. Kuehlmann, A., McMillan, K.L., Brayton, R.K.: Probabilistic state space search. In: Proc. of Computer-Aided Design (CAD 1999), pp. 574–579. IEEE Press, Los Alamitos (1999)
48. Kurshan, R.P., Levin, V., Yenigün, H.: Compressing transitions for model checking. In: Brinksma, E., Larsen, K.G. (eds.) *CAV 2002*. LNCS, vol. 2404, pp. 569–581. Springer, Heidelberg (2002)
49. Larsen, K.G., Larsson, F., Pettersson, P., Yi, W.: Efficient verification of real-time systems: Compact data structure and state-space reduction. In: Proc. of Real-Time Systems Symposium (RTSS 1997), pp. 14–24. IEEE Computer Society Press, Los Alamitos (1997)
50. Lerda, F., Sisto, R.: Distributed-memory model checking with SPIN. In: Dams, D.R., Gerth, R., Leue, S., Massink, M. (eds.) *SPIN 1999*. LNCS, vol. 1680, p. 22. Springer, Heidelberg (1999)
51. Lerda, F., Visser, W.: Addressing dynamic issues of program model checking. In: Dwyer, M.B. (ed.) *SPIN 2001*. LNCS, vol. 2057, pp. 80–102. Springer, Heidelberg (2001)
52. Lin, F., Chu, P., Liu, M.: Protocol verification using reachability analysis: the state space explosion problem and relief strategies. *Computer Communication Review* 17(5), 126–134 (1987)
53. Mihail, M., Papadimitriou, C.H.: On the random walk method for protocol testing. In: Dill, D.L. (ed.) *CAV 1994*. LNCS, vol. 818, pp. 132–141. Springer, Heidelberg (1994)
54. Mailund, T., Westergaard, W.: Obtaining memory-efficient reachability graph representations using the sweep-line method. In: Jensen, K., Podelski, A. (eds.) *TACAS 2004*. LNCS, vol. 2988, pp. 177–191. Springer, Heidelberg (2004)
55. Ozdemir, K., Ural, H.: Protocol validation by simultaneous reachability analysis. *Computer Communications* 20, 772–788 (1997)
56. Parreaux, B.: Difference compression in SPIN. In: Proc. of Workshop on automata theoretic verification with the SPIN model checker (SPIN 1998) (1998)
57. Pelánek, R.: Evaluation of on-the-fly state space reductions. In: Proc. of Mathematical and Engineering Methods in Computer Science (MEMICS 2005), pp. 121–127 (2005)
58. Pelánek, R.: Web portal for benchmarking explicit model checkers. Technical Report FIMU-RS-2006-03, Masaryk University Brno (2006)

59. Pelánek, R.: BEEM: Benchmarks for explicit model checkers. In: Bošnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 263–267. Springer, Heidelberg (2007)
60. Pelánek, R., Hanžl, T., Černá, I., Brim, L.: Enhancing random walk state space exploration. In: Proc. of Formal Methods for Industrial Critical Systems (FMICS 2005), pp. 98–105. ACM Press, New York (2005)
61. Pelánek, R., Rosecký, V., Moravec, P.: Complementarity of error detection techniques. In: Proc. of Parallel and Distributed Methods in verifiCation (PDMC) (2008)
62. Pelánek, R., Rosecký, V., Šeděňka, J.: Evaluation of state caching and state compression techniques. Technical Report FIMU-RS-2008-02, Masaryk University Brno (2008)
63. Peled, D.: Combining partial order reductions with on-the-fly model-checking. In: Dill, D.L. (ed.) CAV 1994. LNCS, vol. 818, pp. 377–390. Springer, Heidelberg (1994)
64. Penczek, W., Szreter, M., Gerth, R., Kuiper, R.: Improving partial order reductions for universal branching time properties. *Fundamenta Informaticae* 43(1-4), 245–267 (2000)
65. Penna, G.D., Intrigila, B., Melatti, I., Tronci, E., Zilli, M.V.: Exploiting transition locality in automatic verification of finite state concurrent systems. *Software Tools for Technology Transfer (STTT)* 6(4), 320–341 (2004)
66. Pnueli, A.: In transition from global to modular temporal reasoning about programs. *Logics and models of concurrent systems*, 123–144 (1985)
67. Qian, K., Nymeyer, A.: Guided invariant model checking based on abstraction and symbolic pattern databases. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 497–511. Springer, Heidelberg (2004)
68. Schmidt, K.: Automated generation of a progress measure for the sweep-line method. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 192–204. Springer, Heidelberg (2004)
69. Self, J.P., Mercer, E.G.: On-the-fly dynamic dead variable analysis. In: Bošnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 113–130. Springer, Heidelberg (2007)
70. Sistla, A.P., Godefroid, P.: Symmetry and reduced symmetry in model checking. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, pp. 91–103. Springer, Heidelberg (2001)
71. Stern, U., Dill, D.L.: Using magnetic disk instead of main memory in the murphi verifier. In: Y. Vardi, M. (ed.) CAV 1998. LNCS, vol. 1427, pp. 172–183. Springer, Heidelberg (1998)
72. Vardi, M.Y., Wolper, P.: An automata-theoretic approach to automatic program verification. In: Kozen, D. (ed.) Proceedings of the First Annual IEEE Symposium on Logic in Computer Science (LICS 1986), pp. 332–344. IEEE Computer Society Press, Los Alamitos (1986)
73. Visser, W.: Memory efficient state storage in SPIN. In: Proc. of SPIN Workshop, pp. 21–35 (1996)
74. Wahl, T.: Adaptive symmetry reduction. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 393–405. Springer, Heidelberg (2007)