

TOOL PRESENTATION
The SPINJA Model Checker*

Marc de Jonge¹ and Theo C. Ruys²

¹ TNO, The Netherlands, marcdejonge@gmail.com

² RUwise Consultancy, Deventer, The Netherlands, ruys@mac.com

Abstract. SPINJA is a model checker for PROMELA, implemented in Java. SPINJA is designed to behave similarly to SPIN, but to be more easily extendible and reusable. Despite the fact that SPINJA uses a layered object-oriented design and is written in Java, SPINJA's performance is reasonable: benchmark experiments have shown that, in exhaustive mode, SPINJA is about five times slower than the highly optimized SPIN. For bitstate verification runs the difference is only a factor of two.

1 Introduction

The SPIN model checker [2,10] is arguably one of the most powerful and popular (software) model checkers available. Unfortunately, the highly optimized C code of SPIN is difficult to understand and hard to extend. Moreover, the algorithms implemented in SPIN cannot be reused in other tools.

This paper presents SPINJA[†] [11], a model checker for PROMELA, implemented in Java. SPINJA is designed to behave similarly to SPIN, but to be more easily extendible and reusable.

Related Work. Independently from SPINJA, Mordechai Ben-Ari has started the development of the model checker ERIGONE [1,8], a model checker for (a subset of) PROMELA, implemented in Ada. ERIGONE's goal is to facilitate *learning* concurrency and model checking. NIPS [6,9] is another attempt to provide a stand-alone alternative model checker for PROMELA. NIPS uses the virtual machine approach to model checking. Both ERIGONE and NIPS do not (yet) implement any of SPIN's powerful optimization algorithms (e.g., partial order reduction, bitstate hashing, hash compaction) though.

Sec. 2 explains the design and implementation of SPINJA. Sec. 3 reports on a benchmark experiment with SPINJA and SPIN. Sec. 4 concludes the paper and discusses future work.

2 SPINJA

SPINJA can be used to check for the absence of deadlocks, assertions, liveness properties and never claims in PROMELA models. SPINJA's verification mode can

* The first version of SPINJA has been developed as part of a MSc project within the Formal Methods & Tools group of the University of Twente, The Netherlands.

[†] SPINJA stands for 'Spin in Java'. SPINJA rhymes with *ninja*.

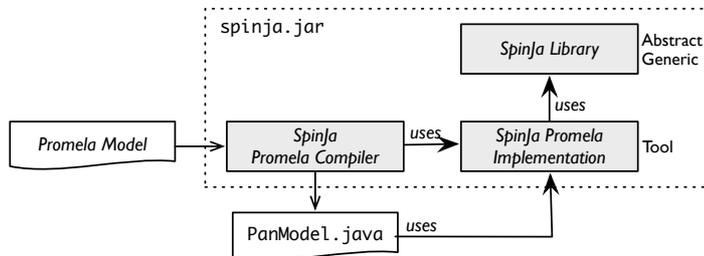


Fig. 1. Architecture of SPINJA.

exploit (nested) depth first search (DFS) or breadth first search (BFS). Bitstate hashing and hash compaction modes are also supported. Furthermore, SPIN’s partial order reduction and statement merging are implemented. SPINJA can also be used for simulation. To enable or disable features, SPINJA uses the same syntax for options as SPIN. And the output of SPINJA also mimics SPIN: seasoned SPIN users should feel right at home.

SPINJA supports a large subset of the PROMELA language. Currently, the following aspects of the PROMELA language are not yet supported: the `unless` statement, `typedef` definitions, non-FIFO communication, and embedded C code.

Design

The research question behind the development of SPINJA was the following: ‘Is it possible to re-implement the core of SPIN in Java using an extendible object-oriented design while being competitive in terms of memory consumption and runtime behaviour?’ Beforehand – given the folklore about Java’s infamous lack of speed, compared to C/C++ – we would have been happy if SPINJA would turn out to be a single order of magnitude slower than SPIN.

From the start we have committed ourselves to Java and a clean object-oriented approach. This section only sketches the design of SPINJA; please refer to [3] for a detailed description.

Fig. 1 shows the high-level architecture of SPINJA. SPINJA follows the same principle as SPIN: given a PROMELA model as input, a verifier program is generated: `PanModel.java`. The SPINJA library, which is used by `PanModel.java`, consists of two parts: a PROMELA specific part and a generic part. The complete SPINJA package (i.e., compiler and library) is supplied as a single Java jar-file.

For the design and implementation of SPINJA we have adopted the conceptual framework for explicit-state model checkers as developed by Mark Katzenbelt et.al. [5]. This layered framework is more flexible than the ‘black box’ approach of conventional model checkers. Fig. 2 shows the layered model as implemented in SPINJA. Each layer extends the layer below it and adds more functionality. This is achieved through inheritance. This way the different parts of the model checker are loosely coupled. This increases the maintainability and reusability of the whole system.

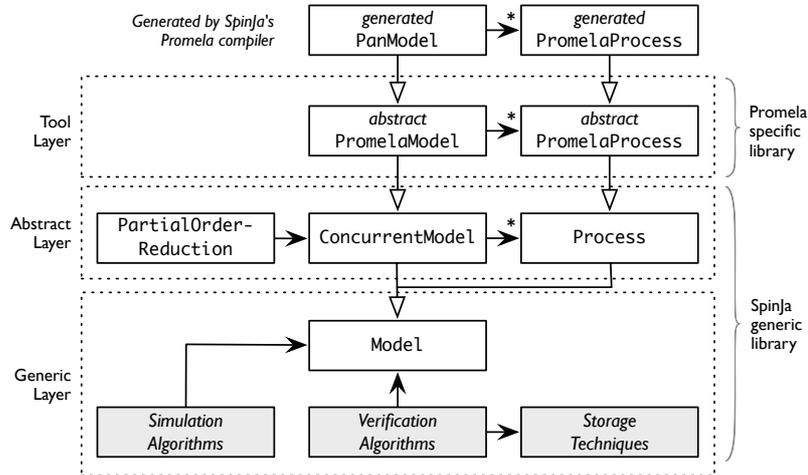


Fig. 2. Layered design of SPINJA.

Generic layer. The Generic layer is the lowest layer. A `Model` is responsible for generating all the states at run-time. The `Model` knows how to take `Transitions` from one state to the next. On the basis of just `State` and `Transition` objects, several important algorithms are implemented within the generic layer: search (DFS, BFS) and simulation algorithms, storage methods, hashing methods.

Abstract layer. The Abstract layer extends on the Generic layer by adding concurrency to the model: a `ConcurrentModel` is a collection of `Processes` which are `Model` objects themselves. The class `PartialOrderReduction` changes the way the set of successor transitions of a `ConcurrentModel` are being computed. The Generic and Abstract layers are completely independent from the PROMELA language and could serve as the basis for any explicit-state model checker.

Tool Layer. Within the Tool layer the abstract `PromelaModel` is defined which forms the blueprint for the generated Java verifier `PanModel`. The `PromelaModel`, for example, knows how to create and terminate PROMELA proctypes and how to add `Channel` objects.

Compiler. The SPINJA PROMELA compiler is responsible for translating the PROMELA specification to the generated Java program: `PanModel`. Much effort has been put into implementing the same semantics for PROMELA as SPIN does: in almost all cases, the resulting automata of the SPINJA compiler are the same as the ones that SPIN produces. Consequently, for most PROMELA models, the number of states and transitions traversed by SPINJA and SPIN are the same.

Implementation. The current version of SPINJA is version 0.9. SPINJA is released as open-source under the Apache 2.0 license and is available from [11]. The development took roughly one man year of work (i.e., [4] and [3]). The code base of SPINJA 0.9 consists of 144 files, 25k loc and its size is 1100Kb. By comparison,

<i>verification mode</i>	\odot <i>states</i>	\odot <i>transitions</i>	\odot SPINJA	\odot SPIN no -02	\odot SPIN with -02	SJA/(S no -02)	SJA/(S with -02)
Exhaustive, with POR	4.2×10^6	19.9×10^6	38.7	14.8	8.63	2.6×	4.5×
Exhaustive, no POR	4.2×10^6	20.2×10^6	38.4	13.9	8.06	2.8×	4.8×
Bitstate Hashing (with POR)	4.2×10^6	19.9×10^6	21.3	17.0	10.67	1.3×	2.0×
Hash Compaction (with POR)	4.2×10^6	19.9×10^6	20.5	14.7	8.56	1.4×	2.4×

Table 1. Summary of benchmark experiments: SPINJA vs. SPIN (*no -02* and *with -02*). Running times are in seconds.

SPIN 5.2.4 consists of 48 files, 48k loc and is 1200Kb and Erigone 2.1.5 consists of 71 files, 8k loc and is 300Kb.

3 Experiments

As mentioned in Sec. 2, we are interested in how a Java verifier generated by SPINJA compares with the C verifier generated by SPIN with respect to time and memory consumption. In this section we report on a benchmark experiment. For this paper, we have used ten medium size PROMELA models from the BEEM benchmark suite [7]: *at.4*, *elevator2.3*, *fischer.6*, *hanoi.3*, *loyd.2*, *mcs.3*, *peterson.4*, *sorter.3*, *szymanski.4* and *telephony.4*.

We compared SPIN against SPINJA in (i) a default exhaustive safety run, (ii) with partial order reduction (POR) disabled, (iii) in bitstate hashing mode (with `-k3 -w31`), and (iv) in hash compaction mode. We considered two versions of SPIN’s `pan` verifier: (i) default compilation without optimizations and (ii) a compilation with `gcc`’s `-02` optimization option set. All runs were limited to 1Gb of memory and a depth of 1.2×10^6 steps. All models could be verified completely except for *hanoi.3* which requires an even larger depth. The experiments were run on an Apple MacBook 2.4Ghz Intel Core 2 Duo, with 4Gb of RAM, running Mac OS X 10.5.8, with SPIN 5.2.4, SPINJA 0.9, `gcc` 4.2.1 and Java 1.6.0_17.

Table 1 summarizes the *averages* (i.e. \odot) of running the verifiers: the number of states, the number of transitions and the running time (i.e., *user time*) of the verifiers in the various verification modes. Generation and compilation time have not been taken into account. The last two columns list the average speed difference factor between SPINJA and SPIN *without -02*, and SPINJA and SPIN *with -02*. Several important observations can be made from Table 1:

- SPIN’s performance profits a lot from `gcc`’s `-02` optimization option.
- Using Java for the implementation of a model checker seems acceptable: without `-02`, SPIN is less than three times faster than SPINJA.
- SPINJA is surprisingly close to SPIN for the two approximate verification modes: the difference in speed is only a factor of two.

We do not report on the memory consumption of the verifiers in Table 1. The reason for this is that there is no substantial difference between SPINJA and SPIN with respect to memory. In most cases, SPINJA needs slightly less memory than SPIN though.

For a more extensive benchmark comparison between SPIN 5.1.6 and SPINJA 0.8 (under Windows XP and using Microsoft's C compiler) we refer to [3].

4 Conclusions

In this paper we presented SPINJA, a model checker for PROMELA, implemented in Java. Due to its layered, object-oriented design we believe that the design and implementation of SPINJA is clear, readable and extensible. Compared to the highly optimized C-based SPIN, the Java-based SPINJA is clearly slower. For bitstate hashing mode, however, the difference is only a factor of two. To extend the usability of SPINJA further, several improvements are planned:

- Support for the missing PROMELA features.
- Support for modelling languages other than PROMELA.
- Implementation of SPIN's `-DREVERSE`, `-DP_RAND`, `-DT_REVERSE` and `-DT_RANDOM` to increase the coverage of bitstate hashing verification runs.
- Improvement of SPINJA's explicit-state storage algorithms.
- Support for embedding Java code (similar to the support for C code in SPIN); this would make it possible to use SPINJA to model check Java applications.
- Development of a code generator for PROMELA on basis of SPINJA's generated verifier code.

References

1. M. Ben-Ari. Teaching Concurrency and Model Checking. In C. S. Pasareanu, editor, *Proc. of SPIN 2009*, LNCS 5578, pages 6–11. Springer, 2009.
2. G. J. Holzmann. *The SPIN Model Checker – Primer and Reference Manual*. Addison-Wesley, Boston, Massachusetts, USA, 2004.
3. M. de Jonge. The SpinJ Model Checker. Master's thesis, University of Twente, Enschede, The Netherlands, September 2008.
4. M. A. Kattenbelt. Towards an Explicit-State Model Checking Framework. Master's thesis, University of Twente, Enschede, The Netherlands, September 2006.
5. M. A. Kattenbelt, T. C. Ruys, and A. Rensink. An Object-Oriented Framework for Explicit-State Model Checking. In *Proc. of VVSS 2007, Eindhoven, The Netherlands, March 2007*, TUE Computer Science Reports, Nr. 07-04, pages 84–92, 2007.
6. M. Weber. An Embeddable Virtual Machine for State Space Generation. In D. Bosnacki and S. Edelkamp, editors, *Proc. of SPIN 2007*, LNCS 4595, pages 168–185. Springer, 2007.
7. BEEM: BEncmarks for Explicit Model checkers.
<http://anna.fi.muni.cz/models/>.
8. ERIGONE. <http://code.google.com/p/erigone/>.
9. NIPS. <http://www.ewi.utwente.nl/~michaelw/nips/>.
10. SPIN. <http://spinroot.com/>.
11. SPINJA. <http://code.google.com/p/spinja/>.

Presentation Plan

SPINJA is a command-line program which is executed from a terminal window. As such, a demo with SPINJA will *not* be a breathtaking visual experience. Still, we believe that a demonstration with SPINJA should be interesting and helpful for people that want to experiment with the tool.

We plan to split the presentation into three parts:

1. (50%) an introduction to the *design* of SPINJA, explaining its high-level, object-oriented architecture,
2. (20%) a *demo* showing how SPINJA works, followed by
3. (30%) an *explanation* of the Java application that is being generated by the SPINJA PROMELA compiler.

1. Architecture of SpinJa (50%)

Seen from the user's perspective, SPINJA might be just an alternative implementation of SPIN. This is not the real strength of SPINJA though.

We believe that the strength and promise of SPINJA lie within its extendible and reusable design. Therefore, the main part of the presentation will be devoted to explaining the (rationale behind the) layered, object-oriented design of SPINJA. We will show that the SPINJA library is general enough to form the basis for any explicit model checker.

2. Demo (20%)

In the demo we will show how SPINJA should be used. We will discuss a small PROMELA model (without errors) and show how to verify this PROMELA model with SPINJA.

We will also demonstrate SPINJA's simulation mode by simulating the PROMELA model using SPINJA's random and interactive simulation modes.

3. Explanation of generated Java application (30%)

After the demo we will dive into the generated Java code (`PanModel.java`) of the PROMELA model of the demo. We will show that this Java code is human readable and easy to understand. We will explain the connection between the generated model and SPINJA's library.

Availability

Both the source and binary distribution of SPINJA are available from [11].