

# BOOM: Taking Boolean Program Model Checking One Step Further\*

Gerard Basler, Matthew Hague, Daniel Kroening, C.-H. Luke Ong,  
Thomas Wahl, and Haoxian Zhao

Oxford University Computing Laboratory, Oxford, United Kingdom

**Abstract.** We present BOOM, a comprehensive analysis tool for Boolean programs. We focus in this paper on model-checking non-recursive concurrent programs. BOOM implements a recent variant of *counter abstraction*, where thread counters are used in a program-context aware way. While designed for bounded counters, this method also integrates well with the Karp-Miller tree construction for vector addition systems, resulting in a reachability engine for programs with *unbounded thread creation*. The concurrent version of BOOM is implemented using BDDs and includes partial order reduction methods. BOOM is intended for model checking system-level code via predicate abstraction. We present experimental results for the verification of Boolean device driver models.

## 1 Introduction

Over the past decade, *predicate abstraction* has evolved into a viable strategy for model checking software, witnessed by the success of device driver verification in Microsoft's SLAM project. The input program is converted into a finite-state *Boolean program*, whose paths overapproximate the original behavior.

Recently, *concurrent* software has gained tremendous stimulus due to the advent of multi-core computing architectures. The software is executed by asynchronous parallel threads, communicating, in the most general case, through fully shared variables. Bugs in such programming environments are known to be subtle and hard to detect by means of testing, strongly motivating formal analysis techniques for concurrent programs.

In this paper, we present BOOM, a model checker for Boolean programs. While BOOM has many features that make it useful for sequential programs [4], we focus here on analyzing the set of reachable states of a *replicated* non-recursive Boolean program. Replication often induces symmetry, which can and must be exploited for the analysis to scale. We present our implementation of a context-aware form of bounded *counter abstraction* [3], and compare its performance to alternative reduction techniques also implemented in BOOM, and to other tools.

Replication materializes in practice as dynamic thread creation. Even without a bound on the number of running threads, the reachability problem for non-recursive concurrent Boolean programs is decidable. We have extended BOOM by

---

\* Supported by the Swiss National Science Foundation (200021-109594) and the Engineering and Physical Sciences Research Council (EP/G026254/1, EP/D037085/1).

a variant of the *Karp-Miller tree*, which operates directly on Boolean programs. We demonstrate that our implementation performs much better in practice than the worst-case complexity of the construction seems to suggest. The result is a practically useful and **exact** reachability analysis for realistic concurrent Boolean programs with arbitrarily many threads.

## 2 Concurrent Boolean Program Analysis with BOOM

BOOM is capable of analyzing the reachable state space of *replicated* programs, where threads may dynamically create other threads during the execution. The Boolean variables are declared either *local* or *shared*. Each thread has its own private copy of the local variables. The shared variables, in contrast, are fully accessible to every thread. Shared-variable concurrency is very powerful and able to simulate many other communication primitives, such as locks.

Concurrent Boolean programs with replicated threads are naturally *symmetric*: the set of transitions of a derived Kripke model is invariant under permutations of the threads. BOOM exploits this property using a form of *counter abstraction*: global states are represented as vectors of counters, one per local state. Each counter tracks the number of threads in the corresponding local state. A transition by a thread translates into an update of the counters for the source and target local state.

The suggested rewriting of program transitions into counter updates can *in principle* be performed at the program text level. In practice, this is usually infeasible due to the *local state explosion problem*: the number of statically determined local states is exponential in the program text size. We recently proposed **context-awareness** as a solution [3]: at exploration time, the context in which a statement is executed is known and exposes the local-state counters that need to be updated. As a natural optimization, a global state in BOOM only keeps counters for *occupied* local states, where at least one thread resides in. Their number is obviously bounded by the number of running threads, which tends to be a tiny fraction of all local states.

### Extending BOOM to Unbounded Thread Creation

If there is no limit to how many threads may be running, the thread counters become unbounded non-negative integers. The induced transition system is an instance of a *vector addition system with [control] states* (VASS); the control state represents the values of the shared program variables. The reachability of a *thread state*  $(s, l)$  (combination of shared and local state) in a concurrent Boolean program is reducible to a VASS *coverability* problem. The latter problem is decidable, using a tree construction proposed by Karp and Miller [8].


BOOM uses the Karp-Miller construction as the starting point for an algorithm to decide thread-state reachability. The local state explosion problem materializes here as the *dimensionality problem* for VASSes. Fortunately, our earlier solution of a context-aware, on-the-fly translation is instrumental in the unbounded

case as well. Our version of the Karp-Miller procedure operates directly on Boolean programs. Bypassing the VASS allows us to avoid the blowup that a static translation into any type of addition system invariably entails. To ameliorate the exponential-space complexity of the Karp-Miller construction, we exploit the special form of vector-addition systems derived from Boolean programs. For example, our implementation keeps a copy of those tree nodes that are *maximal* with respect to the covering relation as partial order. Newly discovered nodes are compared against these maximal nodes only.

### 3 Results

BOOM and our benchmarks are available at <http://www.cprover.org/boom>; we refer the reader to this website for more details on the tool and the benchmarks.

The left chart below compares a plain symbolic exploration of the concurrent Boolean program against BOOM's implementation of **bounded** counter abstraction. Each data point specifies the numbers of threads running. The message of this chart is obvious. The right chart compares plain exploration against BOOM's implementations of partial-order reduction. Comparing left and right, we see that counter abstraction performs somewhat better. In other experiments (not shown), we observed that combining the two gives yet better performance.

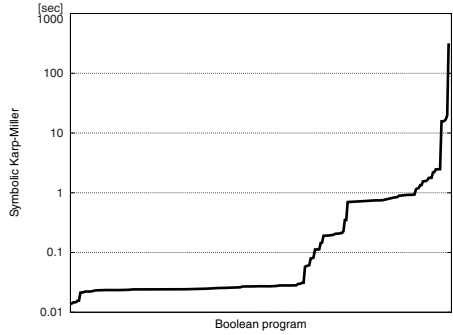


The chart on page 148 (left) compares the bounded version of BOOM with counter abstraction against the *lazy* version of GETAFIX [9] (which performs better than *eager*). GETAFIX targets *recursive* Boolean programs with a bounded number of context-switches. To compare with BOOM, we chose a non-recursive example from the GETAFIX website. The time for GETAFIX to convert the example into a sequential program is tiny and omitted. The table illustrates the time to explore the sequentialized program using MOPED-1, for different context-switch bounds. Note that BOOM explores *all* interleavings.

The graph on page 148 (right) shows our preliminary thread-state analysis of Boolean programs with unbounded thread creation. We see that for many examples, the running times are very small. On the other hand, 301 of 570 cases did not terminate within 60 min. We observed no more than 43 non-zero counters in any global state, despite millions of conceivable local states.

$n$	BOOM	GETAFIX/cont. bd. [sec]					
	[sec]	1	2	3	4	5	6
2	< 0.1	0.1	0.4	2.0	8.7	41	139
3	0.1	0.1	1.0	0.6	4.8	30	187
4	1.2	0.1	1.9	1.2	12.2	146	1318
5	12.1	0.14	2.8	2.3	30.6	426	—
6	88.8	0.2	3.9	3.1	51.7	901	—

Benchmarks on Intel 3GHz, with timeout 60 mins, memory-out 4 GB.



## 4 Related Work and Conclusion

There are a few tools for the analysis of sequential Boolean programs [2,6]. When extended to multiple threads, the problem becomes undecidable. To allow a complete solution, BOOM disallows recursion. There are many tools available for the analysis of VASS. Closest to our work are the applications of these tools to Java [5,7] and to Boolean programs [1]. These tools compile their input into an explicit-state transition system, which will result in a high-dimensional VASS. Our experiments with explicit-state encodings (not shown) indicate that encoding Boolean programs symbolically is mandatory. We believe BOOM to be the first **exact** tool able to analyze non-recursive concurrent Boolean programs with *bounded* replication efficiently, and to extend the underlying technique to the *unbounded* case with encouraging performance.

## References

1. Ball, T., Chaki, S., Rajamani, S.: Parameterized verification of multithreaded software libraries. In: Margaria, T., Yi, W. (eds.) TACAS 2001. LNCS, vol. 2031, p. 158. Springer, Heidelberg (2001)
2. Ball, T., Rajamani, S.: Bebop: A symbolic model checker for Boolean programs. In: Havelund, K., Penix, J., Visser, W. (eds.) SPIN 2000. LNCS, vol. 1885, pp. 113–130. Springer, Heidelberg (2000)
3. Basler, G., Kroening, D., Mazzucchi, M., Wahl, T.: Symbolic counter abstraction for concurrent software. In: Bouajjani, A., Maler, O. (eds.) Computer Aided Verification. LNCS, vol. 5643, pp. 64–78. Springer, Heidelberg (2009)
4. Basler, G., Kroening, D., Weissenbacher, G.: SAT-based summarization for boolean programs. In: Bošnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 131–148. Springer, Heidelberg (2007)
5. Delzanno, G., Raskin, J.-F., Begin, L.V.: Towards the automated verification of multithreaded Java programs. In: Katoen, J.-P., Stevens, P. (eds.) TACAS 2002. LNCS, vol. 2280, p. 173. Springer, Heidelberg (2002)
6. Esparza, J., Schwoon, S.: A BDD-based model checker for recursive programs. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, p. 324. Springer, Heidelberg (2001)

7. Geeraerts, G., Raskin, J.-F., Begin, L.V.: Expand, enlarge and check.. made efficient. In: Etessami, K., Rajamani, S.K. (eds.) CAV 2005. LNCS, vol. 3576, pp. 394–407. Springer, Heidelberg (2005)
8. Karp, R., Miller, R.: Parallel program schemata. *Computer and System Sciences* (1969)
9. Torre, S.L., Madhusudan, P., Parlato, G.: Reducing context-bounded concurrent reachability to sequential reachability. In: Bouajjani, A., Maler, O. (eds.) *Computer Aided Verification*. LNCS, vol. 5643, pp. 477–492. Springer, Heidelberg (2009)