Model Checking, SAT and Bit-Vectors

Armin Biere Johannes Kepler University Linz, Austria

Seminar Theoretical Computer Science

KTH Royal Institute of Technology

Stockholm, Sweden

Monday, 30rd November, 2015







- Symbolic Execution
 - particularly in combination with concolic testing
 - impressive project SAGE at Microsoft (Patrice Godefroid)
- Equivalence Checking
 - first widely adopted formal technique in HW verification
 - originally check combinational equivalence of RTL versus gate-level
 - first use since mid 90'ies, wide-spread adoption since 2000
 - since 10 years applied to sequential equivalence too
 - used for checking last-minute fixes: engineering change orders (ECOs)
 - checking arithmetic circuits, e.g., multipliers, still not completely automatic
- Abstract Interpretation
 - used to check for number / floating-point overflows etc.
 - static analyzer Astrée verified Airbus flight SW (2003)
- Model Checking or Property Checking increasingly used by certain companies
 Model Checking, SAT and Bit-Vectors KTH Stockholm

BurchClarkeMcMillanDillHwang'90: Symbolic Model Checking

CoudertMadre'89: Symbolic Reachability

ClaessenSorensson'12: k-liveness

BiereArthoSchuppan'01: Liveness2Safety

Pnueli'77: Temporal Logic

McMillan'03: Interpolation

McMillan'93: SMV

Bradley'10: IC3

BiereCimattiClarkeZhu'99: Bounded Model Checking

ClarkeEmerson'82: Model Checking

Kurshan'93: Localization

ClarkeEmersonSifakis:

QuielleSifakis'82: Model Checking

BallRajamani'01: SLAM

Turing Award 2007

Holzmann'91: SPIN

GrafSaidi'97: Predicate Abstraction

Holzmann'81: On–The–Fly Reachability

ClarkeGrumbergJahLuVeith'03: CEGAR

SheeranSinghStalmarck'00: k-induction

Peled'94: Partial–Order–Reduction

Symbolic Model Checking without BDDs*

Armin Biere¹, Alessandro Cimatti², Edmund Clarke¹, and Yunshan Zhu¹

¹ Computer Science Department, Carnegie Mellon University 5000 Forbes Avenue, Pittsburgh, PA 15213, U.S.A {Armin.Biere,Edmund.Clarke,Yunshan.Zhu}@cs.cmu.edu ² Istituto per la Ricerca Scientifica e Tecnologica (IRST) via Sommarive 18, 38055 Povo (TN), Italy cimatti@irst.itc.it

Abstract. Symbolic Model Checking [3, 14] has proven to be a powerful technique for the verification of reactive systems. BDDs [2] have traditionally been used as a symbolic representation of the system. In this paper we show how boolean decision procedures, like Stålmarck's Method [16] or the Davis & Putnam Procedure [7], can replace BDDs. This new technique avoids the space blow up of BDDs, generates counterexamples much faster, and sometimes speeds up the verification. In addition, it produces counterexamples of minimal length. We introduce a *bounded model checking* procedure for LTL which reduces model checking to propositional satisfiability. We show that bounded LTL model checking can be done without a tableau construction. We have implemented a model checker **BMC**, based on bounded model checking, and preliminary results are presented.

Bounded Model Checking

[BiereCimattiClarkeZhu-TACAS'99]





simple for <u>safety properties</u> p invariantly true

$$I(s_0) \wedge T(s_0, s_1)) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \bigvee_{i=0}^k \neg p(s_i)$$

harder for liveness properties properties

$$I(s_0) \wedge T(s_0, s_1)) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \bigwedge_{i=0}^k \neg p(s_i) \wedge \bigvee_{l=0}^k T(s_k, s_l)$$

compute and bound k by diameter



Google

Symbolic model checking without BDDs

Authors Armin Biere, Alessandro Cimatti, Edmund Clarke, Yunshan Zhu

- Publication date 1999/1/1
 - Book Tools and Algorithms for the Construction and Analysis of Systems
 - Pages 193-207
 - Publisher Springer Berlin Heidelberg
 - Description Abstract Symbolic Model Checking [3],[14] has proven to be a powerful technique for the verification of reactive systems. BDDs [2] have traditionally been used as a symbolic representation of the system. In this paper we show how boolean decision procedures, like Stälmarck's Method [16] or the Davis & Putnam Procedure [7], can replace BDDs. This new technique avoids the space blow up of BDDs, generates counterexamples much faster, and sometimes speeds up the verification. In addition, it produces counterexamples of minimal ...
 - Total citations Cited by 2076



1999 2000 2001 2002 2003 2004 2003 2006 2007 2006 2009 2010 2011 2012 2013 2014 2013

Scholar articles Symbolic model checking without BDDs A Biere, A Cimatti, E Clarke, Y Zhu - Tools and Algorithms for the Construction and Analysis ..., 1999 Cited by 2076 - Related articles - All 38 versions

[PDF] from cmu.edu

Replacing Testing with Formal Verification in Intel[®] Core[™] i7 Processor Execution Engine Validation

Roope Kaivola, Rajnish Ghughal, Naren Narasimhan, Amber Telfer, Jesse Whittemore, Sudhindra Pandav, Anna Slobodová, Christopher Taylor, Vladimir Frolov, Erik Reeber, and Armaghan Naik

Intel Corporation, JF4-451, 2111 NE 25th Avenue, Hillsboro, OR 97124, USA

Abstract. Formal verification of arithmetic datapaths has been part of the established methodology for most Intel processor designs over the last years, usually in the role of supplementing more traditional coverage oriented testing activities. For the recent Intel[®] Core[™] i7 design we took a step further and used formal verification as the primary validation vehicle for the core execution cluster, the component responsible for the functional behaviour of all microinstructions. We applied symbolic simulation based formal verification techniques for full datapath, control and state validation for the cluster, and dropped coverage driven testing entirely. The project, involving some twenty person years of verification work, is one of the most ambitious formal verification efforts in the hardware industry to date. Our experiences show that under the right circumstances, full formal verification of a design component is a feasible, industrially viable and competitive validation approach.

Introduction 1



තර (D D) වැඩි වැඩි වැඩි හැකින් කරන්න කරන The second s network of the test at a fast at a fast at a second of the second s

6 Formal Verification Value Proposition

cł

The conventional avisdoin about formal verification in industrial context is easy to spell OUTLEDED States and an and a state and a (i) (i) (ii) (iii) (i mal possibilitation and an advantation of a contraction for the second s

a vive execution of the statement of the	
in the second second state of second s	
+ Y may	
	Malap. Usual
A VALVENUE AND A DATA TATATATATATA IS IN A VALVENT A DATATATATATATATATATATATATATATATATATA	ten den etne bischen benden bischen er
of the underlevining gratabatastation on a why his to brit it erecercicos	
State or	
and the second statistic in the index is the index of the international second s	
and a strategiest and a factor of the sector	up plates and
11111111111111111111111111111111111111	
The second	COL
ners of the design, and formal verification	a subsough
and an and the second and a second and as	s that doing a
The statistic field of the statistic field of the statistic design of the statistic statis	ough effort only
a few, in both cases leading to a perceived low return on investment.	The areas where
projects have routinely chosen to do formal verification have then been	limited to those
where an uncaught problem would be so visible and costly that the extr	a effort of doing
formal verification can be justified. As a positive exception, SAT-based	bounded model
checking has been very successfully used as a bug-hunting tool in targe	eted areas.
The third usage model, mixing formal and dynamic techniques on	validating a sin-
gle desire a seconda appealing at a account at the bollow	ménénet atal
province in the second se	at at a fait in the state of th
Thus no warman and a laborated and a second state of the second st	and the second se

aspects of the design and the sets of interesting cases to that the expension sets of

Impact of BMC

- widespread use in industry (EDA)
 - industry embraced bounding part immediately
 - original *industrial* reservations: using SAT vs ATPG
 - original *academic* reservations: incompleteness?
- BMC relies on efficient SAT (SMT) solving
 - breakthroughs in SAT: CDCL '96, VSIDS '01, ...
 - encouraged investment in SAT / SMT research
- extensions to non-boolean domains and SW
 - bounding reduces complexity / decidability
- extensions to completeness
 - diameter checking, k-induction, interpolation





A Short Story on 15 years of Bounded Model Checking

- •1997: interest and capacity of BDDs stalled but there were success stories of other techniques
- Ed Clarke hired Yunshan Zhu & Armin Biere as Post-Docs: Use SAT for Symbolic Model Checking!
- •struggled for 10 months to come up with something that could replace / improve BDDs (mainly looked at QBF then)
- •Alessandro Cimatti came to an AI conference in Pittsburgh and at lunch (at an Indian Restaurant) we realized, that in
- Al Planing they do not care about completeness

What if we apply this to model checking? How to handle temporal logic?

• After one afternoon for the theory and 3 months of implementation and benchmarking later: *TACAS submission*



AWARD

Most influential paper in the first 20 years of TACAS



The Steering Committee of TACAS



SAT Based Model Checking

- BMC
- k-induction
- Abstractions / CEGAR
- Interpolation
- IC3

Abstract Modern satisfiability (SAT) solvers have become the enabling technology of many Model Checkers. In this chapter, we will focus on those techniques most relevant to industrial practice. In *Bounded Model Checking* (BMC), a transition system and a property are jointly unwound for a given number k of steps to obtain a formula that is satisfiable if there is a counterexample for the property up to length k. The formula is then passed to an efficient SAT solver. The strength of BMC is *refutation*: BMC has been used to discover subtle flaws in digital systems. We cover the application of BMC to both hardware and software systems, and to hardware/software co-verification. We also discuss means to make BMC complete, including k-induction, Craig interpolation, abstraction refinement techniques and inductive techniques with iterative strengthening.

SAT Based Model Checking *Armin Biere, Daniel Kröning* Handbook of Model Checking Edmund Clarke, Thomas Henzinger, Helmut Veith, *editors*



Lessons from BMC

- simple but useful ideas are very controversial
 - hard to get accepted (literally)
 - many comments of the sort: we did this before ...
 - main points: make it work, show that it works!
- in retrospective
 - classification considerations might have been useful since we tried to use SAT for symbolic model checking without taking Savitch's theorem into account
 - but might have prevented us going along that route ...



■ P

- problems with polynonmially time-bounded algorithms
- bounds measured in terms of input (file) size
- NP
 - same as P but with non-determininistic choice
 - needs a SAT solver
- PSPACE
 - as P but **space**-bounded
 - QBF and bit-level model checking fall in this class
- NEXPTIME
 - same as NP but with exponential time
- $\blacksquare \ \mathsf{P} \ \subseteq \ \mathsf{NP} \ \subseteq \ \mathsf{PSPACE} \ \subseteq \ \mathsf{NEXPTIME}$
 - usually it is assumed: $P \neq NP$
 - it is further known: NP \neq NEXPTIME



- NP problems
 - anything which can be (polynomially) encoded into SAT
 - combinational equivalence checking, bounded model checking
- PSPACE problems
 - <u>anything</u> which can be encoded (polynomially) into QBF
 - or into (bit-level) symbolic model checking
 - sequential equivalence checking, combinational synthesis or bounded games
- NEXPTIME problems
 - <u>anything</u> which can be encoded **exponentially** into SAT
 - first-order logic Bernays-Schönfinkel class (**EPR**): no functions, $\exists^*\forall^*$ prefix
 - QBF with explicit dependencies (Henkin Quantifiers): DQBF
 - partial observation games, black-box bounded model checking
 - bit-vector logics: QF_BV

NEXPTIME Completeness of Bit-Vectors

joined work with Gergely Kovásznai and Andreas Fröhlich

- QF_BV contained in NEXPTIME
 - bit-blast (exponential)

(set-logic QF_BV)

(assert (= z (bvadd x y)))

- give resulting formula to SAT solver
- we showed QF_BV is NEXPTIME hard by reducing DQBF to QF_BV

 $\forall x_0, x_1, x_2, x_3, x_4 \exists e_0(x_0, x_1, x_2, x_3), e_1(x_1, x_2, x_3, x_4) \varphi$

- polynomially encodes dependencies (for Henkin quantiers)
- my student Andreas has now an (yet unpublished) direct proof
- why are bit-vectors NEXPTIME complete?

(declare-fun x () (_ BitVec 1000000))
(declare-fun y () (_ BitVec 1000000))
(declare-fun z () (_ BitVec 1000000))

x, *y* : **bool**[1000000]

 $y \neq x \land x + y = x \ll 1$

```
(assert (= z (bvshl x (_ bv1 1000000))))
(assert (distinct x y))
Model Checking SA
```

- NP complete: QF_BV_{bw}
 - relate <u>same</u> bits: equality and all bit-wise operators
 - similar to well-known Ackermann reduction
- PSPACE complete: QF_BV_{bw},<<1</p>
 - only allow operators which relate <u>neighbouring</u> bits:
 - base operators: equality, inequality/comparison, bit-wise ops, shift-by-one
 - extended operators: addition, multiplication by constants, single-bit-slices etc.
 - encode in symbolic model checking logarithmically in bit-width
- see our CSR'12, SMT'13 papers and our 2015 journal article in TOCS
- came accross otherwise unsolvable benchmarks from industry!

MODULE main VAR c : boolean; -- carry 'bvadd x y' -- carry 'bvadd y x' d : boolean; -- x0, x1, ... x : boolean; -- y0, y1, ... y : boolean; ASSIGN init (c) := FALSE; init (d) := FALSE; ASSIGN next (c) := c&x | c&y | x&y; -- c + x + y >= 2next (d) := $d_{ky} | d_{kx} | y_{kx}; -- d + y + x >= 2$ DEFINE o := c != (x != y);-- C xor y xor x p := d != (y != x);-- d xor x xor y SPEC AG (o = p)

19/34

Commutativity of Bit-Vector Addition in AIGER



Ripple-Carry-Adder vs Carry-Save-Adder

not really realistic example but shows the fundamental problem of checking arithmetic circuit equivalence



J⊻U

Model Checking, SAT and Bit-Vectors KTH Stockholm

(set- (dec) (dec) (asse (chec)	-logic QF lare-fun z lare-fun y ert (dist: ck-sat)	_BV) x () (_ BitVe y () (_ BitVe inct (bvmul :	ec 12)) ec 12)) x y) (bvmul y x)))	
,	,		12 core	
	1 core	1 core	cube-and-conquer	12 core
bits	Glucose	Lingeling	March iLingeling	Treengeling
01	0.00	0.00	0.00	0.01
02	0.00	0.00	0.00	0.01
03	0.00	0.00	0.00	0.01
04	0.00	0.00	0.02	0.03
05	0.00	0.01	0.05	0.13
06	0.02	0.03	0.36	0.31
07	0.14	0.27	0.63	0.72
08	1.18	1.98	1.38	2.47
09	7.85	10.98	2.63	4.65
10	37.16	41.49	5.02	10.86
11	147.62	214.98	15.72	21.96
12	833.62	649.49	56.57	61.48
13			238.10	263.44

limit of 900 seconds wall clock time

- secret of the success of (combinational) equivalence checking
 - assumption: many internal equivalence points
 - makes BDD and SAT sweeping effective
- problems with arithmetic circuits
 - almost no equivalent internal signals (except for outputs)
 - proof complexity conjectured to be beyond resolution
 - often no "clean" implementation circuit available
- challenges
 - prove conjectured complexity
 - use world-level (bit-vector) information
 - arithmetic reasoning on the bit-level
 - robust integration in SAT and/or SMT solver
- started to collect a large number of such benchmarks





Results of the SAT competition/race winners on the SAT 2009 application benchmarks, 20mn timeout

from Daniel Le Berre



J⊻U

Model Checking, SAT and Bit-Vectors KTH Stockholm

Satisfiability (SAT) related topics have attracted researchers from various disciplines. Logic, applied areas such as planning, scheduling, operations research and combinatorial optimization, but also theoretical issues on the theme of complexity, and much more, they all are connected through SAT.

My personal interest in SAT stems from actual solving: The increase in power of modern SAT solvers over the past 15 years has been phenomenal. It has become the key enabling technology in automated verification of both computer hardware and software. Bounded Model Checking (BMC) of computer hardware is now probably the most widely used model checking technique. The counterexamples that it finds are just satisfying instances of a Boolean formula obtained by unwinding to some fixed depth a sequential circuit and its specification in linear temporal logic. Extending model checking to software verification is a much more difficult problem on the frontier of current research. One promising approach for languages like C with finite word-length integers is to use the same idea as in BMC but with a decision procedure for the theory of bit-vectors instead of SAT. All decision procedures for bit-vectors that I am familiar with ultimately make use of a fast SAT solver to handle complex formulas.

Decision procedures for more complicated theories, like linear real and integer arithmetic, are also used in program verification. Most of them use powerful SAT solvers in an essential way.

Clearly, efficient SAT solving is a key technology for 21st century computer science. I expect this collection of papers on all theoretical and practical aspects of SAT solving will be extremely useful to both students and researchers and will lead to many further advances in the field.

Edmund Clarke

Edmund M. Clarke, FORE Systems University Professor of Computer Science and Professor of Electrical and Computer Engineering at Cornegie Mellon University, is one of the initiators and main contributors to the field of Model Checking, for which he also received the 2007 ACM Turing Award.

In the late 90s Professor Clarke was one of the first researchers to realize that SAT solving has the potential to become one of the most important technologies in model checking.



Editors:
Armin Biere
Marijn Heule
Marijn Walsh

HANDBOOK

of satisfiability

Editors:

Armin Biere

Marijn Heule

Hans van Maaren

Toby Walsh

IOS Press **IOS** Press Frontiers in Artificial Intelligence and Applications

HANDBOOK

•••••f satisfiability

Part I. Theory and Algorithms

🖹 🕹 😤 Armin Biere: Bounded Model Checking. 457-481 🖹 🕹 🤍 Jussi Rintanen: Planning and SAT. 483-504 🖹 🗄 🥰 🕅 Daniel Kroening: Software Verification. 505-532 🖹 🕹 👻 Hantao Zhang: Combinatorial Designs by SAT Solvers. 533-568 🖹 🕹 😤 Fabrizio Altarelli, Rémi Monasson, Guilhem Semerjian, Francesco Zamponi: Connections to Statistical Physics. 569-611 🖹 🗄 🔍 Chu Min Li, Felip Manyà: MaxSAT, Hard and Soft Constraints. 613-631 🖹 🕹 ᅉ Carla P. Gomes, Ashish Sabharwal, Bart Selman: Model Counting. 633-654 🖹 🕹 😤 Rolf Drechsler, Tommi A. Junttila, Ilkka Niemelä: Non-Clausal SAT and ATPG. 655-693 🖹 🕹 🧟 Olivier Roussel, Vasco M. Manquinho: Pseudo-Boolean and Cardinality Constraints. 695-733 ∃ ⊥ ♥ Hans Kleine Büning, Uwe Bubeck: Theory of Quantified Boolean Formulas. 735-760 🖹 🕹 🕅 Enrico Giunchiglia, Paolo Marin, Massimo Narizzano: Reasoning with Quantified Boolean Formulas. 761-780 🖹 立 🔍 Roberto Sebastiani, Armando Tacchella: SAT Techniques for Modal and Description Logics. 781-824 🖹 立 😤 Clark W. Barrett, Roberto Sebastiani, Sanjit A. Seshia, Cesare Tinelli: Satisfiability Modulo Theories. 825-885 🖹 凸 🔍 Stephen M. Majercik:

Stochastic Boolean Satisfiability. 887-925

- 🖹 凸 🔍 John Franco, John Martin: A History of Satisfiability. 3-74
- 目 凸 🔍 Steven David Prestwich: CNF Encodings. 75-97
- 🖹 🗄 😤 🤫 Adnan Darwiche, Knot Pipatsrisawat: Complete Algorithms. 99-130
- 🖹 🕹 ᅉ João P. Marques Silva, Inês Lynce, Sharad Malik: Conflict-Driven Clause Learning SAT Solvers. 131-153
- 🖹 🕹 👻 Marijn Heule, Hans van Maaren: Look-Ahead Based SAT Solvers, 155-184
- 🖹 🕹 👻 Henry A. Kautz, Ashish Sabharwal, Bart Selman: Incomplete Algorithms. 185-203
- 目 凸 🗘 Oliver Kullmann: Fundaments of Branching Heuristics. 205-244
- 🖹 🕹 🤍 Dimitris Achlioptas: Random Satisfiability. 245-270
- 🖹 立 🕅 Carla P. Gomes, Ashish Sabharwal: Exploiting Runtime Variation in Complete Solvers. 271-288
- 目 显 😪 Karem A. Sakallah: Symmetry and Satisfiability. 289-338
- 🖹 🕹 👻 Hans Kleine Büning, Oliver Kullmann: Minimal Unsatisfiability and Autarkies. 339-401
- 🖹 立 😌 Evgeny Dantsin, Edward A. Hirsch: Worst-Case Upper Bounds. 403-424
- 🖹 🕹 🗟 Marko Samer, Stefan Szeider: Fixed-Parameter Tractability. 425-454







File Edit View Document Tools Window Help

👆 5 / 318 💿 🖲 150% -

PREFACE V

Special thanks are due to Armin Biere, Randy Bryant, Sam Buss, Niklas Eén, Ian Gent, Marijn Heule, Holger Hoos, Svante Janson, Peter Jeavons, Daniel Kroening, Oliver Kullmann, Massimo Lauria, Wes Pegden, Will Shortz, Carsten Sinz, Niklas Sörensson, Udo Wermuth, Ryan Williams, and ... for their detailed comments on my early attempts at exposition, as well as to numerous other correspondents who have contributed crucial corrections. Thanks also to Stanford's Information Systems Laboratory for providing extra computer power when my laptop machine was inadequate.

++

Find

* * *

Wow — Section 7.2.2.2 has turned out to be the longest section, by far, in The Art of Computer Programming. The SAT problem is evidently a "killer app," because it is key to the solution of so many other problems. Consequently I can only hope that my lengthy treatment does not also kill off my faithful readers! As I wrote this material, one topic always seemed to flow naturally into another, so there was no neat way to break this section up into separate subsections. (And anyway the format of TAOCP doesn't allow for a Section 7.2.2.2.1.)

I've tried to ameliorate the reader's navigation problem by adding subheadings at the top of each right-hand page. Furthermore, as in other sections, the exercises appear in an order that roughly parallels the order in which corresponding topics are taken up in the text. Numerous cross-references are provided Biere Bryant Buss Eén Gent Heule Hoos Janson Jeavons Kroening Kullmann Lauria Pegden Shortz Sinz Sörensson Wermuth Williams Internet MPR Internet

х

- competitions are used to
 - compare and evaluate implementations and algorithms
 - generate benchmarks used in papers
- SAT competition is one of the largest competitions
 - many solvers, highly competitive
 - portfolio solving, over-tuning issues
 - benchmark selection scheme broken due to competing goals:
 - assess the state-of-the-art
 - high-light new ideas
 - give a fair chance to everybody
- research in SAT solving, verification, etc. in essence empirical science
 - benchmark selection critical
 - how to select benchmarks?
 - for the competition?
 - in your papers?

Conclusion

- what I did not talk about ... (yet)
 - parallel SAT
 - QBF / quantifiers in general
 - huge improvements in local research in recent years
 - how to apply local search to bit-vectors and SMT
 - testing / debugging
 - assertion synthesis
- acknowledgements:

Ed Clarke, all co-authors, collaborators, students and Post-Docs and if would list more names I would struggle with order and probably forget somebody

■ if you have model checking, SMT, or SAT problems you want share let me know ...

looking for Post-Doc's and PhD students too

34/3

Model Checking, SAT and Bit-Vectors KTH Stockholm