Information Processing Letters 135 (2018) 57–61

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

On the complexity of the quantified bit-vector arithmetic with binary encoding

M. Jonáš *, J. Strejček

Faculty of Informatics, Masaryk University, Botanická 68a, 602 00, Brno, Czech Republic

ARTICLE INFO

Article history:
Received 5 September 2017
Received in revised form 21 February 2018
Accepted 23 February 2018
Communicated by Krishnendu Chatterjee

Keywords:
Computational complexity
Satisfiability modulo theories
Fixed-size bit-vectors

ABSTRACT

We study the precise computational complexity of deciding satisfiability of first-order quantified formulas over the theory of fixed-size bit-vectors with binary-encoded bit-widths and constants. This problem is known to be in **EXPSPACE** and to be **NEXPTIME**-hard. We show that this problem is complete for the complexity class **AEXP**(poly) – the class of problems decidable by an alternating Turing machine using exponential time, but only a polynomial number of alternations between existential and universal states.

1. Introduction

The first-order theory of fixed-size bit-vectors is widely used for describing properties of software and hardware. Although most current applications use only the quantifier-free fragment of this logic, there are several use cases that benefit from using bit-vector formulas containing quantifiers [1–5]. Consequently, computational complexity of quantified bit-vector logic has been investigated in recent years. It has been shown that deciding satisfiability of quantified bit-vector formulas is **PSPACE**-complete and it becomes **NEXPTIME**-complete when uninterpreted functions are allowed in addition to quantifiers [6].

However, these results suppose that all scalars in the formula are represented in the unary encoding, which is not the case in practice, because in most of real-world applications, bit-widths and constants are encoded logarithmically. For example, the format SMT-LIB [7], which is an input format for most of the state-of-the-art SMT solvers, represents all scalar values as decimal numbers. Such representation can be exponentially more succinct than the representation using unary-encoded scalars. The satisfiability problem for bit-vector formulas with binary-encoded scalars has been recently investigated by Kovásznai et al. [8]. They have shown that the satisfiability of quantified bit-vector formulas with binary-encoded scalars and with uninterpreted functions is **2-NEXPTIME**-complete. The situation for the same problem without uninterpreted functions is not so clear: deciding satisfiability of quantified bit-vector formulas with binary encoded scalars and without uninterpreted functions (we denote this problem as BV2 satisfiability) is known to be in **EXPSPACE** and to be **NEXPTIME**-hard, but its precise complexity has remained unknown [8].

In this paper, we solve this open problem by identifying the complexity class for which BV2 satisfiability is complete. We use the notion of an alternating Turing machine introduced by Chandra et al. [9] and show that the BV2 satisfiability problem is complete for the class **AEXP**(poly) of problems solvable by an alternating Turing machine using exponential time, but only a polynomial number of alternations.

* Corresponding author.
E-mail addresses: martin.jonas@mail.muni.cz (M. Jonáš), strejcek@fi.muni.cz (J. Strejček).

https://doi.org/10.1016/j.ipl.2018.02.018
0020-0190/

Table 1
Recursive definition of the formula size. Operations include logical connectives, function symbols, and predicate symbols. Each t_i denotes a subterm or a subformula, each i_j denotes a scalar argument of an operation, and $Q \in \{\exists, \forall\}$ [8].

	Expression	Size
Constant	$\lvert c^{[n]}\rvert$	$L(c) + L(n)$
Variable	$\lvert x^{[n]}\rvert$	$1 + L(n)$
Operation	$\lvert o(t_1, \ldots, t_k, i_1, \ldots, i_p)\rvert$	$1 + \sum_{1 \le i \le k} \lvert t_i\rvert + \sum_{1 \le j \le p} L(i_j)$
Quantifier	$\lvert Q\, x^{[n]} \varphi\rvert$	$\lvert x^{[n]}\rvert + \lvert\varphi\rvert$

2. Quantified bit-vector formulas

The *theory of fixed-size bit-vectors* (BV or *bit-vector theory* for short) is a many-sorted first-order theory with infinitely many sorts corresponding to bit-vectors of various lengths. Each bit-vector variable has an explicitly assigned sort, e.g. $x^{[3]}$ is a bit-vector variable of bit-width 3. The BV theory uses only three predicates, namely equality ($=$), unsigned inequality of binary-encoded non-negative integers ($\le_u$), and signed inequality of integers in 2's complement representation ($\le_s$). The signature also contains constants $c^{[n]}$ for each $n \ge 1$ and $0 \le c \le 2^n - 1$, and various interpreted functions, namely addition ($+$), multiplication ($*$), unsigned division ($\div$), bitwise negation ($\sim$), bitwise and ($\&$), bitwise or ($|$), bitwise exclusive or ($\oplus$), left-shift ($\ll$), right-shift ($\gg$), concatenation ($\cdot$), and extraction of a subword starting at the position i and ending at the position j ($extract(\_, i, j)$). Although various sources define the full bit-vector theory with different sets of functions, all such definitions can be polynomially reduced to each other [8]. All numbers occurring in the formula, i.e. values of constants, bit-widths and bounds i, j of extraction, are called *scalars*.

There are more ways to encode scalars occurring in the bit-vector formula: in the *unary encoding* or in a *logarithmic encoding*. In this paper, we focus only on formulas using the *binary encoding*. This covers all logarithmic encodings, since all of them are polynomially reducible to each other. In the binary encoding, $L(n)$ bits are needed to express the number n, where $L(0) = 1$ and $L(n) = \lfloor \log_2 n \rfloor + 1$ for all $n > 0$. The entire formula is encoded in the following way: each constant $c^{[n]}$ has both its value c and bit-width n encoded in binary, each variable $x^{[n]}$ has its bit-width n encoded in binary, and all scalar arguments of functions are encoded in binary. The size of the formula φ is denoted $|\varphi|$. The recursive definition of $|\varphi|$ is given in Table 1. For quantified formulas with binary-encoded scalars, we define the corresponding satisfiability problem:

Definition 1 *([8]).* The BV2 *satisfiability problem* is to decide satisfiability of a given closed quantified bit-vector formula with all scalars encoded in binary.

Similarly to Kovásznai et al. [8], we use an *indexing* operation, which is a special case of the extraction operation that produces only a single bit. In particular, for a term $t^{[n]}$ and a number $0 \le i < n$, the indexing operation $t^{[n]}[i]$ is defined as $extract(t^{[n]}, i, i)$. We assume that bits of bit-vectors are indexed from the least significant. For example, given a bit-vector variable $x^{[6]} = x_5 x_4 x_3 x_2 x_1 x_0$, the value of $x^{[6]}[1]$ refers to x_1. In the following, we use a more general version of the indexing operation, in which the index can be an arbitrary bit-vector term, not only a fixed scalar. This operation can be defined using the indexing operation and the bit-shift operation with only a linear increase in the size of the term:

$$t^{[n]}[s^{[n]}] \stackrel{\text{df}}{\equiv} (t^{[n]} \gg s^{[n]})[0].$$

3. Alternation complexity

We assume a basic familiarity with an *alternating Turing machine* (ATM) introduced by Chandra, Kozen, and Stockmeyer [9], and basic concepts from the complexity theory, which can be found for example in Kozen [10]. We recall that each state of an ATM is either *existential* or *universal*. Existential states behave like states of a non-deterministic Turing machine: a run passing through an existential state continues with one of the possible successors. In contrast to this, a run entering a universal state forks and continues into all possible successors. Hence, runs of an ATM are trees. Such a run is accepting if each branch of the run ends in an accepting state.

This section recalls some complexity classes related to alternating Turing machines. Computations in such complexity classes are bounded not only by time and memory, but also by the number of alternations between existential and universal states during the computation. Although bounding both time and memory is useful in some applications, in this paper we need only complexity classes related to ATMs that are bounded in time and the number of alternations. Therefore, the following definition introduces a family of complexity classes parameterized by the number of steps and alternations used by corresponding ATMs.

Definition 2. Let $t, g\colon \mathbb{N} \to \mathbb{N}$ be functions such that $g(n) \ge 1$. We define the complexity class $\mathbf{ATIME}(t, g)$ as the class of all problems A for which there is an alternating Turing machine that decides A and, for each input of length n, it needs at most $t(n)$ steps and $g(n) - 1$ alternations along every branch of every run. If T and G are classes of functions, let $\mathbf{ATIME}(T, G) = \bigcup_{t \in T, g \in G} \mathbf{ATIME}(t, g)$.

Chandra et al. have observed several relationships between classical complexity classes related to time and memory and the complexity classes defined by ATMs [9]. We recall relationships between alternating complexity classes and the classes **NEXPTIME** and **EXPSPACE**, which are important for this paper. It can easily be seen that the class **NEXPTIME** corresponds to all problems solvable by an alternating Turing machine that starts in an existential state and can use exponential time and no alternations: this yields an inclusion $\mathbf{NEXPTIME} \subseteq \mathbf{ATIME}(2^{\mathcal{O}(n)}, 1)$. On the other hand, results of Chandra et al. imply that **EXPSPACE** is precisely the complexity class $\mathbf{ATIME}(2^{n^{\mathcal{O}(1)}}, 2^{n^{\mathcal{O}(1)}})$ of problems solvable in exponential time and with exponential number of alternations. An interesting class that lies in between those two complexity

classes can be obtained by bounding the number of steps exponentially and the number of alternations polynomially. This class is called **AEXP**(poly).

Definition 3. $\mathbf{AEXP}(\text{poly}) \stackrel{\text{df}}{=} \mathbf{ATIME}(2^{n^{\mathcal{O}(1)}}, n^{\mathcal{O}(1)})$.

The following inclusions immediately follow from the mentioned results.

$$\mathbf{NEXPTIME} \subseteq \mathbf{AEXP}(\text{poly}) \subseteq \mathbf{EXPSPACE}$$

However, it is unknown whether any of the inclusions is strict.

4. Complexity of BV2 satisfiability

In this section, we show that the BV2 satisfiability problem is **AEXP**(poly)-complete. First, we prove that the problem is in the class **AEXP**(poly).

Theorem 1. *The* BV2 *satisfiability problem is in* **AEXP**(poly).

Proof. We describe the alternating Turing machine solving the problem. For a given BV2 formula φ, the machine first converts the formula to the prenex normal form, which can be done in polynomial time without any alternations [11]. The machine then assigns values to all existentially quantified variables using existential states and to all universally quantified variables using universal states. Although this requires exponential time, as there are exponentially many bits whose value has to be assigned, only a polynomial number of alternations is required, because the formula φ can contain only polynomially many quantifiers.

Finally, the machine uses the assignment to evaluate the quantifier-free part of the formula. If the result of the evaluation is true, the machine accepts; it rejects otherwise. The evaluation takes exponential time and no quantifier alternations: the machine replaces all variables by exponentially many previously assigned bits and computes results of all operations from the bottom of the syntactic tree of the formula up. The computation of each of the operations takes time polynomial in the number of bits, which is exponential. □

In the rest of this section, we show that the BV2 satisfiability problem is also **AEXP**(poly)-hard. In particular, we present a reduction of a known **AEXP**(poly)-hard *second-order Boolean formulas satisfiability problem* [12,13] to the BV2 satisfiability.

Intuitively, the *second-order Boolean logic* (SO_2) can be obtained from a quantified Boolean logic by adding function symbols and quantification over such symbols. Alternatively, the SO_2 logic corresponds to the second-order predicate logic restricted to the domain $\{0, 1\}$. Lohrey and Lück have shown that by bounding the number of quantifier alternations in second-order Boolean formulas, problems complete for all levels of the exponential hierarchy can be obtained. Moreover, if the number of quantifier alternations is unbounded, the problem of deciding satisfiability of quantified second-order Boolean formulas is **AEXP**(poly)-complete [12,13].

We now introduce the SO_2 logic more formally. The definitions of the syntax and semantics of SO_2 used in this paper are due to Hannula et al. [14].

Definition 4 (SO_2 *syntax [14]*). Let $\mathcal{F}$ be a countable set of function symbols, where each symbol $f \in \mathcal{F}$ is given an arity $\text{ar}(f) \in \mathbb{N}_0$. The set $\mathsf{SO}_2(\mathcal{F})$ of *quantified Boolean second-order formulas* is defined inductively as

$$\varphi ::= \varphi \wedge \varphi \mid \neg\varphi \mid \exists f\varphi \mid \forall f\varphi \mid f(\underbrace{\varphi, \ldots, \varphi}_{\text{ar}(f) \text{ times}}),$$

where $f \in \mathcal{F}$.

Definition 5 (SO_2 *semantics [14]*). An $\mathcal{F}$*-interpretation* is a function $\mathcal{I}$ that assigns to each symbol $f \in \mathcal{F}$ a Boolean function of the corresponding arity, i.e. $\mathcal{I}(f)\colon \{0,1\}^{\text{ar}(f)} \to \{0,1\}$ for each $f \in \mathcal{F}$. The valuation of a formula $\varphi \in \mathsf{SO}_2(\mathcal{F})$ in $\mathcal{I}$, written $[\![\varphi]\!]_{\mathcal{I}}$, is defined recursively as

$$[\![\varphi \wedge \psi]\!]_{\mathcal{I}} = [\![\varphi]\!]_{\mathcal{I}} * [\![\psi]\!]_{\mathcal{I}},$$
$$[\![\neg\varphi]\!]_{\mathcal{I}} = 1 - [\![\varphi]\!]_{\mathcal{I}},$$
$$[\![f(\varphi_1, \ldots, \varphi_n)]\!]_{\mathcal{I}} = \mathcal{I}(f)([\![\varphi_1]\!]_{\mathcal{I}}, \ldots, [\![\varphi_n]\!]_{\mathcal{I}}),$$
$$[\![\exists f\varphi]\!]_{\mathcal{I}} = \max\left\{ [\![\varphi]\!]_{\mathcal{I}[f\mapsto F]} \mid F\colon \{0,1\}^{\text{ar}(f)} \to \{0,1\} \right\},$$
$$[\![\forall f\varphi]\!]_{\mathcal{I}} = \min\left\{ [\![\varphi]\!]_{\mathcal{I}[f\mapsto F]} \mid F\colon \{0,1\}^{\text{ar}(f)} \to \{0,1\} \right\},$$

where $\mathcal{I}[f \mapsto F]$ is the function defined as $\mathcal{I}[f \mapsto F](f) = F$ and $\mathcal{I}[f \mapsto F](g) = \mathcal{I}(g)$ for all $g \neq f$.

An SO_2 formula φ is *satisfiable* if $[\![\varphi]\!]_{\mathcal{I}} = 1$ for some $\mathcal{I}$.

We call function symbols of arity 0 *propositions* and all other function symbols *proper functions*. An SO_2 formula φ is in the *prenex normal form* if it has the form $\overline{Q}\psi$, where $\overline{Q}$ is a sequence of quantifiers called a *quantifier prefix*, ψ is a quantifier-free formula called a *matrix*, and all proper functions are quantified before propositions. In the following, we fix an arbitrary countable set of function symbols $\mathcal{F}$ and instead of $\mathsf{SO}_2(\mathcal{F})$, we write only SO_2.

Definition 6. The SO_2 *satisfiability problem* is to decide whether a given closed SO_2 formula in the prenex normal form is satisfiable.

Theorem 2 (*[12,13]*). *The* SO_2 *satisfiability problem is* **AEXP**(poly)*-complete.*

We now show a polynomial time reduction of SO_2 satisfiability to BV2 satisfiability and thus finish the main claim of this paper, which states that the BV2 satisfiability problem is **AEXP**(poly)-complete.

Theorem 3. *The* BV2 *satisfiability problem is* **AEXP**(poly)*-hard.*

Proof. We present a polynomial time reduction of SO_2 satisfiability to BV2 satisfiability. Let φ be an SO_2 formula with a quantifier prefix $\overline{Q}$ and a matrix ψ, i.e. $\varphi = \overline{Q}\psi$

Table 2
Completeness results for various bit-vector logics and encodings. This is the table presented by Fröhlich et al. [15] extended by the result proved in this paper.

Encoding	Quantifiers			
	No		Yes	
	Uninterpreted functions		Uninterpreted functions	
	No	Yes	No	Yes
Unary	**NP**	**NP**	**PSPACE**	**NEXPTIME**
Binary	**NEXPTIME**	**NEXPTIME**	**AEXP**(poly)	**2-NEXPTIME**

where ψ is a quantifier-free formula. We construct a bit-vector formula φ^{BV}, such that φ is satisfiable iff the formula φ^{BV} is satisfiable.

In the formula φ^{BV}, each function symbol f of the formula φ is represented by a bit-vector variable x_f of bit-width $2^{\mathrm{ar}(f)}$. Intuitively, the bits of the variable x_f will encode values $f(b_{n-1},\ldots,b_0)$ for all possible inputs $b_0,\ldots,b_{n-1}\in\{0,1\}$. In particular, the value $f(b_{n-1},\ldots,b_0)$ is represented as the bit on the index $\sum_{i=0}^{n-1}(2^i b_i)$ in the bit-vector x_f. Equivalently, this index can be expressed as the numerical value of the bit-vector $b_{n-1}b_{n-2}\ldots b_0$. For example, for a ternary function symbol f, bits of the bit-vector value $x_f = x_7x_6x_5x_4x_3x_2x_1x_0$ will represent values $f(1,1,1)$, $f(1,1,0)$, $f(1,0,1)$, $f(1,0,0)$, $f(0,1,1)$, $f(0,1,0)$, $f(0,0,1)$, and $f(0,0,0)$, respectively.

The reduction proceeds in two steps. First, we inductively construct a bit-vector term ψ^{BV} of bit-width 1, which corresponds to the formula ψ:

- If $\psi \equiv \rho_1 \wedge \rho_2$, we set $\psi^{BV} \equiv \rho_1^{BV}$ & ρ_2^{BV}.
- If $\psi \equiv \neg\rho$, we set $\psi^{BV} \equiv {\sim}\rho^{BV}$.
- If $\psi \equiv f()$ (i.e. f is a proposition), we set $\psi^{BV} \equiv x_f^{[1]}$.
- If $\psi \equiv f(\rho_{n-1},\ldots,\rho_0)$ where $n = \mathrm{ar}(f)$, we set

$$\psi^{BV} \equiv x_f^{[2^n]}\left[0^{[2^n-n]}\cdot\rho_{n-1}^{BV}\cdot\rho_{n-2}^{BV}\cdot\ldots\cdot\rho_0^{BV}\right].$$

Note that because both arguments of the indexing operation have to be of the same sort, $2^n - n$ additional bits have to be added to the index term to get a term of the same bit-width as the term $x_f^{[2^n]}$.

In the second step, we replace each quantifier $Q_i f$ in the quantifier prefix $\overline{Q}$ by a bit-vector quantifier $Q_i x_f^{[2^n]}$, where $n = \mathrm{ar}(f)$, and thus obtain a sequence of bit-vector quantifiers $\overline{Q}^{BV}$. The final formula φ^{BV} is then $\overline{Q}^{BV}(\psi^{BV} = 1^{[1]})$.

Due to the binary representation of the bit-widths, the formula φ^{BV} is polynomial in the size of the formula φ. □

Example 1. Consider an SO_2 formula

$$\exists f \forall p \forall q\,.\,\neg f(p,p,q) \wedge f(p,q\wedge\neg q,q),$$

where f is a ternary function symbol and p, q are propositions. Then the result of the described reduction is the formula

$$\exists x_f^{[8]} \forall x_p^{[1]} \forall x_q^{[1]}\ (\ {\sim}x_f^{[8]}[0^{[5]}\cdot x_p\cdot x_p\cdot x_q]\ \&\ x_f^{[8]}[0^{[5]}\cdot x_p\cdot(x_q\ \&\ {\sim}x_q)\cdot x_q]\ =\ 1^{[1]}).$$

Corollary 1. *The* BV2 *satisfiability problem is* **AEXP**(poly)*-complete.*

5. Conclusions

We have identified the precise complexity class of deciding satisfiability of a quantified bit-vector formula with binary-encoded bit-widths. This paper shows that the problem is complete for the complexity class **AEXP**(poly), which is the class of all problems solvable by an alternating Turing machine that can use exponential time and a polynomial number of alternations. This result settles the open question raised by Kovásznai et al. [8]. Known completeness results for various bit-vector logics including the result proven in this paper are summarized in Table 2.

Acknowledgements

Authors of this work are supported by the Czech Science Foundation, project No. GBP202/12/G061.

References

[1] Sumit Gulwani, Saurabh Srivastava, Ramarathnam Venkatesan, Constraint-based invariant inference over predicate abstraction, in: Verification, Model Checking, and Abstract Interpretation, Proceedings of the 10th International Conference, VMCAI 2009, Savannah, GA, USA, January 18–20, 2009, 2009, pp. 120–135.
[2] Saurabh Srivastava, Sumit Gulwani, Jeffrey S. Foster, From program verification to program synthesis, in: Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17–23, 2010, 2010, pp. 313–326.
[3] Byron Cook, Daniel Kroening, Philipp Rümmer, Christoph M. Wintersteiger, Ranking function synthesis for bit-vector relations, Form. Methods Syst. Des. 43 (1) (2013) 93–120.
[4] Daniel Kroening, Matt Lewis, Georg Weissenbacher, Under-approximating loops in C programs for fast counterexample detection, in: Computer Aided Verification – 25th International Conference, CAV 2013, in: Lect. Notes Comput. Sci., vol. 8044, Springer, 2013, pp. 381–396.
[5] Jan Mrázek, Petr Bauch, Henrich Lauko, Jiří Barnat SymDIVINE, Tool for control-explicit data-symbolic state space exploration, in: Model Checking Software – Proceedings of the 23rd International Symposium, SPIN 2016, Co-located with ETAPS 2016, Eindhoven, The Netherlands, April 7–8, 2016, 2016, pp. 208–213.
[6] Christoph M. Wintersteiger, Youssef Hamadi, Leonardo de Moura, Efficiently solving quantified bit-vector formulas, Form. Methods Syst. Des. 42 (1) (2013) 3–23.
[7] Clark Barrett, Pascal Fontaine, Cesare Tinelli, The SMT-LIB Standard: Version 2.5, Technical report, Department of Computer Science, The University of Iowa, 2015, available at http://www.SMT-LIB.org.
[8] Gergely Kovásznai, Andreas Fröhlich, Armin Biere, Complexity of fixed-size bit-vector logics, Theory Comput. Syst. 59 (2) (2016) 323–376.
[9] Ashok K. Chandra, Dexter Kozen, Larry J. Stockmeyer, Alternation, J. ACM 28 (1) (1981) 114–133.
[10] Dexter Kozen, Theory of Computation, Texts Comput. Sci., Springer, 2006.
[11] John Harrison, Handbook of Practical Logic and Automated Reasoning, Cambridge University Press, 2009.
[12] Markus Lohrey, Model-checking hierarchical structures, J. Comput. Syst. Sci. 78 (2) (2012) 461–490.

[13] Martin Lück, Complete problems of propositional logic for the exponential hierarchy, arXiv:1602.03050 [cs.CC], 2016. (Accessed July 2017).

[14] Miika Hannula, Juha Kontinen, Martin Lück, Jonni Virtema, On quantified propositional logics and the exponential time hierarchy, in: Proceedings of the Seventh International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2016, Catania, Italy, 14–16 September 2016, 2016, pp. 198–212.

[15] Andreas Fröhlich, Gergely Kovásznai, Armin Biere, More on the complexity of quantifier-free fixed-size bit-vector logics with binary encoding, in: Computer Science – Theory and Applications – Proceedings of the 8th International Computer Science Symposium in Russia, CSR 2013, Ekaterinburg, Russia, June 25–29, 2013, 2013, pp. 378–390.