



CoqQFBV: A Scalable Certified SMT Quantifier-Free Bit-Vector Solver

Xiaomu Shi¹, Yu-Fu Fu², Jiayang Liu¹(✉), Ming-Hsien Tsai³,
Bow-Yaw Wang³, and Bo-Yin Yang³

¹ Shenzhen University, Shenzhen, China

² Georgia Institute of Technology, Atlanta, USA

³ Academia Sinica, Taipei City, Taiwan



Abstract. We present a certified SMT QF_BV solver CoqQFBV built from a verified bit blasting algorithm, KISSAT, and the verified SAT certificate checker GRATCHK in this paper. Our verified bit blasting algorithm supports the full QF_BV logic of SMT-LIB; it is specified and formally verified in the proof assistant Coq. We compare CoqQFBV with CVC4, BITWUZLA, and BOLECTOR on benchmarks from the QF_BV division of the single query track in the 2020 SMT Competition, and real-world cryptographic program verification problems. CoqQFBV surprisingly solves more program verification problems with certification than the 2020 SMT QF_BV division winner BITWUZLA without certification.

1 Introduction

Satisfiability Modulo Theories (SMT) solvers for the Quantifier-Free Bit-Vector (QF_BV) logic have been used to verify programs with bit-level accuracy [9, 10]. In such applications, a program verification problem is reformulated as an SMT QF_BV query. An SMT QF_BV solver is then invoked to compute a query result. The query result in turn decides the answer to the program verification problem. For cryptographic assembly programs, a missing carry or borrow flag will result in incorrect computation. Bit-accurate verification is thus necessary for cryptographic programs. SMT QF_BV solvers in fact have been employed to verify such programs [8, 25]. These solvers nonetheless are very complex programs with possibly unknown bugs [7, 18]. Since bugs in SMT QF_BV solvers may induce incorrect query results, program verification cannot be taken without a grain of salt when SMT QF_BV solvers are employed.

In order to check SMT QF_BV query results independently, SMT QF_BV solvers can generate certificates to validate their answers. In the LFSC certificates [14, 23], for instance, an SMT QF_BV query result is certified by correct bit blasting and Boolean Satisfiability (SAT) solving. Such certificates demonstrate that the SMT QF_BV query is reduced to a Boolean SAT query correctly *and* the corresponding SAT query is solved correctly. Although one can certify SAT query results with certificates from SAT solvers [24], it is not always easy to certify correct bit blasting due to complex arithmetic operations in SMT QF_BV

queries. Developing correct and efficient checkers for SMT QF_BV certificates can be very challenging. Indeed, an LFSC certificate checker based on the proof assistant COQ has been developed to improve confidence [12]. Yet the COQ-based certificate checker does not fully support arithmetic operations and thus cannot certify results of SMT QF_BV queries with complicated arithmetic operations. Consequently, the correctness of cryptographic programs still relies on the correctness of SMT QF_BV solvers or their unverified certificate checkers.

In this paper, we take a more direct approach to ensure the correctness of SMT QF_BV query results. Instead of certifying correct bit blasting for every SMT QF_BV query, we specify a bit blasting algorithm and prove its correctness in the proof assistant COQ. In order to formalize the correctness of our bit blasting algorithm, we develop a formal bit-vector theory in COQ. Naturally, the formal theory has to support all arithmetic functions (addition, subtraction, multiplication, division, and remainder) for both signed and unsigned representations as needed in SMT-LIB [3]. Based on our new bit-vector theory, we give a formal semantics for SMT QF_BV queries in COQ. Our semantics follows the SMT-LIB semantics carefully. Particularly, division and remainder are total arithmetic operations even when the divisor is zero. Using our COQ bit-vector theory and semantics, we prove that our bit blasting algorithm always returns a corresponding Boolean formula correctly on any SMT QF_BV query. Since our algorithm has been formally verified, bit blasting is always correct and need not be certified. Through the OCAML program extracted from our verified bit blasting algorithm, a corresponding SAT query is obtained for each SMT QF_BV query and sent to a SAT solver. A SAT certificate checker suffices to validate SAT query results and hence the correctness of answers to SMT QF_BV queries. Since neither complicated SMT QF_BV solvers nor their certificate checkers are trusted, our work can improve the confidence of SMT QF_BV query results.

To our knowledge, our bit-vector theory is the first COQ formalization designed for bit blasting queries from the QF_BV logic of SMT-LIB. Our semantics is the first COQ formalization for full SMT QF_BV queries. We are not aware of any verified bit blasting algorithm or program for full SMT QF_BV queries of SMT-LIB at the time of writing. Even the correctness of its results could be ensured, our certified SMT QF_BV solver CoQQFBV would not be very useful if it were extremely inefficient. In order to evaluate its performance, we run CoQQFBV on benchmarks from the QF_BV division of the single query track in the 2020 SMT Competition. With the same memory and time limits in the competition, our solver successfully finishes 88.72% of the 6861 queries with certification. In comparison, CVC4 with its certificate checker solves 55.97% with certification, and the division winner BITWUZLA solves 98.22% of the benchmarks without certification. Our certified solver outperforms CVC4 with certification significantly. Generating and checking certificates make our certified solver finish about 10% of the queries less than the division winner. The price of accuracy perhaps is not unacceptable for the benchmarks in the competition. To further evaluate CoQQFBV, the certified solver is used to verify linear arithmetic assembly programs from various cryptography libraries such as OpenSSL [30].

CoQQFBV gives certified answers to 96.88% out of the 96 SMT QF_BV queries from real-world cryptographic program verification. CVC4 with its certificate checker certifies 19.79%. Compared with efficient SMT QF_BV solvers without certification, BOOLECTOR is able to solve 100% and BITWUZLA solves 91.67% of the queries. Intriguingly, our certified SMT QF_BV solver outperforms the 2020 division winner BITWUZLA in queries from real-world verification problems. Our certified solver is probably useful for real-world verification problems.

Related Work. As mentioned, SMT certificate generating and checking are challenging. There are few efforts developing SMT QF_BV certificate checkers, let alone verified ones. CVC4 is able to produce unsatisfiability certificates for QF_BV queries, and also equipped with an (unverified) certificate checker [14]. SMTCoQ [12] is proposed to check certificates from SMT solvers VERIT and CVC4. It supports fragments of several logics including the QF_BV logic. Moreover, its correctness is formally proved in COQ. However, the QF_BV logic is not fully supported by SMTCoQ. Z3 also supports certificate generation for the QF_BV logic [19]. The proofs can be reconstructed, thus checked, within proof assistants HOL4 and ISABELLE [6]. But the lack of details in Z3’s generated certificates makes proof reconstruction particularly challenging.

With a similar approach in this paper, GL is a framework for bit blasting finitely bounded ACL2 theorems into SAT queries [28]. Its bit blasting algorithm is formally verified in ACL2. Though it is not designed for SMT-LIB, most of the operations defined in the QF_BV logic are supported, except division and concatenation for instance. A bit blasting algorithm is defined and verified in HOL4 as well [13]. Neither [28] nor [13] aims to develop a scalable SMT QF_BV solver. CoQQFBV accepts SMT-LIB inputs with fully supported QF_BV logic while adopting performance optimizations such as caches.

In ISABELLE and HOL4, one can use the bit-vector libraries to conform SMT-LIB operations, see [17] for example. Under the frame of COQ, coq-bits is a formalization of logical and arithmetic operations on bit-vectors [15]. The library provides the mapping between bit-vector operations and abstract number operations. Different from our theory, it does not support division/remainder or signed operations. WHY3 [11] provides a bit-vector theory which is formalized in COQ too. It defines the division by zero in a different way from SMT-LIB. Moreover, the operations are defined based on integer operations. Our new bit-vector theory instead defines bit-vector operations through bit manipulation. It is more suitable for the correctness proof of bit blasting algorithms.

We have the following organization. After the introduction, an overview is given in Sect. 2. Section 3 reviews preliminaries. Our formal bit-vector theory is presented in Sect. 4. It is followed by the formal semantics of SMT QF_BV queries (Sect. 5). The correctness of our bit blasting algorithm is established in Sect. 6. Section 7 outlines the construction of our certified SMTQF_BV solver. Experiments are presented in Sect. 8. Section 9 concludes our presentation.

2 Methodology Overview

Given an SMT QF_BV query, a bit blasting algorithm computes a Boolean formula such that the SMT QF_BV query is satisfiable if and only if the Boolean formula is satisfiable. The QF_BV logic contains arithmetic operations for bit-vectors. Computing an equi-satisfiable Boolean formula for an arbitrary SMT QF_BV query can be very complicated and susceptible to errors. Our goal is to construct a correct bit blasting program for every SMT QF_BV query. The correctness of the program moreover is verified by the proof assistant COQ to minimize gaps or even errors in hand-written proofs.

Our construction is based on a new formal bit-vector theory `coq-nbits` (Sect. 4). In `coq-nbits`, we define bit-vectors and their functions on top of the COQ data type for Boolean sequences. In order to support the QF_BV logic of SMT-LIB fully, five arithmetic bit-vector functions (addition, subtraction, multiplication, division, and remainder) are defined in our formal theory. To establish the correctness of our definitions, formal proofs are provided to relate bit-vector functions with their arithmetic counterparts. For instance, we show the number represented by the output of the bit-vector negation function is indeed the arithmetic negation of the number represented by the input bit-vector.

Using our `coq-nbits` theory, we then give a formal semantics for SMTQF_BV queries as defined in SMT-LIB (Sect. 5). In our formalization, a QF_BV predicate denotes a Boolean value; and a QF_BV expression denotes a bit-vector. An SMT QF_BV query is formalized as a Boolean combination of QF_BV predicates on QF_BV expressions over QF_BV variables and bit-vector constants. In order to demonstrate the correctness of our formal semantics for SMT QF_BV queries, formal proofs are provided to show that our formal semantics coincides with those defined in SMT-LIB.

Our bit blasting algorithm is given in COQ (Sect. 6). It extends Tseitin transformation for Boolean formulae to SMT QF_BV queries. More precisely, a QF_BV predicate is transformed to a literal with a Boolean formula; a QF_BV expression is transformed to a literal sequence with a Boolean formula. Using our formalization of SMT QF_BV queries, the correctness of bit blasting algorithm is established in COQ by mutual induction. To improve efficiency, our bit blasting algorithm is further optimized with more economic transformations and a cache. The optimized bit blasting algorithm is also verified with formal COQ proofs.

Our formally verified bit blasting algorithm is written in the COQ specification language. It is not yet a program compilable into executable binary codes. Using the code extraction mechanism in COQ, an OCAML program is extracted from our verified bit blasting algorithm. The OCAML program takes expressions in our formal SMT QF_BV query syntax as inputs and returns expressions in our formal syntax for Boolean formulae as outputs. SAT solvers can be employed to decide satisfiability of output Boolean formulae. Their certificates can be validated by SAT certificate checkers independently (Sect. 7).

3 Preliminaries

Let v be a Boolean *variable* with values ff and tt . A *literal* is of the form v or $\neg v$. A *clause* is a disjunction $l_0 \vee l_1 \vee \dots \vee l_k$ of literals l_0, l_1, \dots, l_k . A Boolean formula in the *conjunctive normal form (CNF)* is a conjunction $c_0 \wedge c_1 \wedge \dots \wedge c_m$ of clauses c_0, c_1, \dots, c_m . A *SAT query* is a Boolean CNF formula. An *environment* maps Boolean variables to their values. Given a SAT query, the *Boolean satisfiability problem* is to decide if the query evaluates to tt on some environments.

A bit-vector of *width* w is written as $\#bb_{w-1}b_{w-2} \dots b_0$ with $b_i \in \{0, 1\}$ for $0 \leq i < w$. In the *unsigned* representation, the bit-vector $\#bb_{w-1}b_{w-2} \dots b_0$ denotes the natural number (non-negative integer) $\sum_{0 \leq i < w} b_i 2^i$; in *two's complement (signed)* representation, it denotes the integer $\sum_{0 \leq i < w-1} b_i 2^i - 2^{w-1} b_{w-1}$. For instance, $\#b1010$ denotes 10 and -6 in the unsigned and two's complement representations respectively. We use $bv2nat(bv)$ for the natural number denoted by the bit-vector bv in the unsigned representation; and $nat2bv(w, i)$ stands for the bit-vector of width w representing the natural number i modulo 2^w .

Let $bv = \#bb_{w-1}b_{w-2} \dots b_0$ and $cv = \#bc_{u-1}c_{u-2} \dots c_0$ be bit-vectors of widths w and u respectively. The following QF_BV operations are defined in the QF_BV logic of SMT-LIB: *concat* $bv\ cv \triangleq \#bb_{w-1}b_{w-2} \dots b_0c_{u-1}c_{u-2} \dots c_0$ is the concatenation of bv and cv ; *extract* $i\ j\ bv \triangleq \#bb_i b_{i-1} \dots b_j$ extracts bits from bv where $0 \leq j \leq i < w$; *bvnot* bv , *bvand* $bv\ cv$, and *bvor* $bv\ cv$ are the bitwise complement, and, or operations respectively. Additionally, *bvneg* $bv \triangleq nat2bv(w, 2^w - bv2nat(bv))$ is the arithmetic negation operation; *bvadd* $bv\ cv \triangleq nat2bv(w, bv2nat(bv) + bv2nat(cv))$ is the arithmetic addition operation; and *bvmul* $bv\ cv \triangleq nat2bv(w, bv2nat(bv) \times bv2nat(cv))$ is the arithmetic multiplication operation. The arithmetic division and remainder operations are

$$bvudiv\ bv\ cv \triangleq \begin{cases} nat2bv(w, 2^w - 1) & \text{if } bv2nat(cv) = 0 \\ nat2bv(w, bv2nat(bv) \div bv2nat(cv)) & \text{otherwise} \end{cases}$$

$$bvurem\ bv\ cv \triangleq \begin{cases} bv & \text{if } bv2nat(cv) = 0 \\ nat2bv(w, bv2nat(bv) \bmod bv2nat(cv)) & \text{otherwise.} \end{cases}$$

Note that the arithmetic division and remainder operations are defined even when the divisor represents the number zero. Finally, the operations *bvshl* $bv\ cv \triangleq nat2bv(w, bv2nat(bv) \times 2^{bv2nat(cv)})$ shifts the bit-vector bv to the left by $bv2nat(cv)$ bits; *bvshr* $bv\ cv \triangleq nat2bv(w, bv2nat(bv) \div 2^{bv2nat(cv)})$ shifts the bit-vector bv to the right by $bv2nat(cv)$ bits. In addition to bit-vector operations, the QF_BV logic of SMT-LIB defines QF_BV *predicates* on bit-vectors. The predicate *bveq* $bv\ cv$ is true when the bit-vectors bv and cv are equal; *bvult* $bv\ cv$ is true if $bv2nat(bv) < bv2nat(cv)$. In the QF_BV logic of SMT-LIB, both operands of binary operations and predicates must have the same width. Overall, seventeen bit-vector operations and predicates are defined in the QF_BV logic of SMT-LIB. Particularly, arithmetic division and remainder operations with operands in both unsigned and two's complement signed representations are defined in SMT-LIB.

A QF_BV *variable* denotes a bit-vector. A QF_BV *expression* is constructed from QF_BV operations over QF_BV variables and bit-vectors. An SMT QF_BV *query* is a Boolean combination of QF_BV predicates on QF_BV expressions. Let *stores* be mappings from QF_BV variables to bit-vectors. Given an SMT QF_BV query, the *satisfiability modulo QF_BV theory problem* is to decide if the query evaluates to *tt* on some stores.

4 Bit-Vector Theory

We present our formal COQ bit-vector theory `coq-nbits` in this section. The `coq-nbits` theory supports bit-vectors in both unsigned and two's complement signed representations. In `coq-nbits`, a bit-vector is represented by a Boolean sequence of the data type `bits` in the least significant bit-first order.

|| Definition `bits : Set := seq bool.`

In the definition, `bool` and `seq` are the data types for Boolean values (`false` and `true`) and sequences in COQ respectively. For instance, the bit-vector `#b100` is represented by `[:: false; false; true]` in `coq-nbits`.

COQ functions defined for sequences are applicable to bit-vectors. Particularly, `size bv` computes the width of the bit-vector `bv` and `bv ++ cv` is the concatenation of the bit-vectors `bv` and `cv`. It is also straightforward to define auxiliary bit-vector functions. For example, `zeros n` returns the bit-vector of `n` `false`'s; `ones n` returns the bit-vector of `n` `true`'s; `extract i j bv` returns the sub-sequence of the bit-vector `bv` with indices from `j` to `i` where $0 \leq j \leq i < \text{size } bv$. Let `a` \triangleq `[:: false; false; true]`. Then `size a = 3` and `extract 2 1 a = [:: false; true]`.

Bitwise functions are defined as easily. For instance, the bitwise inverse function maps each Boolean value to its complement:

|| Definition `invB bv : bits := map (fun b => ~b) bv.`

Other bitwise functions are defined similarly. Specifically, bitwise and `andB`, bitwise or `orB`, logical left shift `shlB`, logical right shift `shrB` are all defined in `coq-nbits`. Let `b` \triangleq `[:: false; true; true]`. We have `invB b = [:: true; false; false]`, `andB a b = [:: false; false; true]`, and `shlB 1 b = [:: false; false; true]`.

Arithmetic bit-vector functions are slightly more complicated. To prove properties about arithmetic functions, `coq-nbits` provides conversion functions between bit-vectors and natural numbers.

|| Definition `to_N (bv : bits) : N := foldr (fun b res => N_of_bool b + res * 2) 0 bv.`

In the definition, `to_N bv` converts the bit-vector `bv` to a natural number where `N_of_bool false = 0` and `N_of_bool true = 1`. The `to_N` function multiplies the previous result by two and adds the least significant bit `b`. For instance, `to_N a = to_N [:: false; false; true] = 4`. The function `from_N w n`, on the other hand, converts any natural number `n` to a bit-vector of width `w`.

```

|| Fixpoint from_N (w : nat) (n : N) : bits :=
|| match w with
|| | 0 => [::]
|| | S w' => (N.odd n)::(from_N w' (N.div n 2))
|| end.

```

The function first checks the width w . If the width is zero, it returns the empty bit-vector. Otherwise, the function returns the bit-vector with the least significant bit `N.odd n` and the remaining $w - 1$ bits representing n divided by two. Observe that two COQ formalizations of natural numbers are used. The `nat` theory uses the unary representation suitable for inductive proofs; `N` uses the succinct binary representation. The following lemma is proved in COQ:

Lemma 1. *The following properties hold:*

1. $\forall bv, \text{from_N}(\text{size } bv) (\text{to_N } bv) = bv$.
2. $\forall w n, n < 2^w \implies \text{to_N}(\text{from_N } w n) = n$.

The first property shows that bit-vectors can be converted to natural numbers and back to themselves. The second property shows that natural numbers can be converted to bit-vectors with sufficient widths and back to themselves. To see how they are used to prove properties about bit-vector functions in `coq-nbits`, consider the definition of the successor bit-vector function.

```

|| Fixpoint succB (bv : bits) : bits :=
|| match bv with
|| | [::] => [::]
|| | hd::tl => if hd then false::(succB tl) else true::tl
|| end.

```

If the input is the empty bit-vector, the function returns the empty bit-vector. Otherwise, `succB` checks the least significant bit of the input bit-vector. If the bit is `true`, the function computes the successor of the remaining bits and appends `false` as the least significant bit. If the least significant bit of the input is `false`, the function simply changes the least significant bit to `true` and copies the remaining bits. Using the conversion functions, the bit-vector successor is related to the arithmetic successor in the following lemma:

Lemma 2. $\forall bv, \text{succB } bv = \text{from_N}(\text{size } bv) ((\text{to_N } bv) + 1)$.

Lemma 2 says that `succB bv` does compute the bit-vector representing the arithmetic successor of the natural number represented by the bit-vector bv . Observe that the successor bit-vector function is correct when the input bit-vector is empty. It is also correct when there is overflow. Indeed, both sides are zeros of width `size bv` when overflow occurs.

Other arithmetic bit-vector functions are defined and proved in `coq-nbits` similarly. Specifically, the arithmetic negation `negB`, addition `addB`, subtraction `subB`, unsigned multiplication `mulB`, unsigned division `divB`, and unsigned remainder `remB` functions are supported by `coq-nbits`. We give properties to relate the arithmetic functions for bit-vectors and natural numbers.

Lemma 3. *The following properties hold:*

1. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies \text{to_N}(\text{addB } bv\ cv) = (\text{to_N } bv + \text{to_N } cv) \bmod 2^{\text{size } bv}$.
2. $\forall bv\ cv, \text{to_N}(\text{mulB } bv\ cv) = (\text{to_N } bv \times \text{to_N } cv) \bmod 2^{\text{size } bv}$.
3. $\forall bv\ n, \text{divB } bv\ (\text{zeros } n) = \text{ones } (\text{size } bv)$.
4. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies cv \neq \text{zeros } (\text{size } cv) \implies \text{to_N}(\text{divB } bv\ cv) = (\text{to_N } bv) \text{ div } (\text{to_N } cv)$.
5. $\forall bv\ n, \text{remB } bv\ (\text{zeros } n) = bv$.
6. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies cv \neq \text{zeros } (\text{size } cv) \implies \text{to_N}(\text{remB } bv\ cv) = (\text{to_N } bv) \bmod (\text{to_N } cv)$.
7. $\forall bv\ n, \text{to_N}(\text{shlB } n\ bv) = ((\text{to_N } bv) \times 2^n) \bmod 2^{\text{size } bv}$
8. $\forall bv\ n, \text{to_N}(\text{shrB } n\ bv) = (\text{to_N } bv) \text{ div } 2^n$.

Let bv, cv be bit-vectors of width w . Lemma 3 shows that the natural number represented by the bit-vector $\text{addB } bv\ cv$ is equal to the modular sum of the natural numbers represented by bv and cv . Similarly, the natural number represented by $\text{mulB } bv\ cv$ is equal to the modular product of the natural numbers represented by bv and cv . The division and remainder functions in `coq-nbits` follow the SMT-LIB semantics. Specifically, the quotient of any bit-vector divided by zero is equal to the bit-vector of all `true`'s; the remainder of a bit-vector divided by zero is the bit-vector itself. For non-zero divisors, the division and remainder functions behave as expected. The natural number represented by the bit-vector $\text{divB } bv\ cv$ is the quotient of the number represented by bv divided by the number represented by cv ; and the bit-vector $\text{remB } bv\ cv$ represents the remainder of the number represented by bv divided by the number represented by cv . Last but not least, the logical left (`shlB`) and right (`shrB`) shifts correspond to multiplication and division by powers of two respectively.

`coq-nbits` also provides comparison predicates. In addition to the equality predicate `==` inherited from Boolean sequences, `ltB` $bv\ cv$ and `leB` $bv\ cv$ compare the natural numbers represented by the bit-vectors bv and cv . Properties about comparison predicates have also been proved in `CoQ`.

Lemma 4. *The following properties hold:*

1. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies \text{ltB } bv\ cv = (\text{to_N } bv < \text{to_N } cv)$.
2. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies \text{leB } bv\ cv = (\text{to_N } bv \leq \text{to_N } cv)$.

In addition to arithmetic functions and predicates in the unsigned representation, our formal bit-vector theory moreover defines arithmetic functions and predicates for bit-vectors in two's complement representation. For the signed representation, bit-vectors are converted to integers by the `to_Z` function. Arithmetic bit-vector functions and predicates in the signed representation are related to arithmetic integer functions and predicates as follows.

Lemma 5. *The following properties hold:*

1. $\forall bv, \neg(\text{msb } bv \wedge \text{dropmsb } bv = \text{zeros } (\text{size } bv - 1)) \implies \text{to_Z}(\text{negB } bv) = -\text{to_Z } bv$.

2. $\forall bv\ n, 1 < \text{size } bv \implies \text{to_Z } (\text{sarB } n\ bv) = (\text{to_Z } bv) \text{ quot } 2^n.$
3. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies \text{to_Z } (\text{mulB } (\text{sext } (\text{size } cv)\ bv) (\text{sext } (\text{size } bv)\ cv)) = \text{to_Z } bv \times \text{to_Z } cv.$
4. $\forall bv\ cv, 1 < \text{size } bv \implies \text{size } bv = \text{size } cv \implies [\neg(\text{msb } bv \wedge \text{dropmsb } bv = \text{zeros } (\text{size } bv - 1)) \vee cv \neq \text{ones } (\text{size } cv)] \implies \text{to_Z } (\text{sdivB } bv\ cv) = (\text{to_Z } bv) \text{ quot } (\text{to_Z } cv).$
5. $\forall bv\ cv, 1 < \text{size } bv \implies \text{size } bv = \text{size } cv \implies \text{to_Z } (\text{sremB } bv\ cv) = (\text{to_Z } bv) \text{ rem } (\text{to_Z } cv).$
6. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies \text{sltB } bv\ cv = (\text{to_Z } bv < \text{to_Z } cv).$
7. $\forall bv\ cv, \text{size } bv = \text{size } cv \implies \text{sleB } bv\ cv = (\text{to_Z } bv \leq \text{to_Z } cv).$

In the lemma, `sext` $n\ bv$ extends the bit-vector bv by n bits with the sign bit of bv , `msb` bv returns the sign bit of bv , and `dropmsb` bv drops the sign bit of bv . `quot` and `rem` are the quotient and remainder functions for COQ integers. Consider, for instance, the signed division function `sdivB` $bv\ cv$ in `coq-nbits` (Lemma 5(4)). If the dividend bv is of width > 1 , the widths of bv and the divisor cv are equal, and bv is not of the form `#b100...0` or cv is not of the form `#b11...1`, then the bit-vector `sdivB` $bv\ cv$ represents the quotient of the integers represented by bv and cv . The condition may appear counter-intuitive. To see why it is necessary, consider $bv = \#b100...0$ and $cv = \#b11...1$ both of width w . bv and cv thus represent the integers -2^{w-1} and -1 respectively. Their quotient 2^{w-1} however cannot be represented by bit-vectors of width w in two's complement representation. The corner input case is hence excluded. The corner case is also excluded from the arithmetic negation function (Lemma 5(1)).

The `coq-nbits` theory has several important differences from the prior COQ formalization in [15]. Our formal bit-vector theory supports both unsigned and two's complement signed representations. It also provides the arithmetic division and remainder functions. Since these features are needed in the QF_BV logic of SMT-LIB, they are essential to the formalization of SMT QF_BV queries. Such important features unfortunately are lacking in the prior formalization. Another noted difference is the numeric representations used in theory developments. Since integers are needed for the QF_BV logic, `coq-nbits` naturally uses binary representations for integers and natural numbers in COQ. The prior formalization on the other hand is mainly based on the unary natural number representation but provides conversion to positive integers in the binary representation.

5 Theory for SMT QF_BV Queries

Using `coq-nbits`, we formalize SMT QF_BV queries. Our formalization consists of two parts: a syntactic representation for SMT QF_BV queries in COQ inductive types and a formal semantics in our bit-vector theory `coq-nbits`.

5.1 Syntax of SMT QF_BV Queries

An SMT QF_BV query is a COQ term of the data type `bexp`. It can be constants `Bfalse` or `Btrue`, a unary predicate `Bnot`, or binary predicates `Band` or `Bor` for

Boolean connectives. Additionally, `Bbveq` and `Bbvult` with two arguments of the data type `exp` are binary QF_BV predicates.

```

Inductive bexp : Type := Bfalse : bexp | Btrue : bexp
| Bnot : bexp -> bexp
| Bband : bexp -> bexp -> bexp | Bbor : bexp -> bexp -> bexp
| Bbveq : exp -> exp -> bexp | Bbvult : exp -> exp -> bexp
(* other QF_BV predicates *)
end with exp : Type :=
| Evar : var -> exp | Econst : bits -> exp
| Ebvnot : exp -> exp
| Ebvand : exp -> exp -> exp | Ebvor : exp -> exp -> exp
| Ebvshl : exp -> exp -> exp | Ebvlshr : exp -> exp -> exp
| Ebvneg : exp -> exp
| Ebvadd : exp -> exp -> exp | Ebvmul : exp -> exp -> exp
| Ebvdiv : exp -> exp -> exp | Ebvurem : exp -> exp -> exp
| Eextract : nat -> nat -> exp -> exp
| Econcat : exp -> exp -> exp
(* other QF_BV operations *)
| Ebvsub : exp -> exp -> exp
end.

```

A COQ term of the data type `exp` represents a QF_BV expression. It can be a QF_BV variable `Evar vid` with a variable identifier `vid : var`, a bit-vector constant `Econst bv` with `bv : bits`, a bitwise-not operation `Ebvnot e0`, a bitwise-and operation `Ebvand e0 e1`, a bitwise-or operation `Ebvor e0 e1`, a logical left-shift operation `Ebvshl e0 e1`, or a logical right-shift operation `Ebvlshr e0 e1`. For arithmetic operations, there are `Ebvneg e0` for negation, `Ebvadd e0 e1` for addition, `Ebvmul e0 e1` for multiplication, `Ebvdiv e0 e1` for unsigned division, and `Ebvurem e0 e1` for unsigned remainder with `e0, e1 : exp`. Finally, the extraction `Eextract i j e0` and the concatenation `Econcat e0 e1` operations have the data type `exp` with `i, j : nat` and `e0, e1 : exp`.

5.2 Semantics of SMT QF_BV Queries

In our COQ formalization, an SMT QF_BV query is interpreted on stores. A *store* is a mapping from QF_BV variables to bits. Let σ be a store. The interpretation of $be : \text{bexp}$ on σ is a Boolean value; the interpretation of $e : \text{exp}$ on σ is a bit-vector. Semantic functions `eval_bexp` and `eval_exp` are as follows.

```

Fixpoint eval_bexp (be : bexp) (σ : store) : bool :=
  match be with
| Bfalse => false
| Btrue => true
| Bnot be0 => ~~ (eval_bexp be0 σ)
| Bband be0 be1 => (eval_bexp be0 σ) && (eval_bexp be1 σ)
| Bbor be0 be1 => (eval_bexp be0 σ) || (eval_bexp be1 σ)
| Bbveq e0 e1 => (eval_exp e0 σ) == (eval_exp e1 σ)
| Bbvult e0 e1 => ltB (eval_exp e0 σ) (eval_exp e1 σ)

```

```

(* other QF_BV predicates *)
end with eval_exp (e : exp) (σ : store) : bits :=
match e with
| Evar v => Store.acc v σ
| Econst bv => bv
| Ebvnot e0 => invB (eval_exp e0 σ)
| Ebvand e0 e1 => andB (eval_exp e0 σ) (eval_exp e1 σ)
| Ebor e0 e1 => orB (eval_exp e0 σ) (eval_exp e1 σ)
| Ebvshl e0 e1 => shlB (to_nat (eval_exp e1 σ)) (eval_exp e0 σ)
| Ebvshr e0 e1 => shrB (to_nat (eval_exp e1 σ)) (eval_exp e0 σ)
| Ebvneg e0 => negB (eval_exp e0 σ)
| Ebvadd e0 e1 => addB (eval_exp e0 σ) (eval_exp e1 σ)
| Ebvml e0 e1 => mulB (eval_exp e0 σ) (eval_exp e1 σ)
| Ebvdiv e0 e1 => divB (eval_exp e0 σ) (eval_exp e1 σ)
| Ebvrem e0 e1 => remB (eval_exp e0 σ) (eval_exp e1 σ)
| Eextract i j e0 => extract i j (eval_exp e0 σ)
| Econcat e0 e1 => (eval_exp e1 σ) ++ (eval_exp e0 σ)
(* other QF_BV operations *)
| Ebvsub e0 e1 => subB (eval_exp e0 σ) (eval_exp e1 σ)
end.

```

An SMT QF_BV query denotes a value in the COQ data type `bool`. `Bfalse` and `Btrue` denote `false` and `true` respectively. Boolean negation, conjunction, and disjunction correspond to `~`, `&&`, and `||` in `bool` respectively. For QF_BV predicates, the bit-vector equality `Bbveq` is interpreted by the equality `==` for Boolean sequences. The `coq-nbits` function `ltB` is used to interpret `Bbvlt`.

A QF_BV expression denotes a bit-vector. For basic cases, QF_BV variables are interpreted by corresponding bit-vectors in the store σ through the store access function `Store.acc`; bit-vector constants are interpreted by themselves. Bitwise logical operations `Ebvnot`, `Ebvand`, and `Ebor` are interpreted by corresponding `coq-nbits` functions `invB`, `andB`, and `orB` respectively. For logical shift operations, the offset e_1 is first converted to a natural number through `to_nat (eval_exp e1 σ)` and then passed to the corresponding logical shift functions `shlB` or `shrB` in `coq-nbits`. QF_BV arithmetic operations are interpreted by corresponding `coq-nbits` arithmetic functions as expected. Finally, the extraction `Eextract` and concatenation `Econcat` operations are interpreted by `extract` and `++` in `coq-nbits` respectively.

In an SMT QF_BV query, a QF_BV variable designates a bit-vector of a certain width. An SMT QF_BV query is hence associated with a *signature* Σ mapping QF_BV variables to their respective widths. A store σ *conforms* to a signature Σ if the interpretation of each QF_BV variable on σ has the same width as specified in Σ . Given an SMT QF_BV query $be : \text{bexp}$ with its signature Σ , be is *satisfiable* if there is a store σ conforming to Σ and `eval_bexp be σ = true`.

5.3 Derived QF_BV Operations and Predicates

In the QF_BV logic of SMT-LIB, a number of QF_BV operations and predicates are derived from a small set of core operations and predicates. Consider

the signed comparison predicate $bvslt\ bv\ cv$ in SMT-LIB:

$$\begin{aligned}
 bvslt\ bv\ cv \triangleq & (or\ (and\ (= (extract\ (w-1)\ (w-1)\ bv)\ \#b1) \\
 & (= (extract\ (w-1)\ (w-1)\ cv)\ \#b0)) \\
 & (and\ (= (extract\ (w-1)\ (w-1)\ bv) \\
 & (extract\ (w-1)\ (w-1)\ cv)) \\
 & (bvult\ bv\ cv))).
 \end{aligned}$$

To compare two bit-vectors of width w in two's complement representation, the sign bits are checked. If bv is negative but cv is positive, $bvslt\ bv\ cv$ is true. Otherwise, the signed predicate checks that both operands have the same sign and compares the operands using the unsigned comparison predicate. Interestingly, the arithmetic subtraction operation is actually a derived operation in SMT-LIB: $bvsub\ bv\ cv \triangleq bvadd\ bv\ (bvneg\ cv)$. The arithmetic operation is defined to be the bit-vector sum of minuend and the negation of subtrahend. It is *not*, for instance, defined as $nat2bv(w, bv2nat(bv) - bv2nat(cv))$ because $bv2nat(bv) - bv2nat(cv)$ may not be a natural number.

For derived operations and predicates, there is a subtle yet important difference between our formal semantics and those defined in SMT-LIB. In our formal bit-vector theory `coq-nbits`, most functions and predicates are defined directly. Particularly, the arithmetic subtraction function `subB` is defined by one-bit subtractors in `coq-nbits`. Our formal semantics for the QF_BV arithmetic operation $bvsub$ therefore is defined by the corresponding bit-vector function `subB`. Since our formal semantics did not define $bvsub$ by $bvadd$ and $bvneg$, it could be different from those in SMT-LIB. In order to build a certified solver for the QF_BV logic of SMT-LIB, it is necessary to establish semantic equivalences between both semantic definitions for all derived QF_BV operations and predicates.

To justify our formal semantics, we show the semantics of our definitions and those of SMT-LIB indeed denote the same bit-vector functions or predicates. Consider again the subtraction operation. Recall the semantics of the arithmetic operations $bvadd$ and $bvneg$ are defined by the bit-vector functions `addB` and `negB` respectively. The next lemma is useful to show the semantic equivalence:

Lemma 6. $\forall bv\ cv, size\ bv = size\ cv \implies subB\ bv\ cv = addB\ bv\ (negB\ cv)$.

For all derived QF_BV operations and predicates, we give COQ proofs for the equivalence between our formal semantics and those of SMT-LIB. Particularly, semantics of all QF_BV arithmetic operations and predicates over two's complement representation are equivalent to those in SMT-LIB. Our formal semantics for QF_BV queries is thus certified to be equivalent to SMT-LIB.

6 Certified Bit Blasting

Recall that a SAT query is a Boolean CNF formula. Given an SMT QF_BV query, a bit blasting algorithm computes a SAT query that is satisfiable if and only if the given SMT QF_BV query is satisfiable. Although it is the standard

technique for solving SMT QF_BV queries, bit blasting can be very complex due to arithmetic operations and various optimizations. Bit blasting algorithms therefore can be tedious to construct and thus prone to errors. We verify a bit blasting algorithm for SMT QF_BV queries using our COQ formalization.

Let us start with a simple formalization of Boolean CNF formulae. In our formalization, a clause is represented by a sequence of literals; a CNF formula in turn is represented by a sequence of clauses. Let `bvar` be the data type for Boolean variables. We have the following data types in COQ:

```

|||
| Inductive lit : Set := Pos of bvar | Neg of bvar.
| Definition clause : Set := seq lit.
| Definition CNF : Set := seq clause.
|||

```

Define an *environment* ϵ to be a mapping from `bvar` to `bool`. Given a literal ℓ , a CNF formula f , and an environment ϵ , it is straightforward to define the semantic functions `eval_lit` $\ell \epsilon : \text{bool}$ and `eval_cnf` $f \epsilon : \text{bool}$. A SAT query f is *satisfiable* if there is an environment ϵ such that `eval_cnf` $f \epsilon = \text{true}$.

To illustrate how our COQ proof works, consider Tseitin transformation for the logical negation operation:

```

|||
| Definition bit_blast_Bnot  $\ell$  : lit * CNF :=
|   let  $r$  := a fresh literal in
|   ( $r$ , [:: [::  $r$ ;  $\ell$ ]; [:: ! $r$ ; ! $\ell$ ] ]).
|||

```

Given a literal ℓ , `bit_blast_Bnot` ℓ returns a new literal r and the CNF formula $(r \vee \ell) \wedge (\neg r \vee \neg \ell)$. Tseitin transformation ensures the interpretations of ℓ and r are complementary on any environment ϵ evaluating the CNF formula to true. We give a formal proof using our formalization in COQ:

Lemma 7. $\forall r \text{ cnf } \ell \epsilon, (r, \text{cnf}) = \text{bit_blast_Bnot } \ell \implies \text{eval_cnf } \text{cnf } \epsilon = \text{true} \implies \text{eval_lit } r \epsilon = \sim\sim(\text{eval_lit } \ell \epsilon)$.

The idea is generalized to QF_BV operations naturally. For each QF_BV operation, we construct a literal sequence \vec{r} and a Boolean CNF formula cnf . If cnf evaluates to true on an environment ϵ , the interpretation of \vec{r} on ϵ needs to reflect the semantics of the QF_BV operation. For instance, a COQ proof is given for the QF_BV addition operation:

Lemma 8. $\forall \vec{r} \text{ cnf } \vec{\ell}_0 \vec{\ell}_1 \epsilon, (\vec{r}, \text{cnf}) = \text{bit_blast_Ebvadd } \vec{\ell}_0 \vec{\ell}_1 \implies \text{eval_cnf } \text{cnf } \epsilon = \text{true} \implies \text{eval_lits } \vec{r} \epsilon = \text{addB } (\text{eval_lits } \vec{\ell}_0 \epsilon) (\text{eval_lits } \vec{\ell}_1 \epsilon)$.

Given two literal sequences $\vec{\ell}_0$ and $\vec{\ell}_1$, `bit_blast_Ebvadd` $\vec{\ell}_0 \vec{\ell}_1$ returns a literal sequence \vec{r} and a CNF formula cnf . If cnf evaluates to true on an environment ϵ , then the interpretation of the literal sequence \vec{r} on ϵ is indeed the bit-vector sum of the interpretations of $\vec{\ell}_0$ and $\vec{\ell}_1$ on ϵ . Bit blasting algorithms for other QF_BV operations are given and shown to reflect the semantics of corresponding functions defined in the bit-vector theory `coq-nbits`. Particularly, our bit blasting algorithms for arithmetic division and remainder correctly reflect corresponding arithmetic bit-vector functions in `coq-nbits`.

Recall that the semantics for SMT QF_BV queries is defined over stores for QF_BV variables. In order to prove the correctness of bit blasting algorithms, one has to relate stores for QF_BV variables with environments for Boolean variables. The relation is explicated through literal correspondences. A *literal correspondence* π is a mapping from QF_BV variables to sequences of literals. For each QF_BV variable v , the literal sequence $\pi(v)$ is meant to interpret v on environments for Boolean variables. More formally, let `eval_lits` $\vec{\ell} \epsilon : \text{bits}$ be the bit-vector for the literal sequence $\vec{\ell}$ interpreted on the environment ϵ . The bit-vector `eval_lits` $\pi(v) \epsilon$ is hence the interpretation of the QF_BV variable v on the environment ϵ . Let σ be a store and π a literal correspondence. An environment ϵ is *consistent with σ through π* if the bit-vectors `eval_lits` $\pi(v) \epsilon$ and `Store.acc` $v \sigma$ are equal for every QF_BV variable v in σ . Thus, an environment is consistent with a store if their interpretations of variables coincide.

It is now straightforward to give our bit blasting algorithm for SMT QF_BV queries. For each QF_BV expression, our algorithm first computes literals and CNF formulae for operands recursively. It then invokes an auxiliary bit blasting algorithm to construct result literals and a CNF formula for the QF_BV operation. The literal correspondence is also updated when literals are allocated for QF_BV variables. Finally, the result literals and the updated literal correspondence are returned along with the concatenation of all CNF formulae.

```

Definition bit_blast_bexp  $\Sigma \pi b : \text{lit} * \text{correspondence} * \text{CNF} :=$ 
match  $be$  with
| Bnot  $be_0 =>$ 
  let  $(r_0, \pi', cnf_0) := \text{bit\_blast\_bexp } \Sigma \pi be_0$  in
  let  $(r, cnf) := \text{bit\_blast\_Bnot } r_0$  in
   $(r, \pi', cnf ++ cnf_0)$ 
(* other QF_BV predicates *)
end with bit_blast_exp  $\Sigma \pi e : \text{seq lit} * \text{correspondence} * \text{CNF} :=$ 
match  $e$  with
| Evar  $v =>$ 
  if  $\pi(v)$  is defined then  $(\pi(v), \pi, [::])$ 
  else let  $\vec{r} := \text{fresh literals for } v \text{ according to } \Sigma$  in
  let  $\pi' := \text{update } \pi \text{ with } v \mapsto \vec{r}$  in
   $(\vec{r}, \pi', [::])$ 
| Ebvadd  $e_0 e_1 =>$ 
  let  $(\vec{r}_0, \pi', cnf_0) := \text{bit\_blast\_exp } \Sigma \pi e_0$  in
  let  $(\vec{r}_1, \pi'', cnf_1) := \text{bit\_blast\_exp } \Sigma \pi' e_1$  in
  let  $(\vec{r}, cnf) := \text{bit\_blast\_Ebvadd } \vec{r}_0 \vec{r}_1$  in
   $(\vec{r}, \pi'', cnf ++ cnf_0 ++ cnf_1)$ 
(* other QF_BV operations *)
end.

```

The following COQ theorem establishes the connection between the output literals and the input SMT QF_BV query or expression of the algorithm.

Theorem 1. *Let $be : \text{bexp}$ be an SMT QF_BV query with the signature Σ_{be} , $e : \text{exp}$ a QF_BV expression with the signature Σ_e , and π_0 the empty literal correspondence.*

1. $\forall r \pi \text{ cnf } \sigma \epsilon, (r, \pi, \text{cnf}) = \text{bit_blast_bexp } \Sigma_{be} \pi_0 be \implies \sigma \text{ conforms to } \Sigma_{be} \implies \epsilon \text{ is consistent with } \sigma \text{ through } \pi \implies \text{eval_cnf } \text{cnf } \epsilon = \text{true} \implies \text{eval_lit } r \epsilon = \text{eval_bexp } be \sigma.$
2. $\forall \vec{r} \pi \text{ cnf } \sigma \epsilon, (\vec{r}, \pi, \text{cnf}) = \text{bit_blast_exp } \Sigma_e \pi_0 e \implies \sigma \text{ conforms to } \Sigma_e \implies \epsilon \text{ is consistent with } \sigma \text{ through } \pi \implies \text{eval_cnf } \text{cnf } \epsilon = \text{true} \implies \text{eval_lits } \vec{r} \epsilon = \text{eval_exp } e \sigma.$

Let be be an SMT QF_BV query with the signature Σ_{be} , r and cnf the literal and CNF formula returned by `bit_blast_bexp` respectively. Consider any store conforming to Σ_{be} and any environment consistent with the store. If the environment evaluates the formula cnf to true, Theorem 1 says that the literal r and the SMT QF_BV query be evaluate to the same Boolean value on the environment and store respectively. In other words, the algorithm `bit_blast_bexp` is a generalized Tseitin transformation for SMT QF_BV queries. Particularly, all QF_BV arithmetic operations (addition, subtraction, multiplication, division, and remainder in the unsigned and two's complement representations) are transformed to CNF formulae with formal proofs of correctness in COQ.

A useful corollary to Theorem 1 is the reduction of the satisfiability of SMT QF_BV queries to the satisfiability of SAT queries.

Corollary 1. *Let $be : \text{bexp } be$ be an SMT QF_BV query with the signature Σ_{be} and π_0 the empty literal correspondence. Then*

$$\begin{aligned} \forall r \pi \text{ cnf}, (r, \pi, \text{cnf}) = \text{bit_blast_bexp } \Sigma_{be} \pi_0 be \implies \\ [(\exists \sigma, \sigma \text{ conforms to } \Sigma_{be} \wedge \text{eval_bexp } be \sigma = \text{true}) \iff \\ (\exists \epsilon, \text{eval_cnf } ([:: [:: r]] \text{ ++ cnf}) \epsilon = \text{true})]. \end{aligned}$$

Corollary 1 gives the formal proof of correctness for our bit blasting algorithm `bit_blast_bexp`. Let be be an arbitrary SMT QF_BV query, r and cnf the literal and the CNF formula returned by the algorithm. The corollary shows that the query be is satisfiable if and only if the SAT query $r \wedge \text{cnf}$ is satisfiable. An equi-satisfiable SAT query is indeed obtained from the bit blasting algorithm on every input SMT QF_BV query with a formal proof of correctness.

Recall that several QF_BV operations and predicates are derived from a small number of operations and predicates in SMT-LIB. A naïve bit blasting algorithm could expand derived operations or predicates, and then perform bit blasting on a small set of operations and predicates. Such an algorithm would have a simpler proof of correctness but generate more intermediate literals and clauses. For instance, the naïve algorithm for $bvsub$ would perform bit blasting on $bvneg$ followed by $bvadd$ with intermediate literals and clauses. Our bit blasting algorithm for $bvsub$ on the other hand reflects our semantics defined by the bit-vector function `subB`. Intermediate literals or clauses are not needed. Our bit blasting algorithm hence transforms $bvsub$ more economically than the naïve algorithm.

To improve our bit blasting algorithm further, a cache for QF_BV expressions and predicates is added. In large queries, QF_BV expressions and predicates can

occur a number of times. If a QF_BV expression has several occurrences, our basic bit blasting algorithm will generate result literals and CNF formulae for each occurrence. Consider the SMTQF_BV query

$$(and (bvslt \#b1000 (bvadd x y)) (bvslt (bvadd x y) \#b0111)).$$

The query checks whether the sum of the QF_BV variables x and y can be in a proper range. Since the Boolean predicate *and* has two operands, our basic algorithm invokes the auxiliary bit blasting algorithm for the two comparison predicates. It in turn blasts the same expression *bvadd x y* twice. Repeated bit blasting on the same expression or predicate is redundant. A hash function can detect repeated QF_BV expressions and predicates easily. When an expression or a predicate recurs, the previously computed literals with the empty CNF formula are returned from a cache as the result. More importantly, we give a formal COQ proof of Corollary 1 for the bit blasting algorithm with a cache.

7 A Certified SMT QF_BV Solver

We have so far built a formally verified bit blasting algorithm for SMT QF_BV queries. Using the code extraction mechanism in COQ, an OCAML program corresponding to the verified bit blasting algorithm is obtained. Using a SAT solver and a SAT certificate checker, a certified SMTQF_BV solver can be constructed. Figure 1 gives the flow of our certified solver.




Fig. 1. Certified SMT QF_BV Solver

In the figure, the extracted OCAML program takes an OCAML expression be of the type `bexp` as an input (Sect. 5). The verified program performs bit blasting on the SMT QF_BV query and returns an OCAML expression cnf of the type `lit list list` representing a SAT query (Sect. 6). Precisely, an OCAML term of the type `lit` represents a literal. The OCAML type `lit list` corresponds to the data type for clauses; and the type `lit list list` corresponds to the data type for CNF formulae. The expression cnf is sent to a SAT solver to check satisfiability. If the SAT solver reports SAT, the SMT QF_BV query represented by be is satisfiable. Otherwise, the SAT solver reports UNSAT with a certificate. The certificate is sent to a SAT certificate checker for validation. If it is validated, the SMT QF_BV query be is unsatisfiable with certification.

8 Experiments

In order to evaluate the performance of our verified OCAML bit blasting program, we instantiate our SMT QF_BV solver CoQQFBV based on Fig. 1 as follows. We write an OCAML parser to translate a text file in the SMT-LIB format to an SMT QF_BV query in our formal syntax. The query is sent to the verified OCAML program for bit blasting. We then add an OCAML program to transform the output SAT query to a text file in the DIMACS format. The 2020 SAT Competition winner KISSAT [5] is used to check the satisfiability of the SAT query. If the SAT solver reports UNSAT with a certificate in the DRAT format [31], the certificate is sent to the verified certificate checker GRATCHK [16] for validation. Certificate checkers for SAT solvers use much simpler algorithms than certificate checkers for SMT solvers. They are hence easier to build and prove correct. The correctness of GRATCHK is in fact verified by the proof assistant ISABELLE [22]. We need not trust the certificate checker either.

We ran two experiments to evaluate our certified SMT QF_BV solver. The first experiment is the QF_BV division of the single query track in the 2020 SMT Competition [2]. The second experiment consists of verification problems from various assembly implementations for linear field arithmetic in cryptography libraries such as OpenSSL [30], RELIC [1], and BLST [29]. We compare CoQQFBV against three SMT QF_BV solvers: CVC4 [4] with an LFSC certificate checker [27], the 2020 SMTQF_BV division winner BITWUZLA [20], and the 2019 SMT QF_BV division winner BOOLECTOR [21]. BITWUZLA and BOOLECTOR are designed for efficiency without certification. CVC4 provides an LFSC certificate checker implemented in C [26]. The certificate checker can validate certificates from different theories but is itself not verified. All experiments were run on a Linux machine with a 3.20 GHz CPU and 1 TB memory.¹

8.1 SMT QF_BV Competition

The first experiment is running our certified solver CoQQFBV on tasks from the QF_BV division of the 2020 SMT Competition. We set 60 GB memory limit and 20 min timeout for each task as in the competition. A task solves a single SMT-LIB file sequentially. The SMT QF_BV division contains 6861 files in the SMT-LIB format. All files are marked with *unsat*, *sat*, or *unknown* indicating expected query results. To save running time, we ran 10 tasks concurrently. The experimental results are summarized in Table 1.

In the table, the column N_{SC} indicates the number of solved tasks with certification. O_{SC} is the number of timeouts. E_{SC} shows the number of unsolved tasks due to tool errors. T_{SC} is the average time for solved tasks. CoQQFBV solves 6087 (88.72%) and CVC4 with its certificate checker solves 3840 (55.97%) with certification. We observe three stack overflow errors during bit blasting in CoQQFBV. These errors are induced by deep recursion. Among 328 errors from CVC4, 249 are segmentation faults raised by the LFSC certificate checker.

¹ CoQQFBV is available at <https://github.com/fmlab-iis/coq-qfbv.git>.

Table 1. Experimental results on the 2020 SMT QF_BV division

Tool	N_{SC}		O_{SC}	E_{SC}	T_{SC}	N_S		O_S	E_S	T_S
CoQQFBV	6087	(88.72%)	771	3	119.69	6169	(89.91%)	689	3	81.74
CVC4	3840	(55.97%)	2693	328	74.63	4255	(62.02%)	2544	62	56.87
BITWUZLA	–	–	–	–	–	6739	(98.22%)	122	0	16.09
BOOLECTOR	–	–	–	–	–	6719	(97.93%)	142	0	15.44

Table 2. Experimental results on the 2020 SMT QF_BV division by categories

Tool	N_{SC}		T_{SC}	P_{SU}	N_S		T_S
4238 <i>unsat</i> tasks							
CoQQFBV	3838	(90.56%)	146.72	291.35 MB	3920	(92.50%)	86.51
CVC4	1762	(41.58%)	86.68	266.61 MB	2177	(51.37%)	49.68
BITWUZLA	–	–	–	–	4188	(98.82%)	12.75
BOOLECTOR	–	–	–	–	4180	(98.63%)	11.72
2553 <i>sat</i> tasks							
CoQQFBV	–	–	–	–	2242	(87.82%)	73.26
CVC4	–	–	–	–	2078	(81.39%)	64.41
BITWUZLA	–	–	–	–	2524	(98.86%)	21.08
BOOLECTOR	–	–	–	–	2516	(98.55%)	21.31
70 <i>unknown</i> tasks							
CoQQFBV	5	(7.14%)	173.17	203.52 MB	7	(10.00%)	128.26
CVC4	–	–	–	–	0	(0.00%)	–
BITWUZLA	–	–	–	–	27	(38.57%)	66.36
BOOLECTOR	–	–	–	–	23	(32.86%)	48.58

The same table also compares against efficient but uncertified solvers. To evaluate the overhead from certificate checking, the two certified solvers CoQQFBV and CVC4 still generate certificates but do not validate them. The column N_S gives the number of solved tasks without certification. O_S is the number of timeouts. E_S indicates the number of errors, and T_S is the average time for solved tasks. Our certified solver CoQQFBV finishes 6169 (89.91%) tasks. The CVC4 solver finishes 4255 (62.02%) tasks. CoQQFBV and CVC4 solve 82(= 6169 – 6087) and 415(= 4255 – 3840) more tasks without certification respectively. Since our bit blasting algorithm is verified for all inputs, CoQQFBV does not certify bit blasting on each query and hence induces less overhead. The 2020 and 2019 SMT QF_BV division winners BITWUZLA and BOOLECTOR finish 6739 (98.22%) and 6719 (97.93%) tasks without certification respectively. CoQQFBV solves about 10% less tasks with certification than the 2020 track winner BITWUZLA without certification. It also performs significantly better than CVC4 with a general SMT certificate checker.

Table 2 compares the four solvers by tasks from the three expected query results. Among the 4238 *unsat* tasks, CoQQFBV and CVC4 give certified answers to 3838 (90.56%) and 1762 (41.58%) of them respectively. The column P_{SU} gives the average size of certificates. Efficient solvers BITWUZLA and BOOLECTOR give 4188 (98.82%) and 4180 (98.63%) uncertified answers respectively.

Among the 2553 *sat* tasks, BITWUZLA and BOOLECTOR finish 2524 (98.86%) and 2516 (98.55%) of them respectively. CoQQFBV and CVC4 solve only 2242 (87.82%) and 2078 (81.39%) *sat* tasks respectively. For the 70 tasks marked *unknown*, BITWUZLA and BOOLECTOR respectively answer 27 (38.57%) and 23 (32.86%) of them without certification. Our certified SMT QF_BV solver finds two *sat* and five *unsat* tasks. Answers to the five *unsat* tasks are all certified. CVC4 with its certificate checker fails to solve any *unknown* task. For the benchmarks from the 2020 SMT QF_BV division, our certified solver CoQQFBV appears to be more scalable than CVC4 with its general SMT certificate checker.

Table 3. Average time for CoQQFBV components

Task	T_{BB}	T_{SAT}	T_{Cert}
<i>unsat</i>	41.84	49.92	73.51
<i>sat</i>	37.08	62.09	–
<i>unknown</i>	32.34	121.99	62.86

Table 3 further decomposes the time spent on different components in CoQQFBV. The column T_{BB} gives the average time for our verified OCAML bit blasting program; T_{SAT} gives the average time used by the SAT solver KISSAT; and T_{Cert} contains the average time for the certificate checker GRATCHK. For the tasks in the QF_BV division, the time for SAT solving and certificate checking are comparable. In comparison, the OCAML bit blasting program seems to take an unexpectedly large amount of time and hence can still be improved.

8.2 Linear Field Arithmetic in Cryptography

In this section, we evaluate our certified SMT QF_BV solver on benchmarks from real-world assembly implementations in various cryptography libraries such as OpenSSL [30], RELIC [1], and BLST [29]. In elliptic curve cryptography, arithmetic operations over large finite fields are needed. A field element is typically represented by hundreds of bits. A field arithmetic operation takes two field elements and returns a field element as the result. In the signature scheme Ed25519 used in OpenSSH, for instance, a field element belongs to the residue system modulo the prime number $2^{255} - 19$. Field sum of two field elements is obtained by the arithmetic sum modulo $2^{255} - 19$. Commodity processors however do not

Table 4. Experimental results on cryptographic assembly program verification

Tool	N_{SC}	T_{SC}	P_{SU}	N_S	T_S
CoqQFBV	93 (96.88%)	121.42	168.45 MB	93 (96.88%)	68.96
CVC4	19 (19.79%)	6.66	267.92 MB	46 (47.92%)	40.16
BITWUZLA	–	–	–	88 (91.67%)	16.07
BOOLECTOR	–	–	–	96 (100.00%)	18.25

support arithmetic instructions with operands in hundreds of bits natively. Field arithmetic has to be implemented by 32- or 64-bit instructions. The functional specification of the field addition used in Ed25519 may look as follows.

$$\left\{ \begin{array}{l} \{ \sum_{i=0}^3 a_i \times 2^{64 \times i} < 2^{255} - 19 \wedge \sum_{i=0}^3 b_i \times 2^{64 \times i} < 2^{255} - 19 \} \\ \text{x25519_fe64_add}(r_0, r_1, r_2, r_3, a_0, a_1, a_2, a_3, b_0, b_1, b_2, b_3) \\ \left\{ \begin{array}{l} \sum_{i=0}^3 r_i \times 2^{64 \times i} \equiv \sum_{i=0}^3 a_i \times 2^{64 \times i} + \sum_{i=0}^3 b_i \times 2^{64 \times i} \pmod{2^{255} - 19} \\ \wedge \\ \sum_{i=0}^3 r_i \times 2^{64 \times i} < 2^{255} - 19 \end{array} \right\} \end{array} \right\}$$

Let a_i, b_i, c_i be 64-bit variables (registers) for $0 \leq i \leq 3$. The specification says that the output field element represented by r_i 's computed by the program `x25519_fe64_add` is the field arithmetic sum of the input elements represented by a_i 's and b_i 's. In finite field arithmetic programs, over- or under-flow in assembly instructions lead to incorrect results, and bit-accurate program verification is required. We obtain 46 implementations and generate 96 SMT QFBV queries from verification conditions in order to evaluate our certified solver in this experiment.

Table 4 shows the verification results with the same memory and time limits in the 2020 SMT Competition. All SMT QFBV queries are expected to be unsatisfiable. BOOLECTOR successfully solves all queries (100%) without certification. The 2020 QFBV track winner BITWUZLA finishes 88 queries (91.67%) without certification. Surprisingly, CoqQFBV gives certified answers to 93 queries (96.88%). The verified SAT certificate checker GRATCHK used in CoqQFBV successfully validates all certificates for the real-world cryptographic program verification problems. In comparison, CVC4 solves 46 queries (47.92%) but certifies only 19 (19.79%). The CVC4 certificate checker raises segmentation faults on the 27 (= 46 – 19) solved but uncertified queries. These certificates are perhaps too complicated to be validated by the unverified LFSC certificate checker. For the SMT QFBV queries from real-world program verification problems, our certified solver CoqQFBV seems to perform slightly better than the efficient but uncertified SMT QFBV solver BITWUZLA. Our certified solver is probably scalable enough for certain bit-accurate program verification problems.

9 Conclusion

We combine algorithm design with interactive theorem proving to build a scalable certified SMT QF.BV solver COQQFBV in this work. Our certified solver employs a verified OCAML bit blasting program and the verified certificate checker GRATCHK to improve the confidence in SMT QF.BV query results. Experiments on the QF.BV division of the 2020 SMT Competition and real-world cryptographic program verification suggest that COQQFBV is useful.

For future work, we plan to specify and verify more heuristics to further optimize COQQFBV. Particularly, cryptographic program verification requires more sophisticated range checks. More verified bit blasting algorithms for such checks will undoubtedly improve the confidence of bit-accurate program verification.

Acknowledgements. We thank all the anonymous reviewers for their insightful comments and suggestions. We thank the authors of SSReflect for its powerful language and libraries. We would like to give special thanks to Prof. Moshe Vardi for his encouragement. The work is supported by National Natural Science Foundation of China under the Grant Numbers 62002228, 61802259 and 61836005; the Guangdong Science and Technology Department under the Grant Number 2018B010107004; Ministry of Science and Technology of Taiwan under the Grant Numbers MOST108-2221-E-001-010-MY3 and MOST108-2221-E-001-009-MY2; Academia Sinica for Sinica Investigator Award AS-IA-109-M01; the Data Safety and Talent Cultivation Project AS-KPQ-109-DSTCP.

References

1. Aranha, D.F., Gouvêa, C.P.L., Markmann, T., Wahby, R.S., Liao, K.: RELIC is an Efficient LIBrary for Cryptography. <https://github.com/relic-toolkit/relic>
2. Barbosa, H., Hoenicke, J., Hyvarinen, A.: International Satisfiability Modulo Theories Competition (SMT-COMP) (2020). <https://smt-comp.github.io/2020/>
3. Barrett, C., Stump, A., Tinelli, C.: The SMT-LIB Standard: Version 2.6. Technical report, Department of Computer Science, The University of Iowa (2017). <http://smtlib.cs.uiowa.edu/papers/smt-lib-reference-v2.6-r2017-07-18.pdf>
4. Barrett, C., et al.: CVC4. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 171–177. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_14
5. Biere, A., Fazekas, K., Fleury, M., Heisinger, M.: CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling entering the SAT competition 2020. In: Balyo, T., Froløys, N., Heule, M., Iser, M., Järvisalo, M., Suda, M. (eds.) SAT Competition 2020 - Solver and Benchmark Descriptions. B, vol. B-2020-1, pp. 50–53. University of Helsinki (2020)
6. Böhme, S., Fox, A.C.J., Sewell, T., Weber, T.: Reconstruction of Z3’s bit-vector proofs in HOL4 and Isabelle/HOL. In: Jouannaud, J.-P., Shao, Z. (eds.) CPP 2011. LNCS, vol. 7086, pp. 183–198. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25379-9_15
7. Brummayer, R., Biere, A.: Fuzzing and delta-debugging SMT solvers. In: Dutertre, B., Strichman, O. (eds.) Satisfiability Modulo Theories (SMT), pp. 1–5. ACM (2009)

8. Chen, Y.F., et al.: Verifying Curve25519 software. In: Ahn, G.J., Yung, M., Li, N. (eds.) ACM Computer and Communications Security (CCS), pp. 299–309. ACM (2014)
9. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 168–176. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24730-2_15
10. Dockins, R., Foltzer, A., Hendrix, J., Huffman, B., McNamee, D., Tomb, A.: Constructing semantic models of programs with the software analysis workbench. In: Blazy, S., Chechik, M. (eds.) VSTTE 2016. LNCS, vol. 9971, pp. 56–72. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48869-1_5
11. Dross, C., Fumex, C., Gerlach, J., Marché, C.: High-Level Functional Properties of Bit-Level Programs: Formal Specifications and Automated Proofs. Research Report RR-8821, INRIA Saclay, December 2015. <https://hal.inria.fr/hal-01238376>
12. Ekici, B., Katz, G., Keller, C., Mabsout, A., Reynolds, A.J., Tinelli, C.: Extending SMTCoq, a certified checker for SMT (extended abstract). In: Electronic Proceedings in Theoretical Computer Science, vol. 210, pp. 21–29 (2016)
13. Fox, A.C.J.: LCF-style bit-blasting in HOL4. In: van Eekelen, M., Geuvers, H., Schmaltz, J., Wiedijk, F. (eds.) ITP 2011. LNCS, vol. 6898, pp. 357–362. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22863-6_26
14. Hadarean, L., Barrett, C., Reynolds, A., Tinelli, C., Deters, M.: Fine grained SMT proofs for the theory of fixed-width bit-vectors. In: Davis, M., Fehnker, A., McIver, A., Voronkov, A. (eds.) LPAR 2015. LNCS, vol. 9450, pp. 340–355. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48899-7_24
15. Kennedy, A., Benton, N., Jensen, J.B., Dagand, P.E.: Coq: the world’s best macro assembler? In: Schrijvers, T. (ed.) Principles and Practice of Declarative Programming (PPDP), pp. 13–24. ACM (2013)
16. Lammich, P.: Efficient verified (UN)SAT certificate checking. In: de Moura, L. (ed.) CADE 2017. LNCS (LNAI), vol. 10395, pp. 237–254. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63046-5_15
17. Lochbihler, A.: Fast machine words in Isabelle/HOL. In: Avigad, J., Mahboubi, A. (eds.) ITP 2018. LNCS, vol. 10895, pp. 388–410. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-94821-8_23
18. Mansur, M.N., Christakis, M., Wüstholtz, V., Zhang, F.: Detecting critical bugs in SMT solvers using blackbox mutational fuzzing. In: Devanbu, P., Cohen, M., Zimmermann, T. (eds.) ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE), pp. 701–712. ACM (2020)
19. de Moura, L.M., Bjørner, N.: Proofs and refutations, and Z3. In: Rudnicki, P., Sutcliffe, G., Konev, B., Schmidt, R.A., Schulz, S. (eds.) Proceedings of the LPAR 2008 Workshops, Knowledge Exchange: Automated Provers and Proof Assistants, and the 7th International Workshop on the Implementation of Logics. CEUR Workshop Proceedings, vol. 418. CEUR-WS.org (2008). <http://ceur-ws.org/Vol-418/paper10.pdf>
20. Niemetz, A., Preiner, M.: Bitwuzla at the SMT-COMP 2020. CoRR abs/2006.01621 (2020). <https://arxiv.org/abs/2006.01621>
21. Niemetz, A., Preiner, M., Biere, A.: Boolector 2.0. J. Satisfiability Boolean Modeling Comput. **9**(1), 53–58 (2014)
22. Nipkow, T., Wenzel, M., Paulson, L.C. (eds.): Isabelle/HOL – A Proof Assistant for Higher-Order Logic. LNCS, vol. 2283. Springer, Heidelberg (2002). <https://doi.org/10.1007/3-540-45949-9>

23. Oe, D., Reynolds, A., Stump, A.: Fast and flexible proof checking for SMT. In: Dutertre, B., Strichman, O. (eds.) Satisfiability Modulo Theories (SMT), pp. 6–13. ACM (2009)
24. Ozdemir, A., Niemetz, A., Preiner, M., Zohar, Y., Barrett, C.: DRAT-based bit-vector proofs in CVC4. In: Janota, M., Lynce, I. (eds.) SAT 2019. LNCS, vol. 11628, pp. 298–305. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-24258-9_21
25. Polyakov, A., Tsai, M.H., Wang, B.Y., Yang, B.Y.: Verifying arithmetic assembly programs in cryptographic primitives. In: Schewe, S., Zhang, L. (eds.) Concurrency Theory (CONCUR), pp. 4:1–4:16. LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2018)
26. Reynolds, A., Stump, A.: LFSC checker. <https://github.com/CVC4/LFSC>
27. Stump, A., Oe, D., Reynolds, A., Hadarean, L., Tinelli, C.: SMT proof checking using a logical framework. *Formal Methods Syst. Des.* **42**, 91–118 (2013)
28. Swords, S., Davis, J.: Bit-blasting ACL2 theorems. In: Hardin, D., Schmaltz, J. (eds.) The ACL2 Theorem Prover and its Applications (ACL2). EPTCS, vol. 70, pp. 84–102 (2011)
29. The blst Developers: The blst BLS12-381 signature library. <https://github.com/supranational/blst>
30. The OpenSSL Project: The OpenSSL repository. <https://github.com/openssl/openssl>
31. Wetzler, N., Heule, M.J.H., Hunt, W.A.: DRAT-trim: efficient checking and trimming using expressive clausal proofs. In: Sinz, C., Egly, U. (eds.) SAT 2014. LNCS, vol. 8561, pp. 422–429. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09284-3_31

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

