



Split Gröbner Bases for Satisfiability Modulo Finite Fields

Alex Ozdemir^{1,2(✉)}, Shankara Pailoor², Alp Bassa², Kostas Ferles²,
Clark Barrett¹, and Işıl Dillig²



¹ Stanford University, Stanford, USA
aozdemir@cs.stanford.edu
² Veridise, Stanford, USA



Abstract. Satisfiability modulo finite fields enables automated verification for cryptosystems. Unfortunately, previous solvers scale poorly for even some simple systems of field equations, in part because they build a full Gröbner basis (GB) for the system. We propose a new solver that uses multiple, simpler GBs instead of one full GB. Our solver, implemented within the cvc5 SMT solver, admits specialized propagation algorithms, e.g., for understanding bitsums. Experiments show that it solves important bitsum-heavy determinism benchmarks far faster than prior solvers, without introducing much overhead for other benchmarks.

1 Introduction

Finite fields are critical to many cryptosystems. They underlie the AES-GCM cipher and ECDH key-exchange, which are used in over 80% of web requests [2, 42]. They also underlie zero-knowledge proof systems (ZKPs) and multi-party computation protocols that are used in billion-dollar private cryptocurrencies [27, 28, 40, 46], private DNS filters [34], agricultural auctions [8], discrimination studies [5], and US inter-agency data sharing [3].

Since (finite-)field-based cryptosystems are so prevalent, bugs in their implementations can have serious consequences. Furthermore, such bugs are not hypothetical. They routinely cause CVEs in OpenSSL [18, 19, 48] and compromise cryptocurrencies [1, 57, 62].

Motivated by this problem, recent research has explored automated verification for field-based computations [50, 53]. However, these techniques inherit scalability challenges from the field-solving capabilities of current Satisfiability Modulo Theories (SMT) solvers. The best SMT solver [50] for fields of cryptographic size ($\approx 2^{256}$) uses Gröbner bases (GBs) [10]. A GB can answer many questions about a system of equations, but the GB itself must first be computed.

Unfortunately, computing a GB has high theoretical complexity: doubly exponential in the worst case [45]. In practice, computing a GB *can* be feasible for some systems [50], but it is intractable for others, even simple ones. For example, consider a prime field—representable as the integers modulo a prime p . Suppose that $p \geq 2^b$ and consider the following system in variables $X_1, \dots, X_b, Z$:

$$\bigwedge_{i=1}^b X_i(1 - X_i) = 0 \quad \wedge \quad X_1 + 2X_2 + 4X_3 + \cdots + 2^{b-1}X_b = 0 \quad \wedge \quad X_b Z = 1$$

In some sense, this system is simple: the first equation forces each X_i to be 0 or 1, and the second equation forces every X_i to be 0, which then contradicts the final equation. However, computing a GB for this system using current algorithms takes exponential time. We investigate systems like this in Sect. 3, but essentially there are two conclusions: first, a GB is hard to compute because of the *combination* of the bitsum $\sum_i 2^{i-1}X_i$ and the bit constraints $X_i(1 - X_i) = 0$; second, bitsums and bit constraints are common when verifying systems that use ZKPs. So, the scalability of GB-based reasoning with bitsums is a real problem for ZKP verification.

To overcome this problem, we present a new approach for solving or refuting a system S of finite field equations. The key idea is that of a *split Gröbner basis*. If S is split into (possibly overlapping) subsystems $S_1 \wedge \cdots \wedge S_k = S$, and B_i is a GB for S_i , then we call the sequence $B_1, \dots, B_k$ a split GB for S . A split GB approximates a full GB for S : it gives detailed information about each subsystem S_i , but more limited information about S . In exchange for this approximation, if each S_i is “small” or “simple,” then the split GB might be easier to compute.

In this paper, we present a decision procedure for finite field arithmetic based on the idea of *iteratively refining* a split GB. It starts with some split of S and then refines it as necessary by sharing equations between the S_i ’s. We also add an extensible *propagation* algorithm for deducing new equations. Sharing equations increases the cost of computing the split basis but also improves the approximation that it offers. The key advantage is that the procedure can often solve or refute S before any S_i becomes too hard to compute a basis for.

We implement our approach as a solver for prime fields within the cvc5 SMT solver [4]. Our solver (a) splits bitsums and their bit constraints across two subsystems and (b) includes a specialized propagator for bitsum reasoning. This is particularly effective for important, bitsum-heavy verification problems related to ZKPs. For these problems, experiments show that our solver exponentially improves on prior work; for other problems, it has low overhead.

One application we consider is verifying field blaster ($\mathbb{F}$ -blaster) rules in a ZKP compiler: these rules encode Boolean and bit-vector operations as (conjunctions of) field equations (see Sect. 2). We give a new SMT encoding for rule correctness, prove our encoding is correct, and show that combining it with our new solver improves the state of the art for $\mathbb{F}$ -blaster verification [52]. To summarize, our key contributions are:

1. **Split**: an abstract decision procedure for field solving using a split Gröbner basis instead of a full Gröbner basis.
2. **BitSplit**: an instantiation of **Split**, optimized for bitsums and implemented in cvc5. It is exponentially faster than prior solvers on important benchmarks.
3. An application: a new encoding for $\mathbb{F}$ -blaster verification conditions that improves the state of the art for $\mathbb{F}$ -blaster verification by leveraging **BitSplit**.

The rest of the paper is organized as follows. First, we review related work (Sect. 1.1), give background (Sect. 2), and present a motivating example (Sect. 3). Then, we explain our abstract and concrete decision procedures (Sect. 4) and present experiments (Sect. 5). Last, we apply our solver to the problem of verified $\mathbb{F}$ -blasting (Sect. 6).

1.1 Related Work

There are two prior finite field solvers for SMT: Hader et al. [36, 38, 39] use subresultant regular subchains [58], and Ozdemir et al. [50] use Gröbner bases. As we will see (Sect. 5), only the latter scales to large fields. Our work builds on it.

Other prior works propose verification and linting tools for ZKPs. QED² [53] checks whether an output variable Y in some system is uniquely determined by the values of input variables $X_1, \dots, X_m$. Another project [52] verifies that a ZKP compiler's $\mathbb{F}$ -blaster is correct. These both use satisfiability modulo finite fields and could benefit from our work. Other tools are purely syntactic [20, 59, 60].

Further afield, others consider finite fields in *interactive* theorem provers, applied to mathematics [9, 16, 31, 41], to program correctness [25, 26, 54, 55], and even to ZKPs [13, 15, 29, 43]. In contrast, our work is fully automatic.

2 Background

Here we summarize necessary definitions and facts about finite fields [21, Part IV], computer algebra [17], satisfiability modulo finite fields (SMFF) [47, 50], and applications of SMFF [52, 53]. See the references for further details.

Finite Fields and Polynomials. For naturals $a \geq 1$, $[a]$ denotes $\{1, \dots, a\}$. In general, $\mathbf{x}$ denotes a list of elements $x_1, \dots, x_m$. Let p be a prime. $\mathbb{F}_p$ (abbreviated $\mathbb{F}$ when p is clear) denotes the unique finite field of order p , represented as $\{0, \dots, p-1\}$ with addition and multiplication modulo p . A field of prime order is also called a *prime field*. Let $\mathbf{X}$ be a list of n variables: $(X_1, \dots, X_n)$. $\mathbb{F}[\mathbf{X}]$ is the set of polynomials in $\mathbf{X}$ with coefficients from $\mathbb{F}$. For $f \in \mathbb{F}[\mathbf{X}]$, let $\deg(f)$ be its degree and $\text{vars}(f)$ be the set of variables appearing in it.

Ideals and Their Zeros. Let $S = \{s_1, \dots, s_m\}$ be a set of polynomials in $\mathbb{F}[\mathbf{X}]$. $\langle S \rangle$ denotes the *ideal* that is *generated* by S : the set $\{\sum_i f_i s_i : f_i \in \mathbb{F}[\mathbf{X}]\}$. Let $\mathbf{S} = (S_1, \dots, S_k)$ be a list of *sets* of polynomials. Then, we define $\langle \mathbf{S} \rangle \triangleq (\cup_i S_i)$.

Let $M : \mathbf{X} \rightarrow \mathbb{F}$ be a map from variables $\mathbf{X}$ to values in $\mathbb{F}$. For $f \in \mathbb{F}[\mathbf{X}]$, denote the evaluation of f on M by $f[M]$; a *zero* of f is an M with $f[M] = 0$. The common zeros of S are denoted $\mathcal{V}_{\mathbb{F}}(S)$ (abbreviated $\mathcal{V}(S)$). Note that $\mathcal{V}(S) = \mathcal{V}(\langle S \rangle)$. When studying polynomial systems, one generally considers the system given by the ideal it generates, as it has more structure and has the same set of zeros. For any $f \in \mathbb{F}[\mathbf{X}]$, if $f \in \langle S \rangle$, then $\mathcal{V}(\{f\}) \supseteq \mathcal{V}(S)$. One implication of this is that $1 \in \langle S \rangle$ implies that $\mathcal{V}_{\mathbb{F}}(S)$ is empty. However, the converse does not hold: for example, the polynomial $X^2 + 1$ has no zero in $\mathbb{F}_3$, but $1 \notin \langle X^2 + 1 \rangle$.

Gröbner Bases. A Gröbner basis (GB) is a kind of polynomial set that is often used for solving polynomial systems. Two facts about GBs are relevant to this paper. First, there is an algorithm, `GB`, that for any polynomial set S , computes a GB B such that $\langle B \rangle = \langle S \rangle$. In this case, we say that B is a GB for S or for $\langle S \rangle$. (But: note that in this paper, B does not always refer to a GB!) Second, there is an algorithm `InIdeal`(f, B) that determines whether $f \in \langle B \rangle$ for polynomial f and GB B .¹ Thus, if `InIdeal`(1, `GB`(S)) returns true, this shows that $\mathcal{V}(S)$ is empty. Moreover, `InIdeal`(1, B) is computable in polytime if B is a GB since 1 reduces by B iff B contains a non-zero constant [17].

Satisfiability Modulo Finite Fields (SMFF). Previous work [38, 50] defines the theory of finite fields, which we summarize here using the usual terminology of many-sorted first order logic with equality [24]. For every finite field $\mathbb{F}$, let the signature Σ include: sort $\mathbb{F}$, binary function symbols $+\mathbb{F}$ and $\times\mathbb{F}$, constants $n \in \{0, \dots, |\mathbb{F}| - 1\} \subset \mathbb{N}$, and the inherited equality symbol $\approx_{\mathbb{F}}$. The theory of finite fields requires that any Σ -interpretation interprets $\mathbb{F}$ as $\mathbb{F}$, n as the n^{th} element of $\mathbb{F}$, and $+$, $\times$, and $\approx$ as addition, multiplication, and equality in $\mathbb{F}$. Previous work reduces the satisfiability problem for this theory to the problem of finding an element of $\mathcal{V}(S)$ given S or determining that there is no such element [50]. In this work, we consider the latter problem.

Applying SMFF to ZKPs. Prior work applies SMFF to verification for zero-knowledge proof systems (ZKPs) [50, 52, 53]. Practical ZKPs [11, 30, 33] allow one to prove knowledge of a solution to a system of *field equations* $\Phi(\mathbf{X}, \mathbf{Y})$, while keeping all or part of the solution secret. Since Φ is usually meant to encode a function from $\mathbf{X}$ to $\mathbf{Y}$, recent tools attempt to verify *determinism*: that the value of $\mathbf{X}$ uniquely determines the value of $\mathbf{Y}$ [53, 56, 59, 60]. Determinism can be written as a single satisfiability query solved with SMFF:

$$\Phi(\mathbf{X}, \mathbf{Y}) \wedge \Phi(\mathbf{X}', \mathbf{Y}') \wedge \mathbf{X} = \mathbf{X}' \wedge \mathbf{Y} \neq \mathbf{Y}' \quad (1)$$

The formula (1) is satisfiable if and only if Φ is **nondeterministic**. Determinism is important for two reasons. First, constructing (1) only requires identifying the inputs and outputs, making the specification task trivial and automatable. Second, determinism violations are frequent; one caused the Tornado Cash bug [57], and they are part of over half of the bugs in the ZK Bug Tracker [1]. Third, determinism violations cause real vulnerabilities. A recent survey of ZKP vulnerabilities concludes that insufficient constraints (which typically manifest as non-determinism) account for 95% of constraint-system-level vulnerabilities [12]. In Sect. 6, we give another reason why determinism is important: it can imply stronger properties.

3 Motivating Example

In this section, we explore a class of problems that is both important and challenging for existing SMFF solvers. First (Sect. 3.1), we explain the source

¹ The definition of GB and these algorithms depends on a *monomial order*. Throughout the paper, we use *grexlex* order. We discuss monomial orders in Appendix A.

```

1  template Num2Bits(b) { // split 'in' into 'b' bits.
2      signal input in;
3      signal output out[b];
4      var bitSum = 0;
5      for (var i = 1; i <= b; i++) {
6          out[i] * (out[i] - 1) == 0; // 'out[i]' is 0 or 1
7          bitSum += out[i] * 2 ** (i - 1); // add a term to the accumulating bitsum
8      }
9      bitSum == in; // 'in' is the bitsum of 'out'
10 }

```

Fig. 1. Num2Bits: a widely-used circomlib library function. It converts a prime field element into an b -bit binary representation (assuming this is possible).

and prevalence of these problems—determinism queries with bit-splitting. Second (Sect. 3.2), we explore why they are hard for GB-based reasoning, and we present evidence that the core challenge is the combination of bitsums and bit-constraints. Third (Sect. 3.3), we sketch the design of a decision procedure that can meet this challenge.

3.1 Verifying the Determinism of Num2Bits

The circom language is used to synthesize field equations for ZKPs. Figure 1 shows a slice of the circom program Num2Bits. It relates an input signal in to its binary representation as an array of signals out . The code generates a set of field equations that encode this relationship. The `==` operator generates equations. Line 6 generates the equation forcing $out[i]$ to be either 1 or 0, line 7 adds $out[i]$ to the expression that is accumulating terms in the bitsum, and line 9 generates the equation equating the bitsum to in . Thus, the equations are:

$$\Phi(in, out) := \left(in = \sum_{i=1}^b 2^{i-1} out[i] \right) \wedge \bigwedge_{i=1}^b out[i](out[i] - 1) = 0 \quad (2)$$


Here, b is constant. For any $j \in [b]$, the output $out[j]$ is deterministic if the following SMFF query is unsatisfiable:

$$\exists in, in', out, out'. \quad \Phi(in, out) \wedge \Phi(in', out') \wedge in = in' \wedge out[j] \neq out'[j] \quad (3)$$

Importance. Nearly every circom project uses Num2Bits or similar templates that bit-split field elements. This is because bit encodings are a natural way to encode common operations like range-checks ($x \in \{l, \dots, u\}$) and comparisons ($<, >$) as field equations. In fact, in a crawl of all public circom Github projects, we found that 98% of projects use Num2Bits or other circuits with bitsums. Furthermore, bitsums are *very* common in many programs; for example, in circomlib’s SHA2 implementation, 64% of the variables appear in some bitsum. We describe our methodology for these measurements in Appendix B.

Table 1. Different ideal families with bitsums and bit-constraints.

Ideal Family	Generators
$I_{2,\det}(b)$	$\text{B}\Sigma\text{P}(Y, \mathbf{X}) \cup \text{B}\Sigma\text{P}(Y', \mathbf{X}') \cup \{Y - Y'\} \cup \{(X_b - X'_b)Z - 1\}$
$I_2(b)$	$\text{B}\Sigma\text{P}(Y, \mathbf{X}) \cup \text{B}\Sigma\text{P}(Y', \mathbf{X}') \cup \{Y - Y'\}$
$I_1(b)$	$\text{B}\Sigma\text{P}(Y, \mathbf{X})$
$I_{1,\text{val}}(b)$	$\text{B}\Sigma\text{P}(Y, \mathbf{X}) \cup \{Y\}$

**Fig. 2.** GB computation time for different systems at different bitsum lengths.

3.2 The Challenge of Bit-Splitting

Unfortunately, state-of-the-art SMFF solvers struggle with (3). The solver of Hader et al. [38] scales poorly with field size (Sect. 5), and ZKP security typically requires $|\mathbb{F}| \approx 2^{255}$. It fails for (3), even when $b = 1$. The GB-based solver of Ozdemir et al. [50] scales better with $|\mathbb{F}|$, but poorly with b . It can handle many large-field benchmarks, but it cannot solve (3) for $b = 32$, even in a week.

To understand the problem, consider how a GB-based solver handles (3). First, it computes a polynomial set S such that $\mathcal{V}(S)$ encodes solutions to (3):

$$\begin{aligned}
 S = \{ & Y - Y', \quad Y - \sum_{i=1}^b 2^{i-1} X_i, \quad Y' - \sum_{i=1}^b 2^{i-1} X'_i, \\
 & X_1^2 - X_1, \dots, X_b^2 - X_b, \quad X_1'^2 - X_1', \dots, X_b'^2 - X_b', \\
 & (X_j' - X_j)Z - 1 \}
 \end{aligned} \tag{4}$$

In this system, in , in' , out , and out' are represented by variables Y , Y' , $\mathbf{X}$, and $\mathbf{X}'$ respectively. The inequality $X_j \neq X'_j$ becomes the polynomial $(X'_j - X_j)Z - 1$ (for fresh Z) which can be zero only if $X_j \neq X'_j$. Next, the solver attempts to compute a GB for (4). But this takes time exponential in b , as we will see.

To empirically investigate the cause of the slowdown, we consider other families of ideals generated by sets similar to (4). Table 1 shows four ideal families of increasing simplicity that all include bit-splitting. The polynomials are in variables $(X_1, \dots, X_b, X'_1, \dots, X'_b, Y, Y', Z)$, and we define the set $\text{B}\Sigma\text{P}(Y, \mathbf{X})$ as:

$$\text{B}\Sigma\text{P}(Y, (X_1, \dots, X_b)) \triangleq \{Y - \sum_{i=1}^b 2^{i-1} X_i, X_1^2 - X_1, \dots, X_b^2 - X_b\}.$$

The first family, $I_{2,\det}(b)$, is exactly (4), for $j = b$. The second, I_2 , removes the polynomial that enforces disequality. The third, I_1 , removes one of the bitsum




Fig. 3. High-level information flow in **BitSplit**: our concrete decision procedure.

and bit-constraint sets. The fourth, $I_{1,\text{val}}$, fixes the lone bitsum to a specific value ($Y = 0$). Computing a GB for *any of these families* takes time exponential in b .²

Figure 2 shows the times (using Singular [32]; others are similar). $I_{1,\text{val}}$ is easiest to compute a GB for, and I_2 is the hardest, but all take exponential time.

Interestingly, the singleton set of just the bitsum $\{Y - \sum_{i=1}^b 2^{i-1} X_i\}$ and the set of bit-constraints without the bitsum $\{X_1^2 - X_1, \dots, X_b^2 - X_b\}$ are both already GBs. It appears that the **combination** of the bitsum and the bit-constraints is what makes computing a GB hard.

Translation to Bit-Vectors: a Dead End. Since ZKPs process finite-field equations, the system (2) has coefficients in a *finite field*. Yet, the appearance of the bitsum pattern makes it tempting to attempt some kind of translation into the bit-vector domain. After all, in that domain, bit-decomposition is easy to reason about! However, this intuitive appeal is misleading. In practice, the approach is not trivial, since (in the general case) the system Φ includes other (non-bitsum) equations too. In fact, previous attempts to solve finite-field equations by translation to bit-vectors have been shown to be very ineffective [50]. Thus, performing some finite-field reasoning seems crucial.

3.3 Cooperative Reasoning: A Path Forward

We have seen that verifying **Num2Bits** is hard with only GBs. Yet, **Num2Bits** is easy to verify when we combine GBs with other kinds of reasoning. Consider the following inferences about $\langle S \rangle$ (Eq. 4): Since $\mathbf{X}, \mathbf{X}'$ are bit representations of Y, Y' respectively and $Y - Y'$ is in $\langle S \rangle$, every $X'_i - X_i$ must be too. This is the *congruence* rule for the function from a number to its bit representation. Then, since $f = X'_j - X_j$ and $g = (X'_j - X_j)Z - 1$ are both in $\langle S \rangle$ a GB shows that $1 = fZ - g$ is also in $\langle S \rangle$. But, if $1 \in \langle S \rangle$, then S can have no common zeros. So, (3) is UNSAT, and **Num2Bits** is deterministic. The key here is to use GB-based reasoning *and* non-GB-based reasoning (congruence for bit representations).

² For Fig. 2, we work in $\mathbb{F}_p$, where p is the smallest prime greater than $2^b - 1$. However, the results are similar for other values of p as well.

Our decision procedure **BitSplit** mixes GB-based and non-GB-based reasoning to understand the contents of an ideal $\langle S \rangle$. Figure 3 illustrates its architecture. There are three modules: each learns new polynomials in $\langle S \rangle$ and potentially shares them with other modules. The *sparse* module computes a GB for all polynomials except *bitsum polynomials* (or *bitsums*): those of form $Y - \sum_i 2^{i-1} X_i$. Its name refers to the fact that bitsums are dense: they have many terms. The *linear* module computes a GB for all linear polynomials (including all bitsums). The *unique bit representation* module infers bit equalities using congruence.

This architecture has three key features. First, it includes non-GB-based reasoning. Second, every polynomial is handled by *some* GB-based module (either the sparse or linear module); this will play a role in correctness. Third, by splitting bitsums (which go into the linear module) and bit-constraints (which go into the sparse module), it avoids computing a GB for both simultaneously.

4 Approach

In this section, we present our decision procedure. Given a set of polynomials G , our procedure either finds a common zero $M \in \mathcal{V}(G)$ or determines that none exists. Recall from Sect. 2 that satisfiability modulo $\mathbb{F}$ reduces to this problem.

To explain our decision procedure, we first introduce a *split Gröbner basis* (Sect. 4.1), which can be easier to compute than a full GB, but can also be less useful when deciding satisfiability. Next, we present our abstract decision procedure **Split**, which manipulates split Gröbner bases (Sect. 4.2). **Split** is parameterized by the number of bases k and also by some subroutines. We show that if the subroutines meet suitable conditions, then **Split** is sound and terminating (Theorem 3). Finally, we instantiate **Split** with $k = 2$ by defining the necessary subroutines (Sect. 4.3). The result is a concrete decision procedure **BitSplit** which is optimized for reasoning about bitsums.³ We evaluate **BitSplit** experimentally in Sect. 5.

4.1 Split Gröbner bases

Definition 1 (Split Gröbner basis). A *split Gröbner basis* for ideal I is a sequence $(B_1, \dots, B_k)$ of Gröbner bases such that $I = \langle \mathbf{B} \rangle$.

We make a few relevant observations about this definition.

1. A split GB generalizes a GB: that is, $(\text{GB}(S))$ is always a split GB for $\langle S \rangle$.
2. Split GBs for an ideal I are not unique.
3. The split GB definition relaxes the GB definition: while GBs can be hard to compute, split GBs need not be. For example, the ideal $\langle f_1, \dots, f_n \rangle$ has split GB $(\{f_1\}, \dots, \{f_n\})$.

³ We use the name “**BitSplit**” because the procedure is optimized for bitsums (used in bit-splitting) and because the name suggests an instantiation of the “**Split**” procedure.

<pre> 1 Function Monolithic: In: $G \subset \mathbb{F}[\mathbf{X}]$ Out: A zero $M \in \mathcal{V}(G)$ or $\perp$ 2 3 $B \leftarrow \text{GB}(G)$; 4 if $1 \in \langle B \rangle$ then return $\perp$; 5 return FindZero(B) </pre>	<pre> 1 Function Split: In: $G \subset \mathbb{F}[\mathbf{X}]$ Out: A zero $M \in \mathcal{V}(G)$ or $\perp$ 2 $\mathbf{G} \leftarrow (\{p \in G : \text{init}(i, p)\})_{i=1}^k$; 3 $\mathbf{B} \leftarrow \text{SplitGB}(\mathbf{G})$; 4 if $\exists i. 1 \in \langle B_i \rangle$ then return $\perp$; 5 return SplitFindZero($\mathbf{B}$) </pre>
(a) The prior decision procedure [50].	(b) Our abstract procedure Split.

Fig. 4. The prior decision procedure (Monolithic) [50] and our framework (Split).

Informally, a split GB allows one to navigate a trade-off between the computational expense of computing GBs and the power of their ideal membership tests. Generally, a smaller split GB where each individual GB represents more of I makes $\text{InIdeal}(\cdot, B_i)$ more informative. On the other hand, a bigger split GB where each GB represents less of I makes the split basis easier to compute. Section 3 gave an example of this: it is hard to compute a GB for $\langle \sum_{i=1}^b 2^{i-1} X_i, X_1^2 - X_1, \dots, X_b^2 - X_b \rangle$, but $(\{\sum_{i=1}^b 2^{i-1} X_i\}, \{X_1^2 - X_1, \dots, X_b^2 - X_b\})$ is *already* a split GB.

```

1 Function SplitGB:
  | In:  $\mathbf{G} = (G_i)_{i=1}^k$ : a list of generator sets
  | Out:  $\mathbf{B} = (B_i)_{i=1}^k$ : a split GB; initially each  $B_i$  is empty.
2   while  $\cup_i G_i$  is not empty do
3     | for  $i \in [k]$  do  $B_i \leftarrow \text{GB}(G_i \cup B_i)$ ;  $G_i \leftarrow \emptyset$ ;
4     | for  $p \in (\cup_j B_j) \cup \text{extraProp}(\mathbf{B}), i \in [k]$  do
5     |   | if  $\text{admit}(i, p) \wedge p \notin \langle B_i \rangle$  then  $G_i \leftarrow G_i \cup \{p\}$ ;
6   return  $\mathbf{B}$ 

```

Algorithm 1: SplitGB computes a split Gröbner basis, with propagation.

4.2 Abstract Procedure: Split

Our starting point is a prior solver based on Gröbner bases [50]. Figure 4a shows the prior procedure, which we call Monolithic, and Fig. 4b shows our new procedure, which is named Split. Monolithic begins by computing a GB B and returning $\perp$ if $1 \in \langle B \rangle$. Recall that $1 \in \langle B \rangle$ implies $\mathcal{V}(G)$ is empty, but the converse does not hold; thus, this is a sound but incomplete test for unsatisfiability. If the problem remains unsolved, then Monolithic proceeds to FindZero, which is a (complete) backtracking search over elements of $\mathbb{F}$.

The key difference in Split is that it works with a split GB $\mathbf{B}$ for $\langle G \rangle$. First (line 2), we split G into subsets $G_1 \cup \dots \cup G_k = G$; these may overlap. Second (line 3), we compute a Gröbner basis B_i for each subset G_i (and perform additional propagations, discussed later). If some $\langle B_i \rangle$ contains 1, we return $\perp$. Third (line

5), we fall back to a (complete) backtracking search based on $\mathbf{B}$. We will now discuss each phase in more detail.

Splitting. Splitting is done with a function $\text{init}(i, p)$ that decides whether polynomial p should initially be included in basis i . The function init is a parameter of Split . The only requirement of init is that no polynomial can be ignored:

Definition 2 (Covering init). *The function init is **covering** when for all $p \in \mathbb{F}[\mathbf{X}]$, there exists an $i \in [k]$ such that $\text{init}(i, p) = \top$.*

Computing a Split GB and Propagating. In the second stage, we compute a split GB $\mathbf{B}$ using SplitGB (Algorithm 1). To start, SplitGB sets each B_i to be a GB for $\langle G_i \rangle$. However, SplitGB also adds to each B_i additional polynomials called *propagations*. Propagations can be *inter-basis* (from a different B_j) or *extra* (from a subroutine extraProp). Through extraProp , one can extend SplitGB with specialized reasoning (e.g., for bitsums). Whether a propagation p is admitted into B_i is controlled by a subroutine $\text{admit}(i, p)$. Through admit , a basis can reject a polynomial p that would slow down future GB computations.

Now, we explain SplitGB in detail. In each iteration of the outer loop, B_i is a current basis and G_i is a set of polynomials that will be added in the next round. First, B_i is computed from the previous G_i and B_i . Then, polynomials from each B_j are added to each G_i if $\text{admit}(i, \cdot)$ accepts them and $\langle B_i \rangle$ doesn't contain them already. Any propagations from $\text{extraProp}(\mathbf{B})$ are added in the same way. The loop iterates until there are no new additions.

The correctness of SplitGB depends on extraProp , but not admit . As captured by Definition 3, $\text{extraProp}(\mathbf{B})$ must only return polynomials in $\langle \mathbf{B} \rangle$. If extraProp obeys this requirement, then SplitGB terminates and preserves the generated ideal, as stated in Theorem 1. The proof is in Appendix C; correctness is straightforward, and termination follows from the same theory that guarantees termination for Buchberger's algorithm [10]. We discuss efficiency later.

Definition 3 (Sound extraProp). *The function extraProp is **sound** when for all $\mathbf{B} \in (2^{\mathbb{F}[\mathbf{X}]})^k$, $\text{extraProp}(\mathbf{B}) \subseteq \langle \mathbf{B} \rangle$.*

Theorem 1. *If extraProp is sound, then $\text{SplitGB}(\mathbf{G})$ terminates and returns a split Gröbner basis $\mathbf{B}$ such that $\langle \mathbf{B} \rangle = \langle \mathbf{G} \rangle$ and $\langle B_i \rangle \supseteq \langle G_i \rangle$ for all i .*

Backtracking Search. SplitFindZero (Algorithm 2) is our conflict-driven search. Given a split basis $\mathbf{B}$, it returns $M \in \mathcal{V}(\langle \mathbf{B} \rangle)$ if possible, and $\perp$ if $\mathcal{V}(\langle \mathbf{B} \rangle)$ is empty. It uses a subroutine $\text{SplitZeroExtend}(\mathbf{B})$ which searches for an $M \in \mathcal{V}(\langle \mathbf{B} \rangle)$ by focusing on B_1 , as we explain below. SplitZeroExtend returns one of three possibilities: an $M \in \mathcal{V}(\langle \mathbf{B} \rangle)$; $\perp$, indicating that $\mathcal{V}(\langle \mathbf{B} \rangle)$ is empty; or a *conflict polynomial* $p \in (\cup_i B_i) \setminus \langle B_1 \rangle$ that it failed to account for in its B_1 -focused search. In the last case, SplitFindZero adds p to B_1 and tries SplitZeroExtend again. Each conflict is new information that is added to B_1 from some other B_i .

SplitZeroExtend is based on the FindZero algorithm of prior work [50]. FindZero is a backtracking search based on a GB B . In each recursive step, it assigns a single variable to a single value. Rather than doing an exhaustive case

1 Function SplitFindZero:

In: $\mathbf{B} = (B_i)_{i=1}^k$: a split GB
Out: A zero $M \in \mathcal{V}(\langle \mathbf{B} \rangle)$ or $\perp$

2 while conflict $p \leftarrow \text{SplitZeroExtend}(\mathbf{B})$ **do**

3 $\mathbf{B} \leftarrow \text{SplitGB}(B_1 \cup \{p\}, B_2, \dots, B_k);$

4 **return** $\text{SplitZeroExtend}(\mathbf{B})$

5 Function SplitZeroExtend:

In: $\mathbf{B} = (B_i)_{i=1}^k$: the current split GB

In: $G \subset \mathbb{F}[\mathbf{X}]$: the original generators; if omitted, equal to $\cup_i B_i$

In: A partial map $M : \mathbf{X} \rightarrow \mathbb{F}$; if omitted, empty

Out: A total map M or a conflict polynomial p or $\perp$

6 if $\exists i. 1 \in \langle B_i \rangle$ **then**

7 **if** $\exists p \in G \setminus \langle B_1 \rangle, \text{vars}(p) \subseteq \text{vars}(M) \wedge p[M] \neq 0$ **then return** p ;

8 **else return** $\perp$;

9 **if** $|M| = n$ **then return** M ;

10 for $(X_{j_i} \mapsto z_i) \in \text{ApplyRule}(B_1, M)$ **do**

11 $r \leftarrow \text{SplitZeroExtend}(\text{SplitGB}((B_j \cup \{X_{j_i} - z_i\})_{j=1}^k), G, M \cup \{X_{j_i} \mapsto z_i\});$

12 **if** $r \neq \perp$ **then return** r ;

13 **return** $\perp$

Algorithm 2: SplitFindZero finds zeros using split Gröbner bases.

split for each variable, a subroutine **ApplyRule** analyzes B and constructs a list (an implicit disjunction) of single-variable assignments $X_{j_1} \mapsto z_1, \dots, X_{j_\ell} \mapsto z_\ell$ that *cover* $\mathcal{V}(B)$. That is, for each $M \in \mathcal{V}(B)$, there exists i such that $M[X_{j_i}] = z_i$. Thus, we know that if a solution exists, it must agree with at least one of these assignments. For example, with $B = \{X_1^2 - X_2, X_1(X_2 - 1)\}$, every solution must assign X_1 to 0 or X_2 to 1, so any set of assignments including these would do. **ApplyRule** might, for instance, return exactly $\{X_1 \rightarrow 0, X_2 \rightarrow 1\}$. For each i , **FindZero** recurses on $B \leftarrow \text{GB}(B \cup \{X_{j_i} - z_i\})$. It backtracks if $1 \in \langle B \rangle$ and succeeds if every variable has been assigned.

SplitZeroExtend adapts **FindZero** to a split GB, essentially by running **FindZero** on B_1 and using **SplitGB** instead of **GB**. It also uses a limited notion of conflicts to prune the search space. It is given a split basis $\mathbf{B}$ (that changes in each recursion), a generator set G (that is fixed across recursions and is initially equal to $\cup_i B_i$), and a partial map M from variables to values. First (lines 6–8), it checks whether 1 is in any $\langle B_i \rangle$. There are two cases here. If some polynomial $p \in G \setminus \langle B_1 \rangle$ fully evaluates to a non-zero value, p is returned as a conflict. Otherwise, $\perp$ is returned. Second (line 9), if M is total, then it is returned as a common zero. Third (lines 10–12), **SplitZeroExtend** uses **ApplyRule** (from [50]) to obtain a list of single-variable assignments that cover $\mathcal{V}(B_1)$. For each assignment in the list, it attempts to construct a solution by adding that assignment to M and to each B_i and recursing. If no branch succeeds, it returns $\perp$.

For each conflict that **SplitZeroExtend** returns, **SplitFindZero** will call it again with a new starting split basis. Theorem 2 states the correctness of **SplitFindZero**. The correctness of **Split** (Theorem 3) is a corollary. The proofs are in Appendix D (Table 2).

Table 2. The functions that parameterize **Split**.

Function signature	Semantics
$\text{init}(i \in [k], p \in \mathbb{F}[\mathbf{X}]) \rightarrow \{\top, \perp\}$	whether to initialize basis B_i with p
$\text{admit}(i \in [k], p \in \mathbb{F}[\mathbf{X}]) \rightarrow \{\top, \perp\}$	whether to accept p into B_i during propagation
$\text{extraProp}(\mathbf{B} \in (2^{\mathbb{F}[\mathbf{X}]})^k) \rightarrow 2^{\mathbb{F}[\mathbf{X}]}$	additional polynomials to propagate

Table 3. Which polynomials our bases accept. The linear basis accepts linear polynomials. The sparse basis accepts non-bitsums initially, and then equalities.

Basis # (i)	Name	$\text{init}(i, p)$ definition	$\text{admit}(i, p)$ definition
1	Sparse	$\neg \text{isBitsum}(p)$	$\text{isEq}(p)$
2	Linear	$\deg(p) \leq 1$	$\deg(p) \leq 1$

Theorem 2. Let $\mathbf{B}$ be a split GB. If **extraProp** is sound then **SplitFindZero**($\mathbf{B}$) terminates and returns an element of $\mathcal{V}_{\mathbb{F}}(\langle \mathbf{B} \rangle)$ iff one exists.

Theorem 3. Let G be a polynomial set. If **extraProp** is sound and **init** is covering, then **Split**(G) terminates and returns an element of $\mathcal{V}_{\mathbb{F}}(G)$ iff one exists.

4.3 Concrete Procedure: BitSplit

Bases. To construct **BitSplit**, we instantiate **Split** with $k = 2$. We call B_1 the *sparse* basis and B_2 the *linear* basis, and we define **init** and **admit** as shown in Table 3. We explain **extraProp** later.

We carefully avoid allowing a bitsum $X - \sum_{i=0}^k 2^i X_i$ and its bit constraints $(X_i^2 - X_i)_{i=1}^k$ in the same basis. Initially, the sparse basis rejects only bitsums ($\text{isBitsum}(p)$ is defined as $\exists \ell > 1, \exists Y, X_1, \dots, X_\ell \in \mathbf{X}, p = Y - \sum_{i=0}^\ell 2^i X_i$). During propagation, the sparse basis accepts polynomials that encode equalities ($\text{isEq}(p)$ is defined as $\exists X, Y \in \mathbf{X}, z \in \mathbb{F}, p = X - Y \vee p = X - z$). The linear basis accepts (in initialization and propagation) any linear polynomial. Our definition of **admit** is quite narrow (to accelerate calls to **GB**), but we ensure that both ideals accept equalities, since **extraProp** generates these. In our experiments, we consider some other definitions of **admit**, but they do not improve performance.

Extra Propagation. Our **extraProp** subroutine simply implements congruence for bitsums. That is, consider the following polynomials, with $m < \log_2 |\mathbb{F}|$:

$$Y - \sum_{i=1}^m 2^{i-1} X_i \qquad Y' - \sum_{i=1}^m 2^{i-1} X'_i$$

If all X_i and X'_i are known to have value zero or one (because $X_i^2 - X_i$ is in some $\langle B_j \rangle$) and Y and Y' are known to be equal ($Y - Y'$ is in some $\langle B_j \rangle$), then it propagates $X_i - X'_i$ for all i . Similarly, if Y is known to be a constant c ($Y - c$ is in some $\langle B_j \rangle$), then each X_i must be equal to the i^{th} bit of c as an unsigned integer. Soundness for **extraProp** follows from bit representation uniqueness.

Inter-basis Interactions. SplitGB treats each B_i as a source of polynomials that might be added to other B_j . It does not use $\langle B_i \rangle$ as the source; this would be sound, but enumerating the infinite set $\langle B_i \rangle$ is impossible. The natural question is whether inter-basis propagation within SplitGB is nevertheless complete, that is, whether all polynomials $p \in \langle B_i \rangle$ that are admissible to B_j are in the ideal generated by the polynomials actually added to B_j .

We have both positive and negative results for BitSplit: Lemma 1 shows that propagation from the sparse basis to the linear basis **is** complete. The proof is in Appendix E. Example 1 shows that propagation from the linear basis to the sparse basis **is not** complete. There is a natural way to fix this: enumerate each variable pair X, Y , and propagate $X - Y$ to the sparse basis if $X - Y$ is in the ideal generated by the linear basis. However, our experiments (Sect. 5) show that this doesn't empirically improve solver performance for our benchmarks.

Lemma 1. *Let B be a Gröbner basis under a graded order (a degree compatible order, i.e., for all monomials p, q , $\deg(p) < \deg(q) \implies p < q$); then, every linear $p \in \langle B \rangle$ is in the ideal generated by the linear elements of B .*

Example 1. Consider $\mathbb{F}_5[W, X, Y, Z]$ in grevlex order. Then $B_1 = \{W - X - Y + Z, Y - Z\}$ is a GB. The only polynomial in B_1 that is admissible to the sparse basis is $Y - Z$. Now consider $W - X$. It is in $\langle B_1 \rangle$ (it is the sum of B_1 's elements) and it is admissible to the sparse basis. However, it is not in $\langle Y - Z \rangle$; i.e., it **is not** generated by the subset of B_1 that is admissible to the sparse basis.

Connections. In some respects, our $\mathbb{F}$ -solver resembles two prior SMT ideas: theory combination and portfolio solving with clause sharing. As in theory combination [6], we reduce a problem (a system of field equations) to sub-problems (subsets of the original system) that are handled by loosely-coupled sub-solvers (bases and propagators), each using different reasoning. As in portfolio solving with clause sharing [44, 61], each sub-solver derives lemmas in a common language (not clauses, but polynomials) that they share with one another. Our work also resembles a prior combination of algebraic and propositional reasoning for preprocessing Boolean formulas by sharing $\mathbb{F}_2$ equations between algebraic and propositional modules [14]. However, our focus is on solving equations in a very large finite field with constraints of different structure.

Efficiency. In the worst case, BitSplit builds a GB for the full system (similar to Monolithic). A GB for degree- d polynomials in n variables can have size d^{2^n} [45], so the worst-case complexity of BitSplit (and Monolithic) is doubly exponential.

However, in the next section we will see that BitSplit is efficient on a number of problems of practical interest. For these problems it improves exponentially on Monolithic. Here, we give intuition for the source of the advantage. Consider a bitsum-heavy determinism problem. As discussed in Sect. 3, computing a full GB is hard, so Monolithic performs poorly. However, BitSplit can use extraProp to reason about the uniqueness of the bit-splitting and use its split GB to reason about other parts of the system. This might allow it to refute the system of equations without ever directly computing a GB for the full system.

5 Experiments

Now we present our experiments, which answer three empirical questions:

1. How does BitSplit perform when solving bitsum-heavy determinism queries?
(*Exponentially better than the prior state of the art.*)
2. How does BitSplit perform when solving other queries?
(*Similar to the prior state of the art.*)
3. How do BitSplit’s components impact its performance? (*Propagation is key.*)

We implement BitSplit in `cvc5` [4] as a solver for the theory of finite fields. This includes preprocessing that identifies bitsums in larger polynomials and isolates them for use in BitSplit. Our test bed is a cluster with Intel Xeon E5-2637 v4 CPUs. Each run gets one CPU, 8GB memory, and a time limit of 300 s. After presenting the benchmarks, we compare BitSplit to prior SMT $\mathbb{F}$ -solvers `ffsat` [38]⁴ and `Monolithic` [50], and we compare BitSplit to variants of itself.

5.1 Benchmarks

Table 4 shows our benchmarks, most of which concern the correctness of ZK libraries (`circomlib` [7]) and compilers (`ZoKrates` [23] and `CirC` [49]). There are six families. The `CirC-D` benchmarks verify the determinism of operator encoding rules in `CirC`, at bitwidths up to 32. As we discuss in the next section (Sect. 6), these benchmarks are important to `CirC`’s correctness, but are hard to solve. The `Seq` benchmarks verify the determinism of constraint systems with

Table 4. Our benchmark families. `QED2` [53], `Small` [38], `TV` [50], and `CirC-S` [52] are from prior work. `CirC-D` is a set of large determinism benchmarks based on prior work [52]; see Sect. 6. `Seq` is a set of determinism benchmarks for computations that perform a sequence of bit-splits; see Appendix F.

Family	#	Description
<code>CirC-D</code>	640	Determinism for <code>CirC</code> $\mathbb{F}$ -blaster rules of bitwidth ≤ 32 (Sect. 6)
<code>Seq</code>	100	Determinism for sequenced bit-splits (Appendix F)
<code>QED²</code>	100	Determinism for <code>circomlib</code> , generated by <code>QED²</code> [53]
<code>CirC-S</code>	100	Soundness for <code>CirC</code> $\mathbb{F}$ -blaster rules of bitwidth ≤ 4 [52]
<code>TV</code>	100	Translation validation for ZKP compilers on boolean programs [50]
<code>Small</code>	100	Randomly generated with a small field: $ \mathbb{F} \leq 211$ [38]

⁴ At the time of our experiments, `ffsat` was a Sage-based Python tool for solving conjunctions of equations [35]. We wrapped it with a simple SMT-LIB parser that invokes `ffsat` if the query is sufficiently simple. Since then, `ffsat` has been re-implemented in Yices [22, 37]; future work should compare against that implementation.

sequences of bit-splits. We discuss them further in Appendix F. The QED² benchmarks are determinism queries for circomlib generated by QED² [53]. The CirC-S benchmarks are soundness tests for CirC’s operator rules, at bitwidths up to 4 [52]. The TV benchmarks are translation validation queries for ZoKrates and CirC, as applied to boolean functions [50]. Finally, the Small benchmarks are random, small-field (i.e., $|\mathbb{F}| < 2^8$) benchmarks from the evaluation of `ffsat` [38]. To keep the benchmark set from being too big, all families from prior work are sampled at random from that work’s benchmarks.

5.2 Comparison to Prior Solvers

First, we compare BitSplit against prior solvers Monolithic [50] and `ffsat` [38]. Table 5 shows the number of solved benchmarks by family and result. `ffsat` is successful only when the field is small. BitSplit improves on Monolithic on families that test determinism (QED² and CirC-D) but suffers slightly on other benchmarks. BitSplit is slightly worse on SAT instances but better at UNSAT ones. Figure 5 presents the same results as cactus plots for the determinism families and the other families.

To better understand BitSplit’s advantage, we focus on the CirC-D family. Each CirC-D benchmark tests the determinism of an operator rule at a specific bitwidth. We consider how the solve time scales with bitwidth. Figure 6 shows the results for arithmetic, shift, and comparison operators. Monolithic’s solve time grows exponentially for all of these, while BitSplit’s time is generally insignificant.

Table 5. Solved benchmarks, by family and result. BitSplit’s gains are on determinism queries (the QED² and CirC-D families) and unsatisfiable benchmarks.

Solver	Solved	By Family					By Result		
		CirC-D	Seq	QED ²	CirC-S	TV	Small	SAT	UNSAT
BitSplit	969	582	100	59	92	70	66	88	881
Monolithic	475	191	13	38	94	72	67	90	385
<code>ffsat</code>	67	0	0	0	0	0	67	54	13




Fig. 5. On determinism benchmarks, BitSplit dominates Monolithic; on other benchmarks, they perform similarly.

BitSplit struggles only with division and remainder; verifying their determinism would require understanding that integer division is deterministic, as encoded in field constraints. We omit bitwise operators (e.g., `bvor`) from this experiment. Their operator rules assume that the input bit-vectors are *already represented as bits*, so their benchmarks do not include any bitsums. To summarize, BitSplit can verify many operators exponentially faster than Monolithic.

5.3 Comparison to Variants

To better understand BitSplit, we compare it against six variants of itself:

- BS-LinFirst: make the linear basis (not the sparse basis) B_1
- BS-NoIntProp disable inter-basis propagation
- BS-NoExtProp disable extraProp
- BS-FullIntProp: complete linear-to-sparse propagation (Sect. 4.3, fixes Example 1)
- BS-DenseProp for the sparse basis, use $\text{admit}(p) = \deg(p) \leq 1 \wedge |\text{vars}(p)| \leq 16$.
- BS-QuadProp for the linear basis, use $\text{admit}(p) = \deg(p) \leq 2$.




Fig. 6. Solve time for CirC-D benchmarks for different operators. Monolithic’s solve time grows exponentially, while BitSplit’s solve time usually does not.

Table 6. BitSplit v. variants of itself. Weaker propagation (BS-NoExtProp, BS-NoIntProp) gives worse results, but other changes have less impact.

Solver	Solved	By Family					By Result			
		CirC-D	Seq	QED ²	CirC-S	TV	Small	SAT	UNSAT	
BitSplit	969	582	100	59	92	70	66	88	881	
BS-LinFirst	959	576	100	58	92	69	64	84	875	
BS-NoIntProp	877	576	24	58	86	70	63	84	793	
BS-NoExtProp	344	131	0	34	45	69	65	85	259	
BS-FullIntProp	953	576	97	56	92	69	63	83	870	
BS-DenseProp	898	580	33	58	92	70	65	85	813	
BS-QuadProp	898	580	32	59	92	71	64	86	812	
Monolithic	475	191	13	38	94	72	67	90	385	

Table 6 shows how many benchmarks each variant solves, with both `BitSplit` and `Monolithic` for comparison. First, changing the basis order (`BS-LinFirst`) has little effect. Second, disabling propagation (`BS-NoIntProp` and `BS-NoExtProp`) significantly hurts performance. Third, making inter-basis propagation complete (`BS-FullIntProp`) actually hurts performance slightly, perhaps because it takes quadratic time. Finally, defining `admit` more admissibly (`BS-DenseProp` and `BS-QuadProp`) makes little difference for many families, but significantly hurts performance on sequential bit-splits.

These results justify the key role that propagation plays in `BitSplit`. They also suggest that `BitSplit` would be a good choice for `cvc5`'s default field solver.

6 Application

Prior work uses `Monolithic` to do bounded verification for a zero-knowledge proof (ZKP) compiler pass [52]. In this section, we improve their results using `BitSplit`. Thus, this section is a case study that shows the utility of `BitSplit` for a downstream verification task. Our improvement relies not just on a new solver (`BitSplit`), but also on a new verification strategy. First (Sect. 6.1), we give background on the verification task. Second (Sect. 6.2), we state our new strategy, prove it is correct, and show that it is more efficient—when using `BitSplit`.

6.1 Background on Verifiable Field-Blasting

We consider the *finite field blaster* in a ZKP compiler: its responsibilities include encoding bit-vector operations as field equations [52]. At a high level, the field blaster is a collection of *encoding rules*. Each rule is a small algorithm that is specific to some operator (e.g., `bvadd`). It is given field variables that encode the operator's inputs according to some *encoding scheme*. A rule defines new variables, creates equations, and ultimately returns a field variable that encodes the output of the rule's operator.

As an example, we describe an encoding scheme for bit-vectors and a rule for bit-vector addition. The scheme encodes a length- b bit-vector x as a field variable x' with value in $\{0, \dots, 2^b - 1\} \subseteq \mathbb{F}$ (assuming $|\mathbb{F}| \gg 2^b$). If x' and x have the same (unsigned) integer value, we say that *valid*(x', x) holds. Suppose our rule applies to the addition of x and y , encoded as x' and y' . Our rule defines the following field variables. First, for each $i \in \{1, \dots, b+1\}$, it defines z'_i to 1 if the i^{th} bit of the integer sum of the unsigned values of x' and y' is one, and zero otherwise. Second, it defines $z' = \sum_{i=1}^b 2^{i-1} z'_i$. Then, it enforces these equations:

$$x' + y' = \sum_{i=1}^{b+1} 2^{i-1} z'_i \quad \wedge \quad z' = \sum_{i=1}^b 2^{i-1} z'_i \quad \wedge \quad \bigwedge_{i=1}^{b+1} z'_b (z'_b - 1) = 0$$

Finally, it returns z' . Informally, the idea of this rule is to bit-decompose the sum $x' + y'$ and then use the bit-decomposition to reduce that sum modulo 2^b . For example, if $b = 2$, $x' = 3$, and $y' = 1$, then the unique solution for the z'_i is $z'_1 = 0$, $z'_2 = 0$, $z'_3 = 1$, and then z' must be 0.

In general, an encoding rule for operator o maps a sequence of input encodings (field variables) $\mathbf{e}$ to three outputs: F , A , and e .⁵ Each field variable e_i encodes some bit-vector variable t_i . The first output, $F = \{z_1 \mapsto s_1, \dots, z_\ell \mapsto s_\ell\}$, is a mapping that defines ℓ fresh field variables: $z_1, \dots, z_\ell$. Variable z_i is mapped to a term s_i (in variables $\mathbf{e}$) that defines what value z_i is intended to take. The second output, A , is conjunction of field equations in variables $\mathbf{e}$ and $\mathbf{z}$. The final output is e : a distinguished variable that encodes the rule's output $o(\mathbf{t})$.

Prior work defines *correctness* for encoding rules as the conjunction of two properties: *completeness* and *soundness*. If all rules are correct, then they constitute a correct F-blaster [52]. Completeness says that if each e_i validly encodes t_i and the z_i take the values prescribed by F , then e validly encodes $o(\mathbf{t})$ and A holds. That is, completeness requires the following formula to be valid:

$$((\bigwedge_i \text{valid}(e_i, t_i)) \implies (A \wedge \text{valid}(e, o(\mathbf{t})))) [F]$$

Soundness says that if each e_i validly encodes t_i and A holds, then e validly encodes $o(\mathbf{t})$. That is, the following must be valid:

$$(A \wedge \bigwedge_i \text{valid}(e_i, t_i)) \implies \text{valid}(e, o(\mathbf{t}))$$

Verifier Performance. After fixing the sorts of the t_i (e.g., to bit-vectors of size 4), one can encode soundness and completeness as SMT queries. This enables automatic, bounded verification: one checks these properties up to some input bitwidth bound b using an SMT solver. However, the soundness query is especially challenging for the SMT solver. In prior work, some soundness queries for $b = 4$ could not be solved in 5 min with Monolithic. More generally, solving time grew exponentially with bit-width for most operators [52].

6.2 A New Strategy for Verifying Operator Rules

We propose a different strategy for automatically verifying operator rules. We define *determinism* for operator rules. It says that an operator rule applied to equal inputs should yield equal outputs. That is, if (A, e) and (A', e') are rule outputs for inputs $\mathbf{e}$ and $\mathbf{e}'$ respectively, then the following must be valid:

$$(A \wedge A' \wedge \mathbf{e} = \mathbf{e}') \implies e = e'$$

We prove the following theorem in Appendix G:

Theorem 4. *An operator rule that is deterministic and complete is also sound.*

⁵ Actually, in prior work [52] and in our implementation, encodings are *type-tagged sequences* of field *terms*. In this paper we treat them as single variables to simplify the exposition. Generalization is straightforward, but notationally tedious.




Fig. 7. The best way to verify that CirC rules are fully correct is to prove completeness using `exhaust` and prove determinism (D) using `BitSplit`.

Thus, to verify rule correctness, it suffices to verify completeness and **determinism**. This approach is promising because `BitSplit` is very effective on determinism queries (they were the CirC-D benchmarks in Sect. 5). So, a verification strategy comprises two choices: whether to prove soundness (S) or determinism (D) and whether to use `BitSplit` or `Monolithic`. In all cases, we prove completeness using `exhaust` (a specialized approach from prior work) [52]. For each strategy, we try to verify every bit-vector rule up to width 32. We limit SMT queries to 5 min each, using the same test bench as before.

Figure 7 shows verification time using different strategies. The best strategy is our new one. This approach verifies 66% more rule-bitwidth pairs than the next best strategy: proving soundness with `Monolithic`. More importantly, in our new strategy, verifying determinism (using `BitSplit`) is not the bottleneck: the bottleneck is proving completeness (using `exhaust`). Whereas, when proving soundness with `Monolithic`, `Monolithic` is the bottleneck. Further improvements will require new ideas for proving completeness.

7 Conclusion

We have presented a new approach for $\mathbb{F}$ -solving in SMT. Our contributions are three-fold. First, we proposed an abstract decision procedure `Split` that avoids computing a full Gröbner basis. Second, we described an instantiation of it (`BitSplit`) that is highly effective for bitsum-heavy determinism queries. Third, we applied `BitSplit` to a problem in ZKP compiler verification.

There are many directions for future work. First, we believe other instantiations of `Split` (beyond `BitSplit`) might be useful, for example, by considering other kinds of propagations (`extraProp`) and other conditions under which propagation is allowed (`admit`). Second, `Split` makes very limited use of CDCL(T) features that are known to improve performance: it acts only once a full propositional assignment is available; it constructs no theory lemmas; and it propagates no literals. Third, in this paper, we focus on applications of the theory of finite fields to ZKPs. Finite fields should also be relevant to many other kinds of cryptosystems, including algebraic multi-party computation and those based on elliptic curves. We leave these opportunities to future work.

Acknowledgements. We appreciate the help, support, and advice of Cesare Tinelli, Daniela Kaufmann, Haniel Barbosa, Mathias Preiner, Matthew Sotoudeh, Thomas Hader, the CAV reviewers, and all of the cvc5 developers.

This work was funded in part by NSF grant number 2110397, the Stanford Center for Automated Reasoning, and the Simons Foundation.

A Additional Background

This appendix is available in the full version of the paper [51].

B Computing Bitsum Usage in Real World Projects

This appendix is available in the full version of the paper [51].

C Proof of Theorem 1

This appendix is available in the full version of the paper [51].

D Proof of Theorems 2 and 3

This appendix is available in the full version of the paper [51].

E Proof of Lemma 1

This appendix is available in the full version of the paper [51].

F The Seq Benchmark Family

This appendix is available in the full version of the paper [51].

G Proof of Theorem 4

This appendix is available in the full version of the paper [51].

References

1. 0xPARC. ZK bug tracker. <https://github.com/0xPARC/zk-bug-tracker>. Accessed 5 Sept 2023, via archive.org
2. Anderson, B., McGrew, D.: TLS beyond the browser: Combining end host and network data to understand application behavior. In: IMC (2019)
3. Archer, D., O’Hara, A., Issa, R., Strauss, S.: Sharing sensitive department of education data across organizational boundaries using secure multiparty computation (2021)
4. Barbosa, H., et al.: cvc5: a versatile and industrial-strength SMT solver. In: TACAS (2022)
5. Barlow, R.: Computational thinking breaks a logjam (2015). <https://www.bu.edu/cise/computational-thinking-breaks-a-logjam/>
6. Barrett, C., Tinelli, C.: Satisfiability modulo theories. In: Handbook of Model Checking, pp. 305–343. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-10575-8_11
7. Bellés-Muñoz, M., Isabel, M., Muñoz-Tapia, J.L., Rubio, A., Baylina, J.: Circom: a circuit description language for building zero-knowledge applications. IEEE Trans. Dependable Secure Comput. (2022)
8. Bogetoft, P., et al.: Secure multiparty computation goes live. In: FC (2009)
9. Braun, D., Magaud, N., Schreck, P.: Formalizing some “small” finite models of projective geometry in coq. In: International Conference on Artificial Intelligence and Symbolic Computation (2018)
10. Buchberger, B.: A theoretical basis for the reduction of polynomials to canonical forms. SIGSAM Bulletin (1976)
11. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: IEEE S&P (2018)
12. Chaliasos, S., Ernstberger, J., Theodore, D., Wong, D., Jahanara, M., Livshits, B.: Sok: what don’t we know? understanding security vulnerabilities in snarks (2024). <https://arxiv.org/abs/2402.15293>
13. Chin, C., Wu, H., Chu, R., Coglio, A., McCarthy, E., Smith, E.: Leo: a programming language for formally verified, zero-knowledge applications (2021). Preprint at <https://ia.cr/2021/651>
14. Choo, D., Soos, M., Chai, K.M.A., Meel, K.S.: Bosphorus: Bridging anf and cnf solvers. IEEE, In DATE (2019)
15. Coglio, A., McCarthy, E., Smith, E., Chin, C., Gaddamadugu, P., Dellepère, M.: Compositional formal verification of zero-knowledge circuits (2023). <https://ia.cr/2023/1278>
16. Cohen, C.: Pragmatic quotient types in coq. In: ITP (2013)
17. Cox, D., Little, J., O’Shea, D.: Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra. Springer Science & Business Media (2013)
18. CVE-2014-3570. <https://nvd.nist.gov/vuln/detail/CVE-2014-3570>
19. CVE-2017-3732. <https://nvd.nist.gov/vuln/detail/CVE-2017-3732>
20. Dahlgren, F.: It pays to be Circomspect (2022). <https://blog.trailofbits.com/2022/09/15/it-pays-to-be-circomspect/>. Accessed 15 Oct 2023
21. Dummit, D.S., Foote, R.M.: Abstract algebra, vol. 3. Wiley Hoboken (2004)
22. Dutertre, B.: Yices 2.2. In: CAV (2014)
23. Eberhardt, J., Tai, S.: ZoKrates—scalable privacy-preserving off-chain computations. In: IEEE Blockchain (2018)

24. Enderton, H.B.: A mathematical introduction to logic. Elsevier (2001)
25. Erbsen, A., Philipoom, J., Gross, J., Sloan, R., Chlipala, A.: Systematic generation of fast elliptic curve cryptography implementations. Technical report, MIT (2018)
26. Erbsen, A., Philipoom, J., Gross, J., Sloan, R., Chlipala, A.: Simple high-level code for cryptographic arithmetic: With proofs, without compromises. *ACM SIGOPS Operating Syst. Rev.* **54**(1) (2020)
27. Y. Finance. Monero quote (2023). <https://finance.yahoo.com/quote/XMR-USD/>. Accessed 13 Oct 2023
28. Y. Finance. Zcash quote (2023). <https://finance.yahoo.com/quote/ZEC-USD/>. Accessed 13 Oct 2023
29. Fournet, C., Keller, C., Laporte, V.: A certified compiler for verifiable computing. In: CSF (2016)
30. Gabizon, A., Williamson, Z.J., Ciobotaru, O.: Plonk: permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge (2019). <https://ia.cr/2019/953>
31. Gonthier, G., et al.: A machine-checked proof of the odd order theorem. In: ITP, pp. 163–179 (2013)
32. Greuel, G.-M., Pfister, G., Schönemann, H.: Singular-a computer algebra system for polynomial computations. In: Symbolic Computation and Automated Reasoning, pp. 227–233. AK Peters/CRC Press (2001)
33. Groth, J.: On the size of pairing-based non-interactive arguments. In: EURO-CRYPT (2016)
34. Grubbs, P., Arun, A., Zhang, Y., Bonneau, J., Walfish, M.: Zero-knowledge middleboxes. In: USENIX Security (2022)
35. Hader, T.: Ffsat. <https://github.com/Ovascos/ffsat>, commit 67fecde
36. Hader, T.: Non-linear SMT-reasoning over finite fields (2022). MS Thesis (TU Wein)
37. Hader, T., Kaufmann, D., Irfan, A., Graham-Lengrand, S., Kovács, L.: Mcsat-based finite field reasoning in the yices2 smt solver (2024)
38. Hader, T., Kaufmann, D., Kovács, L.: SMT solving over finite field arithmetic. In: LPAR (2023)
39. Hader, T., Kovács, L.: Non-linear SMT-reasoning over finite fields. In: SMT (2022). Extended Abstract
40. Hopwood, D., Bowe, S., Hornby, T., Wilcox, N.: Zcash protocol specification (2013). <https://raw.githubusercontent.com/zcash/zips/master/protocol/protocol.pdf>
41. Komendantsky, V., Konovalov, A., Linton, S.: View of computer algebra data from coq. In: International Conference on Intelligent Computer Mathematics (2011)
42. Kotzias, P., Razaghpanah, A., Amann, J., Paterson, K.G., Vallina-Rodriguez, N., Caballero, J.: Coming of age: a longitudinal study of TLS deployment. In: IMC (2018)
43. Liu, J., et al.: Certifying zero-knowledge circuits with refinement types (2023). <https://ia.cr/2023/547>
44. Marescotti, M., Hyvärinen, A.E.J., Sharygina, N.: Clause sharing and partitioning for cloud-based SMT solving. In: Artho, C., Legay, A., Peled, D. (eds.) ATVA 2016. LNCS, vol. 9938, pp. 428–443. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46520-3_27
45. Mayr, E.W., Meyer, A.R.: The complexity of the word problems for commutative semigroups and polynomial ideals. *Adv. Math.* **46**(3), 305–329 (1982)
46. Monero technical specs (2022). <https://monerodocs.org/technical-specs/>

47. Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Solving SAT and SAT Modulo Theories: From an abstract davis–putnam–logemann–loveland procedure to DPLL(T). J. ACM (2006)
48. OpenSSL bug 1953. <https://www.mail-archive.com/openssl-dev@openssl.org/msg23869.html>
49. Ozdemir, A., Brown, F., Wahby, R.S.: CirC: compiler infrastructure for proof systems, software verification, and more. In: IEEE S&P (2022)
50. Ozdemir, A., Kremer, G., Tinelli, C., Barrett, C.: Satisfiability modulo finite fields. In: CAV (2023)
51. Ozdemir, S., Pailoor, A., Bassa, A., Ferles, K., Barrett, C., Dillig, I.: Split Gröbner bases for satisfiability modulo finite fields (2024). <https://ia.cr/2024/572>. Full version
52. Ozdemir, A., Wahby, R.S., Brown, F., Barrett, C.: Bounded verification for finite-field-blasting. In: CAV (2023)
53. Pailoor, S., et al.: Automated detection of under-constrained circuits in zero-knowledge proofs. In: PLDI (2023)
54. Philipoom, J.: Correct-by-construction finite field arithmetic in Coq. Ph.D. thesis, Massachusetts Institute of Technology (2018)
55. Schwabe, P., Viguier, B., Weerwag, T., Wiedijk, F.: A coq proof of the correctness of x25519 in tweetnacl. In: CSF (2021)
56. Soureshjani, F.H., Hall-Andersen, M., Jahanara, M., Kam, J., Gorzny, J., Ahmadvand, M.: Automated analysis of halo2 circuits (2023). <https://ia.cr/2023/1051>
57. Tornado.cash got hacked. by us (2019). <https://tornado-cash.medium.com/tornado-cash-got-hacked-by-us-b1e012a3c9a8>. Accessed 13 Oct 2023
58. Wang, D.: Elimination methods. Springer Science & Business Media (2001)
59. Wang, F.: Ecne: automated verification of zk circuits (2022). <https://0xparc.org/blog/ecne>
60. Wen, H., et al.: Practical security analysis of zero-knowledge proof circuits (2023)
61. Wintersteiger, C.M., Hamadi, Y., de Moura, L.: A concurrent portfolio approach to SMT solving. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 715–720. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02658-4_60
62. Zcash counterfeiting vulnerability successfully remediated (2019). <https://electriccoin.co/blog/zcash-counterfeiting-vulnerability-successfully-remediated/>. Accessed 13 Oct 2023

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.

