



Bounded verification for finite-field-blasting in a compiler for zero knowledge proofs

Alex Ozdemir¹ · Riad S. Wahby² · Fraser Brown² · Clark Barrett¹

Received: 16 March 2024 / Accepted: 30 March 2025 / Published online: 3 May 2025

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2025

Abstract

Zero Knowledge Proofs (ZKPs) are cryptographic protocols by which a prover convinces a verifier of the truth of a statement without revealing any other information. Typically, statements are expressed in a high-level language and then compiled to a low-level representation on which the ZKP operates. Thus, *a bug in a ZKP compiler can compromise the statement that the ZK proof is supposed to establish*. This paper takes a step towards ZKP compiler correctness by partially verifying a *field-blasting* compiler pass, a pass that translates Boolean and bit-vector logic into equivalent operations in a finite field. First, we define correctness for field-blasters and ZKP compilers more generally. Next, we describe the specific field-blaster using a set of encoding rules and define verification conditions for individual rules. Finally, we connect the rules and the correctness definition by showing that if our verification conditions hold, the field-blaster is correct. We have implemented our approach in the CirC ZKP compiler and have proved bounded versions of the corresponding verification conditions using an SMT solver. We show that our partially verified field-blaster does not hurt the performance of the compiler or its output; we also report on four bugs uncovered during verification.

Keywords Zero-knowledge proofs · Finite fields · Compilers · Verification · SMT

1 Introduction

Zero-Knowledge Proofs (ZKPs) are powerful tools for building privacy-preserving systems. They allow one entity, the *prover* \mathcal{P} , to convince another, the *verifier* \mathcal{V} , that some secret data satisfies a public property, *without revealing anything else about the data*. ZKPs underlie a large (and growing!) set of critical applications, from billion-dollar private cryptocurrencies, like Zcash [25, 55] and Monero [2], to research into auditable sealed court orders [21], private gun registries [27], privacy-preserving middleboxes [24], and zero-knowledge proofs of exploitability [12]. This breadth of applications is possible because of the generality of ZKPs. In general, \mathcal{P} knows a secret *witness* w , whereas \mathcal{V} knows a *property*

✉ Alex Ozdemir
aozdemir@stanford.edu

¹ Stanford, Stanford, USA

² Carnegie Mellon, Pittsburgh, USA

ϕ and a public *instance* x . \mathcal{P} must show that $\phi(x, w) = \top$. Typically, x and w are vectors of variables in a finite field \mathbb{F} , and ϕ can be any system of equations over the variables, using operations $+$ and \times . Because ϕ itself is an input to \mathcal{P} and \mathcal{V} , and because of the expressivity of field equations, a single implementation of \mathcal{P} and \mathcal{V} can serve many different purposes.

Humans find it difficult to express themselves directly with field equations, so they use *ZKP compilers*. A ZKP compiler converts a high-level predicate ϕ' into an equivalent system of field equations ϕ . In other words, a ZKP compiler *generalizes* a ZKP: by compiling ϕ' to ϕ and then using a ZKP for ϕ , one obtains a ZKP for ϕ' . There are many industrial [3, 5, 6, 15, 22, 46, 57, 68] and academic [4, 19, 29, 30, 47, 49, 52, 56, 65] ZKP compilers.

The correctness of a ZKP compiler is critical for security—a bug in the compiler could admit proofs of false statements—but verification is challenging for three reasons. First, the definition of correctness for a ZKP compiler is non-trivial; we discuss later in this section. Second, ZKP compilers span multiple domains. The high-level predicate ϕ' is typically expressed in a language with common types such as Booleans and fixed-width integers, while the output ϕ is over a large, prime-order field. Thus, any compiler correctness definition must span these domains. Third, ZKP compilers are evolving and performance-critical; verification must not inhibit future changes or degrade compiler performance.

In this work, we develop tools for automatically verifying the *field-blaster* of a ZKP compiler. A ZKP compiler’s field-blaster is the pass that converts from a formula over Booleans, fixed-width integers, and finite-field elements, to a system of field equations; as a transformation from bit-like types to field equations, the field-blaster exemplifies the challenge of cross-domain verification.

Our paper makes three contributions. First, we formulate a precise correctness definition for a ZKP compiler. Our definition ensures that a correct compiler preserves the completeness and soundness of the underlying ZK proof system.¹ More specifically, given a ZK proof system where statements are specified in a low-level language L , and a compiler from a high-level language H to L , if the compiler is correct by our definition, it extends the ZK proof system’s soundness and completeness properties to statements in H . Further, our definition is preserved under sequential composition, so proving the correctness of each compiler pass individually suffices to prove correctness of the compiler itself.

Second, we give an architecture for a verifiable field-blaster. In our architecture, a field-blaster is a set of “encoding rules.” We give verification conditions (VCs) for these rules, and we show that if the VCs hold, then the field-blaster is correct. Our approach supports *automated* verification because (bounded versions of) the VCs can be checked automatically. This reduces both the up-front cost of verification and its maintenance cost.

Third, we do a case study. Using our architecture, we implement a new field-blaster for CirC [47] (“SIR-see”), an infrastructure used by state-of-the-art ZKP compilers. We verify bounded versions of our field-blaster’s VCs using SMT-based finite-field reasoning [48], and show that our field-blaster does not compromise CirC’s performance. We also report on four bugs that our verification effort uncovered, including a soundness bug that allowed the prover to “lie” about the results of certain bit-vector comparisons. We note that the utility of our techniques is not limited to CirC: our ideas would apply equally well to other compilers to field-based ZKPs [15, 46, 68, 70].

¹ Roughly speaking, a ZK proof system is complete if it is possible to prove every true statement, and is sound if it is infeasible to prove false ones.

In the next sections, we discuss related work (Sect.1.1), give background on ZKPs and CirC (Sect.2), present a field-blasting example (Sect.3), describe our architecture (Sect.4), give our verification conditions (Sect.5), and present the case study (Sect.6).

An abridged version of this paper was published in the 2023 proceedings of CAV [50]. In this expanded version, we give additional details. First, we discuss field-blaster optimization (Sect. 4.4) and give proofs of the main theorems (Thms. 1 and 3). Second, we include more detailed experimental results on how VC verification time (Figs. 7 and 8), and the performance of the verified field-blaster (Fig. 9). Finally, we explain all the bugs found in CirC (Sect. 6.3). The artifact for this paper is at <https://doi.org/10.5281/zenodo.7922914>.

1.1 Related work

Verified compilers There is a rich body of work on verifying the correctness of traditional compilers. We focus on compilation for ZKPs; this requires different correctness definitions that relate bit-like types to prime field elements. In the next paragraphs, we discuss more fine-grained differences.

Compiler verification efforts fall into two broad categories: *automated*—verification leveraging automated reasoning solvers—and *foundational*—manual verification using proof assistants (e.g., Coq [9] or Isabelle [45]). CompCert [37], for example, is a Coq-verified C compiler with verified optimization passes (e.g., [41]). Closest to our work is backend verification, which proves correct the translation from an intermediate representation to machine code. CompCert’s lowering [38] is verified, as is CakeML’s [32] lowering to different ISAs [20, 59]. While such foundational verification offers strong guarantees, it imposes a heavy proof burden; creating CompCert, for example, took an expert team eight years [58], and any updates to compiler code require updates to proofs.

Automated verification, in contrast, does not require writing and maintaining manual proofs.² Cobalt [35], Rhodium [36], and PEC [33] are domain-specific languages (DSLs) for writing automatically-verified compiler optimizations and analyses. Most closely related to our work is Alive [40], a DSL for expressing verified peephole optimizations, local rewrites that transform snippets of LLVM IR [1] to better-performing ones. Alive addresses transformations over fixed types (while we address lowering to finite field equations) and formulates correctness in the presence of undefined behavior (while we formulate correctness for ZKPs). Beyond Alive, Alive2 [39] provides translation validation [42, 53] for LLVM [34], and VeRA [11] verifies range analysis in the Firefox JavaScript engine.

There is also work on verified compilation for domains more closely related to ZKPs. The Porcupine [16] compiler automatically synthesizes representations for fully-homomorphic encryption [64], and Gillar [60] proves that optimization passes in the Qiskit [62] quantum compiler are semantics-preserving. While these works compile from high-level languages to circuit representations, the correctness definitions for their domains do not apply to ZKP compilers.

Verified compilation to cryptographic proofs Prior works on verified compilation for ZKPs (or similar) take the foundational approach (with attendant proof maintenance burdens), and they do not formulate a satisfactory definition of compiler correctness.

² Automated verification generally leverages solvers. This is a particularly appealing approach in our setting, since CirC (our compiler infrastructure of interest) already supports compilation to SMT formulas.

PinocchioQ [19] builds on CompCert [37]. The authors formulate a correctness definition that preserves the *existential soundness* of a ZKP but does not consider completeness, knowledge soundness, or zero-knowledge (see Sect. 2.2). Leo [15] is a ZKP compiler that produces (partial) ACL2 [28] proofs of correct compilation; work to emit proofs from its field-blaster is ongoing.

Recent work defines security for *reductions of knowledge* [31]. These let \mathcal{P} convince \mathcal{V} that it knows a witness for an instance of relation \mathcal{R}_1 by proving it knows a witness for an instance of an easier-to-prove relation \mathcal{R}_2 . Unlike ZKP compilers, \mathcal{P} and \mathcal{V} *interact* to derive \mathcal{R}_2 using \mathcal{V} 's randomness (e.g., proving that two polynomials are nonzero w.h.p. by proving that a random linear combination of them is), whereas ZKP compilers run ahead of time and non-interactively.

Further afield, Ecne [67] is a tool that attempts to verify that the input to a ZKP encodes a *deterministic* computation. It does not consider any notion of a specification of the intended behavior. A different work [26] attempts to automatically verify that a “widget” given to a ZKP meets some specification. They consider widgets that could be constructed manually or with a compiler. Our focus is on verifying a compiler pass.

2 Background

2.1 Logic

We assume usual terminology for many-sorted first-order logic with equality ([18] gives a complete presentation). We assume every signature includes the sort `Bool`, constants `True` and `False` of sort `Bool`, and symbol family \approx_σ (the equality operator, abbreviated \approx) with sort $\sigma \times \sigma \rightarrow \text{Bool}$ for each sort σ . We also assume a family of conditionals: symbols ite_σ (“if-then-else”, abbreviated *ite*) of sort $\text{Bool} \times \sigma \times \sigma \rightarrow \sigma$.

A *theory* is a pair $\mathcal{T} = (\Sigma, \mathbf{I})$, where Σ is a signature and \mathbf{I} is a class of Σ -interpretations. A Σ -*formula* is a term of sort `Bool`. A Σ -formula ϕ is *satisfiable* (resp., *unsatisfiable*) in \mathcal{T} if it is satisfied by some (resp., no) interpretation in \mathbf{I} . We focus on two theories. The first is \mathcal{T}_{BV} , the SMT-LIB theory of bit-vectors [54, 63], with signature Σ_{BV} including a bit-vector sort $\text{BV}_{[n]}$ for each $n > 0$ with bit-vector constants $c_{[n]}$ of sort $\text{BV}_{[n]}$ for each $c \in [0, 2^n - 1]$, and operators including $\&$ (bitwise and, or) and $+_{[n]}$ (addition modulo 2^n). We write $t[i]$ to refer to the i 'th bit of bit-vector t , where $t[0]$ is the least-significant bit. The other theory is \mathcal{T}_{F_p} , which is the theory corresponding to the finite field of order p , for some prime p [48]. This theory has signature Σ_{F_p} containing the sort FF_p , constant symbols $0, \dots, p - 1$, and operators $+$ and \times .

In this paper, we assume all interpretations interpret sorts and symbols in the same way. We write $\text{dom}(v)$ for the set interpreting the sort of a variable v . We assume that `Bool`, `True`, and `False` are interpreted as $\{\top, \perp\}$, \top , and \perp , respectively; Σ_{BV} -interpretations follow the SMT-LIB standard; and Σ_{F_p} -interpretations interpret symbols as the corresponding elements and operations in \mathbb{F}_p , a finite field of order p (for concreteness, this could be the integers modulo p). Note that only the values of variables can vary between two interpretations.

For a signature Σ , let t be a Σ -term of sort σ , with free variables x_1, \dots, x_n , respectively of sort $\sigma_1, \dots, \sigma_n$. We define the function $\hat{t} : \text{dom}(x_1) \times \dots \times \text{dom}(x_n) \rightarrow \text{dom}(t)$ as follows. Let $\mathbf{x} \in \text{dom}(x_1) \times \dots \times \text{dom}(x_n)$. Let \mathcal{M} be an interpretation that interprets each x_i as \mathbf{x}_i . Then $\hat{t}(\mathbf{x}) = t^{\mathcal{M}}$ (i.e., the interpretation of t in \mathcal{M}). For example, the term $t = a \wedge \neg a$ defines $\hat{t} : \text{Bool} \rightarrow \text{Bool} = \lambda x. \perp$. In the following, we follow the convention used above in using the

standard font (e.g., x) for logical variables and a sans serif font (e.g., \mathbf{x}) to denote meta-variables standing for values (i.e., elements of $\sigma^{\mathcal{M}}$ for some σ and \mathcal{M}). Also, abusing notation, we'll conflate single variables (of both kinds) with vectors of variables when the distinction doesn't matter. Note that a formula ϕ is *satisfiable* if there exist values \mathbf{x} such that $\hat{\phi}(\mathbf{x}) = \top$. It is *valid* if for all values \mathbf{x} , $\hat{\phi}(\mathbf{x}) = \top$.

For terms s, t and variable x , $t[x \mapsto s]$ denotes t with all occurrences of x replaced with s . For a sequence of variable-term pairs, $S = (x_1 \mapsto s_1, \dots, x_n \mapsto s_n)$, $t[S]$ is defined to be $t[x_1 \mapsto s_1] \cdots [x_n \mapsto s_n]$.

2.2 Zero knowledge proofs

As mentioned above, Zero-knowledge proofs (ZKPs) make it possible to prove that some secret data satisfies a public property—without revealing the data itself. See [61] for a full presentation; we give a brief overview here, and then describe how general-purpose ZKPs are used.


Overview and definitions In a cryptographic proof system, there are two parties: a *verifier* \mathcal{V} and a *prover* \mathcal{P} . \mathcal{V} knows a public *instance* x and asks \mathcal{P} to show that it has knowledge of a secret *witness* w satisfying a public *predicate* $\phi(x, w)$ from a predicate class Φ (a set of formulas) (i.e., $\hat{\phi}(x, w) = \top$). Figure 1 illustrates the workflow. First, a trusted party runs an efficient (i.e., polytime in an implicit security parameter λ) algorithm $\text{Setup}(\phi)$ which produces a *proving key* pk and a *verifying key* vk . Then, \mathcal{P} runs an efficient algorithm $\text{Prove}(\text{pk}, x, w) \rightarrow \pi$ and sends the resulting *proof* π to \mathcal{V} . Finally, \mathcal{V} runs an efficient verification algorithm $\text{Verify}(\text{vk}, x, \pi) \rightarrow \{\top, \perp\}$ that accepts or rejects the proof. A zero-knowledge argument of knowledge for class Φ is a tuple $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$ with three informal properties for every $\phi \in \Phi$ and every $x \in \text{dom}(x), w \in \text{dom}(w)$:

- *Perfect completeness*: if $\hat{\phi}(x, w)$ holds, then $\text{Verify}(\text{vk}, x, \pi)$ holds;
- *Computational knowledge soundness* [10]: an efficient adversary that does not know w cannot produce a π such that $\text{Verify}(\text{vk}, x, \pi)$ holds; and
- *Zero-knowledge* [23]: π reveals nothing about w , other than its existence.

Technically, the system is an “argument” rather than a “proof” because soundness only holds against efficient adversaries. Also note that knowledge soundness requires that an entity must “know” a valid w' to produce a proof; it insufficient w' to simply exist. See the full version for precise definitions [51, App. A].

Representations for ZKPs As mentioned above, ZKP applications are manifold (Sect. 1)—from cryptocurrencies to private registries. This breadth of applications is possible because ZKPs support a broad class of predicates. Most commonly, these predicates are expressed as *rank-1 constraint systems* (R1CSs). Recall that \mathbb{F}_p is a prime-order finite field (also called a *prime field*). We will drop the subscript p when it is not important. In an R1CS, x and w are vectors of elements in \mathbb{F} ; let $z \in \mathbb{F}^m$ be their concatenation. The function $\hat{\phi}$ can be defined by three matrices $A, B, C \in \mathbb{F}^{n \times m}$; $\hat{\phi}(x, w)$ holds when $Az \circ Bz = Cz$, where \circ is the element-wise product. Thus, ϕ can be viewed as n conjoined *constraints*, where each constraint i is of the form $(\sum_j a_{ij}z_j) \times (\sum_j b_{ij}z_j) \approx (\sum_j c_{ij}z_j)$ (where the a_{ij}, b_{ij} and c_{ij} are constant symbols from

Fig. 1 The information flow for a zero-knowledge proof.



Σ_{F_p} , and the z_j are a vector of variables of sort FF_p). That is, each constraint enforces a single non-linear multiplication.

2.3 Compilation targeting zero knowledge proofs

To write a ZKP about a high-level predicate ϕ , that predicate is first compiled to an RICS. A *ZKP compiler* from class Φ (a set of Σ -formulas) to class Φ' (a set of Σ' -formulas) is an efficient algorithm $\text{Compile}(\phi \in \Phi) \rightarrow (\phi' \in \Phi', \text{Ext}_x, \text{Ext}_w)$.³ Given a predicate $\phi(x, w)$, it returns a predicate $\phi'(x', w')$ as well as two efficient and deterministic algorithms, instance and witness *extenders*: $\text{Ext}_x : \text{dom}(x) \rightarrow \text{dom}(x')$ and $\text{Ext}_w : \text{dom}(x) \times \text{dom}(w) \rightarrow \text{dom}(w')$.⁴ For example, CirC [47] can compile a Boolean-returning C function (in a subset of C) to an RICS.

At a high-level, ϕ and ϕ' should be “equisatisfiable”, with Ext_x and Ext_w mapping satisfying values for ϕ to satisfying values for ϕ' . That is, for all $x \in \text{dom}(x)$ and $w \in \text{dom}(w)$ such that $\hat{\phi}(x, w) = \text{T}$, if $x' = \text{Ext}_x(x)$ and $w' = \text{Ext}_w(x, w)$, then $\hat{\phi}'(x', w') = \text{T}$. Furthermore, for any x , it should be impossible to (efficiently) find w' satisfying $\hat{\phi}'(\text{Ext}_x(x), w') = \text{T}$ without knowing a w satisfying $\hat{\phi}(x, w) = \text{T}$. In Sect. 5.1, we precisely define correctness for a predicate compiler.

One can build a ZKP for class Φ from a compiler from Φ to Φ' and a ZKP for Φ' . Essentially, one runs the compiler to get a predicate $\phi' \in \Phi'$, as well as Ext_x and Ext_w . Then, one writes a ZKP to show that $\hat{\phi}'(\text{Ext}_x(x), \text{Ext}_w(x, w)) = \text{T}$. In the full version we detail this construction and prove it is secure [51, App. A].

Optimization The primary challenge when using ZKPs is cost: typically, Prove is at least three orders of magnitude slower than checking ϕ directly [66]. Since Prove’s cost scales with n (the constraint count), it is *critical* for the compiler to minimize n . The space of optimizations is large and complex, for two reasons. First, the compiler can introduce fresh variables. Second, only equisatisfiability—not logical equivalence—is needed. Compilers in this space exploit equisatisfiability heavily to efficiently represent high-level constructs (e.g., Booleans, bit-vectors, arrays, ...) as an RICS.

As a (simple!) example, consider the Boolean computation $a \approx (c_1 \vee \dots \vee c_k)$. Assume that c'_1, \dots, c'_k are variables of sort FF and that we add constraints $c'_i(1 - c'_i) \approx 0$ to ensure that c'_i has to be 0 or 1 for each i . Assume further that $(c'_i \approx 1)$ encodes c_i for each i . How can one additionally ensure that a' (also of sort FF) is also forced to be equal to 0 or 1 and that $(a' \approx 1)$ is a correct encoding of a ? Given that there are $k - 1$ ORs, natural approaches use $\Theta(k)$ constraints. One clever approach is to introduce variable x' and enforce constraints $x'(\sum_i c'_i) \approx a'$ and $(1 - a')(\sum_i c'_i) \approx 0$. In any interpretation where any c_i is true, the corresponding interpretation for a' must be 1 to satisfy the second constraint; setting x' to the sum’s inverse satisfies the first. If all c_i are false, the first constraint ensures a' is 0. This technique assumes the sum does not overflow; since ZKP fields are typically large (e.g., with p on the order of 2^{255}), this is usually a safe assumption.

CirC CirC [47] is an infrastructure for building compilers from high-level languages (e.g., a C subset), to RICSs. It has been used in research projects [4, 13], and in industrial R&D. Figure 2 shows the structure of an RICS compiler built with CirC. First, the front-end of the compiler converts the source program into CirC-IR. CirC-IR is a term IR based

³ While in Sect. 2.2, ϕ was a system of field equations, here—and for the rest of the paper— ϕ denotes a high-level predicate and ϕ' is the (low-level) field equations.

⁴ For technical reasons, the runtime of Ext_x and the size of its description must be $\text{poly}(\lambda, |x|)$ —not just $\text{poly}(\lambda)$ [51, App. A]




Fig. 2 The architecture of CirC.

on SMT-LIB that includes: Booleans, bit-vectors, fixed-size arrays, tuples, and prime fields.⁵ Second, the compiler optimizes and simplifies the IR so that the only remaining sorts are Booleans, bit-vectors, and the target prime field. Third, the compiler lowers the simplified IR to an R1CS predicate over the target field. For ZKPs built with CirC, *the completeness, soundness, and zero-knowledge of the end-to-end system depend on the correctness of CirC itself.*

3 Overview and example

To start, we view CirC’s lowering pass as two passes (Fig. 2). The first pass, “(finite-)field-blasting,” converts a many-sorted IR (representable as a $(\Sigma_{BV} \cup \Sigma_F)$ -formula) to a conjunction of field equations (Σ_F -equations). The second pass, “flattening,” converts this conjunction of field equations to an R1CS.

Our focus is on verifying the first pass. We begin with a worked example of how to field-blast a small snippet of CirC-IR (Sect. 3.1). This example will illustrate four key ideas (Sect. 3.2) that inspire our field-blaster’s architecture.

3.1 An example of field-blasting

We start with an example CirC-IR predicate expressed as a $(\Sigma_{BV} \cup \Sigma_F)$ -formula:

$$\phi \triangleq (x_0 \oplus w_0) \wedge (w_1 +_{[4]} x_1 \approx w_1) \wedge (x_2 \& w_1 \approx x_2) \wedge (x_3 \approx w_2 \times w_2) \tag{1}$$

The predicate includes: the XOR of two Booleans (“ \oplus ”), a bit-vector sum, a bit-vector AND, and a field product. x_0 and w_0 are of sort `Bool`, x_1 , x_2 , and w_1 are of sort `BV[4]`, and x_3 and w_2 are of sort `FFp`. We’ll assume that $p \gg 2^4$. Table 1 summarizes the new variables and assertions we create during field-blasting; we describe the origin of each assertion and new variable in the next paragraphs.

Lowering clause one (Booleans) We begin with the Boolean term $(x_0 \oplus w_0)$. We will use 1 and 0 to represent \top and \perp . We introduce variables x'_0 and w'_0 of sort `FFp` to represent x_0 and w_0 respectively. To ensure that w'_0 is 0 or 1, we assert: $w'_0(w'_0 - 1) \approx 0$.⁶ $x_0 \oplus w_0$ is then represented by the expression $1 - x'_0 - w'_0 + 2x'_0w'_0$. Setting this equal to 1 enforces

⁵ We list all CirC-IR operators for Booleans, bit-vectors, and prime fields in the full version [51, App. C] Almost all are from SMT-LIB.

⁶ Later (Sect. 5), we will see that “well-formedness” constraints like this are unnecessary for instance variables, such as x_0 .

Table 1 New variables and assertions when compiling the example ϕ

Clause	Term from ϕ	Assertions	New variables	Notes
1	x_0		x'_0	
	w_0	$w'_0(w'_0 - 1) \approx 0$	w'_0	
	$x_0 \oplus w_0$	$1 \approx 1 - w'_0 - x'_0 + 2w'_0x'_0$		
2	x_1		$x'_{1,u}$	
	w_1	$w'_{1,i}(w'_{1,i} - 1) \approx 0$	$w'_{1,i}$	$i \in [0, 3]$
	$x_1 +_{[4]} w_1$	$s' \approx x'_{1,u} + \sum_{i=0}^3 2^i w'_{1,i}$	s'	
		$s'_i(s'_i - 1) \approx 0$	s'_i	$i \in [0, 4]$
		$s' \approx \sum_{i=0}^4 2^i s'_i$		
$x_1 +_{[4]} w_1 \approx w_1$	$s'_i \approx w'_{1,i}$		$i \in [0, 3]$	
3	x_2		$x'_{2,u}$	
	x_2 (bits)	$x'_{2,i}(x'_{2,i} - 1) \approx 0$	$x'_{2,i}$	$i \in [0, 3]$
		$x'_{2,u} \approx \sum_{i=0}^3 2^i x'_{2,i}$		
$x_2 \& w_1 \approx x_2$	$x'_{2,i}w'_{1,i} \approx x'_{2,i}$		$i \in [0, 3]$	
4	x_3, w_2		x'_3, w'_2	
	$x_3 \approx w_2 \times w_2$	$x'_3 \approx w'_2 \times w'_2$		

that $x_0 \oplus w_0$ must be true. These new assertions and fresh variables are reflected in the first three rows of the table.

Lowering clause two and three (bit-vectors) Before describing how to bit-blast the second and third clauses in ϕ , we discuss bit-vector representations in general. A bit-vector t can be viewed as a sequence of b bits or as a non-negative integer less than 2^b . These two views suggest two natural representations in a prime-order field: first, as b elements t'_0, \dots, t'_{b-1} , that encode the bits of t as 0 or 1 (in our encoding, t'_0 is the least-significant bit and t'_{b-1} is the most-significant bit); second, as one field element t'_u , whose unsigned value agrees with t (assuming the field's size is at least 2^b). The first representation is simple, but with it, some field values (e.g., 2^b) don't corresponding to any possible bit-vector. With the second approach, by including equations $t'_i(t'_i - 1) \approx 0$ in our system, we ensure that any satisfying assignment corresponds to a valid bit-vector. However, the extra b equations increase the size of our compiler's output.

We represent ϕ 's w_1 bit-wise: as $w'_{1,0}, \dots, w'_{1,3}$, and we represent the instance variable x_1 as $x'_{1,u}$.⁷ For the constraint $w_1 +_{[4]} x_1 \approx w_1$, we compute the sum in the field and bit-decompose the result to handle overflow. First, we introduce new variable s' and set it equal to $x'_{1,u} + \sum_{i=0}^3 2^i w'_{1,i}$. Then, we bit-decompose s' , requiring $s' \approx \sum_{i=0}^4 2^i s'_i$, and $s'_i(s'_i - 1) \approx 0$ for $i \in [0, 4]$. Finally, we assert $s'_i \approx w'_{1,i}$ for $i \in [0, 3]$. This forces the lowest 4 bits of the sum to be equal to w_1 .

The constraint $x_2 \& w_1 \approx x_2$ is more challenging. Since x_2 is an instance variable, we initially encode it as $x'_{2,u}$. Then, we consider the bit-wise AND. There is no obvious way

⁷ We represent w_1 bit-wise so that we can ensure the representation is well-formed with constraints $w'_{1,i}(w'_{1,i} - 1) \approx 0$. As previously noted, such well-formedness constraints are not needed for an instance variable like x_1 .⁸

Table 2 Encodings for each term sort. Only bit-vectors have two encoding kinds

	Variant contents			Semantics
	encoded_term	Kind	Terms	Validity condition
t : Bool	bit	f	$f \approx ite(t, 1, 0)$	
t : BV _[b]	uint	f	$f \approx \sum_i ite(t[i] \approx 1_{[1]}, 2^i, 0)$	
t : BV _[b]	bits	f_0, \dots, f_{b-1}	$\bigwedge_i f_i \approx ite(t[i] \approx 1_{[1]}, 1, 0)$	
t : FF	field	f	$t \approx f$	

to encode a bit-wise operation, other than bit-by-bit. So, we convert $x'_{2,u}$ to a bit-wise representation: We introduce witness variables $x'_{2,0}, \dots, x'_{2,3}$ and equations $x'_{2,i}(x'_{2,i} - 1) \approx 0$ as well as equation $x'_{2,u} \approx \sum_{i=0}^3 2^i x'_{2,i}$. Then, for each i we require $x'_{2,i} w'_{1,i} \approx x'_{2,i}$.

Lowering the final clause (field elements) Finally, we consider the field equation $x_2 \approx w_2 \times w_2$. Our target is also field equations, so lowering this is straightforward. We simply introduce primed variables and copy the equation.

3.2 Key ideas

This example highlights four ideas that guide the design of our field-blaster:

1. *Fresh variables and assertions*: Field-blasting uses two primitive operations: creating new variables in ϕ' (e.g., w'_0 to represent w_0) and adding new assertions to ϕ' (e.g., $w'_0(w'_0 - 1) \approx 0$).
2. *Encodings*: For a term t in ϕ , we construct a field term (or collection of field terms) in ϕ' that represent the value of t . For example, the Boolean w_0 is represented as the field element w'_0 that is 0 or 1.
3. *Operator rules*: if t is an operator applied to some arguments, we can encode t given encodings of the arguments. For example, if t is $x_0 \oplus w_0$, and x_0 is encoded as x'_0 and w_0 as w'_0 , then t can be encoded as $1 - x'_0 - w'_0 + 2x'_0 w'_0$.
4. *Conversions*: Some sorts can be represented by encodings of different kinds. If a term has multiple possible encodings, the compiler may need to convert between them to apply some operator rule. For example, we converted x_2 from an unsigned encoding to a bit-wise encoding before handling an AND.

4 Architecture

In this section, we present our field-blaster architecture. To compile a predicate ϕ to a system of field equations ϕ' , our architecture processes each term t in ϕ using a post-order traversal. Informally, it represents each t as an “encoding” in ϕ' : a term (or collection of terms) over variables in ϕ' . Each encoding is produced by a small algorithm called an “encoding rule”.

Below, we define the type of encodings Enc (Sect. 4.1), the five different types of encoding rules (Sect. 4.2), and a calculus that iteratively applies these rules to compile all of ϕ (Sect. 4.3).

4.1 Encodings

Table 2 presents our tagged union type Enc of possible term encodings. Each variant comprises the term being encoded, its tag (the *encoding kind*), and a sequence of field terms. The encoding kinds are `bit` (a Boolean as 0/1), `uint` (a bit-vector as an unsigned integer), `bits` (a bit-vector as a sequence of bits), and `field` (a field term trivially represented as a field term). Each encoding has an intended semantics: a condition under which the encoding is considered valid. For instance, a bit-encoding of Boolean t is valid if the field term f is equal to $\text{ite}(t, 1, 0)$.

4.2 Encoding rules

An encoding rule is an algorithm that takes and/or returns encodings, in order to represent some part of the input predicate as field terms and equations.

4.2.1 Primitive operations

A rule can perform two primitive operations: creating new variables and emitting assertions. In our pseudocode, the primitive function $\text{fresh}(\text{name}, t, \text{isInst}) \rightarrow \text{Enc}$ creates a fresh variable. Argument `isInst` is a Boolean indicating whether the result x' is an instance variable (as opposed to a witness). Argument t is a field term (over variables from ϕ and previously defined primed variables) that expresses how to compute a value for x' . For example, to create a field variable w' that represents Boolean witness variable w , a rule can call $\text{fresh}(w', \text{ite}(w, 1, 0), \perp)$. The compiler uses t to help create the Ext_x and Ext_w algorithms. A rule asserts a formula t' (over primed variables) by calling $\text{assert}(t')$.

4.2.2 Rule Types

There are five types of rules: (1) Variable rules $\text{variable}(t, \text{isInst}) \rightarrow \text{Enc}$ take a variable t and its instance/witness status and return an encoding of that variable made up of fresh variables. (2) Constant rules $\text{const}(t) \rightarrow \text{Enc}$ take a constant term t and produce an encoding of t comprising terms that depend only on t . Since t is a constant, the terms in the result e can be evaluated to field constants (see the calculus in Sect. 4.3).⁸ The `const` rule cannot call `fresh` or `assert`. (3) Equality rules $\text{assertEq}(e, e')$ take two encodings of the same kind and emit assertions that equate the underlying terms. (4) Conversion rules $\text{convert}(e, \text{kind}') \rightarrow \text{Enc}$ take an encoding and convert it to an encoding of a different kind. Conversions are only non-trivial for bit-vectors, which have two encoding kinds: `uint` and `bits`. (5) Operator rules apply to terms t of form $o(t_1, \dots, t_n)$. Each operator rule takes t , o , and encodings of the child terms t_i and returns an encoding of

⁸ Having `const(t)` return terms that depend on t (rather than directly returning constants) is useful for constructing verification conditions for `const`.

```

fn variable(t, isInst) → Enc :
  if isInst:
    t' ← fresh(name(t) || 'u',
                ∑i ite(t[i] ≈ 1[1], 2i, 0), ⊤)
    return t, uint, t'
  else:
    for i in [0, size(sort(t)) - 1]:
      t'_i ← fresh(name(t) || i,
                  ite(t[i] ≈ 1[1], 1, 0), ⊥)
      assert(t'_i(t'_i - 1) = 0)
    return t, bits, t'_0, ..., t'_{size(sort(t))-1}

fn const(t) → Enc :
  for i in [0, size(sort(t)) - 1]:
    t'_i ← ite(t[i] ≈ 1[1], 1, 0)
  return t, bits, t'_0, ..., t'_{size(sort(t))-1}

fn assertEq(e : Enc, e' : Enc) :
  if kind(e) = bits:
    for i in [0, size(terms(e)) - 1]:
      assert(terms(e)[i] ≈ terms(e')[i])
  elif kind(e) = uint:
    assert(terms(e)[0] ≈ terms(e')[0])

fn convert(e : Enc, kind' : Kind) → Enc :
  t ← encoded_term(e)
  if kind(e) = bits and kind' = uint:
    return t, uint, ∑i 2iterms(e)[i]
  elif kind(e) = uint and kind' = bits:
    e' ← variable(t, ⊥)
    assert(terms(e)[0] ≈ ∑i 2iterms(e')[i])
  return e'

```

Fig. 3 Pseudocode for some bit-vector rules. “||” denotes string concatenation. `variable`’s implementation depends on `isInst`: it uses a small `uint` encoding for instances, but it must bit-split witnesses to ensure they have a value in $\{0, \dots, 2^{\text{size}(\text{sort}(t))} - 1\}$. `const` bit-splits the constant it’s given, `assertEq` asserts unsigned or bit-wise equality, and `convert` either does a bit-sum or bit-split.

```

fn bvZeroExt(t, o : Op, e : Enc) :
  if kind(e) = bits:
    w ← size(terms(e))
    for i in [0, w - 1]:
      t'_i ← terms(e)[i]
    for i in [0, o.newBits - 1]:
      t'_{w+i} ← 0
    return t, bits, t'_0, ..., t'_{w+o.newBits-1}
  else:
    return t, kind(e), terms(e)

fn bvMulUInt(t, o : Op, e' : [Enc]) :
  w ← size(sort(encoded_term(e[0])))
  W ← size(e') × w
  assume W < ⌊log2 p⌋
  s' ← ∏i terms(ei)[0]
  b ← ff2bv(W, s')
  for i in [0, W - 1]:
    t'_i ← fresh(i, ite(b[i], 1, 0), ⊥)
    assert(t'_i(t'_i - 1) ≈ 0)
  assert(s' ≈ ∑i=0W-1 2it'_i)
  return t, bits, t'_0, ..., t'_{w-1}

```

Fig. 4 Pseudocode for some bit-vector operator rules. `bvZeroExt` zero-extends a bit-vector; for bit-wise encodings, it adds zero bits, and for unsigned encodings, it simply copies the original encoding. `bvMulUInt` multiplies bit-vectors, all assumed to be unsigned encodings. We show only the case where the multiplication cannot overflow in the field: in this case the rule performs the multiplication in the field, and bit-splits the result to implement reduction modulo 2^b . If field overflow is possible, the field-blaster uses associativity to split the multiplication into smaller ones. If even a binary multiplication could overflow, the field-blaster throws an error. The rules use `ff2bv`, which converts from a field element to a bit-vector (discussed in Sect. 6.1).

t . Some operator rules require specific kinds of encodings; before using such an operator rule, our calculus (Sect. 4.3) calls the `convert` rule to ensure the input encodings are the correct kind. Figure 3 gives pseudocode for the first four rule types, as applied to bit-vectors. Figure 4 gives pseudocode for two bit-vector operator encoding rules. A field-blaster uses many operator rules: in our case study (Sect. 6) there are 46.

4.3 Calculus

We now give a non-deterministic calculus describing how our field-blaster applies rules to compile a predicate $\phi(x, w)$ into a system of field equations.

A calculus state is a tuple of three items: (E, A, F) . The *encoding store* E is a (multi-) map from terms to sets of encodings. The *assertions formula* A is a conjunction of all field equations asserted via `assert`. The *fresh variable definitions sequence* F is a sequence consisting of pairs, where each pair (v, t) matches a single call to `fresh(v, t, ...)`.

Figure 5 shows the transitions of our calculus. We denote the result of a rule as $A', F', e' \leftarrow r(\dots)$, where A' is a formula capturing any new assertions, F' is a sequence of pairs capturing any new variable definitions, and e' is the rule's return value. We may omit one or more results if they are always absent for a particular rule. For encoding store E , $E \cup (t \mapsto e)$ denotes the store with e added to t 's encoding set.

There are five kinds of transitions. The `Const` transition adds an encoding for a constant term. The `construle` returns an encoding e whose terms depend on the constant c ; e' is a new encoding identical to e , except that each of its terms has been evaluated to obtain a field constant. The `Var` transition adds an encoding for a variable term. The `Conv` transition takes a term that is already encoded and re-encodes it with a new encoding kind. The `kinds` operator returns all legal values of kind (Table 2) for encodings of a given sort. (That is, `uintand bitsfor` for bit-vectors, `bitfor` for Booleans, and `fieldfor` for field elements.) The `Opr` transition applies operator rule r . This transition is only possible if r 's operator kind agrees with o , and if its input encoding kinds agree with e . The `Finish` transition applies when ϕ has been encoded. It uses `constand assertEq` to build assertions that hold when $\phi = \top$. Rather than producing a new calculus state, it returns the outputs of the calculus: the assertions and the variable definitions.

To meet the requirements of the ZKP compiler, our calculus must return two extension function: `Extx` and `Extw` (Sect. 2.3). Both can be constructed from the fresh variable definitions F . One subtlety is that `Extx(x)` (which assigns values to fresh instance variables) is a function of x only—it cannot depend on the witness variables of ϕ . We ensure

$$\begin{array}{c}
 \frac{\text{constant term } c \quad e \leftarrow \text{const}(c) \quad e' \leftarrow \text{map}(\text{eval}, e)}{E := E \cup (c \mapsto e')} \text{Const} \\
 \\
 \frac{\text{variable term } v \quad A', F', e \leftarrow \text{variable}(v, \text{isInst}(v))}{E := E \cup (v \mapsto e), \quad A := A \wedge A', \quad F := F \parallel F'} \text{Var} \\
 \\
 \frac{(t \mapsto e) \in E \quad \text{kind} \in \text{kinds}(\text{sort}(t)) \quad A', F', e' \leftarrow \text{convert}(e, \text{kind})}{E := E \cup (t \mapsto e'), \quad A := A \wedge A', \quad F := F \parallel F'} \text{Conv} \\
 \\
 \frac{(t_i \mapsto e_i) \in E \quad t = o(\vec{t}) \quad A', F', e' \leftarrow r(t, o, \vec{e})}{E := E \cup (t \mapsto e'), \quad A := A \wedge A', \quad F := F \parallel F'} \text{Op}_r \\
 \\
 \frac{(\phi \mapsto e) \in E \quad e_{\top} \leftarrow \text{const}(\top) \quad A', F' \leftarrow \text{assertEq}(e, e_{\top})}{\text{return } (A \wedge A', \quad F \parallel F')} \text{Finish}
 \end{array}$$

Fig. 5 The transition rules of our rewriting calculus.

this by allowing fresh instance variables to only be created by the `variable` rule, and only when it is called with `isInst = T`.

4.4 Implementation and Optimization

4.4.1 Policy

Our calculus is non-deterministic because in some situations, multiple transitions are applicable. For instance, a conversion can apply at virtually any time. Some policy is needed to decide which transitions to apply. The strategy that decides which transition to apply affects field-blaster performance but *not* correctness.

Our implementation’s policy follows three rules. First, if an operator rule for an un-encoded term can be applied (i.e., if the necessary input encoding types are available), then apply it. Second, if only one rule matches the operator for an un-encoded term, then perform the necessary conversions to meet the constraints on that rule’s input encoding types. Third, if multiple rules match the operator of an un-encoded term, then call the `choose` function to determine which rule to apply (and, which input encoding conversions to apply first).

We implement a function `choose(t, E) → id` that takes the current encoding store and the term to encode, and returns an identifier for which operator rule to apply. `choose` only needs to handle terms that can match multiple operator rules. In our field-blaster, the only ambiguity is between different rules for encoding bit-vector extensions. There is one rule when the input is bit-wise encoded and another when the input’s unsigned value is encoded. For brevity, Fig. 4 shows pseudocode for both rules in a single function: `bvZeroExt`.

4.4.2 Equality Assertions

As an optimization, our implemented field-blaster deviates slightly from our calculus as described in Sect. 4.3. Essentially, it “pushes down” equality assertions. That is, for an input predicate $\phi = \bigwedge_i (t_i = t'_i) \wedge \bigwedge_i b_i$, the field-blaster *does not* encode ϕ . Instead, the field-blaster encodes all t_i , t'_i , and b_i and then asserts that the encodings for each t_i and t'_i are equal, and that each $b_i = T$. Since asserting equalities is often cheaper (i.e. uses fewer field multiplications) than encoding their result, this is an optimization.

It is straightforward—but tedious—to modify the calculus and its proof of correctness to show that this optimization yields a correct field-blaster. Thus, we do not treat this optimization formally in this paper.

5 Verification conditions

In this section, we first define correctness for a ZKP compiler (Sect. 5.1). Then, we give verification conditions (VCs) for each type of encoding rule (Sect. 5.2). Finally, we show that if these VCs hold, our calculus is a correct ZKP compiler (Sect. 5.3).

5.1 Correctness definition

Definition 1 (*Correctness*) A ZKP compiler $\text{Compile}(\phi) \rightarrow (\phi', \text{Ext}_x, \text{Ext}_w)$ is **correct** if it is demonstrably complete and demonstrably sound.

- *Demonstrable completeness*: For all $x \in \text{dom}(x)$, $w \in \text{dom}(w)$ such that $\hat{\phi}(x, w) = \top$,

$$\hat{\phi}'(\text{Ext}_x(x), \text{Ext}_w(x, w)) = \top$$

- *Demonstrable soundness*: There exists an efficient algorithm $\text{Inv}(x', w') \rightarrow w$ such that for all $x \in \text{dom}(x)$, $w' \in \text{dom}(w')$ such that $\hat{\phi}'(\text{Ext}_x(x), w') = \top$,

$$\hat{\phi}(x, \text{Inv}(\text{Ext}_x(x), w')) = \top$$

Demonstrable completeness (respectively, soundness) requires the existence of a witness for ϕ' (resp., ϕ) when a witness exists for ϕ (resp., ϕ'); this existence is *demonstrated* by an efficient algorithm Ext_w (resp., Inv) that computes the witness. While one might imagine a requirement that Inv uniquely inverts Ext_w (something like, $\forall x, \forall w, \text{Inv}(\text{Ext}_w(x), \text{Ext}_w(x, w)) = w$), this is not necessary for ZKP security (Thm. 2), so we omit it, to be more general.

Correct ZKP compilers are important for two reasons. First, since sequential composition preserves correctness, one can prove a multi-pass compiler is correct pass-by-pass. Figure 6 shows how two compilers $\text{Compile}'$ and $\text{Compile}''$ can be composed into $\text{Compose}(\text{Compile}', \text{Compile}'')$. The composition is provably correct:

Theorem 1 (Compiler Composition) If $\text{Compile}'$ and $\text{Compile}''$ are correct, then the compiler $\text{Compose}(\text{Compile}', \text{Compile}'')$ (Fig. 6) is correct.

Proof We argue demonstrable completeness and demonstrable soundness.

Demonstrable Completeness (Def. 1, *Demonstrable Completeness*). Fix x, w such that $\hat{\phi}(x, w) = \top$. Let $x' = \text{Ext}'_x(x)$, $w' = \text{Ext}'_w(x, w)$, $x'' = \text{Ext}''_x(x')$, and $w'' = \text{Ext}''_w(x', w')$. We must show that $\hat{\phi}''(\text{Ext}_x(x), \text{Ext}_w(x, w)) = \top$; substituting, we must show that

$$\hat{\phi}''(x'', w'') = \top \tag{2}$$

From the demonstrable correctness of $\text{Compile}'$, $\hat{\phi}(x, w) = \top$ implies that:

$$\hat{\phi}'(x', w') = \top$$

From the demonstrable correctness of $\text{Compile}''$, this implies our goal (2).

Demonstrable Soundness (Def. 1, *Demonstrable Soundness*). Fix x, w'' such that $\hat{\phi}''(\text{Ext}_x(x), w'') = \top$. Let Inv' be the algorithm guaranteed to exist by the demonstrable soundness of $\text{Compile}'$, and let Inv'' be the algorithm guaranteed to exist by the demonstrable soundness of $\text{Compile}''$. Further, let $x' = \text{Ext}'_x(x)$, $w' = \text{Inv}''(x', w'')$ and $w = \text{Inv}'(x, w')$.

```
def Compose(Compile', Compile'')(phi in Phi) -> (phi'' in Phi'', Ext_x, Ext_w):
  phi', Ext_x, Ext'_w -<- Compile'(phi)
  phi'', Ext'_x, Ext''_w -<- Compile''(phi')
  def Ext_x(x): return Ext''_x(Ext'_x(x))
  def Ext_w(x, w): return Ext''_w(Ext'_x(x), Ext'_w(x, w))
  return (phi'', Ext_x, Ext_w)
```

Fig. 6 The composition of compilers $\text{Compile}'$ and $\text{Compile}''$.

Table 3 VCs related to encoding uniqueness

Property	Condition
Valid encoding uniqueness	$(valid(e, t) \wedge valid(e', t)) \rightarrow equal(e, e')$
Valid encoding uniqueness	$(valid(e, t) \wedge valid(e, t')) \rightarrow t \approx t'$
fromTermcorrectness	$valid(fromTerm(t, kind), t)$
toTermcorrectness	$valid(e, toTerm(e))$

Define $Inv(x, w'') = w$; we must show that $\hat{\phi}(x, Inv(x, w'')) = T$; substituting, we must show that

$$\hat{\phi}(x, w) = T \tag{3}$$

By the demonstrable soundness of $Compile''$, $\hat{\phi}''(Ext_x(x), w'') = T$ implies that

$$T = \hat{\phi}'(Ext_x(w), w') = \hat{\phi}'(x', w')$$

By the demonstrable soundness of $Compile'$, this implies our goal (3). □

Second, a correct ZKP compiler from Φ to Φ' is useful because it can be used to generalize a ZKP for Φ' to one for Φ . Stating this property precisely requires tools from cryptography. We state it informally here, as Theorem 2.

Theorem 2 (ZKP Generalization) (informal) Given a correct ZKP compiler $Compile$ from Φ to Φ' and a ZKP for Φ' , we can construct a ZKP for Φ .

Proof We state and prove this theorem properly in the full version [51, App. A], along with the necessary cryptographic definitions.

Here, we sketch the idea.

Completeness follows from the demonstrable completeness of $Compile$ and the completeness of the ZKP for Φ' .

For knowledge soundness, the knowledge soundness of the ZKP for Φ' is used to extract a valid w' , then $Invis$ used to compute a valid w .

For zero-knowledge, it suffices that x' depends only on x (via Ext_x) and not on w . □

5.2 Rule VCs

Recall (Sect. 4) that our language manipulates encodings through five types of encoding rules. We give verification conditions for each type of rule. Intuitively, these capture the correctness of each rule in isolation. Next, we'll show that they imply the correctness of a ZKP compiler that follows our calculus.

Our VCs quantify over valid encodings. That is, they have the form: “for any valid encoding e of term t , ...” We can quantify over an encoding e by making each $t_i \in terms(e)$ a fresh variable, and quantifying over the t_i . Encoding validity is captured by a predicate $valid(e, t)$, which is defined to be the validity condition in Table 2. Each

Table 4 VCs for encoding rules

Rule	Property	Condition
Operator	Sound	$(A \wedge \bigwedge_i \text{valid}(e_i, t_i)) \rightarrow \text{valid}(e', o(\mathbf{t}))$
$e' \leftarrow r_o(\mathbf{e})$	Complete	$((\bigwedge_i \text{valid}(e_i, t_i) \rightarrow (A \wedge \text{valid}(e', o(\mathbf{t})))))[F]$
Equality	Sound	$(A \wedge \bigwedge_i \text{valid}(e_i, t_i)) \rightarrow (t_1 \approx t_2)$
$r_=(e_1, e_2)$	Complete	$((t_1 \approx t_2) \wedge \bigwedge_i \text{valid}(e_i, t_i) \rightarrow A)[F]$
Conversion	Sound	$(A \wedge \text{valid}(e, t)) \rightarrow \text{valid}(e', t)$
$e' \leftarrow r_-(e)$	Complete	$(\text{valid}(e, t) \rightarrow (A \wedge \text{valid}(e', t)))[F]$
Variable	Sound ($t \in w$)	$A \rightarrow \exists t'. \text{valid}(e', t')$
	Sound ($t \in x$)	$(A \rightarrow \text{valid}(e', t))[F_x]$
$e' \leftarrow r_v(t)$	Complete	$(A \wedge \text{valid}(e', t))[F]$
Constant	–	$\text{valid}(e, t)$
$e \leftarrow r_c(t)$		

VC containing encoding variables \mathbf{e} implicitly represents a conjunction of instances of that VC, one for each possible tuple of kinds of \mathbf{e} , which is fixed for each instance. If a VC contains $\text{valid}(e, t)$, the sort of t is constrained to be *compatible* with $\text{kind}(e)$. For a kind and a sort to be compatible, they must occur in the same row of Table 2. We define the equality predicate $\text{equal}(e, e')$ as $\bigwedge_i \text{terms}(e)[i] \approx \text{terms}(e')[i]$.

Encoding Uniqueness First, we require the uniqueness of valid encodings, for any fixed encoding kind. Table 3 shows the VCs that ensure this. Each row is a formula that must be valid, for all compatible encodings and terms. The first two rows ensure that there is a bijection from terms to their valid encodings (in the first row, we consider only instances for which $\text{kind}(e) = \text{kind}(e')$). The function $\text{fromTerm}(t, \text{kind}) \rightarrow e$ maps a term and an encoding kind to a valid encoding of that kind, and the function $\text{toTerm}(e) \rightarrow t$ maps a valid encoding to its encoded term. The third and fourth rows ensure that fromTerm and toTerm are correctly defined. We will use toTerm in our proof of calculus soundness (Thm. 3) and we will use fromTerm to optimize VCs for faster verification (Sect. 6.1).

For an example of the *valid*, *fromTerm*, and *toTerm* functions, consider a Boolean b encoded as an encoding e with kind bit and whose terms consist of a single field element f . Validity is defined as $\text{valid}(e, b) = f \approx \text{ite}(b, 1, 0)$, $\text{toTerm}(f)$ is defined as $f \approx 1$, and $\text{fromTerm}(b, \text{bit})$ is $(b, \text{bit}, \text{ite}(b, 1, 0))$.

VCs for encoding rules. Table 4 shows our VCs for the rules of Fig. 5. For each rule application, A and F denote, respectively, the assertions and the variable declarations generated when that rule is applied.⁹ Now, we explain some of the VCs in detail.

First, consider a rule r_o for operator o applied to inputs t_1, \dots, t_k . The rule takes input encodings e_1, \dots, e_k and returns an output e' . It is sound if the validity of its inputs and its assertions imply the validity of its output. It is complete if the validity of its inputs implies its assertions and the validity of its output, after substituting fresh variable definitions.

⁹ As a warning, note that the completeness VC **do not** imply the soundness ones, because of the substitutions F . The completeness VCs have form $(P \rightarrow (A \wedge Q))[F]$ and the soundness VCs have form $((P \wedge A) \rightarrow Q)$. As an example of how the first does not imply the second, let P and A be \top , x be an integer variable, Q be $x = 0$, and F assign $x = 0$. Then, $(P \rightarrow (A \wedge Q))[F]$ reduces to \top , which does not imply $x = 0$ (which is what soundness reduces to).

Second, consider a variable rule. Its input is a variable term t , and it returns e' , a putative encoding thereof. Note that e' does not actually contain t , though the substitutions in F may bind the fresh variables of e' to functions of t . For the rule to be sound when t is a witness variable ($t \in w$), the assertions must imply that e' is valid for *some* term t' . For the rule to be sound when t is an instance variable ($t \in x$), the assertions must imply that e' is valid for t , when the instance variables in e' are replaced with their definition (F_x denotes F , restricted to its declarations of instance variables).¹⁰ For the variable rule to be complete (for an instance or a witness), the assertions and the validity of e' for t must follow from F .

Third, consider a constant rule. Its input is a constant term t , and it returns an encoding e . Recall that the terms of e are always evaluated, yielding e' which only contains constant terms. Thus, correctness depends only on the fact that e is always a valid encoding of the input t . This can be captured with a single VC.

5.3 A correct field-blasting calculus

Given rules that satisfy these verification conditions, we show that the calculus of Sect. 4.3 is a correct ZKP compiler.

Theorem 3 (Correctness) With rules that satisfy the conditions of Sect. 5.2, the calculus of Sect. 4.3 is demonstrably complete and sound (Def. 1).

Proof We begin with a note on notation. Then, we prove demonstrable correctness. Then, we prove demonstrable soundness.

Notation. If $\phi(x)$ is a formula in variable x , and x is a value, by $\phi[x \mapsto x]$ we denote ϕ with its variable x replaced with a constant term of value x . Note that for our theories of interest (bit-vectors and prime fields), all values have a corresponding constant. When the variable x and the value being substituted x are denoted with the same letter, we will abbreviate $\phi[x \mapsto x]$ as $\phi[x]$. Finally, for a term ϕ which contains no variables, we denote the fact that $\hat{\phi}$ evaluates to z by $\phi \Downarrow z$. Thus, \Downarrow denotes the evaluation relation for variable-free terms and values. Evaluation for variable-free terms is unique because we allow interpretations to differ only in variables (Sect. 2.1).

Demonstrable soundness. *Sketch.* At a high level, we use *toTerm* to define *Inv* that transforms a satisfying witness z' for ϕ' into witness z for ϕ . Then, we induct on a trace of the field-blaster’s execution to show that, under z , every encoding that the compiler produces is valid. Then, a simple analysis of the final rule (equality assertion) shows that $\hat{\phi}(z) = T$. Now, we give the proof in full.

Let the compiler take $\phi(x, w)$ as input and produce as output: $\phi'(x', w')$, Ext_x , and Ext_w . Fix an instance x for ϕ , and let $x' \leftarrow \text{Ext}_x(x)$. Fix a w' such that $\hat{\phi}'(x', w') = T$, and let $z' = (x', w')$. Let \mathcal{T} (the “trace”) be the sequence of transitions that the compiler takes. For each witness variable $w_i \in w$, there is a transition in \mathcal{T} that encodes w_i as some e_i . We define a function $\text{Inv}(z') \rightarrow w$. For variable $w_i \in w$, it outputs the evaluation of $\text{toTerm}(e_i)[z']$. Let w denote the output of *Inv* and let $z = (x, w)$. It suffices to show that $\hat{\phi}(z) = T$.

¹⁰ The different soundness conditions for instance and witness variables play a key role in the proof of Theorem 3. Essentially: since the condition for instances replaces variables with their definitions, the validity of the encodings of instance variables need not be explicitly enforced in A. This is why some constraints could be omitted in our field-blasting example.⁸

By inducting over \mathcal{T} , we will show that for each $(t \mapsto e) \in E$ (where E is the final encoding store in \mathcal{T}), $\text{valid}(e[z'], t[z]) \Downarrow \top$. Initially, E is empty, so the property holds. We have one inductive case for each transition rule (except equality).

Operator rules are our first inductive case. Consider a transition for operator o applied to arguments \mathbf{t} , with operator rule $r: A', F', e' \leftarrow r(o, \mathbf{e})$. By the inductive hypothesis, we have that for all i , $\text{valid}(e_i[z'], t_i[z]) \Downarrow \top$. Furthermore, we have that $A'[z'] \Downarrow \top$. Operator soundness requires that $(A \wedge \bigwedge_i \text{valid}(e_i, t_i)) \rightarrow \text{valid}(e', o(\mathbf{t}))$ is valid (i.e., holds for all values of variables). Then, instantiation and evaluation give the desired result:

$$\begin{aligned} (A \wedge \bigwedge_i \text{valid}(e_i, t_i)) &\rightarrow \text{valid}(e', o(\mathbf{t})) \text{ is valid } (\text{operator soundness}) \\ (A \wedge \bigwedge_i \text{valid}(e_i, t_i)) &\rightarrow \text{valid}(e', o(\mathbf{t}))[z, z'] \Downarrow \top (\text{inst\&eval}) \\ (A[z'] \wedge \bigwedge_i \text{valid}(e_i[z'], t_i[z])) &\rightarrow \text{valid}(e'[z'], o(\mathbf{t}))[z] \Downarrow \top (\text{move subs to vars}) \\ (\top \wedge \bigwedge_i \top) &\rightarrow \text{valid}(e'[z'], o(\mathbf{t}))[z] \Downarrow \top (\text{eval}) \\ &\text{valid}(e'[z'], o(\mathbf{t}))[z] \Downarrow \top (\text{eval}) \end{aligned}$$

Variables are our second inductive case. For a witness variable w_i , encoded as e_i , the value of w_i in \mathbf{w} is that of $\text{toTerm}(e_i)[z']$. So, we must show that $\text{valid}(e_i[z'], \text{toTerm}(e_i)[z']) \Downarrow \top$. Variable soundness (for witnesses) implies that there exists a term constant t of value \mathbf{t} such that $\text{valid}(e_i[z'], t) \Downarrow \top$. Instantiating the correctness of toTerm , we have that $\text{valid}(e_i[z'], \text{toTerm}(e_i[z'])) \Downarrow \top$. Instantiating the uniqueness of valid encodings, the validity of two different terms for the same encoding implies that the terms are equal, i.e., $t = \text{toTerm}(e_i[z']) \Downarrow \top$, so $\text{toTerm}(e_i[z']) \Downarrow \mathbf{t}$. Using this equivalence, we eliminate t from our previous valid term, giving that $\text{valid}(e_i[z'], \text{toTerm}(e_i[z'])) \Downarrow \top$, which is the desired conclusion.

For an instance variable x_i , encoded as e_i with definitions $F', z' \mapsto z'$ agrees with F' on the instance variables of F' . Thus, instantiating variable soundness (for instances) gives $\text{valid}(e_i[z'], x_i[z]) \Downarrow \top$, which is what we need.

The recursive cases for constants and conversions hold *mutatis mutandis*. These cases complete the induction.

Now, consider the final transition in \mathcal{T} . By construction, it is an equality transition: $e_{\top} \leftarrow \text{const}(\top)$; $A', F' \leftarrow \text{assertEq}(e_{\top}, e_{\phi})$. By the constant completeness condition, $\text{valid}(e_{\top}, \top) \Downarrow \top$. By induction, we have $\text{valid}(e_{\phi}[z'], \phi[z]) \Downarrow \top$. Through instantiating equality soundness, substituting, and evaluating, we find:

$$\begin{aligned} (A \wedge \text{valid}(e_{\phi}, \phi) \wedge \text{valid}(e_{\top}, \top)) &\rightarrow \phi = \top \text{ is valid } (\text{eq. sound}) \\ ((A \wedge \text{valid}(e_{\phi}, \phi) \wedge \text{valid}(e_{\top}, \top)) &\rightarrow \phi = \top)[z, z'] \Downarrow \top (\text{inst\&eval}) \\ ((A[z'] \wedge \text{valid}(e_{\phi}[z'], \phi[z]) \wedge \text{valid}(e_{\top}[z'], \top)) &\rightarrow \phi[z] = \top) \Downarrow \top (\text{move subs}) \\ ((\top \wedge \top \wedge \top) &\rightarrow \phi[z] = \top) \Downarrow \top (\text{eval}) \\ &\phi[z] \Downarrow \top \quad (\text{eval}) \end{aligned}$$

This is exactly what we sought to show. Thus, our compiler is demonstrably sound.

Demonstrable completeness. *Sketch.* At a high level, we consider the substitutions F and a satisfying witness \mathbf{z} for ϕ . We induct on a trace of the field-blaster's execution to

show that, under F and z , every encoding that the compiler produces is valid, and that all assertions are satisfied. Now, we give the proof in full.

Let the compiler take $\phi(x, w)$ as input and produce as output: $\phi'(x', w')$, Ext_x , and Ext_w . Fix x and w such that $\phi(x, w) = \top$. Let $z \leftarrow (x, w)$. We must show that for $z' \leftarrow (\text{Ext}_x(x), \text{Ext}_w(z))$, $\phi'[z'] \Downarrow \top$. Again, let \mathcal{T} be the compiler's transition sequence. Let A be the final assertion set and let F be the final sequence of fresh variable definitions. It suffices to show that $A[F][z] = \top$, since the substitutions z' are equivalent to the substitution sequence defined by F and z , because of z' 's definition in terms of the Ext_x and Ext_w functions.

We proceed by induction on \mathcal{T} . Our inductive hypothesis is that after each transition, $A[F][z] \Downarrow \top$ and for each $(t \mapsto e) \in E$, $\text{valid}(e[F], t)[z] \Downarrow \top$. Initially, A and E are empty, so this holds. There is one inductive case for each transition type (except equality). The transition begins with the calculus in state (E, A, F) . It might create new constraints A' , new definitions F' , and a new encoding e for some term t . It suffices to show that $A'[F][F'][z] \Downarrow \top$ and that $\text{valid}(e[F][F'], t)[z] \Downarrow \top$.

First, consider a transition for operator o applied to arguments \mathbf{t} , with operator rule r : $A', F', e' \leftarrow r(o, \mathbf{e})$. The inductive hypothesis gives that for all i , $\text{valid}(e_i[F], t_i)[z] \Downarrow \top$. We instantiate operator completeness and proceed:

$$\begin{aligned} & ((\bigwedge_i \text{valid}(e_i, t_i)) \rightarrow (A' \wedge \text{valid}(e', t)))[F'] \text{ is valid}(\text{op completeness}) \\ & ((\bigwedge_i \text{valid}(e_i, t_i)) \rightarrow (A' \wedge \text{valid}(e', t)))[F, F', z] \Downarrow \top(\text{inst\&eval}) \\ & ((\bigwedge_i \text{valid}(e_i[F], t_i)[z]) \rightarrow (A'[F][F'] \wedge \text{valid}(e'[F][F'], t)))[z] \Downarrow \top(\text{move subs}) \\ & ((\bigwedge_i \top) \rightarrow (A'[F][F'] \wedge \text{valid}(e'[F][F'], t)))[z] \Downarrow \top(\text{eval}) \\ & A'[F][F'][z] \wedge \text{valid}(e'[F][F'], t)[z] \Downarrow \top(\text{eval}) \end{aligned}$$

which is what we sought to show.

Second, consider a transition for variable z_i . Variable completeness states that the conjunction $(A' \wedge \text{valid}(e', z_i))[F']$ holds for all values of z_i ; thus it evaluates to \top under z , as desired.

The inductive cases for conversion rules and constant rules also hold, *mutatis mutandis*. These cases complete the induction.

Finally, consider the calculus's final equality transition: $e_\top \leftarrow \text{const}(\top)$; $A', F' \leftarrow \text{assertEq}(e_\top, e_\phi)$. By constant completeness, $e_\top \Downarrow \top$. Moreover, $\phi[z] \Downarrow \top$. By the inductive hypothesis, $\text{valid}(e_\phi[F], \phi)[z] \Downarrow \top$. We instantiate equality completeness and proceed:

$$\begin{aligned} & ((\phi = \top \wedge \text{valid}(e_\phi, \phi) \wedge \text{valid}(e_\top, \top)) \rightarrow A')[F'] \text{ is valid}(\text{eq. complete}) \\ & ((\phi = \top \wedge \text{valid}(e_\phi, \phi) \wedge \text{valid}(e_\top, \top)) \rightarrow A')[F, F', z] \Downarrow \top(\text{sub\&eval}) \\ & ((\phi[z] = \top \wedge \text{valid}(e_\phi[F], \phi)[z] \wedge \text{valid}(e_\top, \top)) \rightarrow A'[F][F'][z]) \Downarrow \top(\text{move subs}) \\ & ((\top = \top \wedge \top \wedge \top) \rightarrow A'[F][F'][z]) \Downarrow \top(\text{eval}) \\ & A'[F][F'][z] \Downarrow \top(\text{eval}) \end{aligned}$$

That is exactly what we sought to show. Thus, our compiler is demonstrably complete. □

6 Case study: a verifiable field-blaster for CirC

We implemented and partially verified a field-blaster for CirC [47]. Our implementation is based on a refactoring of CirC’s original field-blaster to conform to our encoding rules (Sect. 4.2) and consists of ≈ 850 lines of code (LOC).¹¹ As described below, we have (partially) verified our encoding rules, but trust our calculus (Sect. 4.3, ≈ 150 LOC) and our flattening implementations (Fig. 2, ≈ 160 LOC).

While porting rules, we found 4 bugs in CirC’s original field-blaster (see Sect. 6.3), including a severe soundness bug. Given a ZKP compiled with CirC, the bug allowed a prover to incorrectly compare bit-vectors. The prover, for example, could claim that the unsigned value of 0010 is greater than *or less than* that of 0001 . A patch to fix all 4 bugs (in the original field-blaster) has been upstreamed, and we are in the process of upstreaming our new field-blaster implementation into CirC.

6.1 Verification evaluation

Our implementation constructs the VCs from Sect. 5.2 and emits them as SMT-LIB (extended with a theory of finite fields [48]). We verify them with *cvc5*, because it can solve formulas over bit-vectors and prime fields [48]. The verification is partial in that it is bounded in two ways. We set $b \in \mathbb{N}$ to be the maximum bit-width of any bit-vector and $a \in \mathbb{N}$ to be the maximum number of arguments to any n -ary operator. In our evaluation, we used $a = 4$ and $b = 4$. These bounds are small, but they were sufficient to find the bugs mentioned above.

Optimizing completeness VCs Generally, *cvc5* verifies soundness VCs more quickly than completeness VCs. This is surprising at first glance. To see why, consider the soundness (S) and completeness (C) conditions for a conversion rule from e to e' that generates assertions A and definitions F :

$$S \triangleq (A \wedge \text{valid}(e, t)) \rightarrow \text{valid}(e', t) \quad C \triangleq (\text{valid}(e, t) \rightarrow (A \wedge \text{valid}(e', t)))[F]$$

In both, t is a variable, e contains variables, and there are variables in e' and A that are defined by F . In C , though, some variables are replaced by their definitions in F —which makes the number of variables (and thus the search space)—seem smaller for C than S . Yet, *cvc5* is slower on C .

The problem is that, while the field operations in A are standard (e.g., $+$, \times , and $=$), the definitions in F use a CirC-IR operator that (once embedded into SMT-LIB) is hard for *cvc5* to reason about. That operator, $(\text{ff2bv } b)$, takes a prime field element x and returns a bit-vector v . If x ’s integer representative is less than 2^b , then v ’s unsigned value is equal to x ; otherwise, v is zero.

The ff2bv operator is trivial to evaluate but hard to embed. *cvc5*’s SMT-LIB extension for prime fields only supports $+$, \times and $=$, so no operator can directly relate x to v . Instead, we encode the relationship through b Booleans that represent the bits of v . To test whether $x < 2^b$, we use the polynomial $f(x) = \prod_{i=0}^{2^b-1} (x - i)$, which is zero only on $[0, 2^b - 1]$. The bit-splitting essentially forces *cvc5* to guess v ’s value; further, f ’s high degree slows down the Gröbner basis computations that form the foundation of *cvc5*’s field solver.

¹¹ Our implementation is in Rust, as is CirC.

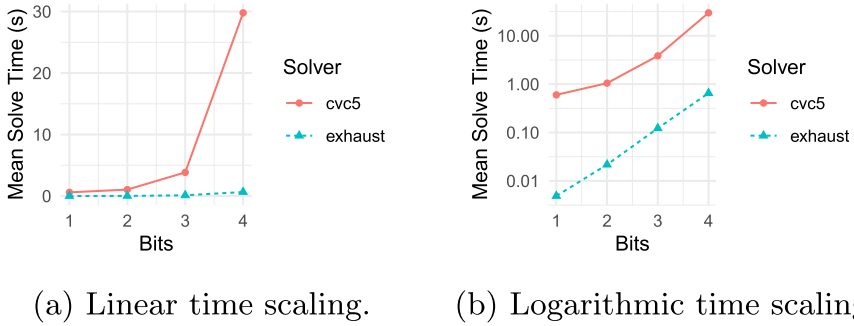


Fig. 7 How average VC verification time depends on bit-width b . Includes only VCs that are verified by the pertinent solver at all bit-widths.

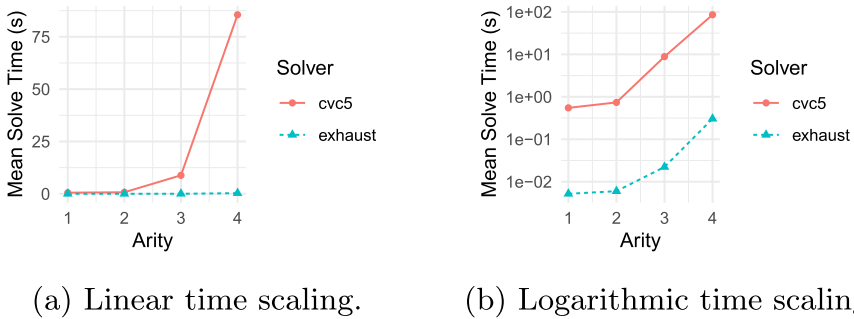


Fig. 8 How average VC verification time depends on operator arity a . Includes only VCs that are verified by the pertinent solver at arity 4

Table 5 VCs verified by different solvers. ‘uniqu’ denotes the VCs of Table 3; others are from Table 4

Type	Prop	VCs	Verified			Unver
			cvc5	Exhaust	Either	
const	–	6	6	5	6	0
conv	C	8	8	8	8	0
conv	S	8	8	4	8	0
eq	C	10	10	9	10	0
eq	S	10	10	9	10	0
op	C	259	247	247	259	0
op	S	263	259	126	259	4
unip	–	40	40	0	40	0
var	C	12	12	10	12	0
var	S	6	6	0	6	0

‘C’ denotes completeness; ‘S’: soundness

Table 6 The performance of CirC with the verified and unverified field-blaster. Metrics are summed over the 61 functions in the Z# standard library

Metric	Unverified	Verified
Time(s)	27.27	25.05
Mem. (GB)	6.56	6.42
Constraints	559,445	559,445

To optimize verification of the completeness VCs, we reason about CirC-IR directly. First, we use the uniqueness of valid encodings and the *fromTerm* function. Since the VC assumes *valid(e, t)*, we know *e* is equal to *fromTerm(t, kind(e))*. We use this equality to eliminate *e* from the completeness VC, leaving:

$$(A \wedge \text{valid}(e', t))[F][e \mapsto \text{fromTerm}(t, \text{kind}(e))]$$

Since *F* defines all variables in *A* and *e'*, the only variable after substitution is *t*. So, when *t* is a Boolean or small bit-vector, an exhaustive search is very effective¹²; we implemented such a solver in 56 LOC, using CirC's IR as a library.

For soundness VCs, this approach is less effective. The *fromTerms* substitution still applies, but if *F* introduces fresh field variables, they are not eliminated and thus, the final formula contains field variables, so exhaustion is infeasible.

Verification results We ran our VC verification on machines with Intel Xeon E5-2637 v4 CPUs.¹³ Each attempt is limited to one physical core, 8GB memory, and 30 min. Table 5 shows the number of VCs verified by *cvc5* and our exhaustive solver. As expected, the exhaustive solver is effective on completeness VCs for Boolean and bit-vector rules, but ineffective on soundness VCs for rules that introduce fresh field variables. There are four VCs that neither solver verifies within 30 min: *bvadd* with (*b* = 4, *a* = 4), and *bvmul* with (*b* = 3, *a* = 4) and (*b* = 4, *a* ≥ 3). Most other VCs verify instantly. Now, we analyze how VC verification time depends on *a* and *b*.

In Fig. 7a, we show how VC verification time scales with the bit-width *b* of the rule that the VC models. Each data point is the average solve time at a particular bit-width over all VC operator rules that the solver verifies at bit-width 4. We see that *cvc5*'s run time grows substantially with the bit-size. The exhaustive solver is far faster on average (though it verifies only completeness VCs). Figure 7b shows the same data with a logarithmic time scale. Both solvers appear to take time exponential in the bit-width.

In Fig. 8a, we show how VC verification time grows with the arity *a* of the operator that the VC is written for. As before, we include only VCs that the solver verifies at arity 4. Note that this excludes binary operators and ternary operators like if-then-else. The trend is similar to that observed with bit-width: solver runtime grows quickly with the arity. Figure 8b shows the same data with a logarithmic scale.

6.2 Performance and output quality evaluation

We compare CirC with our field-baster (“Verified”) against CirC with its original field-blaster (“Unverified”)¹⁴ on three metrics: compiler runtime, memory usage, and the final RICS constraint count. Our benchmark set is the standard library for CirC's Z# input

¹² So long as the exhaustive solver reasons directly about all CirC-IR operators.

¹³ We omit the completeness VCs for *ff2bv*

¹⁴ After fixing the bugs we found. See Sect. 6.3.

File	Seconds		Memory (kB)		Constraint		Rules
	Ver.	Unver.	Ver.	Unver.	Ver.	Unver.	
ecc/edwardsAdd.zok	0.00	0.00	7048	6960	17	17	30
ecc/edwardsCompress.zok	0.09	0.09	21604	21636	511	511	772
ecc/edwardsNegate.zok	0.00	0.00	6792	6840	2	2	7
ecc/edwardsOnCurve.zok	0.00	0.00	6848	6884	6	6	16
ecc/edwardsOrderCheck.zok	0.01	0.00	7304	8988	57	57	72
ecc/edwardsScalarMult.zok	0.55	0.62	127064	100748	9206	9206	10496
ecc/proofOfOwnership.zok	0.57	0.49	126428	100832	9392	9392	10419
EMBED.zok	0.00	0.00	6872	7056	1	1	3
field.zok	0.00	0.00	6760	6940	1	1	3
hashes/mimc7/mimc7R20.zok	0.01	0.01	7840	8052	80	80	124
hashes/mimcSponge/mimcFeistel.zok	0.15	0.17	20080	20004	662	662	1325
hashes/pedersen/512bitBool.zok	0.54	1.14	48664	48672	3573	3573	8617
hashes/pedersen/512bit.zok	0.59	0.55	52540	52744	4083	4083	8424
hashes/sha256/1024bitPadded.zok	3.96	4.59	750852	841956	84727	84727	49749
hashes/sha256/1024bit.zok	2.55	2.99	546200	550216	57382	57382	33242
hashes/sha256/1536bit.zok	3.93	4.78	772816	863696	86253	86253	49953
hashes/sha256/256bitPadded.zok	1.20	1.36	276076	273748	27645	27645	16442
hashes/sha256/512bitPacked.zok	2.65	3.03	551276	541852	55897	55897	34692
hashes/sha256/512bitPadded.zok	2.45	2.85	550252	542844	55856	55856	33036
hashes/sha256/512bit.zok	1.25	1.50	281908	277880	28505	28505	16523
hashes/sha256/embed/IVconstants.zok	0.01	0.02	9068	8804	256	256	514
hashes/sha256/shaRound.zok	1.21	1.55	283432	281228	29057	29057	16879
hashes/utis/256bitsDirectionHelper.zok	0.00	0.00	7456	7188	16	16	65
utils/casts/bool_128_to_u32_4.zok	0.03	0.04	21616	21612	4	4	272
utils/casts/bool_256_to_u32_8.zok	0.05	0.07	21620	21632	8	8	544
utils/casts/field_to_u16.zok	0.03	0.02	21604	21604	17	17	38
utils/casts/field_to_u32.zok	0.02	0.03	21604	21548	33	33	70
utils/casts/field_to_u64.zok	0.03	0.03	21604	21628	66	66	134
utils/casts/field_to_u8.zok	0.02	0.02	21600	21600	9	9	22
utils/casts/u16_from_bits.zok	0.00	0.00	7144	6824	1	1	4
utils/casts/u16_to_bits.zok	0.00	0.00	7144	7056	17	17	50
utils/casts/u16_to_field.zok	0.00	0.00	7112	7108	1	1	3
utils/casts/u16_to_u32.zok	0.00	0.00	8720	6632	1	1	4
utils/casts/u16_to_u64.zok	0.00	0.00	8356	6796	1	1	4
utils/casts/u32_4_to_bool_128.zok	0.06	0.05	21572	21572	132	132	392
utils/casts/u32_8_to_bool_256.zok	0.06	0.08	21620	21620	264	264	784
utils/casts/u32_from_bits.zok	0.01	0.00	6988	6960	1	1	4
utils/casts/u32_to_bits.zok	0.00	0.00	7580	6956	33	33	98
utils/casts/u32_to_field.zok	0.00	0.00	6952	6776	1	1	3
utils/casts/u32_to_u64.zok	0.00	0.00	7184	6976	1	1	4
utils/casts/u64_from_bits.zok	0.00	0.00	7152	7208	1	1	4
utils/casts/u64_to_bits.zok	0.01	0.00	8160	7704	65	65	194
utils/casts/u64_to_field.zok	0.00	0.00	6996	8680	1	1	3
utils/casts/u8_from_bits.zok	0.00	0.00	6584	6876	1	1	4
utils/casts/u8_to_bits.zok	0.00	0.00	7080	7064	8	8	17
utils/casts/u8_to_field.zok	0.00	0.00	8796	6932	1	1	3
utils/casts/u8_to_u16.zok	0.00	0.00	6984	8708	1	1	4
utils/casts/u8_to_u32.zok	0.00	0.00	7064	6700	1	1	4
utils/casts/u8_to_u64.zok	0.00	0.00	6796	6704	1	1	4
utils/multiplexer/lookup1bit.zok	0.00	0.00	6708	8416	1	1	6
utils/multiplexer/lookup2bit.zok	0.00	0.00	7044	7080	3	3	20
utils/multiplexer/lookup3bitSigned.zok	0.00	0.00	9100	6788	4	4	26
utils/pack/bool/nonStrictUnpack256.zok	0.06	0.06	21612	21552	257	257	769
utils/pack/bool/pack128.zok	0.06	0.03	21544	21632	1	1	388
utils/pack/bool/pack256.zok	0.06	0.07	21612	21556	1	1	772
utils/pack/bool/unpack128.zok	0.06	0.06	21604	21660	129	129	386
utils/pack/u32/nonStrictUnpack256.zok	0.08	0.11	21616	21604	255	255	543
utils/pack/u32/pack128.zok	0.12	0.10	21620	21628	129	129	396
utils/pack/u32/pack256.zok	0.06	0.07	21640	21604	257	257	788
utils/pack/u32/unpack128.zok	0.10	0.10	21632	21628	129	129	274

Fig. 9 Compile time, memory usage, and constraint count for all standard library functions.

language (which extends ZoKrates [17, 70] v0.6.2). Our testbed runs Linux with 32GB memory and an AMD Ryzen 2700.

Table 6 shows the total compile time and compiler memory usage of CirC with the verified and unverified field-blaster. These metrics are summed over the entire set of benchmark programs. There is no difference in constraints, but the verified field-blaster slightly improves compiler performance: -8% time and -2% memory.

We think that the small performance improvement is unrelated to the fact that the new field-blaster is verified. The difference is in a downstream compiler pass—the RICS optimizer. We believe the cause of the difference is in how our field-blaster represents *bit-constraints*. A *bit-constraint* for variable $x \in \mathbb{F}$ is an equation that ensures that $x \in \{0, 1\} \subset \mathbb{F}$. Field-blasters emit many bit-constraints, for instance when converting a bit-vector from an unsigned integer encoding to a bitwise encoding. The previous field-blaster emitted bit-constraints of form $x \times x = x$. (Observe that this polynomial's roots are exactly $\{0, 1\}$.) Our field-blaster emits bit-constraints of form $x \times (x - 1) = 0$. This constraint encodes the same polynomial, but with one fewer variable occurrence. It appears that the RICS optimizer maintains an index over variable occurrences, and that this index performs slightly better (in both time and space) when there are fewer occurrences.

In Fig. 9, we show compiler runtime, memory usage, and final constraint count for every function in the Z# standard library, with and without the verified field-blaster. We also show the number of encoding rules applied in the verified field-blaster; this is generally similar to the final number of constraints. The number of constraints is always exactly equal. The slight improvement in compiler runtime and memory usage seems well distributed across most non-trivial benchmarks. The benchmarks where the verified blaster is sometimes slower are those based on elliptic-curve cryptography (e.g., twisted Edwards curve arithmetic [8] and the Pedersen hash function). Note that these circuits include few bit-constraints relative to the overall number of constraints. The number of rule applications (in the verified blaster) varies widely, but is correlated with the number of constraints.

6.3 Bugs found in the CirC field-blaster

In this subsection, we explain the four bugs we found in CirC's field-blaster. We found these bugs by observing VC violations while porting operator rules from CirC's field-blaster to our verifiable field-blaster.

Incompleteness: finite field reciprocal CirC-IR includes a finite field reciprocal operator: $z \leftarrow 1/y$. In the IR evaluator, if y is zero, then z is set to be zero. However, the field-blaster emits the equation $zy = 1$; this is unsatisfiable for $y = 0$. This is an incompleteness bug: it makes the whole output predicate unsatisfiable. The completeness VC for the finite-field division rules catches this bug.

One change would be to emit $zyy = y$. This fixes completeness: now if y is non-zero, then $z = x/y$ is a solution for any x . But, it compromises soundness: now if $y = 0$, any (z, x) pair is a solution.

Our actual fix is to emit $ay = 1 - t \wedge ty = 0 \wedge ta = 0$. In this system t is zero if z is zero and one otherwise. And, a is the inverse of y if y is non-zero and 0 otherwise. Completeness is easy to check. For soundness, observe that if y is non-zero, then the second equation forces $t = 0$, and then a must be the inverse of y (first equation). On the other hand, if y is zero, then t is 1 (first equation), so a is zero (last equation).

Incompleteness: bit-shifts CirC-IR includes bit-vector shifts (left shift, arithmetic right shift, and logical right shift). Following SMT-LIB, overshifting should *saturate*: i.e., shifting a length- b bit-vector by $\geq b$ bits should be equivalent to shifting by $b - 1$

<pre> fn bvUgeBuggy($t_x, t_y, x : \text{Enc}, y : \text{Enc}$) : kind($x$) = uint \wedge kind(y) = uint $b \leftarrow \text{size}(\text{sort}(t_x))$ $\Delta \leftarrow x - y$ $\Delta' \leftarrow \text{terms}(x) - \text{terms}(y)$ for i in $[0, b - 1]$: $\Delta'_i \leftarrow \text{fresh}(i, \text{ite}(\Delta[i], 1, 0), \perp)$ assert($\Delta'_i(\Delta'_i - 1) \approx 0$) return ffEq($\Delta', \sum_{i=0}^{b-1} 2^i \Delta'_i$) </pre>	<pre> fn bvUgeCorrect($t_x, t_y, x : \text{Enc}, y : \text{Enc}$) : kind($x$) = uint \wedge kind(y) = uint $b \leftarrow \text{size}(\text{sort}(t_x))$ $\Delta \leftarrow \text{bvUext}(x, 1) - \text{bvUext}(y, 1)$ for i in $[0, b]$: $\Delta'_i \leftarrow \text{fresh}(i, \text{ite}(\Delta[i], 1, 0), \perp)$ assert($\Delta'_i(\Delta'_i - 1) \approx 0$) assert($\Delta' \approx -2^b \Delta'_b + \sum_{i=0}^{b-1} 2^i \Delta'_i$) return (bit, $1 - \Delta'_b$) </pre>
---	---

Fig. 10 Pseudocode for buggy and correct bit-vector unsigned \geq . All sub-routines have been inlined, except ffEq which produces a boolean bit encoding of whether the inputs are equal, as field elements.

bits. However, the field-blaster required the shift amount to be less than b . This is an incompleteness bug, and it is caught by the completeness VCs for the shift rules.

The fix for this bug is relatively straightforward. The idea is to simply add a test for overshifting. That is, we field-blast $x \gg_{\text{complete}} y$ as $\text{ite}(x \geq b, b - 1, b) \gg_{\text{incomplete}} y$, and likewise for the other shift operators.

Non-determinism: unsigned bit-vector division CirC-IR includes unsigned bit-vector division: $z \leftarrow x/y$. Following SMT-LIB, when $y = 0$, the value of z should be $2^b - 1$. However, when $y = 0$, the field-blaster emits equations that allow z to have any value in $\{0, \dots, 2^b - 1\}$. The soundness VC for the unsigned bit-vector division rule catches this bug. Fixing is again fairly straightforward, through a special-case for when $y = 0$.

Unsoundness: bit-vector comparisons CirC-IR includes operators for signed and unsigned bit-vector comparisons: $z \leftarrow x \bowtie y$. For example, signed \geq and unsigned $<$. For all these operators, the field-blaster uses a utility that computes $\Delta = x - y$ (the difference of signed values) and tests whether $\Delta \geq 0$. This test reduces to testing whether Δ “fits in b unsigned bits”. To test this, the field-blaster emits bit-constrained fresh variables $\Delta_0, \dots, \Delta_{b-1}$ and tests $\Delta = \sum_i 2^i \Delta_i$. This approach is unsound: while the Δ_i variables are supposed to be set to the bits of Δ (if it is non-negative), this is not ensured. By setting the bits to the wrong decomposition, the equality doesn’t hold, even if $0 \leq \Delta < 2^b$. In the context of a ZKP, this allows a malicious prover to equivocate about whether $x \geq y$. We first caught this bug with the soundness VC for the operator rule for signed bit-vector \geq . However, the subroutine “fits in bits” is also used in all bit-vector comparisons, as well as division and remainder.

To fix the bug, observe that Δ always fits in $b + 1$ signed bits. After enforcing (not testing!) the signed bit decomposition, the sign bit indicates whether $\Delta \geq 0$. In Fig. 10, we show the buggy rule and the correct rule.

Significance In most applications of ZKPs, safety properties (e.g., the solvency of an exchange) depend on soundness, while liveness properties (e.g., whether a specific transaction completes) depend on completeness. Thus, the soundness bugs are more serious. The bug in unsigned division only affects predicates that assume SMT-LIB semantics for division by zero. The comparison bug is much more serious: it affects any predicate that compares or divides bit-vectors.

Impact For each bug, we submitted a report and a patch to CirC’s existing field-blaster. All of these patches have been upstreamed. We are still considering whether to upstream the verifiable field-blaster itself.

7 Discussion

In this work, we present the first automatically verifiable field-blaster. We view the field-blaster as a set of rules; if some (automatically verifiable) conditions hold for each rule, then the field-blaster is correct. We implemented a performant and partially verified field-blaster for CirC, finding 4 bugs along the way.

Our approach has limitations. First, we require the field-blaster to be written as a set of encoding rules. Second, we only verify our rules for bit-vectors of bounded size and operators of bounded arity. Third, we assume that each rule is a pure function: for example, it doesn't return different results depending on the time. Future work might avoid the last two limitations through bit-width-independent reasoning [43, 44, 69] and a DSL (and compiler) for encoding rules. It would also be interesting to extend our approach to: a ZKP with a non-prime field [7, 14], a compiler IR with partial or non-deterministic semantics, or a compiler with correctness that depends on computational assumptions.

Acknowledgements We appreciate the help and guidance of Andres Nötzli, Dan Boneh, and Evan Laufer. This material is in part based upon work supported by the DARPA SIEVE program and the Simons foundation. Any opinions, findings, and conclusions or recommendations expressed in this report are those of the author(s) and do not necessarily reflect the views of DARPA. It is also funded in part by NSF grant number 21110397 and the Stanford Center for Automated Reasoning.

References

1. LLVM language reference manual. <https://llvm.org/docs/LangRef.html>
2. Monero technical specs. <https://monerodocs.org/technical-specs/> (2022)
3. Aircscript. <https://github.com/OxPolygonMiden/air-script>
4. Angel S, Blumberg AJ, Ioannidis E, Woods J (2022) Efficient representation of numerical optimization problems for SNARKs. In: USENIX Security
5. Bellés-Muñoz M, Isabel M, Muñoz-Tapia JL, Rubio A, Baylina J (2022) Circom: a circuit description language for building zero-knowledge applications. *IEEE Trans Depend Secure Comput* 20(6):4733–4751
6. Bellman. <https://github.com/zkcrypto/bellman>
7. Ben-Sasson E, Bentov I, Horesh Y, Riabzev M (2019) Scalable zero knowledge with no trusted setup. In: CRYPTO
8. Bernstein DJ, Birkner P, Joye M, Lange T, Peters C (2008) Twisted edwards curves. In: AFRICACRYPT
9. Bertot Y, Castéran P (2013) Interactive theorem proving and program development: Coq'Art: the calculus of inductive constructions. Springer Science & Business Media, UK
10. Blum M, Feldman P, Micali S (1988) Non-interactive zero-knowledge and its applications. In: STOC
11. Brown F, Renner J, Nötzli A, Lerner S, Shacham H, Stefan D (2020) Towards a verified range analysis for JavaScript JITs. In: PLDI
12. Campanelli M, Gennaro R, Goldfeder S, Nizzardo L (2017) Zero-knowledge contingent payments revisited: attacks and payments for services. In: CCS
13. Chen E, Zhu J, Ozdemir A, Wahby RS, Brown F, Zheng W (2023) Silph: a framework for scalable and accurate generation of hybrid MPC protocols
14. Chiesa A, Hu Y, Maller M, Mishra P, Vesely N, Ward N (2020) Marlin: preprocessing zkSNARKs with universal and updatable SRS. In: EUROCRYPT
15. Chin C, Wu H, Chu R, Coglio A, McCarthy E, Smith E (2021) Leo: a programming language for formally verified, zero-knowledge applications, Preprint at <https://ia.cr/2021/651>
16. Cowan M, Dangwal D, Alaghi A, Trippel C, Lee VT, Reagan B. (2021) Porcupine: a synthesizing compiler for vectorized homomorphic encryption. In: PLDI
17. Eberhardt J, Tai S (2018) ZoKrates—scalable privacy-preserving off-chain computations. In: IEEE Blockchain
18. Enderton HB (2001) A mathematical introduction to logic. Elsevier, UK
19. Fournet C, Keller C, Laporte V (2016) A certified compiler for verifiable computing. In: CSF

20. Fox A, Myreen MO, Tan YK, Kumar R (2017) Verified compilation of CakeML to multiple machine-code targets. In: CPP
21. Frankle J, Park S, Shaar D, Goldwasser S, Weitzner D (2018) Practical accountability of secret processes. In: USENIX Security
22. Goldberg L, Papini S, Riabzev M (2021) Cairo—a Turing-complete STARK-friendly CPU architecture, Preprint at <https://ia.cr/2021/0163>
23. Goldwasser S, Micali S, Rackoff C (1985) The knowledge complexity of interactive proof-systems. In: STOC
24. Grubbs P, Arun A, Zhang Y, Bonneau J, Walfish M (2022) Zero-knowledge middleboxes. In: USENIX Security
25. Hopwood D, Bowe S, Hornby T, Wilcox N (2016) Zcash protocol specification. <https://raw.githubusercontent.com/zcash/zips/master/protocol/protocol.pdf>
26. Jiang K, Chait-Roth D, DeStefano Z, Walfish M, Wies T (2023) Less is more: refinement proofs for probabilistic proofs. IEEE S & P
27. Kamara S, Moataz T, Park A, Qin L (2021) A decentralized and encrypted national gun registry. In: IEEE S & P
28. Kaufmann M, Manolios P, Moore JS (2013) Computer-aided reasoning: ACL2 case studies, vol 4. Springer Science & Business Media, USA
29. Kosba A, Papadopoulos D, Papamanthou C, Song D (2020) MIRAGE: succinct arguments for randomized algorithms with applications to universal ZK-SNARKs. In: USENIX Security
30. Kosba A, Papamanthou C, Shi E (2018) xJsnark: a framework for efficient verifiable computation. In: IEEE S & P
31. Kothapalli A, Parno B (2022) Algebraic reductions of knowledge, Preprint at <https://ia.cr/2022/009>
32. Kumar R, Myreen MO, Norrish M, Owens S (2014) CakeML: a verified implementation of ML. In: POPL
33. Kundu S, Tatlock Z, Lerner S (2009) Proving optimizations correct using parameterized program equivalence. In: PLDI
34. Lattner C, Adve V (2004) LLVM: a compilation framework for lifelong program analysis & transformation. In: CGO
35. Lerner S, Millstein T, Chambers C (2003) Automatically proving the correctness of compiler optimizations. In: PLDI
36. Lerner S, Millstein T, Rice E, Chambers C (2005) Automated soundness proofs for dataflow analyses and transformations via local rules. In: POPL
37. Leroy X (2009) Formal verification of a realistic compiler. Commun ACM 52(7):107–115
38. Leroy X (2009) A formally verified compiler back-end. J Autom Reason 43(4):363–446
39. Lopes NP, Lee J, Hur CK, Liu Z, Regehr J (2021) Alive2: bounded translation validation for LLVM. In: PLDI
40. Lopes NP, Menendez D, Nagarakatte S, Regehr J (2015) Provably correct peephole optimizations with Alive. In: PLDI
41. Mullen E, Zuniga D, Tatlock Z, Grossman D (2016) Verified peephole optimizations for CompCert. In: PLDI
42. Necula GC (2000) Translation validation for an optimizing compiler. In: PLDI
43. Niemetz A, Preiner M, Reynolds A, Zohar Y, Barrett C, Tinelli C (2019) Towards bit-width-independent proofs in SMT solvers. In: CADE
44. Niemetz A, Preiner M, Reynolds A, Zohar Y, Barrett C, Tinelli C (2021) Towards satisfiability modulo parametric bit-vectors. J Autom Reason 65(7):1001–1025
45. Nipkow T, Wenzel M, Paulson LC (2002) Isabelle/HOL: a proof assistant for higher-order logic. Springer
46. Noir. <https://noir-lang.github.io/book/index.html>
47. Ozdemir A, Brown F, Wahby RS (2022) Circ: compiler infrastructure for proof systems, software verification, and more. In: IEEE S & P
48. Ozdemir A, Kremer G, Tinelli C, Barrett C (2023) Satisfiability modulo finite fields. In: CAV, Preprint at <https://ia.cr/2023/091>
49. Ozdemir A, Wahby R, Whitehat B, Boneh D (2020) Scaling verifiable computation using efficient set accumulators. In: USENIX Security
50. Ozdemir A, Wahby RS, Brown F, Barrett C (2023) Bounded verification for finite-field-blasting. In: CAV, (conference version)
51. Ozdemir A, Wahby RS, Brown F, Barrett C (2023) Bounded verification for finite-field-blasting. Cryptology ePrint Archive, <https://ia.cr/2023/778>, (Full version)

52. Parno B, Howell J, Gentry C, Raykova M (2016) Pinocchio: nearly practical verifiable computation. *Commun ACM* 59(2):103–112
53. Pnueli A, Siegel M, Singerman E (1998) Translation validation. In: TACAS
54. Ranise S, Tinelli C, Barrett C (2017) SMT fixed size bit-vectors theory. <https://smtlib.cs.uiowa.edu/theories-FixedSizeBitVectors.shtml>
55. Sasson EB, Chiesa A, Garman C, Green M, Miers I, Tromer E, Virza M (2014) Zerocash: decentralized anonymous payments from Bitcoin. In: IEEE S & P
56. Setty S, Braun B, Vu V, Blumberg AJ, Parno B, Walfish M (2013) Resolving the conflict between generality and plausibility in verified computation. In: EuroSys
57. Snarky. <https://github.com/o1-labs/snarky>
58. Stewart G, Beringer L, Cuellar S, Appel AW (2015) Compositional CompCert. In: POPL
59. Tan YK, Myreen MO, Kumar R, Fox A, Owens S, Norrish M (2019) The verified CakeML compiler backend. *J Funct Programm* 29:e2
60. Tao R, Shi Y, Yao J, Li X, Javadi-Abhari A, Cross AW, Chong FT, Gu R (2022) Giallar: push-button verification for the Qiskit quantum compiler. In: PLDI
61. Thaler J (2022) Proofs, arguments, and zero-knowledge. Manuscript
62. The Qiskit authors and maintainers: Qiskit: an open-source framework for quantum computing (2021). <https://doi.org/10.5281/zenodo.2573505>, The Qiskit maintainers request that the full list of Qiskit contributors be included in any citation. Regretfully, we cannot comply, as the list is two pages long
63. Tinelli C (2015) SMT core theory. <https://smtlib.cs.uiowa.edu/theories-Core.shtml>
64. Viand A, Jattke P, Hithnawi A (2021) SoK: fully homomorphic encryption compilers. In: IEEE S & P
65. Wahby RS, Setty S, Howald M, Ren Z, Blumberg AJ, Walfish M (2015) Efficient RAM and control flow in verifiable outsourced computation. NDSS
66. Walfish M, Blumberg AJ (2015) Verifying computations without reexecuting them. *Commun ACM* 58(2):74–84
67. Wang F (2022) Ecne: automated verification of ZK circuits <https://0xparc.org/blog/ecne>
68. Zinc. <https://zinc.matterlabs.dev/>
69. Zohar Y, Irfan A, Mann M, Niemetz A, Nötzli A, Preiner M, Reynolds A, Barrett C, Tinelli C (2022) Bit-precise reasoning via int-blasting. In: CADE
70. ZoKrates. <https://zokrates.github.io/>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.