ZKSMT: A VM for Proving SMT Theorems in Zero Knowledge

Daniel Luick, John C. Kolesar, and Timos Antonopoulos, *Yale University;*
William R. Harris and James Parker, *Galois, Inc.;* Ruzica Piskac, *Yale University;*
Eran Tromer, *Boston University;* Xiao Wang and Ning Luo, *Northwestern University*

This paper is included in the Proceedings of the 33rd USENIX Security Symposium.

August 14–16, 2024 • Philadelphia, PA, USA

978-1-939133-44-1

ZKSMT: A VM for Proving SMT Theorems in Zero Knowledge

Daniel Luick
Yale University
daniel.luick@yale.edu

John C. Kolesar
Yale University
john.kolesar@yale.edu

Timos Antonopoulos
Yale University
timos.antonopoulos@yale.edu

William R. Harris
Galois, Inc.
bll.hrris@gmail.com

James Parker
Galois, Inc.
james@galois.com

Ruzica Piskac
Yale University
ruzica.piskac@yale.edu

Eran Tromer
Boston University
tromer@bu.edu

Xiao Wang
Northwestern University
wangxiao@northwestern.edu

Ning Luo
Northwestern University
ning.luo@northwestern.edu

Abstract

Verification of program safety is often reducible to proving the unsatisfiability (i.e., validity) of a formula in Satisfiability Modulo Theories (SMT): Boolean logic combined with theories that formalize arbitrary first-order fragments. Zero-knowledge (ZK) proofs allow SMT formulas to be validated without revealing the underlying formulas or their proofs to other parties, which is a crucial building block for proving the safety of proprietary programs. Recently, Luo et al. (CCS 2022) studied the simpler problem of proving the unsatisfiability of pure Boolean formulas but does not support proofs generated by SMT solvers. This work presents ZKSMT, a novel framework for proving the validity of SMT formulas in ZK. We design a virtual machine (VM) tailored to efficiently represent the verification process of SMT validity proofs in ZK. Our VM can support the vast majority of popular theories when proving program safety while being complete and sound. To demonstrate this, we instantiate the commonly used theories of equality and linear integer arithmetic in our VM with theory-specific optimizations for proving them in ZK. ZKSMT achieves high practicality even when running on realistic SMT formulas generated by Boogie, a common tool for software verification. It achieves a three-order-of-magnitude improvement compared to a baseline that executes the proof verification code in a general ZK system.

1 Introduction

Formal verification is the process of using mathematical reasoning to prove the correctness of programs. It has been used to verify large-scale real-world programs like compilers [40], operating systems [29, 36], and the Transport Layer Security (TLS) protocol [9]. To confirm that a program adheres to some property, both are translated into some mathematical formalism so that the problem of reasoning about programs is reduced to reasoning about mathematical objects.

Boolean logic is the simplest formalism used for verification, but almost all formal verification tasks need something beyond pure Boolean logic. Satisfiability Modulo Theories (SMT) is a well-explored formalism that extends the concept of Boolean satisfiability with theories such as equality with uninterpreted functions and linear integer arithmetic. Tools known as SMT solvers [4, 14, 16, 19, 32, 46] are among the most widely used verification tools. SMT solvers reason about SMT formulas automatically: they can generate both proofs for valid SMT formulas and counterexamples for invalid ones.

Standard techniques for program verification require all relevant information to be completely public: both the program and the proof must be available to everyone who wants to check whether the program is safe. In practice, the owner of the program and the verifier of the program are not always the same entity, and the two may not trust each other. If a program contains sensitive intellectual property, the owner of the program cannot demonstrate the program's safety without revealing the program itself. This limitation results in real-world situations where vendors are forced to reveal their software. For example, cryptographic modules must be FIPS 140-2 [3] certified to be allowed for use in US government systems. The certification process requires that auditors inspect the cryptographic software and run a series of test vectors on the software. Instead of requiring vendors to share their proprietary software for certification, a better approach would enable vendors to prove compliance while keeping their intellectual property secret.

Zero-knowledge (ZK) proofs are a cryptographic primitive that could make this approach a reality. ZK proofs enable a prover to demonstrate that they know a witness w that satisfies a public predicate P without revealing anything about the value of w. In a secrecy-preserving program verifier built on ZK, the predicate will include (1) a formal but easily understood specification of *all* of a program's execution, defined over concepts in a high-level language (including, e.g., tuples, arrays, and classes), and represented as an SMT formula, and (2) a binary executable that the vendor is prepared to distribute (perhaps because it is sufficiently obfuscated). The witness would include (a) a high-level program that satisfies the specification, (b) a proof that it does so, and (c) a proof that the secret high-level program is observationally equivalent to the public binary executable. Implementing a complete

framework that provides such a proof is a massive undertaking that will require several major conceptual and engineering achievements. This paper focuses on proving component **(b)**, which depends primarily on validating SMT deductions.

To instantiate such a system, one could take an existing tool that can convert any C-like program (e.g., [7, 15, 27, 31]) and apply it to execute the program that verifies the validation of the SMT formula in ZK. However, we observe that such an approach does not sufficiently scale. When using a state-of-the-art ZK toolchain [15] to run on a short benchmark with only 6 steps, the end-to-end running time is almost *two hours* (Sec. 7.3)! Since typical SMT proofs often need hundreds of steps, this approach is clearly impractical, due to a few reasons:

1. To prove arbitrary programs in ZK, all tools adopt the Von Neumann architecture, providing an execution environment that resembles the cleartext computation. However, translating the SMT proof verification program to such a format (e.g., TinyRAM [7]) incurs a huge overhead, a necessary cost to achieve the highest expressiveness.

2. Each SMT theory has its own verification techniques, which in turn means different optimization opportunities in ZK. Using a generic tool essentially prohibits theory-specific optimizations.

3. Supporting random access memory (RAM) in ZK protocols is often the most costly component [8, 17, 20, 30, 33, 44, 48, 49]. Although SMT verification features unique access patterns in how and when it reads from RAM, they cannot be captured in a generic toolchain.

Essentially, we need a framework that allows modular support of new theories (like cleartext SMT solvers) and the flexibility to introduce customized protocols for different rules. This framework should be compact, general, and compatible with common ZK optimizations simultaneously. The first two require reasonable proof size and the ability to express reasoning on first-order theories. Achieving a level of usability for these two features in SMT in cleartext took decades [4, 14, 16, 19, 32, 46]. Introducing a layer of ZK to the SMT-proof system should ideally maintain the same level of compactness and generality while being efficient; this requires a ZK mindset from the outset. Finally, the whole framework should allow incremental development, meaning that the support of different SMT theories could be added over time.

Our Contributions and Technical Novelty In this paper, we introduce ZKSMT, an efficient and extendable framework for zero-knowledge proofs of SMT formulas. Unlike ZKUN-SAT [42], which only supports propositional logic, the goal of ZKSMT is to support first-order logic, which is more versatile and more commonly used for software verification (See Section 8 for a full comparison with ZKUNSAT). ZKSMT represents and validates complex SMT proofs involving a

dynamic set of rules, whereas ZKUNSAT's proofs only require repeated execution of a single rule. Ideally, we want a system that can be extended with more rules easily, can be encoded in ZK with high efficiency, and is not tied to any specific cryptographic backend. To this end, we adopt a VM approach.

- We introduce ZKSMT, a virtual machine that allows efficient encoding and validation of refutation proofs of SMT formulas in the context of zero-knowledge proofs. Compared to VMs that encode generic computation, ZKSMT is specifically designed for SMT validation and arithmetization at the same time. In particular, our VM can be efficiently instantiated with look-up tables and polynomial operations, both of which are efficiently supported in many ZK protocols. In addition, ZKSMT is expressive: it can be instantiated to efficiently validate any SMT formula in first-order logic.

- We instantiate three common theories in ZKSMT: Boolean logic, equality with uninterpreted functions (EUF), and linear integer arithmetic (LIA). To achieve practical efficiency, we design optimized arithmetizations of these theories within our VM. For example, we show that many rules can be checked efficiently as constraints over multisets, which in turn can be efficiently arithmetized as operations over polynomials, resulting in improved complexity.

- We implement ZKSMT in ZK and benchmark it over formulas that are generated by the Boogie verification toolchain [5] and the Wisconsin Safety Analyzer [2] benchmark suite (from the official SMT-LIB benchmark set [1]). The results for Boogie show that ZKSMT achieves a speed-up of more than three orders of magnitude compared to a state-of-the-art system [15]. ZKSMT can also verify an "ultra-large" proof instance from the Wisconsin Safety Analyzer set with 200,000 proof steps in about 3 hours.

Information Leakage Our system does reveal some size parameters of the proof (e.g., the number of proof steps). We also made some privacy-efficiency trade-offs by revealing the number of occurrences (but not the order) of each proof rule. Together with other techniques, our trade-offs enable the impressive improvement mentioned above.

2   Preliminaries

2.1   Quantifier-Free First-Order Logic

Formula Structure Logical formulas are mathematical statements that assert a property of functions and predicates; the class of formulas that we consider in this work have the following structure. A set of function symbols is a set in which each element has an arity, denoted $|f|$ for $f \in \mathcal{F}$. The arity of a function may be any natural number, including 0. The set of terms over function symbols $\mathcal{F}$ and variables $\mathcal{V}$, denoted $T_{\mathcal{F}, \mathcal{V}}$, is the smallest set containing $\mathcal{V}$ and $f(t_0, \ldots, t_{|f|-1})$ for all function symbols $f \in \mathcal{F}$ and terms $t_i \in T_{\mathcal{F}, \mathcal{V}}$. For instance,

the function symbols for linear integer arithmetic include all integer literals n, with $|n| = 0$, the negation operator $-$, with $|-| = 1$, and the addition operator $+$, with $|+| = 2$. An example of a term over these function symbols and the variable x is $-(x + 3)$.

Predicate symbols, similar to function symbols, are a set equipped with arities. The set of atoms over variables $\mathcal{V}$, function symbols $\mathcal{F}$, and predicate symbols $\mathcal{P}$ is the set of all $P(t_0, \ldots, t_{|P|})$ for predicate symbols $P \in \mathcal{P}$ and terms $t_i \in T_{F,V}$. The formulas over $\mathcal{F}$, $\mathcal{P}$, and $\mathcal{V}$ are all Boolean combinations of atoms over $\mathcal{F}$, $\mathcal{P}$, and $\mathcal{V}$, i.e. all objects built from atoms using the distinguished formulas True and False and the constructors negation, conjunction, disjunction, and implication. For example, linear integer arithmetic has the predicate symbols $=$, $\leq$, and $<$, with $|=| = |\leq| = |<| = 2$. A formula over these function and predicate symbols and the variable x is $x = 0 \vee 10 \leq x + 2$.

The definitions of terms and formulas can be described by the following BNF grammar for terms t and formulas b:

$$t ::= v \mid f(t_0, \ldots, t_n)$$
$$b ::= \text{True} \mid \text{False} \mid P(t_0, \ldots, t_n) \mid \neg b_0 \mid \bigwedge\{b_0, \ldots, b_n\} \mid$$
$$\bigvee\{b_0, \ldots b_n\} \mid b_0 \to b_1$$

Formula Validity and Proofs One approach for assigning meaning to function and predicate symbols is to specify which of the formulas defined over them are conclusions of a given set of assumed formulas. Evidence that a formula is a conclusion of some assumptions $\mathcal{A}$ is represented as a proof: a tree-shaped argument whose nodes are formulas, each derived from its children by a step of inference.

More precisely, a theory is a set of automatically recognizable proof steps, each of which consists of **(1)** a symbol, referred to as the rule identifier, which has a finite arity, **(2)** a set of formulas known as the premises, and **(3)** a formula referred to as the conclusion. The proofs of a formula φ in theory $\mathcal{T}$ under assumed formulas $\mathcal{A}$ are the smallest set such that **(1)** each assumption $\varphi \in \mathcal{A}$ is a proof of itself, and **(2)** if $P_0, \ldots, P_n$ are proofs of formulas $\varphi_0, \ldots, \varphi_n$, and R is a proof step with φ' as its conclusion and $\varphi_0, \ldots, \varphi_n$ as its premises, then R and $P_0, \ldots, P_n$ form a proof of φ'. If φ has a proof in $\mathcal{T}$ under $\mathcal{A}$, then φ is derived in $\mathcal{T}$ from $\mathcal{A}$. A refutation of formula φ in theory $\mathcal{T}$ is a proof of False in $\mathcal{T}$ under assumption φ. Multiple theories can be combined into a single theory by combining the programs that recognize applications of their proof rules. Thus, when convenient, we may consider either individual theories in isolation (to explain facts that they can derive) or a combination of multiple theories (when describing benchmarks that use many theories in combination).

Defining a formal theory for a previously unformalized domain of interest, and obtaining assurance that it proves exactly the formulas of interest, can be non-trivial. Our work is applicable in a setting where each theory of interest is accompanied by a public definition of the theory as a set of

inference rules that the prover and verifier have agreed allows the derivation of only desired conclusions from assumptions.

2.2 SMT Theories of Interest

We now introduce illustrative examples of inference rules that define three logical theories of central importance: propositional logic, equality with uninterpreted functions, and linear integer arithmetic. Each of these theories is commonly used by program verifiers to verify critical properties of software, and each is supported by the current implementation of our protocol. Each inference rule is presented using a standard notation where the rule's premises occur above a horizontal bar and its conclusion occurs below.

2.2.1 Propositional Logic

Propositional logic rules—i.e., how formulas constructed from conjunction, disjunction, and negation can be proved and used to prove other formulas—include the following.

ExclMid The rule ExclMid formalizes the law of the excluded middle, stating that each proposition or its negation must hold:

$$\overline{a \vee \neg a}$$

Resolution The rule Res formalizes the idea of reasoning by case splitting. If both $p \vee A$ and $\neg p \vee B$ hold, then either A must hold (when $\neg p$ holds) or B must hold (when p holds):

$$\frac{p \vee A \quad \neg p \vee B}{A \vee B}$$

DeDup The de-duplication rule DeDup prunes duplicated disjuncts. A weak form of it (which can be applied n times to prune disjuncts that are repeated n times) is this:

$$\frac{a \vee a \vee B}{a \vee B}$$

Given that resolution alone is complete for proving refutations in propositional logic and there are existing protocols that verify resolution proofs in ZK [42], it may be surprising that we consider a large collection of rules instead of a minimal subset. However, practical SMT theorem provers [14] often generate proofs that use many distinct rules, both to minimize their tool's output and to simplify their implementations. While such proofs could be rewritten to use a more restricted rule set, the consequences for both the size of the resulting proof and the performance of a subsequent ZK proof that verifies it are non-obvious and well beyond the scope of the this work.

2.2.2 Equality with Uninterpreted Functions

The theory of equality with uninterpreted functions (EUF) enables SMT to describe general properties of system operations without explicitly defining their complete behavior, which can be helpful for modeling complex systems that consist of multiple modules. EUF contains three rules—Reflexivity, Symmetry, and Transitivity—that express the fact that equality

is reflexive, symmetric, and transitive (i.e., that it is, unsurprisingly, an equivalence relation); their definitions are straightforward. It also contains an infinite family of rules, $\mathsf{Cong_n}$ for all $n \in \mathbb{N}$, which express that applying an n-ary function f to n equal arguments produces equal results:

$$\frac{a_0 = b_0 \quad \ldots \quad a_{n-1} = b_{n-1}}{f(a_0, \ldots, a_{n-1}) = f(b_0, \ldots, b_{n-1})}$$

2.2.3 Linear Integer Arithmetic

Linear integer arithmetic (LIA) is a commonly used first-order theory of integers that includes addition and multiplication by constants but does not permit multiplication between variables. It is used to model the semantics of both bounded and unbounded arithmetic.

Multiplication Distribution The rule MulDist is the general law that multiplication distributes over addition, specialized to the case of constant left factors. It can be applied to derive, e.g., that $4 * (2x + 3y) = 8x + 12y$. Its general form is this:

$$\overline{c * (\sum_{i=0}^{n} d_i * x_i) = \sum_{i=0}^{n} c * d_i * x_i}$$

Here $x_0, \ldots, x_n$ are arbitrary terms; $c, d_0, \ldots, d_n$ are constants.

Farkas' Lemma Farkas' Lemma derives strict inequalities from the terms in a larger strict inequality. It can be expressed as a family of inference rules, indexed by a term size n:

$$\frac{\sum_{i=0}^{n} c_i * a_i - c_i * b_i > 0}{\bigvee_{i=0}^{n} a_i > b_i}$$

A similar rule can be applied to derive a slightly more constrained disjunction when a linear term of the identical form is given to be equal to 0.

2.3 An Example Formalizing Software Safety

We can use EUF and LIA to model safety properties for numerical type conversion in languages like C. Suppose that we want to prove the safety of a function that overwrites an entry in an array:

```
1 void find_and_replace(int[] a, int x, int y) {
2   int i = get_position(a, x);
3   if (0 <= i) {
4     a[i] = y;
5   }
6 }
```

The helper function `get_position(a, x)` returns the index of the first occurrence of `x` within `a` and returns -1 if `x` does not appear in `a`. `find_and_replace` includes a safeguard for the possibility that `x` is not in `a`: in that event, it does nothing rather than attempting to write to an index of `a`. We can prove the safety of `find_and_replace` using the following SMT formula:

$$g(a, x) = -1 \lor (0 \le g(a, x) \land g(a, x) < l(a)) \tag{1}$$
$$\land \; i = g(a, x) \tag{2}$$
$$\land \; 0 \le i \tag{3}$$
$$\land \; \neg(0 \le i \land i < l(a)) \tag{4}$$

Line 1 represents the behavior of `get_position(a, x)`. We abbreviate the name `get_position` to `g`, and `l` is a function that returns the length of an array. Line 2 represents the assignment of a value to `i`, and line 3 indicates that we are modeling the situation when the conditional `if` is satisfied. Line 4 represents the safety property that we want to establish: the index `i` is within the bounds of the array when the check `0 <= i` passes. Instead of phrasing the safety property as a proof goal, we aim to prove the unsatisfiability of the scenario where the safety property is negated.

We can give this formula as an input to an external SMT solver that produces a refutation proof that ZKSMT can use. In Sec. 3.1, we will discuss the encoding that ZKSMT uses to represent the refutation proof for this formula.

2.4 Zero-Knowledge Proofs

A zero-knowledge proof [24, 26] allows a prover to convince a verifier that it possesses an input w such that $P(w) = 1$ for some public predicate P, while revealing no additional information about w. There have been many lines of work in designing practically efficient ZK protocols under different settings and assumptions (e.g., [25, 28, 34, 35]). ZKSMT uses a special type of ZK protocol commonly referred to as "commit-and-prove" ZK [13], which allows a witness to be committed and later proven over multiple predicates while ensuring consistency of the committed values.

Although ZKSMT can be instantiated with any commit-and-prove ZK, we use the recent VOLE-ZK series for maximum efficiency [6, 18, 51] and, in particular, take advantage of optimizations for polynomials [53] and RAM operations [20]. We also use the permutation check originally proposed by [10].

Note that ZK proofs and refutation proofs are two different concepts, one in cryptography and one in formal methods. The verification procedure of a refutation proof is encoded as a statement proven by two parties using a ZK protocol.

3 ZKSMT Architecture

To verify an SMT refutation proof, ZKSMT examines the whole proof, step by step, in a loop: one such step is depicted in Fig. 1. In each iteration, ZKSMT **(1)** fetches the rule to be applied to the current step, **(2)** fetches the rule's premises, and **(3)** verifies that the derived formula is a valid conclusion of the proof rule. The overall structure resembles the design of a Von Neumann processor that executes only straight-line instructions (i.e., instructions that always transfer control to their successor). The available set of proof rules resembles a CPU's set of supported instructions. The proofs themselves
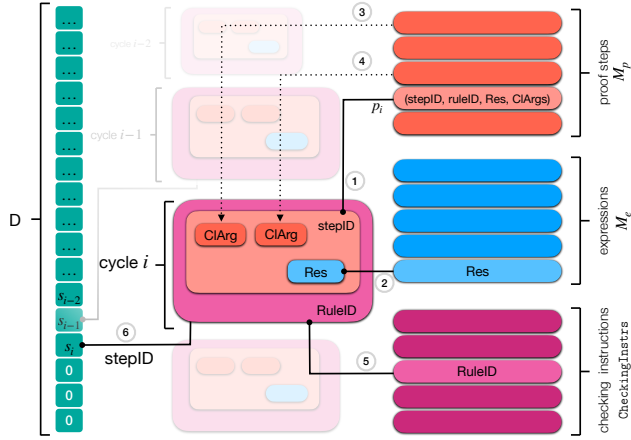
Figure 1: The retrieval and processing of information over one of ZKSMT's steps. Operations are numbered in the order of their occurrence. Data concerning rules applied and proof expressions to the right is used to check step validity, depicted in the middle. The result of the check is written to a storage cell in D, on the left.

are similar to programs composed of sequential instructions. In this analogy, the checking instruction responsible for each individual proof rule can be envisioned as analogous to a CPU's arithmetic-logic unit to handle specific computations.

Meanwhile, the main checker acts as the control unit, orchestrating the overall verification process. For each proof step, ZKSMT relies on a fixed-length array of formulas to store premises associated with the current step, functioning much like instruction operands. Furthermore, temporary storage is needed for a derived conclusion, a pointer to the next proof step, etc. The expression table, similar to the memory in CPU architecture, is read-only in this context. Our main philosophy is to develop a flexible VM that can efficiently encode and verify SMT refutation proofs when the underlying VM is instantiated using ZK protocols. This way, we can plug in any suitable ZK protocol for ZKSMT and bring in optimizations in CPU design. Below we introduce our VM's encoding for formulas and proofs and its execution strategy.

3.1 Encoding Formulas and Proofs

We first explain how ZKSMT represents SMT formulas and checks that particular formulas can be proved from others according to the rules of a logical theory. ZKSMT's encoding of SMT formulas, and the complex terms that they may contain, critically enables it to prove formulas in theories beyond what existing techniques can support.

Encoding Formulas in an Expression Table Every formula is constructed from an operator applied to a finite collection of smaller formulas or terms; thus, it can be represented naturally as an AST. In particular, if we view formulas as being defined by the BNF grammar from Sec. 2.1, we can think of every individual production option used to produce a formula

as a node in the formula's AST. The sub-productions are the node's children. Note that even semantically equivalent formulas can have distinct ASTs (e.g., False and ¬True).

ZKSMT stores the ASTs for all formulas involved in the proof in a read-only table M_e, called the *expression table*. We refer to entries in the table as *expressions*. Each expression represents an individual node within the AST of a formula. A node of an AST has three fields: the node ID (NodeID), the immediate addressing list (ImmAddr), and the indirect addressing list (IndAddr). NodeID specifies the operator being employed, such as Eq or Mul. ImmAddr is used to identify constants and immediate values, like an immediate value of an operand in a CPU. We also consider the names of variables as immediate values. AST nodes with children store the indices of their children within the expression table under the IndAddr field. Most expressions, such as logical negation (Not) and equality (Eq), have a fixed number of children. Others, including Boolean conjunction (And), disjunction (Or), and applications of uninterpreted functions (Apply), can have a variable number of entries within IndAddr.

Note that not all nodes in the table are formulas: some entries simply represent parts of other rows' formulas. A row that encodes a term or formula can have multiple other rows pointing to it if the term/formula appears in different formulas (which can even come from different theories). For example, in Table 1, i has only one entry even though it appears within $i = g(a,x)$, $0 \le i$, and multiple other formulas.

Example 3.1. Table 1 shows a portion of the expression table for the proof in Table 2. The entry with address &14 in Table 1 represents the formula $i < l(a)$, whose NodeID is Lt (less than). The values indicated within the IndAddr field represent the indices for the sub-expression children of $i < l(a)$; specifically, the indices of i (entry &2) and $l(a)$ (entry &5). The sub-expression $l(a)$ (entry &5) is a term rather than a formula and has one sub-expression child a (&3) and the label l as an immediate value stored in ImmAddr.

Encoding Proof Steps A proof step in a theory T consists of an application of a rule, labeled with an identifier with a fixed arity n to formulas $\varphi_0, \ldots, \varphi_n$ to conclude a formula φ (Sec. 2.1). The steps of a theory T of interest are checked in ZKSMT by a finite set of step *checking instructions*, each one checking steps identified by a corresponding rule of T. An occurrence p of a checking instruction has four fields:

- StepID: the position of the step in the execution order.
- RuleID: the identifier of the applied theory rule. The rule identifier r of theory T is identified as the pair (R, T).
- Premises: a list of the StepID's of $\varphi_0, \ldots, \varphi_n$. Each StepID points to a previous step, whose derived formula is a premise of p.
- Result: an index into the expression table to identify the conclusion φ of the current proof step.

Addr.	NodeID	ImmAddr	IndAddr	Meaning
&1	Var	ONE	{}	ONE
&2	Var	i	{}	i
&3	Var	a	{}	a
&4	Var	x	{}	x
&5	Apply	l	{&3}	$l(a)$
&6	Apply	g	{&3,&4}	$g(a,x)$
&7	Mul	0	{&1}	$0 * \mathtt{ONE}$
&8	Mul	-1	{&1}	$-1 * \mathtt{ONE}$
&9	Eq		{&2,&6}	$i = g(a,x)$
&10	Eq		{&7,&7}	$0 = 0$
&11	Leq		{&7,&8}	$0 \le -1$
&12	Not		{&11}	$\neg(0 \le -1)$
&13	Leq		{&7,&2}	$0 \le i$
&14	Lt		{&2,&5}	$i < l(a)$
&15	And		{&13,&14}	$0 \le i \wedge i < l(a)$
&16	Not		{&15}	$\neg(0 \le i \wedge i < l(a))$

Table 1: Part of the expression table M_e for the proof of the safety of `find_and_replace`. We use & to denote the addresses of expressions.

StepID	RuleID	Premises	Result
#2	Assume		**&9**: $\quad i = g(a,x)$
#3	Assume		**&13**: $0 \le i$
#4	Assume		**&16**: $\neg(0 \le i \wedge i < l(a))$
…			
#7	Res	{#5,#6}	$0 - (-1) = 1$
#8	Farkas	{#7}	**&12**: $\neg(0 \le -1)$
…			
#10	Cong		$\neg((0 \le i) = (0 \le -1))$ $\vee \neg(0 = 0) \vee \neg(i = -1)$
#11	Res	{#9,#10}	$\neg(0 = 0) \vee \neg(i = -1)$
#12	Refl		**&10**: $0 = 0$
#13	Res	{#11,#12}	$\neg(i = -1)$
…			
#16	Res	{#14,#15}	**&15**: $0 \le i \wedge i < l(a)$
#17	Res	{#4,#16}	False

Table 2: Part of the proof step array M_p for the proof of the safety of `find_and_replace`. Not all conclusions' addresses are shown. We use # to denote the IDs of proof steps.

A set of instructions T is a *T-logical unit* if there is a bijection from rule identifiers of T to instructions in T such that each instruction succeeds if and on if it is executed in a machine state in which it points encoding of premises and a conclusion that can be derived using its corresponding rule in T.

Size Parameters Five parameters bound the resources used by a ZKSMT instance. **(1)** π is the maximum number of proof steps. It parallels the concept of program size in CPU design and defines the extent of the proof structure that can be examined in a manner similar to how the size of a program in a CPU determines the number of instructions it can execute. **(2)** χ is the maximum number of expressions the proof can use, analogous to the size of the CPU's memory. **(3)** μ is the maximum number of premises in any rule, analogous to the number of registers in a CPU. **(4)** α is the maximum argument list size of any expression, where the argument list size of an expression e is defined as $|e.\mathsf{ImmAddr}| + |e.\mathsf{IndAddr}|$; it is analogous to the bit width of memory entries. **(5)** ρ is the number of distinct rules used in the proof, analogous to the size of the architecture's instruction set. The set of ZKSMT machines over checking instructions T on particular size parameters is denoted $\mathsf{ZKSMT}[\mathtt{T}](\pi, \chi, \mu, \alpha, \rho)$.

To define the necessary components of the machine, we often use the bit widths of these numbers: $\ell_p = \lceil \log(\pi) \rceil$, $\ell_e = \lceil \log(\chi) \rceil$, and $\ell_r = \lceil \log(\rho) \rceil$.

Example 3.2. Some of the entries of M_e and M_p for the refutation of the formula in Sec. 2.3 are shown in Table 1 and Table 2, respectively. The proof applies rules from EUF and LIA as well as rules for Boolean connectives. Most of the 1,041 steps in the proof are omitted, and some of the steps that we show are simplified. For example, we do not show the steps for adding and removing singleton disjunctions.

3.2 Machine Specification and Execution

Once the encodings are specified, we can build the VM on top of them. We show the overall architecture in Fig. 1.

Machine Specification ZKSMT has five main components:
- pc: the *proof counter*, an ℓ_p-bit integer.

- $\{r_i\}, \{t_i\}$: the list of *registers* that store information for the proof step currently being examined. The machine has $2\mu + 2$ registers in total: r_0 stores the conclusion, $r_1, \ldots, r_\mu$ store the premises, and r_{rule} stores the rule ID. The first $\mu + 1$ registers are of size ℓ_e, and r_{rule} is of size ℓ_r. The registers $\{t_1, \ldots, t_\mu\}$ store the addresses of $r_1, \ldots, r_\mu$. The main checker uses them when fetching the premises of a proof step from M_e. Each t_i is of size ℓ_e.

- M_e: the *expression table*, a read-only array of size χ that contains all expressions used in the proof, using the encoding system that we explained in Sec. 3.1.

- M_p: the *step table*, a read-only array of size π that contains all the proof steps used in the proof.

- D: the checking order of the proof. The checking order is the order in which proof steps are validated during execution. If $\mathsf{D}[i] = j$, then the validation of the j^{th} proof step occurs on the i^{th} iteration of the main verification loop.

Machine Execution to Validate a Proof As mentioned above, ZKSMT's proof validation process closely resembles how a machine program is executed in the Von Neumann architecture (using a CPU, memory, etc.). To provide more flexibility in VM execution, we distinguish two orderings: the logical ordering and the checking ordering. The logical ordering is the original ordering of the proof as outlined in Sec. 2: a proof step should not use a result proven in a step that occurs

Algorithm 1: ZKSMT[T]$(\pi, \chi, \mu, \alpha, \rho)$'s execution

Output: True, False

1   D $\leftarrow [0, \ldots, 0]$;
2   **for** pc = 0 *to* $\pi - 1$ **do**
3   **Proof Step Fetch:**
4   (StepID, RuleID, Res, ClArgs) $\leftarrow M_p[\text{pc}]$;
5      $r_{\text{rule}} \leftarrow$ RuleID;
6   **Conclusion Fetch:**
7      $r_0 = M_e[\text{Res}]$;
8   **Premise Fetch:**
9      $t_1, \cdots, t_\mu \leftarrow M_p[\text{ClArgs}_0], \cdots, M_p[\text{ClArgs}_{\mu-1}]$;
10      $r_1, \cdots, r_\mu \leftarrow M_e[t_1.\text{Res}], \cdots, M_e[t_\mu.\text{Res}]$;
11   **Rule Checking:**
12   CheckingInstrs$[r_{\text{rule}}](r_0, \{r_1, \cdots, r_\mu\})$;
13   **Cycle Checking:**
14   **for** $j = 1$ *to* μ **do**
15   assert($t_j.\text{StepID} <$ StepID);
16   D$[i] \leftarrow$ StepID ;
17   **assert**(PermuteCheck(D, $[0, \ldots, \pi - 1]$));
18   TypeCheck(M_e);

later in the logical ordering. The StepID of each proof step is its logical ordering. However, the checking order, which is the order in which proof steps are validated during the VM's execution, does not need to have any relationship with the logical ordering other than the former being a permutation of the latter. Allowing the two orderings to differ helps us to hide the structure of a proof when we instantiate the VM in ZK (Sec. 5.2).

Algorithm 1 provides an overview of ZKSMT's algorithm, which iterates over the set of all proof steps (line 2). Each proof step is verified over five phases: proof step fetching, conclusion fetching, premise fetching, rule checking, and cycle checking. In the fetching phases (lines 3–10), ZKSMT fetches the relevant elements for that step from the tables M_p and M_e based on the values in the fields Result and Premises and stores them in $r_0, \ldots, r_\mu$. Next, the checker determines the checking instruction to execute by examining the value specified in RuleID (lines 11–12). It delegates the responsibility of validating the proof step to the selected instruction and asserts the success of the validation (line 12). Formulas must be proven before being used as premises. Since StepID represents the logical ordering of a derived formula, we can confirm that in cycle checking that the StepID of every formula in Premises for a rule is strictly smaller than the StepID of the rule's conclusion. This is checked by iterating over all rule premises (lines 13–15). To conclude the iteration, the checker assigns StepID to D$[i]$ (line 16).

To address the potential discrepancies between orderings, we need to perform one more check. Every proof step needs to be verified at some point. The array D keeps track of which proof steps have been validated. When the main loop finishes

execution, the main checker verifies that D is a permutation of the list $[0, \cdots, \pi - 1]$ (line 17). If it is, then every step in the refutation proof has been verified.

Well-Formed Expressions The soundness of ZKSMT also relies on the well-formedness of expressions in the table M_e. This can be ensured by a process analogous to proof validation. In particular, we type-check each expression according to a set of type rules, which work similarly to proof rules and are provided as public configurations of ZKSMT. To forbid cyclic expressions, ZKSMT also checks for cycles in M_e, similarly to the check for cycles in proof steps.

3.3 Soundness and Completeness

The following are key properties of ZKSMT that establish that it produces exactly valid SMT formulas. Both theorems are defined over an arbitrary theory T and T-logical unit T, formula φ, and size parameters $\pi, \chi, \mu, \alpha, \rho$ (Sec. 3.1).

In this context, we say that φ is *boundedly verifiable* if it has a derivation in $\mathcal{T}$ containing at most π steps, χ distinct expressions with at most α arguments, and using ρ rules which all have at most μ premises.

Theorem 1 (Soundness). *A VM in* ZKSMT$[T](\pi, \chi, \mu, \alpha, \rho)$ *validates* φ *only if* φ *is boundedly verifiable.*

Theorem 2 (Completeness). *If* φ *is boundedly verifiable, then some VM in* ZKSMT$[T](\pi, \chi, \mu, \alpha, \rho)$ *validates it.*

Proofs of Thm. 1 and Thm. 2 appear in App. B.

4 Instantiating ZKSMT on Practical Theories

In this section, we explain how to instantiate ZKSMT on propositional logic, equality with uninterpreted functions (EUF), and linear integer arithmetic (LIA). We discuss **(1)** the encoding of expressions in each theory, **(2)** the theories' proof rules, and **(3)** the implementations of the checking instructions for an illustrative selection of each theory's rules. Table 3 shows all of the rules that we cover in this section, along with a few others that we discuss later.

4.1 Checking Propositional Logic

We have implemented in ZKSMT an instruction unit that checks applications of the rules of propositional logic. We now describe implementations of checking instructions for selected example rules (Sec. 2.2.1).

ExclMid When the unit instruction that checks applications of ExclMid (the rule formalizing the law of the excluded middle) receives the conclusion expression r_0 from the main checker, it first confirms that r_0's NodeID is Or. Next, the checking instruction retrieves the first two entries a_0 and a_1 from list $r_0.\text{IndAddr}$ and confirms that **(1)** the NodeID of a_0 is Not; and **(2)** the expression table index of a_0's child is the same as the expression table index of a_1. In general, the same technique is implemented by all checking instructions that must check that two expressions are identical: the instructions

RuleID	Side Condition	Premises	Conclusion
Boolean			
Resolution	$\exists p. p \in \langle\!\langle A \rangle\!\rangle, \neg p \in \langle\!\langle B \rangle\!\rangle,$ $\langle\!\langle A \rangle\!\rangle \subseteq \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle p \rangle\!\rangle, \langle\!\langle B \rangle\!\rangle \subseteq \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle \neg p \rangle\!\rangle$	$\bigvee A, \bigvee B$	$\bigvee C$
DeDup	$\forall a \in \langle\!\langle A \rangle\!\rangle.\ a \in \langle\!\langle B \rangle\!\rangle$	$\bigvee A$	$\bigvee B$
ExclMid			$\bigvee \{\neg a, a\}$
EUF			
Congruence	$\exists A, B, f. (fA = fB) \in C, |A| = |B|,$ $\forall i \in \{0, \ldots, |A| - 1\}. \neg(A_i = B_i) \in C$		$\bigvee C$
LIA			
MulDist			$c * (\sum_{i=0}^{n} d_i * x_i) = \sum_{i=0}^{n} c d_i * x_i$
Flatten	$\exists \langle\!\langle C \rangle\!\rangle, \langle\!\langle D \rangle\!\rangle.\ \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle \sum D \rangle\!\rangle = \langle\!\langle A \rangle\!\rangle, \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle D \rangle\!\rangle = \langle\!\langle B \rangle\!\rangle$		$\sum A = \sum B$
Farkas	$\forall i \in \{0, \ldots, n\}.\ m_i \geq 0$ either $c > 0$, or $c = 0$ and $\exists j. \leq_j$ is $<$	$\sum_{i=0}^{n}(m_i * a_i) +$ $(-m_i * b_i) = c$	$\bigvee_{i=0}^{n}\{\neg(a_i \leq_i b_i)\}$

Table 3: A selection of ZKSMT's rules for Boolean logic, EUF, and LIA that we cover in Sec. 4, Sec. 5, and Sec. 7. Tables 4 and 5 in the appendix show all of the proof rules omitted here.

check the equality of indices in the expression table, instead of traversing the expressions' complete ASTs.

Many rules of propositional logic, as in the case of ExclMid, do not have premises. Instead of interacting with the results of previous steps, they introduce simple tautologies that other Boolean rules can use as premises later; the checking instructions for such rules need only to pattern-match the rules' conclusions. However, in general, an instruction may need to validate non-trivial *side conditions* imposed by a rule on the terms matched to its conclusion and premises (similar to the LFSC framework [46]). One example of such a rule is Resolution, whose checking instruction we now describe.

Res The checking instruction for Res (the formalization of unit resolution) checks properties of the *multisets* of propositions that may be in each of its premise clauses. To describe the instruction's implementation, we employ the notation $\langle\!\langle A \rangle\!\rangle$ (or $\langle\!\langle a \rangle\!\rangle$) to represent the multiset containing the elements of a list A (or single element a). Also, $\uplus$ is multiset union.

The checking instruction interprets r_0.IndAddr, from the conclusion r_0, as a multiset $\langle\!\langle C \rangle\!\rangle$ and interprets r_1.IndAddr and r_2.IndAddr, from the premises r_1 and r_2, as multisets $\langle\!\langle A \rangle\!\rangle$ and $\langle\!\langle B \rangle\!\rangle$, respectively. After checking that r_0, r_1, and r_2 are Or nodes, the instruction identifies the expression p, locates p within $\langle\!\langle A \rangle\!\rangle$, and locates $\neg p$ within $\langle\!\langle B \rangle\!\rangle$. Finally, the instruction checks the side conditions $\langle\!\langle A \rangle\!\rangle \subseteq \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle p \rangle\!\rangle$ and $\langle\!\langle B \rangle\!\rangle \subseteq \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle \neg p \rangle\!\rangle$. Note that p can be provided as an *extended witness* so that the checking instruction does not need to search for it.

In general, checking instructions for all propositional rules that have premises, as in the case of Res, must perform pattern matching on both the conclusion r_0 *and* the premises $r_1, \ldots, r_k$ that they receive from the main checker.

Remark 4.1 (Extended witnesses). In the context of zero-knowledge proofs, determining the value of p for the checking instruction for Resolution can be computationally expensive.

To reduce the runtime cost, the proof itself can cache the value of p and provide it for the checking instruction directly. This value serves as an *extended witness*. When it receives an extended witness, the checking instruction only needs to test the side condition on the cached value of p rather than checking all possible options. Multiple other rules use extended witnesses for the same purpose.

DeDup It is straightforward to implement a checking instruction for applications of the de-duplication rule DeDup as presented in Sec. 2.2.1: the instruction simply checks that its conclusion and premise are Or nodes, that the children of the conclusion's node occur in the premise, and that the children of the premise at corresponding positions are identical. However, checking DeDup strictly as presented would unfortunately require a proof to apply it multiple times to remove disjuncts that occur more than twice, and apply another rule formalizing the associativity of disjunction to arrange the premise in an expected form.

Instead, DeDup's actual checking instruction effectively checks repeated applications of such a rule in one step by checking that each distinct element in $\langle\!\langle A \rangle\!\rangle$ is also in $\langle\!\langle B \rangle\!\rangle$, where A is the argument list for the proof step's premise and B is the argument list for its conclusion.

4.2 Checking Equalities with Functions

We have instantiated ZKSMT to refute proofs that use the theory of equality with uninterpreted functions (EUF; Sec. 2.2.2). To check applications of a Congruence$_i$ rule, we model an alternative formulation, easily shown to be logically equivalent to the standard formalization, which derives a disjunction from no premises. The checking instruction for Congruence begins by confirming that the NodeID of the conclusion r_0 is Or. Next, it retrieves the set of expressions indexed by r_0.IndAddr, identifies the pair of function applications, and verifies that the other disjuncts match the corresponding argu-

ments of the function applications.

4.3 Checking Linear Integer Arithmetic

We now describe ZKSMT's representation of expressions from linear integer arithmetic (LIA; Sec. 2.2.3). Then, we discuss implementations of checking instructions in ZKSMT for two LIA rules: MulDist and Farkas.

Expression Representation In ZKSMT's representation of LIA, addition is an n-ary operation, just like $\wedge$ and $\vee$. Singleton sums are allowed, and so are empty sums. The entries in a sum can be arbitrary integer-valued expressions, including other sums. Multiplication in LIA is shorthand for the repeated addition of an expression to itself. A multiplication node always has exactly one child, which can be an arbitrary integer-typed expression. It stores its scaling factor in the ImmAddr field, as we show in entries &7 and &8 in Table 1. The value of the scaling factor can be any integer, positive or negative. We store integer constants in M_e as multiples of a special variable ONE that represents 1. This representation enables checking instructions for rules such as MulDist to assume that the sums in their conclusions contain only Mul nodes rather than having a separate case for integer constants.

Multiplication Distribution MultDist's checking instruction can validate an application of MultDist by combining a bounded AST traversal and simple numerical computations with expression equality checks, implemented as checks for reference equality. Specifically, it first checks that the conclusion node r_0 is an Eq node whose children are **(1)** a Mul node with scaling factor denoted **(1.1)** and child denoted **(1.2)** and **(2)** an Add node. It then iterates over the children of nodes **(1.2)** and **(2)** in lockstep, checking that each child of node **(2)** is a Mul node with the same child as the corresponding Mul node in **(1.2)** and a scaling factor that is the product of **(1.1)** and the scaling factor of the same Mul node.

Farkas' Lemma Although Farkas' Lemma formalizes a somewhat subtle law of linear arithmetic, its application as a formal rule can be checked efficiently within ZKSMT's design. The instruction checks that **(1)** its conclusion operand is a node with operation Or whose children are negated inequalities, **(2)** its premise operand is a node with operation Eq whose children are a linear term matching the pattern given in the conclusion and a nonnegative constant, and **(3)** the sub-expressions of the linear term in the premise match the children of the inequalities in the disjuncts of the conclusion.

5 Zero-Knowledge Support

In this section we describe the technical details of ZKSMT's instantiation in ZK. Recall that the prover needs to demonstrate to the verifier that it knows a refutation proof of a formula without revealing the proof (or even the formula) to the verifier. We first explain how to commit ZKSMT's encoding of a refutation proof in Sec. 5.1. We discuss the details of how ZKSMT validates a committed refutation proof in

ZK in Sec. 5.2. Finally, in Sec. 5.3, we explain the checking instruction protocols that have some non-trivial design component for the ZK setting, continuing our focus on the theories covered in Sec. 4.

5.1 Refutation Proof Commitment

Recall that a refutation proof consists of a set of clauses and a sequence of proof steps. Both the clauses and proof steps can be committed as fixed-length vectors of integers. In detail, for a k-bit integer, we commit each bit individually (i.e., $\mathbb{F}_2^k$) and they can be converted to an extension binary field element (i.e., $\mathbb{F}_{2^k}$) for free thanks to the structure of the VOLE commitment [20]. Let I, $\mathcal{A}_{imm}$, and $\mathcal{A}_{ind}$ denote the set of all possible NodeID values, elements in ImmAddr, and elements in IndAddr, respectively. Given three injective functions $\varepsilon_I : I \to \mathbb{N}$, $\varepsilon_{\mathcal{A}_{imm}} : \mathcal{A}_{imm} \to \mathbb{N}_{>0}$, and $\varepsilon_{\mathcal{A}_{ind}} : \mathcal{A}_{ind} \to \mathbb{N}_{>0}$, an expression e specified by the tuple (NodeID, ImmAddr, IndAddr) can be mapped to the following vector of integers:

$$\{\varepsilon_I(\text{NodeID})\} \| \{\varepsilon_{\mathcal{A}_{imm}}(\text{ImmAddr})\} \| \{\varepsilon_{\mathcal{A}_{ind}}(\text{IndAddr})\}$$

Here, $\varepsilon_{\mathcal{A}_{imm}}$ and $\varepsilon_{\mathcal{A}_{ind}}$ are applied element-wise on the two respective lists. Given concrete encoding schemes ε_I, $\varepsilon_{\mathcal{A}_{imm}}$, and $\varepsilon_{\mathcal{A}_{imm}}$, an expression can be committed by committing its integer vector element-wise. These encoding schemes are made known by both the prover and the verifier.

An expression's NodeID should be kept private. Different NodeIDs take different numbers of operands. To avoid revealing an expression's NodeID from the size of its ImmAddr and IndAddr, we can pad both ImmAddr and IndAddr to the length α that is the upper bound of their size (Sec. 3.1).

Each proof step can be committed in a similar way. Recall that a proof step is encoded by four fields: StepID, RuleID, Result, and Premises which are either integers or lists of integers serving as pointers. The list Premises has its size bounded by μ. Hence, any proof step can be committed as a list of $\mu + 3$ integers.

5.2 Machine Execution in Zero Knowledge

We discuss how machine execution, i.e., the main checker, can be instantiated in ZK. Recall that the main checker performs three key operations:

1. **Fetching essential clauses and expressions.** To verify SMT proofs, we need to read entries from M_e and M_p using committed addresses. We can achieve this by instantiating M_e and M_p with any read-only memory (ROM) protocol [17, 20, 30] in ZK that is compatible with the commitment scheme we use.

2. **Guaranteeing the proof is acyclic.** The proof can be regarded as a DAG with proof steps ordered by their logical order. Proving a graph is a DAG reduces to proving magnitude relationships between pairs of committed integers.

3. **Invoking the corresponding checking instructions.** To ensure the privacy of a proof, the proof rule employed by each proof step should be kept private. This can be achieved generically by multiplexing all checks, but that incurs a high cost and leads to a large overhead. Instead, ZKSMT uses *group checking*, as we will explain next.

Group Checking ZKSMT groups the verification of the proof steps with the same proof rule, where the real checking instruction is the *only* checking instruction that will be called. There is no multiplexing, and no other checking instructions are executed. For instance, all proof steps employing the Resolution rule are verified consecutively, and only the checking instruction of Resolution is invoked on them.

The RuleID of a proof step, which identifies the specific step being validated within a particular checking group, is private to the prover. The StepID of every step that has been verified so far appears in D. The array D is append-only, and at the end of each proof step, the step's committed StepID is appended to it. D can be implemented using a standard array containing commitments when ZKSMT is instantiated in ZK.

The soundness of ZKSMT relies on the permutation checking between D and $\{0, 1, \ldots, \pi - 1\}$ (see Algorithm 1, line 17). Permutation checking ensures that every proof step is validated. When D contains committed values, the Schwartz-Zippel lemma allows for efficient permutation checking.

Remark 5.1 (Leakage and Optimization). By group checking, we reveal the number of applications of each proof rule in the input proof. On the other hand, grouping checking for identical proof rules over different premises and conclusions offers a chance for optimization by using a ZK protocol optimized for batch proofs (i.e., single instruction multiple data (SIMD) optimizations), such as [52].

5.3 Checking Instructions in Zero Knowledge

Some checking instructions for Boolean, EUF, and LIA rules consist of only reading operations over the expression table and comparisons, such as ExclMid (Sec. 2.2.1). All necessary ZK operations are already needed by the main checker, and the same operations suffice for handling these simple rules.

The instantiation of checking instructions becomes complex when the side condition of the proof rule involves traversing the IndAddr. In Sec. 4, we explain how these side conditions can be represented using the language of multisets. This level of abstraction further enables us to leverage the polynomial commitment scheme when instantiating these checking instructions in ZK. Next, we explain how to check two relations, subset and subset$_d$, between multisets using a polynomial commitment scheme. Following this, we will illustrate our implementations of the DeDup and Resolution checking instructions as examples.

Checking Multiset Relations To enable compact representation and efficient operations simultaneously, our protocol encodes multisets as polynomials over a finite field.

For the checking instructions we consider, we focus on two relations: subset and subset up to the number of occurrences (subset$_d$). The subset relation takes multiplicities into account. The multiset $\langle\!\langle A \rangle\!\rangle$ is a subset of the multiset $\langle\!\langle B \rangle\!\rangle$ if the multiplicities of all elements in $\langle\!\langle A \rangle\!\rangle$ are less than or equal to their multiplicities in $\langle\!\langle B \rangle\!\rangle$. On the other hand, $\langle\!\langle A \rangle\!\rangle$ is a subset$_d$ of $\langle\!\langle B \rangle\!\rangle$ if all distinct elements of $\langle\!\langle A \rangle\!\rangle$ also appear in $\langle\!\langle B \rangle\!\rangle$.

Checking the subset relation between two multisets is based on encoding multisets as univariate polynomials. Let Σ be a finite set, and $\mathbb{F}$ a finite field such that $|\mathbb{F}| > |\Sigma|$. Let $\langle\!\langle \Sigma^* \rangle\!\rangle$ be the set of all possible multisets over Σ. Given an injective function $\psi : \Sigma \to \mathbb{F}$, we define an encoding $\gamma_\psi : \langle\!\langle \Sigma^* \rangle\!\rangle \to \mathbb{F}[X]$ of a multiset as univariate polynomials over $\mathbb{F}$ such that for each multiset $\langle\!\langle \ell \rangle\!\rangle$, the images under ψ of the Σ-elements ℓ_i in $\langle\!\langle \ell \rangle\!\rangle$ are the roots of the image of $\langle\!\langle \ell \rangle\!\rangle$ under γ_ψ:

$$\gamma_\psi(\langle\!\langle \{\ell_0, \ldots, \ell_d\} \rangle\!\rangle) = (X - \psi(\ell_0)) \ldots (X - \psi(\ell_d))$$

To check the subset relation between two multisets $\langle\!\langle \ell^{\mathrm{sub}} \rangle\!\rangle$ and $\langle\!\langle \ell^{\mathrm{sup}} \rangle\!\rangle$, the prover commits their polynomial encodings, and the verifier checks that $\gamma_\psi(\langle\!\langle \ell^{\mathrm{sub}} \rangle\!\rangle)$ divides $\gamma_\psi(\langle\!\langle \ell^{\mathrm{sup}} \rangle\!\rangle)$ by attesting that $\gamma_\psi(\langle\!\langle \ell^{\mathrm{sub}} \rangle\!\rangle) \cdot W = \gamma_\psi(\langle\!\langle \ell^{\mathrm{sup}} \rangle\!\rangle)$. Here, W is a private polynomial committed by the prover as an extended witness. We use bivariate polynomials to verify the subset$_d$ relation between two multisets, leveraging an observation from [21]. Let $\bar{\ell}^{\mathrm{sub}}$, $\bar{\ell}^{\mathrm{sup}}$ and $\bar{\ell}$ be permuted versions of ℓ^{sub}, ℓ^{sup} and $\ell = \ell^{\mathrm{sub}} \uplus \ell^{\mathrm{sup}}$ respectively with the d' being the size of $\bar{\ell}^{\mathrm{sub}}$ and d being the size of $\bar{\ell}^{\mathrm{sup}}$. Given the same ψ we use for subset checking, define the following two polynomials:

$$\alpha_\psi(\langle\!\langle \bar{\ell}^{\mathrm{sub}} \rangle\!\rangle, \langle\!\langle \bar{\ell}^{\mathrm{sup}} \rangle\!\rangle) := (1 + X)^{d'} \cdot \Pi_{i=0}^{d'-1} (Y + \psi(\bar{\ell}_i^{\mathrm{sub}}))$$
$$\cdot \Pi_{i=0}^{d-2} (Y \cdot (1 + X) + \psi(\bar{\ell}_i^{\mathrm{sup}}) + X \cdot \psi(\bar{\ell}_{i+1}^{\mathrm{sup}}))$$
$$\beta_\psi(\langle\!\langle \bar{\ell} \rangle\!\rangle) := \Pi_{i=0}^{d'+d-1} ((1 + X) \cdot Y + \psi(\bar{\ell}_i) + \psi(\bar{\ell}_{i+1}) \cdot X)$$

It is proved that $\alpha_\psi(\langle\!\langle \bar{\ell}^{\mathrm{sub}} \rangle\!\rangle, \langle\!\langle \bar{\ell}^{\mathrm{sup}} \rangle\!\rangle)(X, Y)$ equals $\beta_\psi(\langle\!\langle \bar{\ell} \rangle\!\rangle)(X, Y)$ if and only if **(1)** $\langle\!\langle \bar{\ell}^{\mathrm{sup}} \rangle\!\rangle$ is a subset$_d$ of $\langle\!\langle \bar{\ell}^{\mathrm{sup}} \rangle\!\rangle$; and **(2)** $\bar{\ell}^{\mathrm{sub}}$, $\bar{\ell}^{\mathrm{sup}}$ and $\bar{\ell}$ are order-consistent[1]. A set of lists is order-consistent if values appear in the same order across all lists in the set. Putting it all together, to check if the subset$_d$ relation between $\langle\!\langle \ell^{\mathrm{sup}} \rangle\!\rangle$ and $\langle\!\langle \ell^{\mathrm{sub}} \rangle\!\rangle$ holds, the verifier attests the following relation between polynomials:

$$\alpha_\psi(\langle\!\langle \bar{\ell}^{\mathrm{sub}} \rangle\!\rangle, \langle\!\langle \bar{\ell}^{\mathrm{sup}} \rangle\!\rangle) = \beta_\psi(\langle\!\langle \bar{\ell} \rangle\!\rangle)$$
$$\gamma_\psi(\langle\!\langle \bar{\ell}^{\mathrm{sub}} \rangle\!\rangle) = \gamma_\psi(\langle\!\langle \ell^{\mathrm{sub}} \rangle\!\rangle)$$
$$\gamma_\psi(\langle\!\langle \bar{\ell}^{\mathrm{sup}} \rangle\!\rangle) = \gamma_\psi(\langle\!\langle \ell^{\mathrm{sup}} \rangle\!\rangle)$$
$$\gamma_\psi(\langle\!\langle \bar{\ell} \rangle\!\rangle) = \gamma_\psi(\langle\!\langle \ell^{\mathrm{sup}} \rangle\!\rangle) \cdot \gamma_\psi(\langle\!\langle \ell^{\mathrm{sub}} \rangle\!\rangle)$$

Here, the prover computes and commits $\bar{\ell}$, $\bar{\ell}^{\mathrm{sub}}$ and $\bar{\ell}^{\mathrm{sup}}$ using some proper order over Σ.

Resolution Recall that the side condition of Resolution on premise clauses $\bigvee A$, $\bigvee B$ and conclusion clause $\bigvee C$ is

---

[1]See Claim 3.1 [21] and its proof.

that $\langle\!\langle A\rangle\!\rangle \subseteq \langle\!\langle C\rangle\!\rangle \uplus \langle\!\langle p\rangle\!\rangle$ and $\langle\!\langle B\rangle\!\rangle \subseteq \langle\!\langle C\rangle\!\rangle \uplus \langle\!\langle \neg p\rangle\!\rangle$. Here, A, B, and C are lists of addresses of the expression table and p is an address. Given that the size of the expression table is bounded by χ, we can restrict the co-domain of ε_I to $\mathbb{N}_{\leq\chi}$, i.e., $\varepsilon_I : I \to \mathbb{N}_{\leq\chi}$. We further fix an injective function $\psi_I : \mathbb{N}_{\leq\chi} \to \mathbb{F}$ for a given proof. Then the checking instruction of the resolution rule can be implemented by verifying the subset relation between multisets $\varepsilon_I(\langle\!\langle A\rangle\!\rangle)$, $\varepsilon_I(\langle\!\langle C\rangle\!\rangle \uplus \langle\!\langle p\rangle\!\rangle)^2$ and between $\varepsilon_I(B)$ and $\varepsilon_I(\langle\!\langle C\rangle\!\rangle \uplus \langle\!\langle \neg p\rangle\!\rangle)$ using the approach mentioned above, with ψ concretized by ψ_I. By applying ε_I to the multisets, we mean element-wise application.

DeDup The side condition of DeDup asserts that for all $a \in \langle\!\langle A\rangle\!\rangle$ it holds that $a \in \langle\!\langle B\rangle\!\rangle$ given the premise clause $\bigvee A$ and the conclusion clause $\bigvee B$. This side condition can be validated by checking if $\langle\!\langle A\rangle\!\rangle$ is a subset$_d$ of $\langle\!\langle B\rangle\!\rangle$. Using the same encoding scheme as is used for the Resolution rule, we can implement the checking instruction of the DeDup rule by checking the subset$_d$ relation between $\varepsilon_I(\langle\!\langle A\rangle\!\rangle)$ and $\varepsilon_I(\langle\!\langle B\rangle\!\rangle)$. This relation checking can be achieved using the protocol we explain at the beginning of this section.

6 Implementation

We implement our protocol using the EMP-toolkit [50] for ZKP operations (circuits, polynomials, read-only memory access). We instantiated the arithmetic field as the extension field $\mathbb{F}_{2^{128}}$, under which field operations (and their ZK counterparts) can be efficiently implemented. The indices of proof steps are 32-bit integers, which support refutation proofs with more than one billion steps. In addition, as a performance optimization, we use an array M_a known as the *expression list table* to store argument lists for expressions that can take variable numbers of children. Expressions that take a fixed number of children (which is always 1 or 2 for the theories that we cover) store pointers to their children directly in M_e, but nodes that take variable numbers of children store a pointer to an entry in M_a that contains pointers to the expression's children. It allows us to keep the individual entries of M_e small and to avoid the cost of scanning a variable-length argument list for nodes like Not and Eq. We use η to denote the number of lists in M_a. This is not to be confused with α, which is the maximum length of an individual list within M_a.

7 Evaluation

We evaluate ZKSMT to compare our protocol with the prior state of the art. We intend to answer three key questions:

Q1 Does our protocol efficiently validate SMT formulas that formalize the safety and security of practical software?

Q2 Does our protocol scale well in response to increases in proof size?

Q3 Is ZKSMT more efficient than a zkVM running a commodity SMT proof validator?

The results of our experiments allow us to report an affirmative answer for all three questions. For all benchmarks, we ran ZKSMT on AWS instances of type `r5b.4xlarge` with 128 GB of memory, 16 vCPUs, and a 10 Gbps network connection between the prover and verifier. However, the underlying ZK protocols that we use only consume about 100 Mbps bandwidth. We also configured ZKSMT to use 8 threads. Our methodology and results for Q1, Q2, and Q3 appear in Sec. 7.1, Sec. 7.2, and Sec. 7.3, respectively.

7.1 Verifying Practical Software

To answer Q1, we collected a set of SMT formulas whose validity formalizes program correctness. Specifically, the SMT formulas were generated by the Boogie verification toolchain [5]. The Boogie toolchain contains an intermediate language for expressing low-level programs annotated with function requirements and guarantees, along with compilation passes to an intermediate language from high-level languages including C, Spec#, and Dafny [39]. Boogie generates verification conditions from the annotated intermediate programs in the SMT-LIB 2.0 format, which can be validated by SMT solvers like Z3 [16]. We ran Boogie on its test suite to collect the corresponding SMT formulas, and we validated the SMT formulas using the solver SMTInterpol [14, 32] to generate proof certificates that ZKSMT can process.

Fig. 2 shows the runtime of ZKSMT versus the number of proof steps used in each of the SMT statements in the Boogie test suite. ZKSMT is able to verify most of the test suite SMT statements in ZK within a few seconds, but the largest benchmark takes 39 seconds. We also observe a general linear trend between the running time and the number of steps, which is expected. The fluctuation is due to the use of different rules in each instance, since some rules are more costly than others.

7.2 Scalability

To determine how our protocol scales in response to increases in proof size (Q2), we microbenchmark various aspects of ZKSMT while varying the size of input SMT statements.

Proof Breakdown To assess the relative time consumption of different parts of our protocol, we run three of our Boogie benchmarks and separate the timing results into three phases: type checking, Resolution, and all other proof rules. We place Resolution in a phase of its own because, for each of the examples, it takes more time than checking all of the other rules combined. Fig. 3 shows the performance decomposition for the three benchmarks. All of them are related to program safety verification: Lock is a Boogie benchmark for verification of a lock, Houdini is a benchmark on modular contract checking [37], and McCarthy is an adaptation of a standard benchmark for verification of recursive functions [43].

---

[2]When verifying equivalences that involve the combination of multisets through union operations, we compute the product of the two corresponding polynomials.
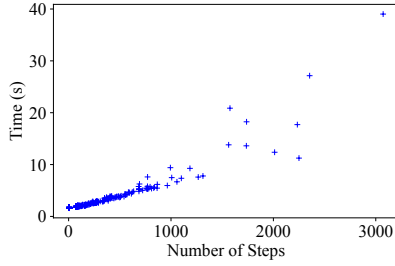
Figure 2: Step counts and time costs for validating SMT statements from the Boogie test suite in ZKSMT.
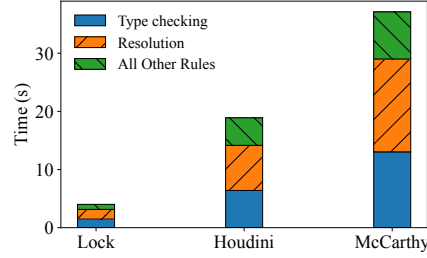

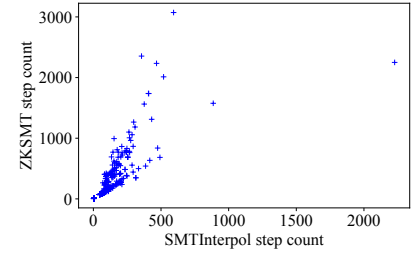Figure 3: Time cost decomposition for Lock, Houdini, and McCarthy benchmarks.

Figure 4: Relative step counts of ZKSMT's proofs and the original SMTInterpol proofs.

We observe that type checking can be as time-consuming as the main checking loop itself. This is due to the fact that ZKSMT needs to fetch every entry of every list in M_a at least once to confirm that its type fits with the list's type.

Max List Size To understand how α, the maximum list size, affects the running time of our rules, we benchmark the running time of our individual proof rules in isolation. To find the amortized cost of each rule, we run the rule 1,000 times and average the result. Most of our rules are simple, so we present the results for only our four most performance-intensive rules: Resolution, Consolidate, Farkas, and Flatten.

The results of varying α with values ranging from 10 to 50 are presented in Fig. 5a. We ran a linear regression and determined that all four rules scale linearly, with R^2 values above 0.99. This occurs because all four rules contain loops or procedures which iterate $O(\alpha)$ times.

Many of ZKSMT's rules are affected by the size of the longest list in the proof because all argument lists are padded to be the same size. The worst-case scenario for ZKSMT would be a proof that operates mainly on short lists but contains one extra-long list that forces all list-traversing rules to perform a large number of iterations. Fortunately, our benchmarks demonstrate that this degenerate case does not appear in practice. In the future we plan to mitigate the effect of α on a proof's overall running time by breaking down list-based rules into smaller pieces, which will improve runtime even more by eliminating the impact of large maximum list sizes.

Table Size Next, we consider how χ, the size of M_e, affects the running time of our four main rules. For each trial, we ran 10,000 instances of a rule with an α value of 10, which was a common value among our benchmarks, an η value of 1,000, and a π value of 10 while varying the value of χ to between 1,000 and 4,000. The results are plotted in Fig. 5b. Unlike our results for α, the running time does not change appreciably. This is because the main operation in these rules that is affected by a change in table size is the cost of accessing an element from ROM, for which the amortized access time does not depend on the number of elements. For similar reasons, changing the value of η does not change the

running time significantly.

Rule Breakdown We also consider which operations make up the running time of our four main rules. With α values of 5 and 21, we divide the running times for each rule into the time taken for memory operations (retrieving entries from M_e or M_a) and the time taken for arithmetic operations (everything else). 5 is a small but still realistic value for α, and 21 is the highest value of α that appears in our Boogie benchmarks. The results appear in Fig. 5c. Arithmetic operations dominate the running time for Resolution, Consolidate, and Flatten, but memory operations dominate the running time for Farkas. This makes sense because, unlike the other three rules, Farkas does not perform any multiset equivalence or containment checks. Multiset checks can work directly with expressions' addresses, but Farkas needs to fetch every entry in its premise and conclusion to pattern-match their NodeIDs and arguments.

Original Proof Size Our work uses a compiler to convert the output of SMTInterpol to the format accepted by ZKSMT. To enable evaluation in zero knowledge, some rules in SMTInterpol, particularly the LIA rules, must be broken down into simpler rules. This increases the proof size. A comparison between the number of proof steps in SMTInterpol and ZKSMT is given in Fig. 4 for the Boogie test suite. The proof size increases by a factor from 1 to 7, which is not problematic because ZKSMT is still vastly more efficient than the generic zkVM solution (Sec. 7.3).

Stress Test To stress test ZKSMT, we ran it against a series of larger tests from the Wisconsin Safety Analyzer [2] benchmark suite found in the official SMT-LIB benchmarks repository [1]. The benchmarks from the Wisconsin Safety Analyzer represent correctness and security properties for commercial off-the-shelf software. The resulting running times are plotted in Fig. 6a. The largest test which passed uses 200K steps, 380K expressions, and a maximum list size α of 97. This verified in about 3 hours, requiring more than 22.9 billion $\mathbb{F}_2$ multiplications and 336 million $\mathbb{F}_{2^{128}}$ multiplications. Larger tests ran out of memory. This demonstrates that ZKSMT can scale up to proofs of a larger size, and gives insight into ZKSMT's current limitations.
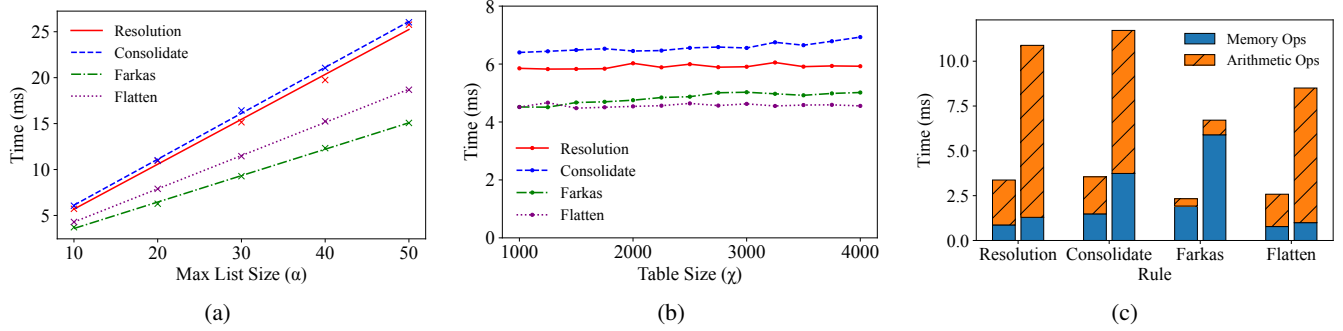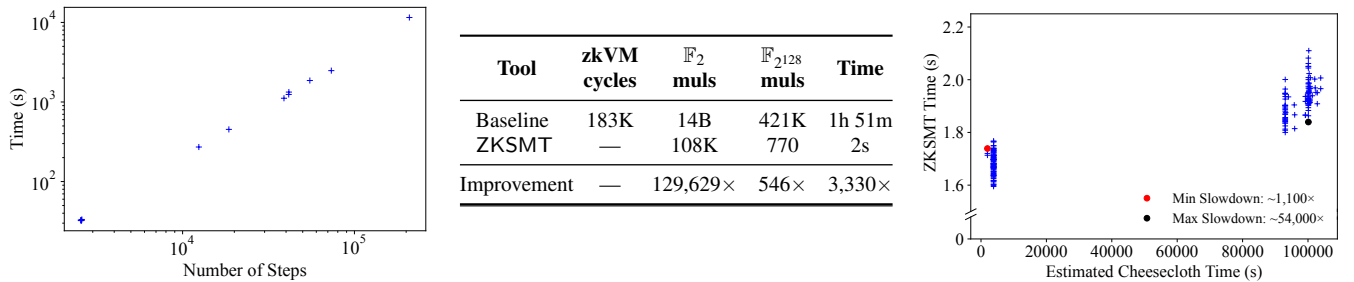
Figure 5: Scalability and rule breakdown of our protocol. Fig. 5a and Fig. 5b contain the running times of a single proof step across different rules for changing values of max list size α and expression table size χ, respectively. Fig. 5c shows the time cost decomposition across the four rules with max list size $\alpha = 5, 21$.

Tool	zkVM cycles	$\mathbb{F}_2$ muls	$\mathbb{F}_{2^{128}}$ muls	Time
Baseline	183K	14B	421K	1h 51m
ZKSMT	—	108K	770	2s
Improvement	—	$129{,}629\times$	$546\times$	$3{,}330\times$

(a) Step counts and time costs for validating SMT statements from the Wisconsin Safety Analyzer, plotted on a logarithmic scale.

(b) Comparison of ZKSMT's performance against Cheesecloth and Diet Mac'n'cheese on the shortest Boogie benchmark.

(c) Cheesecloth's estimated running times for our Boogie benchmarks, compared to the actual running times of ZKSMT.

Figure 6: More experimental results.

7.3 Comparison with Alternative Protocols

Instead of developing a custom ZK protocol to validate SMT formulas, a simpler approach would be to take a commodity SMT proof validator and compile it to a ZK statement using a ZK virtual machine (zkVM). We benchmark the performance of ZKSMT against such a zkVM to determine whether the benefits of a custom ZK protocol are worth the effort (Q3).

In our evaluation, we used Cheesecloth [15] and Diet Mac'n'cheese [6, 22] as the baseline zkVM. Cheesecloth is a general-purpose tool for generating zero-knowledge proof statements that verify the execution of LLVM programs. Diet Mac'n'cheese is an interactive VOLE-based zero-knowledge proof backend, capable of verifying ZK statements.

We developed a cleartext C++ version of ZKSMT that verifies SMT statements, and we used Cheesecloth and Diet Mac'n'cheese to verify the shortest Boogie benchmark (with only 6 steps) in ZK. The results in Fig. 6b demonstrate that ZKSMT is significantly faster than the baseline, taking seconds instead of hours to verify it. With a $3{,}330\times$ reduction in runtime, it is clear that ZKSMT provides a significant improvement over the zkVM approach in enabling SMT validation for program verification in ZK.

For a fuller comparison, we ran Cheesecloth to output the number of zkVM cycles required for 190 of our 440 Boo-

gie benchmarks. Cheesecloth failed to process 28 of the benchmarks, and the remaining 222 benchmarks timed out after 5,000,000 zkVM cycles. From the zkVM cycle count, we estimate the ZK running time of Cheesecloth and Diet Mac'n'cheese, which scales linearly relative to cycle count[3]. Fig. 6c compares this estimated running time with ZKSMT's. At best, Cheesecloth's estimated running time is approximately 1,100 times slower than ZKSMT's running time for the same benchmark. At worst, Cheesecloth's estimated running time is more than 54,000 times slower. We highlight the two extremes in the figure.

The wide gap in running time improvements is due to the fact that Cheesecloth can handle only the smaller benchmarks in our suite. Because we focus on small benchmarks, ZKSMT's running times in Figure 6c are dominated by the fixed startup time of EMP-toolkit, which ranges from 1.5 to 1.7 seconds on our evaluation machine. On the largest benchmarks, this startup time is less than 0.1 percent of the overall running time. Cheesecloth does not have a comparable startup cost, but it is significantly slower than ZKSMT in general. We estimate that, if Cheesecloth could run on larger benchmarks, ZKSMT's speedup rate would stabilize at

---

[3]A single zkVM cycle corresponds to approximately 0.021 seconds of running time.

an even larger number. Still, ZKSMT is much faster than a conservative estimation of Cheesecloth's performance.

8   Related Work

Vulnerability Proofs Prior research on ZK proofs has focused on protocols for proving the existence of bugs and vulnerabilities in programs [15, 27, 31]. With ZKSMT, we work towards the opposite goal of proving that a program is free of bugs and vulnerabilities. Specifically, ZKSMT efficiently proves the unsatisfiability of SMT formulas, which is the first step in a pipeline for safety proofs about programs.

Safety Proofs ZKSMT is not the first ZK protocol to encode refutation proofs. ZKUNSAT [42] is a ZK protocol for validating proofs in propositional logic. Among existing ZK protocols, ZKUNSAT is the most similar to ZKSMT, but ZKSMT is more versatile than ZKUNSAT. ZKUNSAT requires SAT formulas to be in conjunctive normal form, while ZKSMT supports arbitrary AST structures in its more expressive SMT formulas. Also, ZKUNSAT supports only one proof rule, namely resolution. In contrast, ZKSMT's implementation supports dozens of distinct proof rules, and the protocol itself generalizes to any suite of first-order proof rules.

A concurrent work that addresses a similar problem is zkPi [38]. Unlike ZKSMT and ZKUNSAT, zkPi encodes proofs written in interactive theorem provers, primarily Lean. The proofs encoded by zkPi can contain algebraic data types, lambda calculus terms, and induction. ZKSMT differs from zkPi in that it prioritizes efficiency for large-scale first-order proofs that encode formalisms like integer arithmetic directly rather than as part of a broader framework.

General-Purpose ZK Protocols General-purpose ZK protocols emulate program executions in ZK. These protocols are much less efficient than ZKSMT since they must support arbitrary computations while obliviously concealing a program's control flow. When Cheesecloth [15] and TinyRAM [7] execute a program in ZK, they multiplex over all CPU instructions and simulate read, write, malloc, and free operations for the program's memory. Likewise, Pantry [11] and Buffet [49] model programs with memory operations and mutable states. ZKSMT does not incur the same performance costs as Cheesecloth, Pantry, and Buffet because it efficiently checks SMT rules instead of arbitrary computations.

9   Conclusion

This paper introduces ZKSMT, an efficient protocol for validating SMT formulas in ZK. This work sets up exciting future work in multiple directions. First, protocols can be developed for other theories that model practical verification problems but are not currently supported, including the theory of arrays and the theory of bit-vectors [23]. Arrays and bit-vectors are commonly used by symbolic execution engines that execute low-level code [12]. Second, the core logic itself can be

extended to validate formulas that contain universal and existential quantifiers. Prominent program verification toolchains often produce quantified formulas as output [5, 39].

Another direction for further research is the combination of ZKSMT's proofs with other proofs about program safety. ZKSMT is a protocol for validating proofs about programs written in a high-level language. For a full ZK program verification pipeline, the prover would also need to demonstrate that a compiled and distributed binary version of the program is observationally equivalent to the high-level program whose correctness proof is covered by ZKSMT. The construction and representation of such proofs is the subject of translation validation [45, 47] and verified compilers [41]. Validating these proofs in ZK is an exciting direction for future work.

Acknowledgments

References

[1] SMT-LIB: The Satisfiability Modulo Theories Library - Benchmarks. http://smtlib.cs.uiowa.edu/benchmarks.shtml.

[2] WiSA: Wisconsin Safety Analyzer. https://research.cs.wisc.edu/wisa/.

[3] FIPS PUB 140-2. Security requirements for cryptographic modules. *NIST*, 2001.

[4] Haniel Barbosa, Clark Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, et al. cvc5: A versatile and industrial-strength smt solver. In *TACAS*. Springer, 2022.

[5] Mike Barnett, Bor-Yuh Evan Chang, Robert DeLine, Bart Jacobs, and K Rustan M Leino. Boogie: A modular reusable verifier for object-oriented programs. In *FMCO*. Springer, 2005.

[6] Carsten Baum, Alex J. Malozemoff, Marc B. Rosen, and Peter Scholl. Mac'n'cheese: Zero-knowledge proofs for boolean and arithmetic circuits with nested disjunctions. In *Crypto*, 2021.

[7] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Crypto*, 2013.

[8] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *USENIX Security*, 2014.

[9] Karthikeyan Bhargavan, Barry Bond, Antoine Delignat-Lavaud, Cédric Fournet, Chris Hawblitzel, Catalin Hritcu, Samin Ishtiaq, Markulf Kohlweiss, K. Rustan M. Leino, Jay R. Lorch, Kenji Maillard, Jianyang Pan, Bryan Parno, Jonathan Protzenko, Tahina Ramananandro, Ashay Rane, Aseem Rastogi, Nikhil Swamy, Laure Thompson, Peng Wang, Santiago Zanella Béguelin, and Jean Karim Zinzindohoue. Everest: Towards a verified, drop-in replacement of HTTPS. In *SNAPL*, 2017.

[10] Manuel Blum, William S. Evans, Peter Gemmell, Sampath Kannan, and Moni Naor. Checking the correctness of memories. In *32nd FOCS* , 1991.

[11] Benjamin Braun, Ariel J Feldman, Zuocheng Ren, Srinath Setty, Andrew J Blumberg, and Michael Walfish. Verifying computations with state. In *ACM SOSP*, 2013.

[12] Cristian Cadar, Daniel Dunbar, Dawson R Engler, et al. Klee: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *OSDI*, 2008.

[13] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *34th ACM STOC*, 2002.

[14] Jürgen Christ, Jochen Hoenicke, and Alexander Nutz. Smtinterpol: An interpolating smt solver. In *International SPIN Workshop on Model Checking of Software*, 2012.

[15] Santiago Cuéllar, Bill Harris, James Parker, Stuart Pernsteiner, and Eran Tromer. Cheesecloth: Zero-Knowledge proofs of real world vulnerabilities. In *USENIX Security*, 2023.

[16] Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *TACAS*, 2008.

[17] Cyprien Delpech de Saint Guilhem, Emmanuela Orsini, Titouan Tanguy, and Michiel Verbauwhede. Efficient proof of ram programs from any public-coin zero-knowledge system. In *SCN*, 2022.

[18] Samuel Dittmer, Yuval Ishai, and Rafail Ostrovsky. Line-Point Zero Knowledge and Its Applications. In *ITC*, 2021.

[19] Burak Ekici, Alain Mebsout, Cesare Tinelli, Chantal Keller, Guy Katz, Andrew Reynolds, and Clark Barrett. Smtcoq: A plug-in for integrating smt solvers into coq. In *CAV*, 2017.

[20] Nicholas Franzese, Jonathan Katz, Steve Lu, Rafail Ostrovsky, Xiao Wang, and Chenkai Weng. Constant-overhead zero-knowledge for RAM programs. In *ACM CCS*, 2021.

[21] Ariel Gabizon and Zachary J Williamson. plookup: A simplified polynomial protocol for lookup tables. *Cryptology ePrint Archive*, 2020.

[22] Galois, Inc. swanky: A suite of rust libraries for secure computation. https://github.com/GaloisInc/swanky, 2019.

[23] Vijay Ganesh and David L Dill. A decision procedure for bit-vectors and arrays. In *CAV*, 2007.

[24] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3), July 1991.

[25] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In *40th ACM STOC*, 2008.

[26] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, 1985.

[27] Matthew Green, Mathias Hall-Andersen, Eric Hennenfent, Gabriel Kaptchuk, Benjamin Perez, and Gijs Van Laer. Efficient proofs of software exploitability for real-world processors. *PETS*, 2023.

[28] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In *Asiacrypt*, 2010.

[29] Ronghui Gu, Zhong Shao, Hao Chen, Xiongnan (Newman) Wu, Jieung Kim, Vilhelm Sjöberg, and David Costanzo. CertiKOS: An extensible architecture for building certified concurrent os kernels. In *USENIX OSDI*, 2016.

[30] David Heath and Vladimir Kolesnikov. A 2.1 KHz zero-knowledge processor with BubbleRAM. In *ACM CCS*, 2020.

[31] David Heath, Yibin Yang, David Devecsery, and Vladimir Kolesnikov. Zero knowledge for everything and everyone: Fast ZK processor with cached ORAM for ANSI C programs. In *IEEE S&P*, 2021.

[32] Jochen Hoenicke and Tanja Schindler. A simple proof format for smt. In *International Workshop on SMT*, 2022.

[33] Zhangxiang Hu, Payman Mohassel, and Mike Rosulek. Efficient zero-knowledge proofs of non-algebraic statements with sublinear amortized cost. In *Crypto*, 2015.

[34] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *ACM STOC*, 2007.

[35] Marek Jawurek, Florian Kerschbaum, and Claudio Orlandi. Zero-knowledge using garbled circuits: how to prove non-algebraic statements efficiently. In *ACM CCS*, 2013.

[36] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David A. Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. seL4: formal verification of an os kernel. In *ACM SOSP*, 2009.

[37] Shuvendu Lahiri and Julien Vanegue. Explain houdini: Making houdini inference transparent. In *VMCAI*, 2010.

[38] Evan Laufer, Alex Ozdemir, and Dan Boneh. zkpi: Proving lean theorems in zero-knowledge. Cryptology ePrint Archive, Paper 2024/267, 2024. https://eprint.iacr.org/2024/267.

[39] K Rustan M Leino. Dafny: An automatic program verifier for functional correctness. In *LPRA*, 2010.

[40] Xavier Leroy. Formal verification of a realistic compiler. *Communication of ACM*, 2009.

[41] Xavier Leroy, Sandrine Blazy, Daniel Kästner, Bernhard Schommer, Markus Pister, and Christian Ferdinand. Compcert-a formally verified optimizing compiler. In *ERTS*, 2016.

[42] Ning Luo, Timos Antonopoulos, William R. Harris, Ruzica Piskac, Eran Tromer, and Xiao Wang. Proving UNSAT in zero knowledge. In *ACM CCS*, 2022.

[43] Zohar Manna and John McCarthy. *Properties of programs and partial function logic*. Stanford University, 1969.

[44] Payman Mohassel, Mike Rosulek, and Alessandra Scafuro. Sublinear zero-knowledge arguments for RAM programs. In *Eurocrypt*, 2017.

[45] George C Necula. Translation validation for an optimizing compiler. In *ACM SIGPLAN*, 2000.

[46] Duckki Oe, Andrew Reynolds, and Aaron Stump. Fast and flexible proof checking for smt. In *Proceedings of the 7th International Workshop on SMT*, 2009.

[47] Amir Pnueli, Michael Siegel, and Eli Singerman. Translation validation. In *TACAS*, 1998.

[48] Srinath Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In *Crypto*, 2020.

[49] Riad S. Wahby, Srinath T. V. Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *NDSS*, 2015.

[50] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. Emp-toolkit: Efficient multi-party computation toolkit. https://github.com/emp-toolkit, 2016.

[51] Chenkai Weng, Kang Yang, Jonathan Katz, and Xiao Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In *IEEE S&P*, 2021.

[52] Chenkai Weng, Kang Yang, Zhaomin Yang, Xiang Xie, and Xiao Wang. AntMan: Interactive zero-knowledge proofs with sublinear communication. In *ACM CCS*, 2022.

[53] Kang Yang, Pratik Sarkar, Chenkai Weng, and Xiao Wang. QuickSilver: Efficient and affordable zero-knowledge proofs for circuits and polynomials over any field. In *ACM CCS*, 2021.

A Proof Rule Tables

Tables 4 and 5 show our full set of proof rules for Boolean logic, EUF, and LIA. Table 4 contains the simple rules that have no premises or side conditions, and Table 5 contains the more complex rules.

B Proofs

We now prove that ZKSMT is sound and complete. The proofs are independent of the theories on which ZKSMT is instantiated. A proof Π of a formula φ exists in ZKSMT's format under a theory T if and only if a corresponding proof Π' exists for φ in T outside ZKSMT.

B.1 Proof of VM Soundness

Let Π be a proof in ZKSMT's format that derives φ. Assume that Algorithm 1 succeeds for Π. Our goal is to convert Π into a derivation tree Π' that derives φ. Π' needs to follow the structure defined in Section 2.1.

Let T be the theory used for Π. We can use the exact same theory for Π', so the logical structure of the proof can stay the same. We can simply translate all of the proof steps from Π verbatim to create a proof tree for Π'. No new proof steps

RuleID	Conclusion
Boolean	
TruePos	$\bigvee\{\text{True}\}$
FalseNeg	$\bigvee\{\neg\text{False}\}$
ExclMid	$\bigvee\{\neg a, a\}$
ImplPos1	$\bigvee\{a \rightarrow b, a\}$
ImplPos2	$\bigvee\{a \rightarrow b, \neg b\}$
ImplNeg	$\bigvee\{\neg(a \rightarrow b), \neg a, b\}$
EquivPos1	$\bigvee\{a = b, a, b\}$
EquivPos2	$\bigvee\{a = b, \neg a, \neg b\}$
EquivNeg1	$\bigvee\{\neg(a = b), a, \neg b\}$
EquivNeg2	$\bigvee\{\neg(a = b), \neg a, b\}$
EUF	
Refl	$\bigvee\{a = a\}$
Symm	$\bigvee\{a = b, \neg(b = a)\}$
Trans	$\bigvee\{a = c, \neg(a = b), \neg(b = c)\}$
LIA	
Total	$\bigvee\{a \leq b, b < a\}$
Trichotomy	$\bigvee\{a < b, a = b, b < a\}$
AddSingle	$\sum\{a\} = a$
MulSingle	$1 * a = a$
MulDist	$c * (\sum_{i=0}^{n} d_i * x_i) = \sum_{i=0}^{n} cd_i * x_i$

Table 4: ZKSMT's rules that have no premises or side conditions, grouped by theory.

need to be created for Π', but some steps from Π may need to be duplicated. In ZKSMT's format, the conclusion of a proof step can be used as a premise for any number of other proof steps. However, in ordinary proof trees, each step's conclusion can be used only once. If a proof step's conclusion is ever used as a premise for multiple other steps in Π, we can duplicate that proof step in Π' for every place where it is used.

If a proof step is duplicated for Π', then all of its premises need to be duplicated as well. This can lead to a cascading effect where duplicating a proof step requires even more steps to be duplicated as a consequence, but the process is guaranteed to terminate eventually. There are only finitely many steps in Π, each original step takes only finitely many premises, and there are no cycles. An infinite chain of expansions would violate one of those three properties of Π, so we can always find a finite value for π to confirm that φ is boundedly verifiable.

The expression table M_e in Π stores terms and formulas together, but we need to distinguish between terms and formulas for Π'. Recall that ZKSMT's format includes typing constraints (Sec. 3.2). There is a type for formulas, and it does not overlap with the types used for terms in ZKSMT's format, so we can always distinguish between terms and formulas when creating Π'. In Π, the premises and conclusion of every proof step have the type of formulas. We can interpret all expressions of that type as formulas for Π' and treat everything else as a term. We can unfold every term and formula in M_e into a full AST because we forbid M_e to contain cycles. All terms and formulas must be well-formed because they satisfy

the typing rules in Π that are based on T.

One minor issue for the conversion from Π to Π' is the fact that functions in Π can behave as predicates. We need to distinguish between functions and predicates for Π', so we can treat functions from Π as either functions or predicates in Π' depending on whether their output type is the type of formulas. Each function in Π has only one permitted output type, so there can be no ambiguity in the conversion.

Every proof step in Π is checked exactly once, so each step in Π' must be valid according to T. We know that every step in Π is checked once because of the permutation check on line 17 of Algorithm 1. A proof tree is valid if every individual step in the tree is valid, and every step in Π' is a copy of a step in Π, so Π' must be a valid derivation of φ.

All that remains to be shown is that we can fix size parameters for Π'. We already covered π. χ can be the same as it is for Π because no new expressions have been created, only copies of existing ones. μ can be the same as it is for Π because every proof step in Π' takes as many premises as it does in Π. α can be the same as it is for Π because no argument lists have been changed. ρ can be the same as it is for Π because every rule has the same premises in Π' as it does in Π. This covers all five size parameters, so φ is boundedly verifiable, and ZKSMT is sound.

B.2 Proof of VM Completeness

For the reverse direction, we will start with a valid derivation tree Π' and produce a valid proof Π in ZKSMT's format. To convert Π' into Π, we can perform some of the same conversions that we performed for the soundness proof in reverse. We can let Π use the same theory T that Π' uses. Every proof step in Π' can become a valid proof step in Π without being altered. Likewise, every term and formula in Π' that is well-typed becomes a well-typed expression in Π. It is possible that Π' contains groups of identical proof steps, but there is no need to consolidate identical proof steps for Π.

We can construct the expression table M_e by giving every distinct AST node in Π' its own entry. It is important that M_e does not contain any duplicates. Whenever two terms or formulas are identical, they should have the same entry in M_e. This is what allows equality checks to work in instances of ZKSMT. The order of the entries in M_e is unimportant. We can add a type for formulas and preserve all of the distinctions between types for terms. We can treat predicates as formula-typed functions.

To finish the construction of Π, we need to define the size parameters π, χ, μ, α, and ρ. Let π be the number of proof steps in Π'. We copied the structure of Π' without adding or removing any steps, so π is also the number of proof steps in Π. Let χ be the number of distinct AST nodes in the terms and formulas in Π'. This is also the number of expressions in Π because of our definition of M_e. Let μ be the maximum number of premises that any proof step in Π' takes. There must be a maximum because Π' is a finite proof, and this is

RuleID	Side Condition	Premises	Conclusion
Boolean			
Resolution	$\exists p. p \in \langle\!\langle A \rangle\!\rangle, \neg p \in \langle\!\langle B \rangle\!\rangle,$ $\langle\!\langle A \rangle\!\rangle \subseteq \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle p \rangle\!\rangle, \langle\!\langle B \rangle\!\rangle \subseteq \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle \neg p \rangle\!\rangle$	$\bigvee A, \bigvee B$	$\bigvee C$
DeDup	$\forall a \in \langle\!\langle A \rangle\!\rangle.\, a \in \langle\!\langle B \rangle\!\rangle$	$\bigvee A$	$\bigvee B$
OrNil		$\bigvee \{\}$	False
OrSingle		a	$\bigvee \{a\}$
OrSingleRev		$\bigvee \{a\}$	a
AndPos	$\exists A, B. \langle\!\langle \bigwedge A \rangle\!\rangle \uplus \langle\!\langle B \rangle\!\rangle = \langle\!\langle C \rangle\!\rangle, \bigwedge A = \bigwedge_{i=0}^{n} a_i, \bigvee B = \bigvee_{i=0}^{n} \neg a_i$		$\bigvee C$
AndNeg	$a \in \langle\!\langle A \rangle\!\rangle$		$\bigvee \{\neg \bigwedge A, a\}$
OrPos	$a \in \langle\!\langle A \rangle\!\rangle$		$\bigvee \{\bigvee A, \neg a\}$
OrNeg	$\exists A. \langle\!\langle \neg \bigvee A \rangle\!\rangle \uplus \langle\!\langle A \rangle\!\rangle = \langle\!\langle B \rangle\!\rangle$		$\bigvee B$
EUF			
Congruence	$\exists A, B, f. (fA = fB) \in C, |A| = |B|,$ $\forall i \in \{0, \ldots, |A| - 1\}. \neg(A_i = B_i) \in C$		$\bigvee C$
LIA			
TotalInt	$i_0 = m * \mathtt{ONE}, i_1 = (m+1) * \mathtt{ONE}$		$\bigvee \{a \leq i_0, i_1 \leq a\}$
Consolidate	$\exists a, A_a, B_a, C.\ \langle\!\langle A_a \rangle\!\rangle \uplus \langle\!\langle C \rangle\!\rangle = \langle\!\langle A \rangle\!\rangle, \langle\!\langle B_a \rangle\!\rangle \uplus \langle\!\langle C \rangle\!\rangle = \langle\!\langle B \rangle\!\rangle,$ $A_a = \{\alpha_0 * a, \ldots, \alpha_{t-1} * a\}, B_a = \{\beta_0 * a, \ldots, \beta_{t'-1} * a\},$ $\alpha_0 + \cdots + \alpha_{t-1} = \beta_0 + \cdots + \beta_{t'-1}$		$\sum A = \sum B$
Flatten	$\exists \langle\!\langle C \rangle\!\rangle, \langle\!\langle D \rangle\!\rangle.\ \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle \sum D \rangle\!\rangle = \langle\!\langle A \rangle\!\rangle, \langle\!\langle C \rangle\!\rangle \uplus \langle\!\langle D \rangle\!\rangle = \langle\!\langle B \rangle\!\rangle$		$\sum A = \sum B$
Farkas	$\forall i \in \{0, \ldots, n\}.\, m_i \geq 0$ either $c > 0$, or $c = 0$ and $\exists j. \leq_j$ is $<$	$\sum_{i=0}^{n}(m_i * a_i) +$ $(-m_i * b_i) = c$	$\bigvee_{i=0}^{n} \{\neg(a_i \leq_i b_i)\}$

Table 5: ZKSMT's rules that have premises or side conditions, grouped by theory. Capital letters represent argument lists for n-ary operations, and lowercase letters represent individual expressions.

also the maximum number of premises for Π because we did not modify the proof steps. Let α be the maximum number of arguments taken by any AST node in Π'. Again, this maximum must exist because Π' is finite, and it applies equally well to Π. Lastly, let ρ be the number of distinct proof rules used in Π'. Π uses the same suite of proof rules as Π', so ρ is also the number of distinct proof rules in Π.

We need to confirm that Π passes the checks in Algorithm 1. The rule check on line 12 of Algorithm 1 will always pass because Π and Π' use the same theory. The cycle checks in Algorithm 1 must pass because Π' does not contain any cyclic terms, formulas, or proof steps. The type checking must pass because every term and formula in Π' is well-formed. The permutation check on line 17 must succeed because π is defined as the number of steps in Π' and every step from Π' is copied into Π. There are no other checks that can cause Algorithm 1 to fail, so Π must be a valid instance of ZKSMT. Therefore, ZKSMT is complete.