# Proving UNSAT in Zero Knowledge

https://eprint.iacr.org/2022/206

In 2022, the White House issued two reports on the security of open-source software and supply chain, respectively: both reports put forth an outline to provide a nationwide high-resilience cyberinfrastructure. Computer program security is among the most pressing issues and the gold standard for achieving it is formal verification. This is typically achieved by proving the unsatisfiability (UNSAT) of a formula based on the program and the desired property.

Such proofs of security, however, may themselves reveal sensitive methods or intellectual property, especially when applied to critical system components. Verification techniques that require access to sensitive materials can be problematic because the entities typically tasked with performing the verification, such as centralized curated marketplaces, may not be trusted fully by either software owners or software users. This issue has gained significant attention and prompted multiple enforcement actions and legislative proposals in the EU and the US.

This paper initiates a new line of research to facilitate proving the safety properties of privileged and open-source software on critical cyberinfrastructure. The new verification protocol enables developers to prove the safety of their software to users or even the public without revealing the implementation details. This secure, private solution eliminates concerns around third-party involvement in the verification process.

The core idea of this paper is to apply zero-knowledge proofs (ZKP) to the verification of logical formulas that express program safety properties. ZKP is a cryptographic primitive that allows one party (the prover) to prove a claim to a second party (the verifier) without revealing the secret evidence for the claim. ZKP eliminates the need to share source code during verification and thus reduces the attack surface of the verification process.

To make this approach practical, the paper introduces an efficient method for encoding resolution proofs. Resolution proofs are certificates of unsatisfiability for Boolean propositions. A resolution proof consists of a sequence of clauses derived from the input formula, with a contradiction being derived at the end of the resolution as evidence of the formula's unsatisfiability. The authors observed that resolution proofs typically contain clauses with only a small number of literals. Consequently, each clause can be encoded efficiently as a low-degree polynomial over a finite field. This encoding allows for the implementation of a fast ZKP protocol for verifying the validity of the clause derivations. The protocol verifies certain relations between the polynomials while preserving the confidentiality of the refutation proof and input formula.

The authors also implemented the proposed cryptographic system and evaluated its performance using safety property formulas generated from real-world programs that are similar to ones needed by typical cyberinfrastructures. The results indicate that the proposed system could be made readily deployable to offer high resilience in cyberinfrastructures.

This paper received the Distinguished Paper Award at the prestigious ACM CCS 2022 conference. Followup work is underway, extending the expressive power and efficiency of the approach to make it practically applicable to additional systems, languages and security specifications.