

Automated Reasoning for Indistinguishability in the CCSA

Simon Jeanteur

TU Wien

Vienna, Austria

simon.jeanteur@tuwien.ac.at

Matteo Maffei

TU Wien

Vienna, Austria

matteo.maffei@tuwien.ac.at

Laura Kovács

TU Wien

Vienna, Austria

laura.kovacs@tuwien.ac.at

Michael Rawson

University of Southampton

Southampton, United Kingdom

michael@rawsons.uk

Abstract

Cryptographic protocols are the foundation of secure digital communication, yet their design remains error-prone, as evidenced by the vulnerabilities that have plagued even the most widely adopted protocols throughout history. Security properties are typically formalized using either trace properties or indistinguishability, each addressing distinct security guarantees, such as agreement and authenticity for the former and anonymity and strong secrecy for the latter. Formal verification of cryptographic protocols spans both symbolic and computational models. While symbolic techniques enable automation and scalability, they do not provide computational security guarantees. Computational models, though robust, are harder to formalize and automate. Recent advances, such as the *Computationally Complete Symbolic Attacker* (CCSA) model and its logic, the *Bana-Comon Logic* (BC Logic), bridge this gap by supporting both trace properties and indistinguishability. However, despite significant progress in proof assistants, automating indistinguishability remains a challenge due to its combination of unstructured equality theories, complex non-classical calculus, and partially inductive reasoning—all requiring expert knowledge in both cryptography and logic.

This paper introduces a novel approach to automate indistinguishability proofs in the CCSA model, implemented in the automated prover CRYPTOVAMPIRE2. We extend CRYPTOVAMPIRE to support indistinguishability by designing GOLGGE, a PROLOG-inspired backtracking engine over equality graphs (e-graphs), which provides strong, rewrite-driven equational reasoning capabilities. We adapt the BC Logic rules to this new framework, yielding semantically compatible statements. The effectiveness of our approach is demonstrated by automating all indistinguishability goals in the SQUIRREL repository.

CCS Concepts

• Security and privacy → Logic and verification; Formal security models.

Keywords

Security Protocols, Indistinguishability, Formal Methods, Computational Security, Automated Theorem Proving

This is the technical report.

1 Introduction

Cryptographic protocols serve as the foundation enabling secure communication among the components of our digital infrastructure. However, their design remains a challenging and error-prone endeavor. Historically, vulnerabilities have been uncovered in widely adopted protocols, ranging from the classic attack on the Needham-Schroeder protocol [56] to repeated breaches in the TLS protocol [3], and even in emerging blockchain protocols [58]. These incidents underscore the need for rigorous verification methods.

Formal methods have emerged as a robust approach to ensure the correctness of cryptographic protocols. Notably, they have played a pivotal role in the development of modern standards such as TLS 1.3 [15], WireGuard [54], and 5G-AKA [52]. Security properties for these protocols are typically formalized using either *trace properties* [57] or *indistinguishability relations* [25]. Trace properties focus on guarantees for individual execution traces, encompassing secrecy, authentication, and constraints on the order of security-relevant events. Indistinguishability, on the other hand, ensures that an attacker cannot distinguish between two protocol configurations, such as identifying which of two secrets is in use.

1.1 Related Work

Formal verification of cryptographic protocols is broadly categorized into *symbolic* and *computational* models. Symbolic techniques abstract away from the algebraic properties of cryptography, modeling attackers as entities with infinite computational resources but limited to constructing symbolic terms from their existing knowledge (the “Dolev-Yao” attacker [36]). This abstraction simplifies proof techniques and has led to the development of automated verifiers like PROVERIF [17] and TAMARIN [60]. However, symbolic security does not guarantee computational security; protocols proven secure in the symbolic model may still be vulnerable to attacks exploiting weaknesses in their cryptographic implementations [16].

Conversely, the computational model uses probabilistic and computational notions to mirror real-world implementations. Here, attackers are probabilistic polynomial-time Turing machines, and proofs reduce protocol security to mathematical assumptions. While providing strong guarantees, this model is harder to formalize and automate. Nonetheless, several tools have emerged, attempting automation. For instance, EASYCRYPT [14] is a cryptographic proof checker simplifying reasoning about probabilistic computations, typically used for cryptographic schemes. Additionally, CRYPTOVERIF [18] automates game-based proofs for protocols but

lacks support for states. Finally, OWL [39] uses types for integrity and confidentiality properties.

Recent efforts have focused on bridging the gap between symbolic and computational models, seeking to combine the ease of symbolic proofs with the strong guarantees of the computational model. A pivotal development in this direction is the *Computationally Complete Symbolic Attacker* (CCSA) model and its associated logic, the Bana-Comon (BC) Logic [9–11], which has since been formalized as a fragment of higher-order logic [30]. The BC Logic supports both trace properties and indistinguishability (Def. 1), enabling computationally sound formal verification. While early proofs in the BC Logic were conducted manually, recent advancements have led to the development of the interactive proof assistant SQUIRREL [5] and the automated prover CRYPTOVAMPIRE [47], the latter supporting only trace properties. The growing interest in CCSA has further extended its capabilities to include stateful protocols [6], post-quantum security [29], and proof composition [27]. SQUIRREL, as an interactive theorem prover, has made significant progress in usability through custom security games [8], automation of non-cryptographic goals [7], and strong typing [35]. However, it still demands expertise in both cryptography and logic, and full automation, in particular of indistinguishability proofs, remained until now an open problem.

Automating indistinguishability goals poses a number of challenges. Proofs in this setting typically require complex equality reasoning in a non-classical logic, necessitating non-trivial choices from cryptographers when applying cryptography-related rules. These challenges can in principle be addressed individually: congruence closure algorithms [70] handle equational reasoning; some theorem provers can reason in non-classical logics [64]; undirected choice can be modeled as SAT and related approaches [31, 59]; and PROLOG [26] and DATALOG [22] execute rule systems. However, integrating all these requirements into one system is an open challenge. *Satisfiability Modulo Theory* (SMT) solvers [12, 34, 53] are natural candidates, given their ability in principle to handle all these requirements (including non-classical logic [37]), but they already proved to require external assistance for trace properties [47] and remain incapable of handling indistinguishability goals independently, especially in quantified theories. A fundamentally novel approach is thus needed for machine-checked, automated reasoning about indistinguishability properties.

1.2 Our Contributions

This paper introduces the first methodology for automating indistinguishability proofs in the BC Logic. Our approach is implemented in new version of CRYPTOVAMPIRE¹, called CRYPTOVAMPIRE2, and delivers the following key contributions:

(A) Leveraging automated equality reasoning in the form of EGG [70] and inspired by EGGLOG [71], we introduce GOLGGE, a PROLOG-like [26] backtracking engine reasoning modulo equality.

(B) We adapt SQUIRREL’s rules to the GOLGGE engine, rethinking the subterm relations ubiquitous in the CCSA to align with equality reasoning. We also extend the concept of *bi-deducibility* [6] to avoid expensive rules. We ensure compatibility with SQUIRREL’s first-order semantics and prove the soundness of our approach.

¹The old version of the tool will be referred to by its paper, i.e., [47].

(C) We conduct an *experimental evaluation* of CRYPTOVAMPIRE2 on challenging, state-of-the-art protocols from [5, 6], in particular, we consider all indistinguishability queries in the SQUIRREL protocol library that are expressible as bi-processes (i.e., pairs of protocols sharing a structure, but differing on terms), which is the format that our tool supports at present.² CRYPTOVAMPIRE2 automates all such queries, taking as input the protocol specification and the security goals, expressed as trace properties, in less than 2 min 15 s each, on consumer hardware, showcasing its efficiency as a verifier.

2 Overview

CRYPTOVAMPIRE2 attempts to prove the indistinguishability of user-specified protocols. To that end, CRYPTOVAMPIRE2 expects the following inputs from its users: (i) function declarations, (ii) steps representing atomic inputs/outputs, (iii) their instantiations for each protocol, and (iv) constraints on all the above. These constraints could be cryptographic assumptions, simplification rules, structural dependencies across steps, or supplementary lemmas. CRYPTOVAMPIRE2 then attempts to prove the indistinguishability of the specified protocols by building a proof tree using the BC Logic and SQUIRREL-compatible rules.

EXAMPLE 1 (BASIC HASH). *To illustrate, we introduce the classic CCSA example in the form of the following RFID-like protocol [20]:*

$$\begin{aligned} T_j &\rightarrow R : \langle n, \mathcal{H}(n, k_j) \rangle \\ \leftarrow R &: \text{ok if } k_j \text{ is in } R\text{'s database} \end{aligned} \quad (1)$$

Here, a tag T_j outputs a fresh nonce n both in plaintext and hashed ($\mathcal{H}$) with the key k_j shared with a reader R . For simplicity, we model a single reader R operating over multiple tags T (parameterized by j). This protocol features unlinkability, meaning an attacker cannot trace a tag across multiple executions. This is expressed by observing that (1) is indistinguishable from:

$$\begin{aligned} T_j &\rightarrow R : \langle n, \mathcal{H}(n, k'_{ij}) \rangle \\ \leftarrow R &: \text{ok if } k'_{ij} \text{ is in } R\text{'s database} \end{aligned} \quad (2)$$

where i identifies multiple runs of the protocol. In (1), k_j is reused across sessions, while (2) uses a fresh key each time.

We now illustrate the main steps of analyzing Ex. 1 in CRYPTOVAMPIRE2.

2.1 Formalization of the Protocol

Following the CCSA model, cryptography is represented *symbolically* using a term algebra built on a set of honest functions $\mathcal{F}$ and a set of nonce names $\mathcal{N}$.

2.1.1 *Honest Functions.* The functions in $\mathcal{F}$ become *uninterpreted functions* in CRYPTOVAMPIRE2, representing deterministic polynomial-time Turing machines.

EXAMPLE 2 (FUNCTIONS). *To encode Ex. 1, we use the following functions in $\mathcal{F}$: $\langle \_, \_ \rangle$, $\pi_1(\_)$, $\pi_2(\_)$ for pair construction and projection; ok, ko for success and failure of tag authentication; $\mathcal{H}(\_, \_)$ for hash computation; and *if _ then _ else _* for conditionals.*

²This notably includes *all* examples from the original [5].

2.1.2 Honest Randomness. The only source of honest randomness in CRYPTOVAMPIRE2 is *nonces*, representing pairwise independent uniformly random bitstrings of length η , the security parameter. In the logic, nonces are constants used to express cryptographic properties (e.g., collision-resistance). Nonces are indexed to support freshness across an unbounded number of sessions and participants.

EXAMPLE 3 (NONCES). In Ex. 1, $\mathcal{N} = \{n[\_, \_], k[\_, \_], k'[\_, \_]\}$. $n[i, j]$ and $k'[i, j]$ are indexed by the session i and the tag identifier j , while $k[j]$ is indexed only by the tag j . We use nonces to represent keys as they are a source of honest randomness.

2.1.3 Protocol Steps. Within CRYPTOVAMPIRE2, a protocol is a set of atomic steps, and each execution is a valid sequence of those steps. A step takes an *input* and computes an *output* (or *message*). It may be guarded by a *condition*, specifying when its execution succeeds. To consider indistinguishability, protocols must share some common structure to enable lock-step execution similar to PROVERIF's Bi-Processes [19]. In the context of the CCSA we say such protocols are *compatible*, meaning they share their step identifiers.

EXAMPLE 4 (STEPS). We express Ex. 1's (1) and (2) using the following steps. p is the protocol ((1) or (2)), and mk is an alias for the keys: $mk(i, j, (1)) := k[j]$ and $mk(i, j, (2)) := k'[i, j]$.

$T[i, j]$: (i^{th} execution of the j^{th} tag)

Condition: true

Message: $\langle n[i, j], \mathcal{H}(n[i, j], mk(i, j, p)) \rangle$

$Rs[i, j]$: (Authentication of the j^{th} tag on the i^{th} execution)

Condition: $\pi_2(\text{in}) \equiv \mathcal{H}(\pi_1(\text{in}), mk(i, j, p))$

Message: ok

where in stands for $\text{input}@_p Rs[i, j]$.

$Rf[l]$: (No authentication on the l^{th} execution)

Condition: $\neg \exists i, j. \pi_2(\text{in}) \equiv \mathcal{H}(\pi_1(\text{in}), mk(i, j, p))$

Message: ko

where in stands for $\text{input}@_p Rf[l]$.

Protocol steps are executed following an order captured by the relation $a < b$. We also introduce the mutually exclusive relation $\diamond$ to ensure at most one can be executed in the same protocol run.

EXAMPLE 5 (MUTUAL EXCLUSION). In Ex. 1, the reader can only execute one branch per session. Thus, for each session i , the authentication either succeeds or fails, and at most one tag can be authenticated:

$$\text{for all } i, j \neq k, (Rf[i] \diamond Rs[i, j]) \text{ and } (Rs[i, j] \diamond Rs[i, k]) \quad (3)$$

No other ordering constraints apply to this protocol.

2.1.4 Semantic and Cryptographic Properties. The semantics of constants and functions are expressed through rewriting rules, inference rules, and first-order formulas. Following the CCSA model, these are proven against the computational attacker and relevant cryptographic assumptions, capturing what the attacker *cannot* do.

The simplest properties are *directed equality rules*, used by GOLGGE to manipulate proof goals before applying an inference rule.

EXAMPLE 6 (PROJECTION). π_1 is a destructor for $\langle \_, \_ \rangle$, yielding the following rewrite rule:

$$\pi_1(\langle x, y \rangle) \rightsquigarrow x \quad (4)$$

These rules also encode cryptographic and protocol-specific properties. For instance, Lem. 1 describes Ex. 1's authentication.

LEMMA 1 (AUTHENTICATION). Both protocols of Ex. 1 manifest authentication between tags and the reader: when $\text{happens}(\tau) \equiv \text{true}$,

$$\pi_2(\text{input}@_p \tau) \equiv \mathcal{H}(\pi_1(\text{input}@_p \tau), mk(i, j, p)) \rightsquigarrow \left(\begin{array}{l} \exists i'. T[i', j] < \tau \\ \wedge \pi_1(\text{input}@_p \tau) \equiv \pi_1(\text{msg}@_p T[i', j]) \\ \wedge \pi_2(\text{input}@_p \tau) \equiv \pi_2(\text{msg}@_p T[i', j]) \end{array} \right) \quad (5)$$

for any τ . This means that if the condition of a step Rs succeeds (i.e., a reader believes authentication succeeded) then a corresponding step T occurred earlier, and the inputs and outputs of both steps matched (i.e., the corresponding tag indeed tried to authenticate).

Properties that cannot be expressed as equality rules, typically related to cryptographic properties, are *inference rules*. The core mission of CRYPTOVAMPIRE2 is to chain these rules to automate indistinguishability proofs. For Ex. 1, it is relevant to include Prop. 1³.

PROPERTY 1 (PRF). When $\mathcal{H}(\_, \_)$ behaves like a Pseudo Random Function [41], the following rule holds:

$$\frac{u \sim v \vdash C[n] \sim b \quad m, k, n \triangleright_{PRF} C[\_, \_], m}{u \sim v \vdash C[\mathcal{H}(m, k)] \sim b} \text{PRF}$$

where k and n are nonces, $a \sim b$ is the indistinguishability predicate, and $m, k, n \triangleright_{PRF} \_$ is a subterm relation that ensures (i) n is fresh, (ii) the key k was not leaked nor misused, and (iii) $\mathcal{H}(m, k)$ is not used outside of the hole in $C[\_, \_]$.

When these conditions hold, this rule allows the prover to replace $\mathcal{H}(m, k)$ by a fresh nonce in an indistinguishability goal.

PROOF. See Appendix C.5. $\square$

A major contribution of this paper is to merge Prop. 1's (i), (ii) and (iii) into the single predicate $\_, \_ \triangleright_{PRF}$ ensuring each argument search through the same term, leading to soundness modulo rewriting. Finally, properties not expressible as rewrite rules or inferences are given using first-order formulas. For example:

THEOREM 1 (NO GUESSING). For any nonce $n \in \mathcal{N}_{bc}$, BC term t ,

$$n \triangleright t \Rightarrow |n \neq t| \quad (6)$$

where $x \triangleright y$ is also a subterm relation meaning that x is not used in y . That is, honest nonces cannot be guessed: a term that has no dependencies on a nonce ($n \triangleright t$) cannot be equal to it ($|n \neq t|$).

2.2 Proof Search modulo Equality

The presence of inference rules like Prop. 1's PRF naturally suggests PROLOG-like reasoning: a depth-first search in the graph induced by the inference rules to find a proof tree, or in our case, a *Directed Acyclic Graph* (DAG). However, the requirement for robust equality reasoning—due to the many potentially complex equational rules—makes direct use of PROLOG [26] impractical. The closest existing tool to our needs is EGLOG [71], an extension of DATALOG with strong support for equality, but its logical fragment lacks the expressiveness required for our rules. Thus, we developed our own automated reasoning engine: GOLGGE.

³see its proof in Appendix C.5




Figure 1: Applying rewrites over an example e-graph (adapted from [70])

2.2.1 E-Graphs. Both EGGLOG and GOLGGE rely on *e-graphs* [62], efficient data structures for storing terms and a congruence relation (i.e., an equivalence relation $\approx$ such that $\vec{x} \approx \vec{y} \Rightarrow f(\vec{x}) \approx f(\vec{y})$). An e-graph is a collection of *e-classes*, each composed of *e-nodes*. An e-class represents the subset of an equivalence class stored in the e-graph, while an e-node describes one way to construct an element of this e-class by applying a function symbol to other e-classes. Typically, the relation $\approx$ is approximated using *oriented equation rules*. Fig. 1 shows an evolving e-graph under rewriting of $(a * 2)/2$ using the rules $x * 2 \rightsquigarrow y \ll 1$ and $(x * y)/z \rightsquigarrow x * (y/z)$.

2.2.2 Backtracking. EGGLOG’s key insight is to store facts as *e-classes* (subsets of equivalence classes in an e-graph). GOLGGE reuses this insight but flips EGGLOG’s reasoning⁴ from forward (building the DAG from axiom e-classes to the goal e-class, DATALOG-style) to backward (building from the goal e-class toward the axiom e-classes, PROLOG-style). GOLGGE goes further than backward chaining, featuring powerful caching akin to SLG resolution [23].

EXAMPLE 7. Let a and b be two constants (i.e., honest functions with no arguments). To prove $\mathbf{n} \triangleright f(a, \pi_1(\langle b, \mathbf{n} \rangle))$, GOLGGE builds the proof tree (7) transparently using Ex. 6’s rewriting rule (4):

$$\frac{\frac{\mathbf{n} \triangleright c_2 \quad a}{\mathbf{n} \triangleright c_2} \quad \frac{\mathbf{n} \triangleright c_3 \quad b}{\mathbf{n} \triangleright c_3}}{\mathbf{n} \triangleright c_1} f \quad (7)$$

with $c_1 := f(a, \pi_1(\langle b, \mathbf{n} \rangle))$, $c_2 := a$, $c_3 := \pi_1(\langle b, \mathbf{n} \rangle) \approx b$.

This backward reasoning enables dynamic computation of applicable rules for each goal e-class, allowing arbitrarily complex rules. Notably, we can offload some goals to [47], translating them into first-order queries for external SMT solvers [12, 34, 53]. CRYPTO-VAMPIRE2 coordinates Section 2.1 and Section 2.2 together. It translates the protocols and their constraints into a problem for GOLGGE as described in Section 2.1. GOLGGE then attempts to construct a proof, potentially delegating subgoals to [47]. If GOLGGE succeeds, then the protocols are indeed indistinguishable (Th. 4).

3 Preliminaries

We shall use vector notations $\vec{y}$ throughout as a shorthand for $y_1, \dots, y_n$ when n is obvious from the context. We write $\text{fv}(\phi)$ for

⁴hence the name of the tool.

the free variables of ϕ . We also write the *inference rule* $\frac{A_1 \dots A_n}{B}$, with $n \geq 0$, to mean that $A_1 \wedge \dots \wedge A_n \Rightarrow B$ is valid. All probability-related statements are defined in Appendix A.

3.1 Introduction to the CCSA—The BC Logic

CRYPTOVAMPIRE2’s syntax and semantics extend the Bana-Comon Logic (BC Logic) [9]. This section briefly introduces to the BC Logic, while Section 3.2 introduces its extension in CRYPTO-VAMPIRE2 and [47]. We assume the following sets:

- (1) $\mathcal{N}_{\text{bc}}$, a finite set of *nonce* names;
- (2) $\mathcal{F}_{\text{bc}}$, a finite set of *honest* functions;
- (3) $\mathcal{G}_{\text{bc}}$, a finite set of *attacker* functions.

Terms in the BC Logic are built from the following grammar:

$$t := x \mid \mathbf{n} \mid f(\vec{t}) \mid g(\vec{t}) \quad (8)$$

where $\mathbf{n} \in \mathcal{N}_{\text{bc}}$, $f \in \mathcal{F}_{\text{bc}}$, $g \in \mathcal{G}_{\text{bc}}$, and x is a variable.

A computational model $\mathbb{M}$ assigns elements of $\mathcal{N}_{\text{bc}}$, $\mathcal{F}_{\text{bc}}$, and $\mathcal{G}_{\text{bc}}$ to polynomial-time Turing machines, defining an interpretation $\llbracket t \rrbracket_{\text{bc}}^{\mathbb{M}}$ for each BC term t with respect to the security parameter η and a pair $\rho = (\rho_h, \rho_{\mathcal{A}})$ of random tapes (for honest randomness and the adversary, respectively):

- (1) $\llbracket \mathbf{n} \rrbracket_{\text{bc}}^{\mathbb{M}}(1^\eta, (\rho_h, \rho_{\mathcal{A}}))$, with $\mathbf{n} \in \mathcal{N}_{\text{bc}}$, is a slice of length η of ρ_h , pairwise distinct for all elements of $\mathcal{N}_{\text{bc}}$.
- (2) For $f \in \mathcal{F}_{\text{bc}}$,

$$\llbracket f(t_1, \dots, t_n) \rrbracket_{\text{bc}}^{\mathbb{M}}(1^\eta, \rho) = \llbracket f \rrbracket^{\mathbb{M}}(1^\eta, \llbracket t_1 \rrbracket_{\text{bc}}^{\mathbb{M}}(1^\eta, \rho), \dots, \llbracket t_n \rrbracket_{\text{bc}}^{\mathbb{M}}(1^\eta, \rho)) \quad (9)$$

where $\llbracket f \rrbracket^{\mathbb{M}}$ is a polynomial-time Turing machine with no direct access to ρ .

- (3) For $g \in \mathcal{G}_{\text{bc}}$,

$$\llbracket g(t_1, \dots, t_n) \rrbracket_{\text{bc}}^{\mathbb{M}}(1^\eta, \rho) = \llbracket g \rrbracket^{\mathbb{M}}(1^\eta, \rho_{\mathcal{A}}, \llbracket t_1 \rrbracket_{\text{bc}}^{\mathbb{M}}(1^\eta, \rho), \dots, \llbracket t_n \rrbracket_{\text{bc}}^{\mathbb{M}}(1^\eta, \rho)), \quad (10)$$

where $\rho := (\rho_h, \rho_{\mathcal{A}})$ and $\llbracket g \rrbracket^{\mathbb{M}}$ is a polynomial-time Turing machine with direct access to $\rho_{\mathcal{A}}$ but not ρ_h . Thus, terms of this form model attacker computations.

In this paper, we assume $\langle \_, \_ \rangle$ (tuples), π_1, π_2 (projections), $\_ \doteq \_$ (equality), $\text{If } \_ \text{ then } \_ \text{ else } \_$, true , and false are part of $\mathcal{F}_{bc}$. This lets us to define *boolean connectives*, which we distinguish from standard first-order boolean operators by overlining them (e.g., we write $a \bar{\wedge} b := \text{If } a \text{ then } b \text{ else false}$ instead of $a \wedge b$).

DEFINITION 1 (INDISTINGUISHABILITY). Two BC terms a, b are indistinguishable, written $\mathbb{M} \models a \sim_{bc} b$, if for all polynomial-time attackers $\mathcal{A}$,

$$\left| \mathbb{P} \left(\mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \llbracket a \rrbracket_{bc}^{\mathbb{M}}(1^\eta, \rho)) = 1 \right) - \mathbb{P} \left(\mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \llbracket b \rrbracket_{bc}^{\mathbb{M}}(1^\eta, \rho)) = 1 \right) \right| = \text{negl}(\eta) \quad (11)$$

where $\rho := (\rho_h, \rho_{\mathcal{A}})$ with ρ_h and $\rho_{\mathcal{A}}$ independent random tapes. A function h is negligible in η (written $\text{negl}(\eta)$) when $h(\eta) \stackrel{\eta \rightarrow \infty}{=} o(\eta^{-c})$ for all $c \in \mathbb{N}$.

This strong semantic foundation allows us to prove BC axioms and rules using standard cryptographic tools. An important example is Th. 1 from Section 2 stating that it is hard to guess nonces:

THEOREM (1 NO GUESSING). For any nonce $n \in \mathcal{N}_{bc}$, BC term t ,

$$n \blacktriangleright t \Rightarrow |n \neq t| \quad (12)$$

where $x \blacktriangleright y$ is also a subterm relation meaning that x is not used in y .

PROOF SKETCH. Assuming $n \blacktriangleright t$, by induction on t , the random variables $\llbracket n \rrbracket_{bc}^{\mathbb{M}}(1^\eta, \rho)$ and $\llbracket t \rrbracket_{bc}^{\mathbb{M}}(1^\eta, \rho)$ are independent. A detailed proof in the full logic is postponed to Appendix C.2. $\square$

In practice, we restrict ourselves to a class $\mathcal{C}$ of models where function interpretations closely match their real-world counterparts (e.g., $\pi_1(\_)$ as a projection, $\mathcal{H}(\_, \_)$ is a PRF hashing function). The rules and formulas introduced in Section 2.1.4 describe this chosen class of models. We write $\mathcal{C} \models \varphi$ for validity in all models in $\mathcal{C}$.

3.2 Extensions to CRYPTOVAMPIRE2's Logic

3.2.1 The Notion of Protocols. The BC Logic effectively represents symbolic computations and can model a single protocol execution, where network inputs are modeled by applying an attacker function to the history of prior messages (the *trace*). However, protocol verification requires considering all possible executions. CRYPTOVAMPIRE2 adopts SQUIRREL's meta-logic approach, abstracting executions while formally unfolding back into the BC Logic for each interleaving. The result is the Symbolic Logic $\mathcal{S}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C})$ (Fig. 2).

Indices (variables $i, j, \dots$ from $\mathcal{I}$) extend reasoning to unbounded protocols. **Timepoints** T represent when an atomic input/output step is executed, either directly via an indexed identifier a or relatively via the predecessor operator pred . **Predicates** A reason about timepoint ordering ($=, <, \leq$) or step execution (happens). **Terms** t follow the same structure to the BC Logic, with indexed nonce identifiers and function symbols from $\mathcal{N}$ and $\mathcal{F}$, respectively. We restrict ourselves to a single attacker function att . **Binders** (in the form of lookups) allow defining quantifiers:

$$\exists \vec{a}, t := \text{Find } \vec{a} \text{ st. } t \text{ then true else false,} \quad (13)$$

$$\forall \vec{a}, t := \neg \exists \vec{a}, \neg t. \quad (14)$$

Macros relate terms to executions or traces, taking a protocol identifier p and a timepoint T as arguments, representing input

$I :=$	$i, j, \dots$	$i, j \in \mathcal{I}$, indices
$T :=$	τ	timepoints
	$ a[\vec{i}]$	$a[\vec{i}] \in \mathcal{S}$, step
	$ \text{pred}(T)$	previous step
$A :=$	$T = T' \mid T < T' \mid T \leq T'$	order over timepoints
	$ \text{happens}(T) \mid N = N'$	
$t :=$	$u, v, \dots \mid A$	
	$ n[\vec{i}]$	$n[\vec{i}] \in \mathcal{N}$, nonces
	$ f[\vec{i}](\vec{i})$	$f[\vec{i}] \in \mathcal{F}$, honest function
	$ \text{att}(t)$	attacker function
	$ \text{Find } \vec{a} \text{ st. } t \text{ then } t' \text{ else } t''$	lookup
	$ M@_p T$	macro
$P :=$	p	$p \in \mathcal{P}$, protocol identifiers
$M :=$	$\text{input} \mid \text{msg} \mid \text{cond} \mid \text{frame} \mid \text{exec} \mid \text{call}(c[\vec{i}])$	
		$c[\vec{i}] \in \mathcal{C}$, memory cell

Terms built with this grammar form $\mathcal{S}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C})$, shortened to $\mathcal{S}$ when the parameters are obvious. $\vec{a}$ is a vector of variables over timepoints and/or indices.

Figure 2: Grammar of the Symbolic Logic.

(*input*), output (*msg*), memory cell reads (*call*()), execution condition (*cond*), attacker knowledge (*frame*), and the conjunction of conditions of all executed steps (*exec*).

DEFINITION 2 (STEP). A step $a[\vec{i}]$, represented by (c, m, u) , is an atomic input/output operation guarded by a condition:

- (1) c is the condition guarding the step's execution (referred to by $\text{cond}_p(a[\vec{i}])$).
- (2) m is the message output if the condition succeeds (referred to by $\text{msg}_p(a[\vec{i}])$).
- (3) $u(c[\vec{j}])$ is the assignment to cell $c[\vec{j}] \in \mathcal{C}$ (referred to by $\text{cell}_p(c[\vec{j}], a[\vec{i}])$).

where p is the protocol and $\text{fv}(c, m) \subseteq \vec{i}$, $\text{fv}(u(c[\vec{j}])) \subseteq \vec{i} \cup \vec{j}$.

We single out the step *init*, which always has no parameters and its condition set to true , as the first step executed in all protocols. Its output represents the attacker's initial knowledge.

DEFINITION 3 (PROTOCOL). A protocol p is a tuple $(\mathcal{S}, \preceq, \diamond)$ composed of a set $\mathcal{S}$ of steps containing *init* and constraints $\preceq$ (partial order) and $\diamond$ (mutual exclusion) over $\mathcal{S}$. To ensure executability and representability, we require:

- (1) constraints and cell assignments are index-independent (e.g., $a[\vec{i}] \leq b[\vec{j}] \Leftrightarrow a[\vec{k}] \leq b[\vec{l}]$);
- (2) step definitions only reference timepoints smaller than themselves according to $\preceq$ and protocol p and recursive cell writes statically terminate;
- (3) *init* is the smallest element of $\mathcal{S}$ according to $\preceq$.

REMARK 1 (MANY **init** STEPS). When attacker knowledge does not fit in a single step (e.g., publishing a whole sequence $(pk)_i$ of public keys), we introduce extra initialization steps: unguarded steps smaller by $\leq$ than all steps in the main body of the protocol.

3.2.2 Executions. We model protocol execution with a trace:

DEFINITION 4 (TRACE). Given a protocol $p = (\mathcal{S}, \leq, \diamond)$, a trace $\mathbb{T}$ is a tuple $(\mathbb{I}, \mathbb{S}, \sigma_{\mathbb{I}}, \leq_{\mathbb{S}})$ such that:

- (1) $\mathbb{I}$ is a finite set of indices;
- (2) $\sigma_{\mathbb{I}} : \mathcal{J} \rightarrow \mathbb{I}$, idempotent on $\mathbb{I}$, assigns index variables to $\mathbb{I}$;
- (3) $\mathbb{S} \subseteq \{a[\sigma_{\mathbb{I}}(i_1), \dots, \sigma_{\mathbb{I}}(i_n)] \mid a[i_1, \dots, i_n] \in \mathcal{S}\}$, is the downward closed⁵ set by $\leq$ of all executed steps;
- (4) $\leq_{\mathbb{S}}$ is a total order over $\mathbb{S}$ compatible with $\leq$ and $\diamond$:

$$a[\vec{i}] \leq b[\vec{j}] \Rightarrow a[\vec{i}] \leq_{\mathbb{S}} b[\vec{j}], \quad (15)$$

$$a[\vec{i}] \diamond b[\vec{j}] \Rightarrow a[\vec{i}] <_{\mathbb{S}} b[\vec{j}] \wedge b[\vec{j}] <_{\mathbb{S}} a[\vec{i}]. \quad (16)$$

(16) prevents two mutually exclusive steps from being in $\mathbb{S}$ simultaneously, and (15) and downward closure ensures **init** is in $\mathbb{S}$ and is its smallest element.

EXAMPLE 8. Reusing Ex. 1, assume two tags and two rounds. The protocol splits into steps **init**, $\mathbb{T}[i, j]$, $\mathbb{R}s[i, j]$, and $\mathbb{R}f[i]$ as in Ex. 4, with constraints from Equation 3. Let $\mathcal{J} := \mathbb{N}$ and consider:

$$\mathcal{S}_1 = \{\text{init} <_1 \mathbb{T}[1, 2] <_1 \mathbb{R}s[1, 2]\}, \quad (17)$$

$$\mathcal{S}_2 = \{\text{init} <_2 \mathbb{T}[1, 1] <_2 \mathbb{R}s[1, 2] <_2 \mathbb{T}[2, 2]\}, \quad (18)$$

$$\mathcal{S}_3 = \{\text{init} <_3 \mathbb{R}f[1] <_3 \mathbb{T}[1, 1] <_3 \mathbb{R}s[1, 1]\}. \quad (19)$$

(17) describes a valid trace with $\mathbb{I} = \{1, 2\}$, as is (18) despite step $\mathbb{R}s[1, 2]$ failing. However, (19) is not a valid trace as it contradicts $\diamond$ despite the condition of all the steps seemingly holding.

We restrict ourselves to protocols sharing some overall structure enforcing lock step execution by taking inspiration from the notion of Bi-Process [19, 24] used in other automated protocol verifiers supporting indistinguishability like PROVERIF [17].

DEFINITION 5 (COMPATIBLE PROTOCOLS). Two protocols are compatible if they share step identifiers and constraints.

Throughout the remainder of this paper, we will assume $\mathcal{P}$ consists of pairwise compatible protocols.

EXAMPLE 9. In Ex. 1, (1) and (2) are compatible.

DEFINITION 6 (UNFOLDING). Given a trace $\mathbb{T}$ and a term t from the Symbolic Logic, $(t)^{\mathbb{T}}$ is its unfolding into a BC term (Appendix B).

EXAMPLE 10. Continuing Ex. 1, let us unfold $\text{input}_{(1)}(\mathbb{R}s[1, 2])$ according to trace $\mathbb{T}_{(17)}$ described by Ex. 8's (17):

$$\begin{aligned} &(\text{input}_{(1)}(\mathbb{R}s[1, 2]))^{\mathbb{T}_{(17)}} := \\ &\text{att} \left[\begin{array}{l} \langle (\text{cond}_{(1)}(\mathbb{T}[1, 2]))^{\mathbb{T}_{(17)}} \bar{\wedge} (\text{cond}_{(1)}(\text{init}))^{\mathbb{T}_{(17)}} ; \\ \langle \text{If}(\text{cond}_{(1)}(\mathbb{T}[1, 2]))^{\mathbb{T}_{(17)}} \bar{\wedge} (\text{cond}_{(1)}(\text{init}))^{\mathbb{T}_{(17)}} \\ \text{then } (\text{msg}_{(1)}(\mathbb{T}[1, 2]))^{\mathbb{T}_{(17)}} \text{ else fail;} \\ (\text{msg}_{(1)}(\text{init}))^{\mathbb{T}_{(17)}} \rangle \end{array} \right] \quad (20) \end{aligned}$$

⁵i.e., such that if $y \leq x$ and $x \in \mathbb{S}$ then $y \in \mathbb{S}$.

We extend indistinguishability to the Symbolic Logic: Given a computational model $\mathbb{M}$ and a trace $\mathbb{T}$, we write $\mathbb{M}, \mathbb{T} \models a \sim b$ when $\mathbb{M} \models (a)^{\mathbb{T}} \sim_{\text{bc}} (b)^{\mathbb{T}}$. We abbreviate $\mathbb{M}, \mathbb{T} \models a \sim b$ or $\mathbb{C}, \mathbb{T} \models a \sim b$ as $a \sim b$ when context permits. We write $\vdash_{\mathbb{P}} a$ to denote $a \sim \text{true}$ (a trace property). Throughout this paper, we use $\mathbb{C}, \mathbb{T} \models \vdash_{\mathbb{P}} a \equiv b$ as our congruence relation, written $a \equiv b$.

DEFINITION 7 (INDISTINGUISHABLE PROTOCOLS). Two compatible protocols p_1 and p_2 are observationally equivalent or indistinguishable in a class of models $\mathbb{C}$ if for all timepoints τ and traces $\mathbb{T}$,

$$\mathbb{C}, \mathbb{T} \models \text{frame}_{p_1} @_{p_1} \tau \sim \text{frame}_{p_2} @_{p_2} \tau. \quad (21)$$

NOTATIONS 1. We write $\llbracket t \rrbracket_{\mathbb{M}, \mathbb{T}}^{\eta}(\rho)$ for $\llbracket (t)^{\mathbb{T}} \rrbracket_{bc}^{\mathbb{M}}(1^{\eta}, \rho)$, omitting $\mathbb{M}$ and/or $\mathbb{T}$ when obvious.

We link trace properties to the notion of overwhelming sequences (Def. 8) by noting that $\vdash_{\mathbb{P}} P$ means that $(\llbracket P \rrbracket_{\mathbb{M}, \mathbb{T}}^{\eta}(\rho) = 1)_{\eta \in \mathbb{N}}$ is overwhelming.

DEFINITION 8 (OVERWHELMING). A sequence of sets or events $(E_{\eta})_{\eta \in \mathbb{N}}$ is overwhelming if $|1 - \mathbb{P}(E_{\eta})| = \text{negl}(\eta)$.

Thus far, all notions introduced map directly to those of SQUIRREL [5], with the sole exception that macros are tagged with a protocol identifier. This correspondence immediately yields Prop. 2.

PROPERTY 2 (INTERCOMPATIBILITY). Two compatible and executable protocols are indistinguishable in CRYPTOVAMPIRE2 if and only if they are indistinguishable in SQUIRREL [5, 6].

REMARK 2. Note that Prop. 2 only considers the original SQUIRREL [5, 6] and does not account for some of its later extensions such as high-order [30] or quantum soundness [29].

CRYPTOVAMPIRE2 mainly considers the Symbolic Logic. Notably, our use of GOLGGE is constrained to it. However, we sometimes requires more precision than GOLGGE can achieve (e.g., for some trace properties), prompting offloads to [47] which uses a further extension to FOL.

3.2.3 As a First-Order Logic. As noted elsewhere [4, 9, 47]⁶, the BC Logic's semantics do not match classical FOL.

EXAMPLE 11 (COUNTEREXAMPLE TO CLASSICAL SEMANTICS). Consider $\text{first\_bit}(\mathbf{n})$, with $\mathbf{n}$ a nonce. It can be distinguished from both true and false.

Nonetheless, as shown in [47], we can extend the Symbolic Logic to be compatible with classical FOL.

DEFINITION 9 (EVALUATED LOGIC). We define the Evaluated Logic $\mathcal{L}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C})$ as:

$$\begin{aligned} t &:= \dots \quad \text{as in Fig. 2,} \\ \phi &:= |t| \mid P(\vec{t}) \mid \phi \wedge \phi' \mid \phi \vee \phi' \mid \neg \phi \mid \exists \vec{x}. \phi \mid \forall \vec{x}. \phi, \end{aligned} \quad (22)$$

where P denotes predicates (e.g., other subterm relations).

⁶in the technical report of [9]

Given a cryptographic model M and a trace T , we link the Evaluated Logic with the Symbolic Logic—and consequently the rest of the CCSA literature—using models $\mathcal{L}^{M,T}$. These Tarskian models [69] are restricted to be a Herbrand model [43] over $\mathcal{S}$ and evaluate $|t|$ to $\llbracket t \rrbracket_{M,T}^\eta(\rho)$:

$$\begin{aligned} \mathcal{L}^{M,T}(t)(1^\eta, \rho) &:= t \quad \text{when } t \in \mathcal{S}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C}) \\ \mathcal{L}^{M,T}(|t|)(1^\eta, \rho) &:= \begin{cases} \mathbf{0} & \text{if } \llbracket t \rrbracket_{M,T}^\eta(\rho) = \mathbf{0} \\ \mathbf{1} & \text{otherwise} \end{cases} \end{aligned}$$

PROPERTY 3 (RELATION BETWEEN LOGICS).

$$\mathbb{P}(\mathcal{L}^{M,T}(|t|)(1^\eta, \rho) = \mathbf{0}) = \text{negl}(\eta) \text{ iff. } M, T \models_P t. \quad (23)$$

for any $t \in \mathcal{S}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C})$.

THEOREM 2 (WEAK FIRST-ORDER COMPATIBILITY). *Let φ, ϕ be quantifier-free formulas of $\mathcal{L}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C})$ such that:*

- (1) $\vec{x} := \text{fv}(\varphi)$ and $\vec{y} := \text{fv}(\phi)$;
- (2) for all $\vec{t}, \vec{C}, T \models_P \varphi\{\vec{x} \mapsto \vec{t}\}$;
- (3) $(\forall \vec{x}. \varphi) \Rightarrow (\forall \vec{y}. \phi)$ is valid in FOL.

Then $\vec{C}, T \models_P \phi\{\vec{y} \mapsto \vec{t}\}$ for all $\vec{t}$.

Prop. 3 and Th. 2 enable offloading trace properties to classical automated theorem provers for FOL. This is the basis of [47], which we use when the e-graph approach is not complete enough. As the topic is quite orthogonal to indistinguishability, we refer the reader to [47] for more details on the FOL translation specifically and a stronger version of Th. 2.

4 GOLGGE

Automating indistinguishability proofs requires tackling two significant challenges:

(A) Even the simple Ex. 1 requires a wide range of equality rules, from basic semantics (e.g., Ex. 6) to advanced properties like authentication (e.g., Lem. 1) and macro unfolding. While Ex. 6 reduces term size, macro unfolding increases it, and Lem. 1 transforms terms into structurally-unrelated forms. Despite existing research [51], we were not able to complete the set of equations.

(B) The rules of the CCSA calculus can be highly complex: for instance, some encode inductive reasoning (e.g., induction on traces), others require arbitrary trace properties (e.g., step ordering), and some introduce fresh symbols.

This combination of challenges means no existing tool fully meet our needs, necessitating the development of GOLGGE. The closest existing tool to our needs is EGGLOG [71], a powerful equality saturation engine and DATALOG theorem prover. Like classic DATALOG [22], it derives all facts from initial facts using inference rules, but unlike DATALOG, it supports arbitrary rewriting. EGGLOG uses e-graphs [62] to efficiently represent equivalence classes under the congruence relation induced by rewriting. Thus, facts are no longer simple terms, but full equivalence classes instead. However, EGGLOG's DATALOG heritage presents limitations for CRYPTOVAMPIRE2: it requires pre-enumerating all base facts, which is impractical for our calculus where facts like trace properties and introduction of fresh terms emerge dynamically during proofs. Nonetheless, EGGLOG serves as a strong foundation for GOLGGE, particularly its e-graph infrastructure, which we adapt by reversing the reasoning

direction from forward to backward chaining to better handle the unbounded nature of our proof obligations.

This section describes details of GOLGGE. After more formally introducing e-graphs in Section 4.1, we present a simplified version of the engine in Section 4.2 and discuss extensions and optimizations.

4.1 E-Graphs and Equality Saturation

DEFINITION 10 (E-GRAPH). *An e-graph $\mathcal{G}$ over a signature Σ and a congruence relation $\approx$ consists of e-nodes and e-classes:*

- (1) *An e-class $\mathbf{a} := \{\mathbf{n}_1, \dots, \mathbf{n}_m\}$ is a set of e-nodes $\mathbf{n}_i$, each representing a subset of an equivalence class stored in $\mathcal{G}$.*
- (2) *An e-node $\mathbf{n}$ represents a way to construct its e-class $\mathbf{b}$. It takes the form $f(\vec{\mathbf{a}})$ with $f \in \Sigma$ and $\vec{\mathbf{a}}$ a vector of e-classes, expressing that for all term vectors $\vec{x}, \vec{y} \in \vec{\mathbf{a}}, f(\vec{x}) \approx f(\vec{y})$ and both belong to $\mathbf{b}$.*

Fig. 1 illustrates several e-graphs. E-graphs are widely used in automated theorem proving [33, 34, 63] and form a powerful basis for equality saturation engines like EGG [70]. For the latter, the relation $\approx$ is typically approximated by *directed equality rules*.

DEFINITION 11 (DIRECTED EQUALITY RULES). *A directed equality rule has the form*

$$x_1 = t_1, \dots, x_n = t_n \rightsquigarrow y_1 = t'_1, \dots, y_m = t'_m \quad (24)$$

where $x_1, \dots, x_n, y_1, \dots, y_m$ are variables and $t_1, \dots, t_n, t'_1, \dots, t'_m$ are patterns (i.e., terms with variables) such that $\text{fv}(t'_1, \dots, t'_m) \cup \{y_1, \dots, y_m\} \subseteq \text{fv}(t_1, \dots, t_n) \cup \{x_1, \dots, x_n\}$.

A rule is *sound* when for any substitution σ , if $\sigma(x_i) \approx t_i\sigma$ for all i , then $\sigma(y_j) \approx t'_j\sigma$ for all j . We simply write $t \rightsquigarrow t'$ for $x = t \rightsquigarrow x = t'$ with $x \notin \text{fv}(t, t')$. This aligns with the intuitive notion of equality rules.

The orientation of rules helps equality saturation solvers avoid infinite loops by only adding terms in one direction. Rewriting rules underapproximate (often non-tightly) the congruence relation. Given an e-graph $\mathcal{G}$, we write $\equiv_{\mathcal{G}}$ for the induced congruence relation, and, assuming sound equality rules, we get:

PROPERTY 4 (E-GRAPHS UNDERAPPROXIMATE). *For $t, t' \in \mathcal{G}$,*

$$t \equiv_{\mathcal{G}} t' \Rightarrow t \approx t'. \quad (25)$$

DEFINITION 12 (SATURATION). *An e-graph is saturated when it is stable under all its equality rules.*

In particular, once $\mathcal{G}$ is saturated, $\equiv_{\mathcal{G}}$ is a proper congruence relation. Fig. 1 exemplifies the saturation process.

4.2 Base Decision Procedure

GOLGGE expects the following inputs: (i) a set $\mathcal{R}$ of *directed equality rules* (Def. 11), (ii) a set $\mathcal{T}$ of *inference rules* (Def. 13), (iii) a term g to serve as a *goal* to prove.

DEFINITION 13 (INFERENCE RULES). *Given patterns $Q, R_1, \dots, R_n$ such that $\text{fv}(R_1, \dots, R_n) \subseteq \text{fv}(Q)$, the rule*

$$\frac{R_1 \quad \dots \quad R_n}{Q} \quad (26)$$

is sound modulo $\approx$ if for all ground propositions $R'_1, \dots, R'_n, Q'$ where $R'_1, \dots, R'_n$ hold, if we can find a substitution σ such that $R_i\sigma \approx R'_i$ and $Q\sigma \approx Q'$, then Q holds as well.

REMARK 3. *The constraints on free variables in a GOLGGE inference rule are the reverse of those in EGGLOG.*

The tool attempts to build a proof tree for g using the rules in $\mathcal{T}$ and $\mathcal{R}$. Facts, such as the goal, are represented by their e-class (e.g., $\mathbf{g}$ for g) in an EGG e-graph continuously saturated using $\mathcal{R}$. Simultaneously, GOLGGE backtracks over the rules in $\mathcal{T}$, searching for a rule whose conclusion unifies with $\mathbf{g}$ by *e-matching* and then attempts to prove each premise. If it fails, it backtracks. Algorithm 1 shows a simplified implementation of GOLGGE.

Algorithm 1 Simplified version of GOLGGE

```

function PROVE1( $\mathbf{g}$ )
  saturate( $\mathcal{R}$ )
  for  $\frac{R_1 \dots R_n}{Q} \in \mathcal{T}$  do
    if  $Q\sigma \approx \mathbf{g}$  then
       $R_1, \dots, R_n \leftarrow \text{add}(R_1\sigma, \dots, R_n\sigma)$ 
      if  $\forall i, \text{PROVE}_1(R_i\sigma)$  then return true
  return false

```

where saturate($\mathcal{R}$) saturates the underlying e-graph with the rules in $\mathcal{R}$, σ is a substitution found by e-matching, and add adds terms to the e-graph, returning the potentially new e-class.

In the end, the proof generated by the execution of GOLGGE is a proof tree of $\mathcal{T}$, where substitutions by $\equiv_{\mathcal{G}}$ are interleaved between each rule application. This means all predicates must *at least* be compatible with $\approx$ (i.e., if $\tilde{x} \approx \tilde{y}$ then $P(\tilde{x}) \approx P(\tilde{y})$) and potentially more if a predicate can be rewritten into another. Another notable point is that $\mathcal{G}$ only grows during execution. Consequently, $\equiv_{\mathcal{G}}$ also grows monotonically, and no backtracking over it may occur. This property significantly reduces complexity during backtracking. It is a key differentiating factor between GOLGGE and EGGLOG's approach to e-graphs and a typical SMT solver's treatment [34].

THEOREM 3 (SOUNDNESS OF GOLGGE). *If all rules in $\mathcal{T}$ and all equality rules are sound and PROVE₁($\mathbf{g}$) returns true, then for all predicates $P \in \mathbf{g}$, P holds.*

PROOF. By induction on Algorithm 1. $\square$

4.3 Extensions and Optimizations

4.3.1 *Dynamic Inference Rules.* Within CRYPTOVAMPIRE2, restricting ourselves to a predetermined finite set of inference rules is insufficient. For example, the PRF rule of Prop. 1 uses *subterm relations*, requiring proving that some terms are distinct (i.e., $\vdash_P a \neq b$). These queries are impossible for the e-graph, as it only underapproximates equality (Prop. 4). Thus, a natural extension is to determine proof obligations dynamically as we pattern-match on the goal. This allows us to: (i) apply optimized decision procedures for some goals (e.g., subterm reasoning), (ii) trigger transformations on the e-graph that would otherwise be complex (e.g., substitutions in contexts), and (iii) offload queries to [47] (see Ex. 12).

EXAMPLE 12 (TRACE QUERIES). *With dynamic rules, the goal $\vdash_P \varphi$ can be offloaded to [47] and treated as an axiom upon success; otherwise, it behaves as if no rule applies.*

Dynamic rules are more expensive than regular ones, especially offloads to the trace solver, which must often timeout to fail. This makes the caching of Section 4.3.2 yet more critical.

REMARK 4 (EVOLUTION OF $\equiv_{\mathcal{G}}$). *While dynamic rules can modify the e-graph directly (e.g., add terms, merge e-classes), they must not introduce modifications that contradict Prop. 4, as GOLGGE does not backtrack on $\equiv_{\mathcal{G}}$.*

In this new setting, we consider $\mathcal{T}$ as a *function* from e-classes to sets of sets of e-classes representing the next proof obligations.

EXAMPLE 13. $\mathcal{T}(\mathbf{x}) = \{\emptyset\}$ expresses that $\mathbf{x}$ is an axiom, while $\mathcal{T}(\mathbf{x}) = \emptyset$ indicates it is unprovable.

In practice, such sets are lists whose order can be relied on for stateful operations and optimizations (e.g., executing cheap pruning goals before expensive ones).

4.3.2 *Loop Prevention and Memoization.* Loops can arise quickly during the search, as shown in Ex. 14.

EXAMPLE 14 (LOOP). *Recall (4) from Ex. 6 and the PRF rule of Prop. 1. Consider $u \sim v \vdash \pi_1(\langle x, \mathcal{H}(y, k) \rangle) \sim b$. We obtain:*

$$\frac{u \sim v \vdash \pi_1(\langle x, \mathcal{H}(y, k) \rangle) \sim b \quad \text{by (4)}}{u \sim v \vdash \pi_1(\langle x, \mathbf{n} \rangle) \sim b} \text{SUBST} \quad \dots \quad \text{PRF} \quad \frac{}{u \sim v \vdash \pi_1(\langle x, \mathcal{H}(y, k) \rangle) \sim b} \quad (27)$$

thus creating a loop.

To mitigate this, we use *memoization*, tracking currently proven goals and pruning branches that reuse them. This complements *failure and success caching* similar to PROLOG's SLG resolution [23]. Unlike in PROLOG, we cache whole e-classes, significantly reducing the search space, as seen in Ex. 15.

EXAMPLE 15 (CACHING). *Following Ex. 14, caching $u \sim v \vdash x \sim b$ also caches $u \sim v \vdash \pi_1(\langle x, \mathcal{H}(y, k) \rangle) \sim b$, skipping the expensive PRF rule. Caching also lets GOLGGE to detect the loop in Ex. 14.*

Simple cases like Ex. 14 effectively make memoization a requirement for termination.

With all optimizations, GOLGGE behaves as in Algorithm 2, where mem is the memoization hashmap.

Algorithm 2 GOLGGE

```

function PROVE2( $\mathbf{g}$ )
  saturate( $\mathcal{R}$ )
  if  $\mathbf{g} \in \text{mem}$  then return mem}[\mathbf{g}]
  mem[\mathbf{g}]  $\leftarrow$  false
  for  $\{R_1, \dots, R_n\} \in \mathcal{T}(\mathbf{g})$  do
    if  $\forall i, \text{PROVE}_2(R_i)$  then
      mem[\mathbf{g}]  $\leftarrow$  true
      return true
  return false

```

$\mathcal{T}$ can mutate the e-graph and is responsible for adding proof obligations to it.

This implementation of memoization is not tight. It is nonetheless sound (Prop. 5) and empirically sufficient for protocol verification (Section 6).

PROPERTY 5. If $\text{PROVE}_2(\mathbf{g})$ returns true, then so does $\text{PROVE}_1(\mathbf{g})$.

PROOF. By induction in Algorithm 2 on Prop. 5 and $\text{mem}[\mathbf{g}] = \text{true} \Rightarrow \text{PROVE}_2(\mathbf{g}) = \text{true}$. $\square$

5 From Theory to Practice: CRYPTOVAMPIRE2 for Indistinguishability Goals

In this section, we integrate all previously introduced concepts and instantiate them in the form of CRYPTOVAMPIRE2's indistinguishability solver.

5.1 General Shape of the Proofs

Let p_1 and p_2 be two compatible protocols using steps from $\mathcal{S}$, and $\mathbf{T}$ a trace. Recall from Def. 7 that the goal of CRYPTOVAMPIRE2 is to prove:

$$\text{frame}@_{p_1} \tau \sim \text{frame}@_{p_2} \tau \quad (21)$$

This typically requires induction on τ using $\mathbf{T}$'s $\leq_{\mathcal{S}}$ ordering. Induction is notoriously difficult [21, 42], especially in determining when to initiate it. Nonetheless, CRYPTOVAMPIRE2 *always* first applies induction (via the INDUCTION rule of Fig. 3) and adapts the standard closing procedure to address potential missing assumptions of the induction hypothesis (see Section 5.2.4).

This culminates in CRYPTOVAMPIRE2 running GOLGGE once per step, backtracking over the rules defined in this paper, in order to prove the inductive case. Whenever soundness or performance requires it, GOLGGE offloads to [47]'s SMT solving.

5.2 New and Improved Decision Procedures

Throughout this section, we use Prop. 1's PRF as our leading example. Other cryptographic notions typically behave similarly. We remind the reader of the rule:

PROPERTY (1). When $\mathcal{H}(\_, \_)$ behaves like a Pseudo Random Function [41], the following rule holds:

$$\frac{u \sim v \vdash C[\mathbf{n}] \sim b \quad m, \mathbf{k}, \mathbf{n} \triangleright_{\text{PRF}} C[\_], m}{u \sim v \vdash C[\mathcal{H}(m, \mathbf{k})] \sim b} \text{PRF}$$

with the notations introduced in Section 2's Prop. 1.

In this section, we explore the rule sets designed to utilize rules like PRF. We introduce how to efficiently search for the context $C[\_]$ (Section 5.2.1), how to compute $C[\mathbf{n}]$ from $C[\mathcal{H}(m, \mathbf{k})]$ (Section 5.2.2), and how to solve for $m, \mathbf{k}, \mathbf{n} \triangleright_{\text{PRF}}$ (Section 5.2.3). Finally, we extend the DEDUCE rules [6] (Section 5.2.4) to close goals.

$$\begin{array}{c} \text{INDUCTION} \\ \text{frame}@_{p_1} \text{pred}(\mathbf{a}) \sim \text{frame}@_{p_2} \text{pred}(\mathbf{a}) \vdash \\ \text{frame}@_{p_1} \mathbf{a} \sim \text{frame}@_{p_2} \mathbf{a} \\ \text{for all } \mathbf{a} \in \mathcal{S} \setminus \{\text{init}\} \\ \text{true} \sim \text{true} \vdash \text{frame}@_{p_1} \text{init} \sim \text{frame}@_{p_2} \text{init} \\ \hline \text{frame}@_{p_1} \tau \sim \text{frame}@_{p_2} \tau \end{array}$$

Figure 3: Induction rule

5.2.1 *Looking for Candidates.* Searching for surface-level patterns in e-classes is a solved problem addressed via e-matching [33]. However, rules such as PRF require objects of the form $C[\mathbf{f}(\vec{x})]$, an e-class containing a pattern in one of its children. To find contexts of the form $C[m]$ with $\vec{x} := \text{fv}(m)$, we introduce a function marker. This function interacts with the e-graph using rules:

$$m \rightsquigarrow \text{marker}(\vec{x}, m) \quad (28)$$

$$\mathbf{f}(\vec{u}, \text{marker}(\vec{z}, y), \vec{v}) \rightsquigarrow \text{marker}(\vec{z}, \mathbf{f}(\vec{u}, y, \vec{v})) \quad (29)$$

where $\mathbf{f}$ is a constructor of $\mathcal{S}$. We have Prop. 6 by induction.

PROPERTY 6 (MARKER). In a saturated e-graph $\mathcal{G}$, $\text{marker}(\vec{x}, t) \equiv_{\mathcal{G}} t$ and t is of the form $C[m]$.

PROOF. By induction. $\square$

Such markers allow us to trigger rules efficiently, keeping them simple. Overall, Ex. 16 instantiates this concept for PRF.

EXAMPLE 16 (PRF MARKER). We build $\text{candidate}_{\text{PRF}}(m, \mathbf{k}, t)$ as a marker for the PRF rule to find contexts of the form $C[\mathcal{H}(m, \mathbf{k})]$. We then have the following rewriting rules:

$$\mathcal{H}(m, \mathbf{k}) \rightsquigarrow \text{candidate}_{\text{PRF}}(m, \mathbf{k}, \mathcal{H}(m, \mathbf{k})) \quad (30)$$

$$\mathbf{f}(\vec{x}, \text{candidate}_{\text{PRF}}(m, \mathbf{k}, y), \vec{z}) \rightsquigarrow \text{candidate}_{\text{PRF}}(m, \mathbf{k}, \mathbf{f}(\vec{x}, y, \vec{z})) \quad (31)$$

where $\mathbf{k}$ is a nonce and $\mathbf{f}$ is a constructor of $\mathcal{S}$.

Now, we can simply trigger PRF on:

$$u \sim v \vdash \text{candidate}_{\text{PRF}}(m, \mathbf{k}, a) \sim b \quad (32)$$

PRF still requires transforming $C[\mathcal{H}(m, \mathbf{k})]$ into $C[\mathbf{n}]$.

5.2.2 *Context and Substitutions over General Terms.* Many cryptographic rules feature terms of the form $C[\_]$. Such a *context* C represents “a term with holes”, encoding substitution of arbitrary terms. For instance, in PRF, it is used to convey that one can replace all instances of $\mathcal{H}(m, \mathbf{k})$ by $\mathbf{n}$. However, such a notion is not well-defined for e-classes as shown in Ex. 17.

EXAMPLE 17. Consider $C[\mathcal{H}(m, \mathbf{k})] \equiv \mathcal{H}(m, \mathbf{k}) \equiv \text{If } \mathbf{n} = \mathcal{H}(m, \mathbf{k}) \text{ then } 0 \text{ else } \mathcal{H}(m, \mathbf{k})$. What should $C[\mathbf{n}]$ be?

(1) Taking $C[\_] := \_$ yields $C[\mathbf{n}] := \mathbf{n}$

(2) Taking $C[\_] := (\text{If } \mathbf{n} = \_ \text{ then } 0 \text{ else } \_) \text{ yields } C[\mathbf{n}] := 0$

Already those two options point to two different e-classes for $C[\mathbf{n}]$. Moreover, (1) can be used for PRF but subterm checks reject (2).

The difficulty comes from choosing a C simply from what the markers of Section 5.2.1 highlight. This choice affects the related proofs obligations, notably the subterm checks. Flipping the problem around, the proofs of subterms effectively choose a C for us. Thankfully, GOLGGE's memoization effectively saves the proofs of each predicate that is proven; we can use such proofs to build the new term. Gathering everything together, the rule PRF becomes:

$$\frac{u \sim v \vdash \text{subst}_{\text{PRF}}(m, \mathbf{k}, \mathbf{n}; (m, \mathbf{k}, \mathbf{n} \triangleright_{\text{PRF}} a)) \sim b \quad m, \mathbf{k}, \mathbf{n} \triangleright_{\text{PRF}} a, m}{u \sim v \vdash \text{candidate}_{\text{PRF}}(m, \mathbf{k}, a) \sim b} \text{PRF}$$

A relevant subset of the rules for $m, \mathbf{k}, \mathbf{n} \triangleright_{\text{PRF}} t$ is shown in Fig. 4 giving an intuition on how to construct a $C[\_]$ from a proof.

$$\begin{array}{c}
\text{PRFINST} \\
\hline
m, k, n \triangleright_{PRF} \mathcal{H}(m, k) \\
\\
\text{PRFNONCE} \\
\hline
\frac{\vdash_P k \neq n' \quad \vdash_P n \neq n'}{m, k, n \triangleright_{PRF} n'} \\
\\
\text{PRFNEQM} \\
\hline
\frac{\vdash_P m \neq m' \quad m, k, n \triangleright_{PRF} m'}{m, k, n \triangleright_{PRF} \mathcal{H}(m', k)} \\
\\
\text{PRFFA} \\
\hline
\frac{m, k, n \triangleright_{PRF} t_1 \quad \dots \quad m, k, n \triangleright_{PRF} t_n \quad f \in \mathcal{F} \setminus \{\mathcal{H}\}}{m, k, n \triangleright_{PRF} f(t_1, \dots, t_n)} \\
\\
\text{PRFINPUT} \\
\hline
\frac{\text{VAMP}_{PRF, p}(m, k, T) \quad m \in \{\text{exec}, \text{frame}, \text{cell}(\_)\}}{m, k, n \triangleright_{PRF} m@_p \text{pred}(T)}
\end{array}$$

where n is a nonce and $\text{VAMP}_{PRF, p}(m, k, \_)$ is a built-in decision procedure calling [47].

Figure 4: Relevant simplified subset of $m, k, n \triangleright_{PRF}$ rules

5.2.3 Subterms and Equality. In the CCSA, cryptography-related axioms often translate *semantic* observations (e.g., $a \sim b$ or $\vdash_P a \doteq b$) into *syntactic* facts, typically involving subterm searches [5, 47]. These two modes of reasoning often interact poorly, leading to inconsistencies if not handled carefully, especially with CCSA equality [47]. Here, we avoid such syntactic searches and replace them with other predicates. We remark that these subterm searches implicitly encode proofs of existence for polynomial-time algorithms—potentially with access to specific oracles—that an attacker could use to break an underlying cryptographic assumption. This often comes hand in hand with Section 5.2.2’s context retrieval.

EXAMPLE 18 (SUBTERMS AS PROOF). In PRF, $m, k, n \triangleright_{PRF} t$ means that $\llbracket t \rrbracket_{M, T}^\eta$ is a PTIME algorithm computing t ’s value using a hashing oracle but not k nor n (i.e., with no direct access to the corresponding slices of ρ_h outside of negligible events). Thus PRF’s condition turns a distinguisher between $C[n]$ and $C[\mathcal{H}(m, k)]$ into a valid attacker in a PRF sub-security game. Similarly, both PRF and Th. 1, $n \triangleright t$ proves that $\llbracket t \rrbracket_{M, T}^\eta$ is a PTIME algorithm that does not use n . Refer to Appendix C for formal statements for each rule.

In Fig. 4, PRFNONCE, PRFNEQM, and PRFFA already provide a good intuition for Ex. 18. These rules sometimes produce goals of the form $\vdash_P a \neq b$. This stems from Prop. 4’s underapproximation: we need to call [47] to prove overwhelming disequality of terms as we cannot rely on the incomplete $\equiv_{\mathcal{G}}$. Intuitively, this ensures $\llbracket t \rrbracket_{M, T}^\eta$ never uses forbidden terms with non-negligible probability.

To handle terms like $m@_p \text{pred}(T)$, the naive approach would unfold the macro. However, unfolding depends on the current trace, not just m and T , which is unacceptable since we require proofs to hold for any trace. Fortunately, [5, 47] already provide procedures for computing subterms in such cases, which we reuse. Thus, the goal $\text{VAMP}_{PRF, p}(m, k, T)$ simply offloads the search to [47].

In practice, we also extend the $m, k, n \triangleright_{PRF}$ rules to include a guard condition, like to the bi-deduction rule introduced next.

$$\begin{array}{c}
\text{DEDUCEFA} \\
\hline
\frac{n \triangleright u|h_1 \quad n \triangleright v|h_2}{u, v \triangleright_{\#} n, n|h_1, h_2} \\
\\
\text{DEDUCEITE} \\
\hline
\frac{n \triangleright u|h_1 \quad n \triangleright v|h_2}{u, v \triangleright_{\#} n, n|h_1, h_2} \\
\\
\text{DEDUCENONCE} \\
\hline
\frac{n \triangleright u|h_1 \quad n \triangleright v|h_2}{u, v \triangleright_{\#} n, n|h_1, h_2}
\end{array}$$

where $n \triangleright u|h$ is $n \triangleright u$ under the assumption that h holds.

Figure 5: Selected bi-deduction rules

5.2.4 An Improved Bi-Deduction Procedure. CRYPTOVAMPIRE2 can only close a goal using the following rules:

$$\begin{array}{c}
\text{AXIOM} \\
\hline
u \sim v \vdash u \sim v \\
\\
\text{REFL} \\
\hline
u \sim v \vdash a \sim a \\
\\
\text{DEDUCE} \\
\hline
\frac{u, v \triangleright_{\#} a, b | \text{true}, \text{true}}{u \sim v \vdash a \sim b}
\end{array}$$

The rules AXIOM and REFL are standard. REFL differs lightly from [5], allowing for macros in a as ours are tagged by their protocol.

However, the rule DEDUCE leaves us with a new proof obligation for *bi-deduction* [6]. Bi-deduction proofs are relevant to indistinguishability goals of the form $u \sim v \vdash a \sim b$. They express that a and b do not provide more information than u and v in the sense that we can compute a (resp. b) from u (resp. v) using the *same* polynomial-time algorithm $\mathcal{B}$. In [6], $\mathcal{B}$ is restricted to the attacker’s randomness $\rho_{\mathcal{A}}$, meaning it cannot guess fresh nonces. We relax some of this restriction by extending bi-deduction to allow $\mathcal{B}$ limited but sufficient access to the honest tape ρ_h , culminating in the added rule DEDUCENONCE over the typical SQUIRREL rules.

Fig. 5 presents a relevant subset of the bi-deduction rules implemented in CRYPTOVAMPIRE2. These rules are powerful as they do not require additional context: proving bi-deducibility can be done without knowledge of the rest of the protocol, leading to proofs with few choices and much goal reuse.

Let us formally define our new bi-deduction. Let $\pi_H(\rho)$, with $H \subseteq \mathbb{N}$ and ρ a tape, be the projection of ρ induced by taking its value at each index in H . And we write $\text{stable}_{M, T}^\eta(h, H)$ when the term h is *stable* by $H \subseteq \mathbb{N}$, meaning $\llbracket h \rrbracket_{M, T}^\eta((\rho_h, \rho_{\mathcal{A}}))$ keeps its value when changing arbitrarily many bits of ρ_h whose indices are in H . The notion is formalized in Appendix C.2. We can then formally define our new bi-deduction:

DEFINITION 14 (BI-DEDUCTION). Given $u_1^*, u_2^*, v_1^*, v_2^*, h_1^*, h_2^* \in \mathcal{S}$, we say that $u_1^*, u_2^* \triangleright_{\#} v_1^*, v_2^* | h_1^*, h_2^*$ hold when for all h_1, h_2 such that $h_i \equiv h_i^*$, we can find

- (1) u_1, u_2, v_1, v_2 terms
- (2) a sequence $(H_\eta)_{\eta \in \mathbb{N}}$ of subsets of $\mathbb{N}$ of size N_η ;
- (3) $\mathcal{B}$, a polynomial-time Turing machine

such that

$$v_i \equiv v_i^* \text{ and } u_i \equiv \text{If } h_i \text{ then } u_i^* \text{ else } \emptyset \quad (33)$$

$$\text{stable}_{M, T}^\eta(h_i, H_\eta) \Rightarrow \text{stable}_{M, T}^\eta(u_i, H_\eta) \quad (34)$$

$$\mathcal{B}(1^\eta, \rho_{\mathcal{A}}, \pi_{H_\eta}(\rho_h), \llbracket u_i \rrbracket^\eta(\rho)) = \llbracket v_i \rrbracket^\eta(\rho) \text{ over } R_\eta^i \quad (35)$$

where $R_\eta^i = \{\llbracket h_i \rrbracket^\eta(\rho) = 1\}$ and N_η is a polynomial in η

In less formal terms: $(R_{\eta}^i)_{i \in \mathbb{N}}$ restricts the space to where the hypothesis h_i holds. $\pi_{H_{\eta}}$ (Item 2) selects parts of the random tape irrelevant to the computation (Equation 34). H_{η} notably selects the nonces of v_1 and v_2 not used in u_1 and u_2 . $\mathcal{B}$ and (35) allow computing v_1 (resp. v_2) from u_1 (resp. u_2) with access to the attacker tape and the selected part of the *honest* random tape when h_1 (resp. h_2) holds. DEDUCENONCE and (34) ensure that the selected part of ρ_h is independent from the u_i 's computation, letting us share it with the distinguisher $\mathcal{B}$. This notion yields:

PROPERTY 7 (BI-DEDUCTION). *DEDUCE and DEDUCENONCE hold.*

PROOF. See Appendix C.3. $\square$

5.2.5 *Enriching the Frame.* DEDUCE is sometimes insufficient to overcome the limitation of the weak induction hypothesis given by the INDUCTION rule. This notably happens when DEDUCENONCE encounters a nonce n that is not fresh. This is typically overcome by attempting to induct over (36) instead.

$$\langle \text{frame@}_{p_1} \tau, n \rangle \sim \langle \text{frame@}_{p_2} \tau, n \rangle \quad (36)$$

Although CRYPTOVAMPIRE2 does not support goals of this form, we can emulate such a proof technique by proving a weaker protocol which “leaks” n in *init* or using Rmk. 1's extra *init* steps.

On failure, CRYPTOVAMPIRE2 can retry a proof with a frame enriched with some nonces. It will guide its choices by remembering which nonces blocked the previous proof because of DEDUCENONCE. This approach often finds a suitable set of nonces much faster than basic exhaustive search (Table 1).

5.3 Quantifiers and Variables

One major roadblock toward automating ours and SQUIRREL's peculiar fragment of the CCSA comes in the form of *binders*. As a reminder from Section 3.2's Fig. 2, they come in the form of the lookups $\text{Find } \vec{t}, \vec{\tau} \text{ st. } t \text{ then } t' \text{ else } t''$ and can be used to further define objects similar to the first-order quantifiers $\exists$ and $\forall$ as shown in (13) and (14) respectively. These are higher-order terms which neither of our backends (i.e., EGG and SMT-LIB solvers) natively support. This section explores the transformations required to sufficiently capture their semantics throughout CRYPTOVAMPIRE2.

5.3.1 *In First Order Logic.* [47] already presents and uses results to soundly transform $\exists$ and $\forall$ into their first order logic counterparts. Such equivalence does not exist for general lookups. In [47], lookups are each replaced by free functions leveraging Skolemization [68] whose semantics are then described through added axioms.

EXAMPLE 19 (NAMING OF LOOKUPS). *Let $X := \text{Find } \vec{t} \text{ st. } t \text{ then } t' \text{ else } t''$ be a lookup with $\vec{\alpha} := \text{fv}(X)$. We introduce the free function find_X and the skolems $\vec{\text{sk}}_X$ such that:*

$$\text{find}_X(\vec{\alpha}, \vec{\beta}) := \text{If } t \text{ then } t' \text{ else } t'' \quad (37)$$

$$\forall \vec{\alpha}, \vec{\beta}. t \Rightarrow t \mapsto \vec{\text{sk}}_X(\vec{\alpha}) \quad (38)$$

and can now use $\text{find}_X(\vec{\alpha}, \vec{\text{sk}}_X(\vec{\alpha}))$ in lieu of X in first-order formulas.

However, unlike in [47], we do not know ahead of time which lookups are part of the problem. Moreover, as this is a fairly expensive transformation to the SMT solver, we want to reuse as many of the “named” functions as possible. To that end, CRYPTOVAMPIRE2 implements *caching* and *quantifier unification*.

5.3.2 *In the E-Graph.* One could use the same strategy within the e-graph. This, however, prevents rewriting *within* a binder. Instead, we take inspiration from [65] and enable our e-graph to reason about any binder using Modal De Bruijn indices [32] and Peano arithmetic [40]. We introduce the usual two symbols to represent variables 0_λ and S_λ as the zero and successor function respectively. Thus, for instance, $\exists x, \exists y, g(x, y)$ becomes $\exists_\lambda (\exists_\lambda (g(S_\lambda(0_\lambda), 0_\lambda)))$. Moreover, we make S_λ a Modal Operator, in the sense that it applies to arbitrary terms and not just De Bruijn variables, and it intuitively increments all the variables of its arguments. Effectively, we include the following equality rule:

$$S_\lambda(f(t_1, \dots, t_n)) \rightsquigarrow f(S_\lambda(t_1), \dots, S_\lambda(t_n)) \quad (39)$$

Such encoding has the benefit of making free variables explicit. This is important as Section 5.2.1's marker functions can “extract” a term out of its binder, so can rewritings (e.g., $x \equiv_{\mathcal{G}} \pi_1(\langle x, 0_\lambda \rangle)$). In these cases, soundness requires us to treat those free variables as universally quantified (or give up). Encoding bound variables using De Bruijn indices gives us this for free.

5.4 Soundness

The semantics of CRYPTOVAMPIRE2 are directly grounded in traditional cryptography (Section 3). This has the interesting feature that each rule can be proven separately while still retaining Th. 4.

THEOREM 4 (SOUNDNESS OF CRYPTOVAMPIRE2). *Let p_1 and p_2 be two compatible protocols using steps from $\mathcal{S}$. If CRYPTOVAMPIRE2 returns true, then p_1 and p_2 are indistinguishable according to Def. 7 in the computational model under consideration.*

PROOF. Since all inference and equality rules are proven sound, Th. 3 lets us prove Def. 7's (21). $\square$

6 Benchmarks

6.1 Implementation

CRYPTOVAMPIRE2 is available at [48]. It is implemented in around 26 000 lines of RUST [1] code (1 600 for GOLGGE). It contains a modified version of EGG [70] and a reimplement of [47] to fit the new shape of formulas. User input is handled through the STEEL [66] interpreter for SCHEME [2], to which CRYPTOVAMPIRE2-specific bindings were added. This lets the user easily programmatically define their protocol while keeping most of the advantages of a domain-specific language. CRYPTOVAMPIRE2 offloads some goals to SMT solvers (as in [47]) via SMT-LIB [13], using VAMPIRE [53], z3 [34] and cvc5 [12] in parallel.

6.2 Setup and Limitations

6.2.1 *Experimental Setup.* We used a 24-thread machine equipped with an AMD 9900X processor and 64 GB of RAM, less than 1.4 GiB of which was required for all experiments. We used VAMPIRE 4.9, z3 4.15.8 and cvc5 1.3.3. All solver timeouts were set to 150 ms (the

Protocol		Features Enabled	Enriched Nonce Sets	Memoization Hit Rate	Runtime (s)			Longest SMT Run (ms)			
					GOLGGE	SMT	Total	All	VAMPIRE	z3	CVC5
basic_hash [20] Δ		PRF	-	24.0 %	0.1	1.4	1.6	30.8	96.4	31.0	421.3
hash_lock [49] Δ		PRF	-	27.4 %	0.3	2.1	2.3	79.3	137.9	142.1	X
lak_tag [45] Δ		PRF, FA	-	40.8 %	5.3	9.0	14.4	20.9	454.7	21.1	142.1
mw [61] Δ		PRF, XOR, FA	-	44.7 %	16.0	8.7	24.7	22.9	125.5	27.7	451.0
feldhofer [38] Δ		ENC-KP, FA	1	48.3 %	110.6	19.9	130.5	15.0	34.6	18.4	48.8
private_auth [10] Δ		ENC-KP, IND-CCA, FA	-	35.9 %	0.9	3.9	4.9	<u>19.1</u>	11.7	35.4	X
ddh [46]	base Δ	DDH, FA	3	40.4 %	2.3	11.3	13.6	14.6	32.2	37.6	X
	base $\blacksquare$		3	34.6 %	2.6	14.3	17.0	17.7	34.1	35.8	X
	composition $\Delta^\dagger$		1	23.1 %	1.9	11.8	13.6	<u>28.7</u>	16.5	85.8	X
	composition $\blacksquare^\dagger$		3	25.5 %	15.8	94.1	109.9	20.9	33.8	65.7	X
ssh [55]	part 1 $\Delta^\dagger$	DDH, FA	2	31.5 %	2.4	11.5	13.9	22.0	33.9	38.6	402.8
	part 2 $\Delta^\dagger$		2	31.1 %	5.3	24.6	29.9	28.1	39.3	39.7	X
	part 1 $\blacksquare^\dagger$		-	32.8 %	0.5	4.9	5.4	11.8	29.4	23.4	55.2
	part 2 $\blacksquare^\dagger$		-	15.7 %	0.2	1.5	1.7	9.6	23.4	14.0	19.9
lfmtp $\blacksquare^\ddagger$		PRF	-	41.1 %	0.6	4.2	4.8	11.2	30.0	19.6	68.4
toy_state $\blacksquare^\ddagger$		PRF	-	27.5 %	0.1	0.4	0.5	10.3	32.9	14.5	66.8

 $\ddagger$ stateful $\dagger$ using the composition framework [27]**X** fails to prove indistinguishability regardless of timeout Δ unlinkability query $\blacksquare$ strong secrecy query

Table 1: Summary of the Experiments

lowest allowed by VAMPIRE) for the “Runtime” column of Table 1 and 750 ms⁷ for the “Longest SMT Run” column.

6.2.2 Description of the Benchmarks. We evaluated CRYPTOVAMPIRE2 on all supported indistinguishability goals from SQUIRREL’s benchmarks [5, 6]. This means all queries that could be converted and expressed as $\text{frame}@_{p_1}t \sim \text{frame}@_{p_2}t$. This includes stateful protocols and protocols using the composition framework [27], but it notably excludes yubikey-related lemmas. Overall, they consist of unlinkability and anonymity properties in the basic_hash [20], hash_lock [49], lak_tag [45], feldhofer [38], mw [61], private_auth [10], ddh [46] and ssh [55] protocols; strong secrecy of the key in ddh, ssh; and two toyish stateful protocols from SQUIRREL (lfmtp and toy_state). We closely follow the modeling in [5, 6].

All experimental results are summarized in Table 1. It captures the sort of queries tested (unlinkability or strong secrecy) along with the features required of CRYPTOVAMPIRE2 to prove the query. That is, the notable cryptographic assumptions (see Appendix C for detailed rules) and proof techniques required (e.g., usage of Section 5.2.5, and to which extent). We further recorded the running times split between the SMT solver and GOLGGE (i.e., the rest of CRYPTOVAMPIRE2), and the longest time taken by each SMT backend to prove a query.

6.2.3 Limitations. To automate indistinguishability proofs, CRYPTOVAMPIRE2 sometimes requires *trace* lemmas as input. In particular, all stateless protocols except private_auth require one trace lemma capturing the authentication properties expected by their

specification (e.g., Lem. 1 for basic_hash), while both stateful protocols require secrecy lemmas about their state. These lemmas constrain the search space and preserve execution alignment during the proof. They are not arbitrary proof artifacts, but security goals of the protocol specification. While synthesizing such invariants is future work, tools like [47] can already verify the authentication lemmas when provided, and SQUIRREL can verify all of them.

As mentioned above, CRYPTOVAMPIRE2 is currently strictly constrained to bi-processes. Effectively, this restricts our support for stateful protocols, excluding 11 indistinguishability goals from [6]. These goals are typically used as lemmas for broader trace properties.

Furthermore, here we consider only examples from [5, 6]. We do not take into account later extensions made to SQUIRREL.

6.3 Analysis

6.3.1 General Runtime Analysis. Despite the short timeout, we observe that almost all the runtime is spent waiting for the SMT solvers. Deeper analysis reveals that most of it is spent on impossible queries waiting for the solvers to time out. This underscores the importance of the longest SMT query in CRYPTOVAMPIRE2’s runtime, as this dictates the timeout that can be given to the SMT backend. Moreover, we observe that the proportion of time spent on SMT proving reduces as the total runtime increases. This correlates with increased memoization: as time goes on, SMT queries, or goals that would lead to some, are cached.

6.3.2 Backend Comparison. We compare the SMT backends (each solver individually, then all together) on their ability to terminate quickly on all successful queries. This metric is crucial, as it determines the minimum timeout required for a configuration to prove

⁷chosen such that the result of the overall proof is not influenced by the timeout

indistinguishability, which in turn has a significant impact on the total runtime. We manually verified that failing cvc5 runs were not due to timeouts, but rather to incomplete solver strategies. The fastest single solver is marked in **bold** in Table 1.

Interestingly, VAMPIRE alone sometimes outperforms racing all solvers in parallel (including VAMPIRE itself) for a result. This paradoxical outcome is explained by a different resource allocation when all solvers are enabled. Such cases are marked with underlines.

One should note, however, that CRYPTOVAMPIRE2 runs the solvers in close-to-default configurations. Their performance could likely be significantly improved with further tuning.

Overall, VAMPIRE is more often the fastest, while z3 is not far behind with more consistent runtimes. The poor performance of cvc5 contrasts with the results of [47]. This highlights the high volatility of SMT solvers, which justifies why CRYPTOVAMPIRE2 does not rely on a single solver.

7 Conclusion and Future Work

This paper has introduced CRYPTOVAMPIRE2, the first auto-active verifier⁸ for indistinguishability in the computational model using the CCSA framework. To that end, we also introduced GOLGGE—a novel backtracking engine supported by e-graphs—which enables a more relaxed approach to equality, key to CRYPTOVAMPIRE2’s success. CRYPTOVAMPIRE2 was furthermore tested against the examples of the state-of-the-art proof assistant SQUIRREL [5], successfully proving them all.

Future work will focus on extending CRYPTOVAMPIRE2 to advanced CCSA techniques such as quantum soundness [29]. Beyond such extensions, we aim to propagate the new equality reasoning philosophy to [47], alongside various improved heuristics and optimizations for CRYPTOVAMPIRE2 and GOLGGE.

Acknowledgments

This work was partially supported by the European Research Council (ERC) through the Consolidator Grants ARTIST 101002685 and BlockSec 101141432; the TU Wien Doctoral College SecInt; the Austrian Science Fund (FWF) through the SpyCoDe SFB projects F8504 and F8510-N; the Vienna Science and Technology Fund (WWTF) through [ForSmart Grant ID: 10.47379/ICT22007]; and the Austrian Research Promotion Agency (FFG) through the COMET K1 SBA.

This paper was edited for grammar and phrasing using Mistral, Gemini, ChatGPT, and various versions of Qwen. All outputs were inspected by the authors to ensure accuracy and originality.

References

- [1] [n. d.]. *Rust Programming Language*. <https://rust-lang.org/>
- [2] [n. d.]. *The Scheme Programming Language*. <https://www.scheme.org/>
- [3] Nadhem J. Al Fardan and Kenneth G. Paterson. 2013. Lucky Thirteen: Breaking the TLS and DTLS Record Protocols. In *2013 IEEE Symposium on Security and Privacy* (2013-05). 526–540. doi:10.1109/SP.2013.42
- [4] Thibaut Antoine and David Baelde. 2025. Propositional Logics of Overwhelming Truth. In *CSL 2025 - 33rd EACSL Annual Conference on Computer Science Logic* (Amsterdam, Netherlands, 2025-02). Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPICs.CSL.2025.24
- [5] David Baelde, Stéphanie Delaune, Charlie Jacomme, Adrien Koutsos, and Solène Moreau. 2021. An Interactive Prover for Protocol Verification in the Computational Model. In *2021 IEEE Symposium on Security and Privacy (SP)* (San Francisco / Virtual, United States, 2021-05). 537–554. doi:10.1109/SP40001.2021.00078
- [6] David Baelde, Stéphanie Delaune, Adrien Koutsos, and Solène Moreau. 2022. Cracking the Stateful Nut: Computational Proofs of Stateful Security Protocols Using the Squirrel Proof Assistant. In *2022 IEEE 35th Computer Security Foundations Symposium (CSF)* (Haifa, Israel, 2022-08). 289–304. doi:10.1109/CSF54842.2022.9919665
- [7] David Baelde, Stéphanie Delaune, and Stanislas Riou. 2025. SMT-based Automation for Overwhelming Truth. In *38th IEEE Computer Security Foundations Symposium* (Santa Cruz, United States, 2025-06). IEEE, 505–520. doi:10.1109/CSF64896.2025.00019
- [8] David Baelde, Adrien Koutsos, and Justine Sauvage. 2024. Foundations for Cryptographic Reductions in CCSA Logics. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2024-12-09) (CCS ’24). Association for Computing Machinery, 2814–2828. doi:10.1145/3658644.3690193
- [9] Gergei Bana and Hubert Comon-Lundh. 2012. Towards Unconditional Soundness: Computationally Complete Symbolic Attacker. In *Principles of Security and Trust* (Berlin, Heidelberg, 2012) (*Lecture Notes in Computer Science*), Pierpaolo Degano and Joshua D. Guttman (Eds.). Springer, 189–208. doi:10.1007/978-3-642-28641-4_11
- [10] Gergei Bana and Hubert Comon-Lundh. 2014. A Computationally Complete Symbolic Attacker for Equivalence Properties. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (2014). 609–620.
- [11] Gergei Bana, Koji Hasebe, and Mitsuhiro Okada. 2013. Computationally Complete Symbolic Attacker and Key Exchange. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security* (2013) (CCS ’13). Association for Computing Machinery, 1231–1246. doi:10.1145/2508859.2516710
- [12] Haniel Barbosa, Clark W. Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, Alex Ozdemir, Mathias Preiner, Andrew Reynolds, Ying Sheng, Cesare Tinelli, and Yoni Zohar. 2022. Cvc5: A Versatile and Industrial-Strength SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems - 28th International Conference, TACAS 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings, Part I* (2022) (*Lecture Notes in Computer Science*, Vol. 13243), Dana Fisman and Grigore Rosu (Eds.). Springer, 415–442. doi:10.1007/978-3-030-99524-9_24
- [13] Clark Barrett, Pascal Fontaine, and Cesare Tinelli. 2017. The SMT-LIB Standard: Version 2.6.
- [14] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. 2011. Computer-Aided Security Proofs for the Working Cryptographer. In *Advances in Cryptology – CRYPTO 2011* (Berlin, Heidelberg, 2011) (*Lecture Notes in Computer Science*), Phillip Rogaway (Ed.). Springer, 71–90. doi:10.1007/978-3-642-22792-9_5
- [15] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Santiago Zanella-Béguelin. 2014. Proving the TLS Handshake Secure (As It Is). In *Advances in Cryptology – CRYPTO 2014* (2014). Springer, Berlin, Heidelberg, 235–255. doi:10.1007/978-3-662-44381-1_14
- [16] Karthikeyan Bhargavan and Gaetan Leurent. 2016. Transcript Collision Attacks: Breaking Authentication in TLS, IKE, and SSH. In *Proceedings 2016 Network and Distributed System Security Symposium* (San Diego, CA, 2016). Internet Society. doi:10.14722/ndss.2016.23418
- [17] Bruno Blanchet. 2001. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *14th IEEE Computer Security Foundations Workshop (CSFW-14)* (Cape Breton, Nova Scotia, Canada, 2001-06). IEEE Computer Society, 82–96. doi:10.1109/CSFW.2001.930138
- [18] B. Blanchet. 2006. A Computationally Sound Mechanized Prover for Security Protocols. In *2006 IEEE Symposium on Security and Privacy (S&P’06)* (2006-05). 15 pp.–154. doi:10.1109/SP.2006.1
- [19] Bruno Blanchet, Martin Abadi, and Cédric Fournet. 2008. Automated Verification of Selected Equivalences for Security Protocols. 75, 1 (2008), 3–51. doi:10.1016/j.jlap.2007.06.002
- [20] Mayla Brusó, Konstantinos Chatzikokolakis, and Jerry den Hartog. 2010. Formal Verification of Privacy for RFID Systems. In *2010 23rd IEEE Computer Security*

⁸I.e., a largely automated verifier requiring some user input, in our case, some lemmas.

- Foundations Symposium* (2010-07). 75–88. doi:10.1109/CSF.2010.13
- [21] Alan Bundy. 2001. Chapter 13 - The Automation of Proof by Mathematical Induction. In *Handbook of Automated Reasoning*, Alan Robinson and Andrei Voronkov (Eds.). North-Holland, 845–911. doi:10.1016/B978-044450813-3/50015-1
- [22] Stefano Ceri, Georg Gottlob, and Letizia Tanca. 1989. What You Always Wanted to Know about Datalog (and Never Dared to Ask). 1, 1 (1989), 146–166. doi:10.1109/69.43410
- [23] Weidong Chen and David S. Warren. 1996. Tabled Evaluation with Delaying for General Logic Programs. 43, 1 (1996), 20–74. doi:10.1145/227595.227597
- [24] Vincent Cheval and Bruno Blanchet. 2013. Proving More Observational Equivalences with ProVerif. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 226–246. doi:10.1007/978-3-642-36830-1_12
- [25] Michael R. Clarkson and Fred B. Schneider. 2008. Hyperproperties. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium, CSF 2008, Pittsburgh, Pennsylvania, USA, 23-25 June 2008* (2008). IEEE Computer Society, 51–65. doi:10.1109/CSF.2008.7
- [26] Alain Colmerauer, Henri Kanoui, Robert Pasero, and Philippe Rousset. 1973. Un système de communication homme-machine en français. (1973).
- [27] Hubert Comon, Charlie Jaccome, and Guillaume Scerri. 2020. Oracle Simulation: A Technique for Protocol Composition with Long Term Shared Secrets. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (Virtual Event USA, 2020-10-30). ACM, 1427–1444. doi:10.1145/3372297.3417229
- [28] Hubert Comon and Adrien Koutsos. 2017. *Formal Computational Unlinkability Proofs of RFID Protocols*. arXiv:1705.02296 [cs] doi:10.48550/arXiv.1705.02296
- [29] Cas Cremers, Caroline Fontaine, and Charlie Jaccome. 2022. A Logic and an Interactive Prover for the Computational Post-Quantum Security of Protocols. In *2022 IEEE Symposium on Security and Privacy (SP)* (2022-05). 125–141. doi:10.1109/SP46214.2022.9833800
- [30] Baelde David, Adrien Koutsos, and Joseph Lallemand. 2023. A Higher-Order Indistinguishability Logic for Cryptographic Reasoning. In *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)* (2023-06). 1–13. doi:10.1109/LICS56636.2023.10175781
- [31] Martin Davis, George Logemann, and Donald Loveland. 1962. A Machine Program for Theorem-Proving. 5, 7 (1962), 394–397. doi:10.1145/368273.368557
- [32] Nicolaas Govert de Bruijn. 1972. Lambda Calculus Notation with Nameless Dummies, a Tool for Automatic Formula Manipulation, with Application to the Church-Rosser Theorem. 75, 5 (1972), 381–392. doi:10.1016/1385-7258(72)90034-0
- [33] Leonardo de Moura and Nikolaj Bjørner. 2007. Efficient E-Matching for SMT Solvers. In *Automated Deduction – CADE-21* (Berlin, Heidelberg, 2007), Frank Pfenning (Ed.). Springer, 183–198. doi:10.1007/978-3-540-73595-3_13
- [34] Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems* (Berlin, Heidelberg, 2008) (*Lecture Notes in Computer Science*), C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer, 337–340. doi:10.1007/978-3-540-78800-3_24
- [35] Stéphanie Delaune, Clément Herouard, and Joseph Lallemand. 2025. Secrecy by Typing in the Computational Model. 17–32. doi:10.1109/CSF64896.2025.00006
- [36] D. Dolev and A. Yao. 1983. On the Security of Public Key Protocols. 29, 2 (1983), 198–208. doi:10.1109/TIT.1983.1056650
- [37] Clemens Eisenhofer, Ruba Allassaf, Michael Rawson, and Laura Kovács. 2023. Non-Classical Logics in Satisfiability Modulo Theories. In *Automated Reasoning with Analytic Tableaux and Related Methods* (Cham, 2023) (*Lecture Notes in Computer Science*), Revantha Ramanayake and Josef Urban (Eds.). Springer Nature Switzerland, 24–36. doi:10.1007/978-3-031-43513-3_2
- [38] Martin Feldhofer, Sandra Dominikus, and Johannes Wolkerstorfer. 2004. Strong Authentication for RFID Systems Using the AES Algorithm. In *Cryptographic Hardware and Embedded Systems – CHES 2004* (Berlin, Heidelberg, 2004) (*Lecture Notes in Computer Science*), Marc Joye and Jean-Jacques Quisquater (Eds.). Springer, 357–370. doi:10.1007/978-3-540-28632-5_26
- [39] Joshua Gancher, Sydney Gibson, Pratap Singh, Samvid Dharanikota, and Bryan Parno. 2023. Owl: Compositional Verification of Security Protocols via an Information-Flow Type System. In *2023 IEEE Symposium on Security and Privacy (SP)* (2023-05). 1130–1147. doi:10.1109/SP46215.2023.10179477
- [40] Giuseppe Peano. 1889. *Arithmetices principia: nova methodo*. Fratres Bocca.
- [41] Oded Goldreich. 2007. *Foundations of Cryptography*. Vol. 1: Basic Tools. Cambridge Univ. Press.
- [42] Márton Hajdú, Petra Hozzová, Laura Kovács, Johannes Schoisswohl, and Andrei Voronkov. 2020. Induction with Generalization in Superposition Reasoning. In *Intelligent Computer Mathematics* (Cham, 2020) (*Lecture Notes in Computer Science*), Christoph Benz Müller and Bruce Miller (Eds.). Springer International Publishing, 123–137. doi:10.1007/978-3-030-53518-6_8
- [43] Jacques Herbrand. 1930. *Recherches sur la théorie de la démonstration*.
- [44] Edwin Hewitt and Karl Stromberg. 1965. Integration on Product Spaces. In *Real and Abstract Analysis: A Modern Treatment of the Theory of Functions of a Real Variable*, Edwin Hewitt and Karl Stromberg (Eds.). Springer, 377–459. doi:10.1007/978-3-642-88044-5_6
- [45] Lucca Hirschi, David Baelde, and Stéphanie Delaune. 2016. A Method for Verifying Privacy-Type Properties: The Unbounded Case. In *2016 IEEE Symposium on Security and Privacy (SP)* (2016-05). 564–581. doi:10.1109/SP.2016.40
- [46] ISO/IEC 9798-3:2019 2019. IT Security Techniques – Entity Authentication – Part 3: Mechanisms Using Digital Signature Techniques.
- [47] Simon Jeanteur, Laura Kovács, Matteo Maffei, and Michael Rawson. 2024. CryptoVampire: Automated Reasoning for the Complete Symbolic Attacker Cryptographic Model. IEEE Computer Society, 3165–3183. doi:10.1109/SP54263.2024.00246
- [48] Simon Jeanteur, Laura Kovács, Matteo Maffei, and Michael Rawson. 2026. *Crypto-vampire*. <https://github.com/SecPriv/CryptoVampire>
- [49] Ari Juels and Stephen A. Weis. 2009. Defining Strong Privacy for RFID. 13, 1 (2009), 7:1–7:23. doi:10.1145/1609956.1609963
- [50] Olav Kallenberg. 2021. *Foundations of Modern Probability*. Probability Theory and Stochastic Modelling, Vol. 99. Springer International Publishing. doi:10.1007/978-3-030-61871-1
- [51] Donald E. Knuth and Petter B. Bendix. 1970. Simple Word Problems in Universal Algebras. In *Computational Problems in Abstract Algebra*, John Leech (Ed.). Pergamon, 263–297. doi:10.1016/B978-0-08-012975-4.50028-X
- [52] Adrien Koutsos. 2019. The 5G-AKA Authentication Protocol Privacy. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)* (2019-06). 464–479. doi:10.1109/EuroSP.2019.00041
- [53] Laura Kovács and Andrei Voronkov. 2013. First-Order Theorem Proving and Vampire. In *Computer Aided Verification* (Berlin, Heidelberg, 2013) (*Lecture Notes in Computer Science*), Natasha Sharygina and Helmut Veith (Eds.). Springer, 1–35. doi:10.1007/978-3-642-39799-8_1
- [54] Benjamin Lipp, Bruno Blanchet, and Karthikeyan Bhargavan. 2019. A Mechanised Cryptographic Proof of the WireGuard Virtual Private Network Protocol. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)* (2019-06). 231–246. doi:10.1109/EuroSP.2019.00026
- [55] Chris M. Lonvick and Tatu Ylonen. 2006. The Secure Shell (SSH) Transport Layer Protocol. doi:10.17487/RFC4253
- [56] Gavin Lowe. 1995. An Attack on the Needham-Schroeder Public-Key Authentication Protocol. 56, 3 (1995), 131–133. doi:10.1016/0020-0190(95)00144-2
- [57] G. Lowe. 1997. A Hierarchy of Authentication Specifications. In *Proceedings 10th Computer Security Foundations Workshop* (1997-06). 31–43. doi:10.1109/CSFW.1997.596782
- [58] Giulio Malavolta, Pedro Moreno-Sanchez, Clara Schneidewind, Aniket Kate, and Matteo Maffei. 2019. Anonymous Multi-Hop Locks for Blockchain Scalability and Interoperability. In *Proceedings 2019 Network and Distributed System Security Symposium* (San Diego, CA, 2019). Internet Society. doi:10.14722/ndss.2019.23330
- [59] J.P. Marques Silva and K.A. Sakallah. 1996. GRASP-A New Search Algorithm for Satisfiability. In *Proceedings of International Conference on Computer Aided Design* (1996-11). 220–227. doi:10.1109/ICCAD.1996.569607
- [60] Simon Meier, Benedikt Schmidt, Cas Cremers, and David Basin. 2013. The TAMARIN Prover for the Symbolic Analysis of Security Protocols. In *Computer Aided Verification* (Berlin, Heidelberg, 2013) (*Lecture Notes in Computer Science*), Natasha Sharygina and Helmut Veith (Eds.). Springer, 696–701. doi:10.1007/978-3-642-39799-8_48
- [61] David Molnar and David Wagner. 2004. Privacy and Security in Library RFID: Issues, Practices, and Architectures. In *Proceedings of the 11th ACM Conference on Computer and Communications Security* (New York, NY, USA, 2004-10-25) (CCS '04). Association for Computing Machinery, 210–219. doi:10.1145/1030083.1030112
- [62] Charles Gregory Nelson. 1980. *Techniques for Program Verification*. Stanford University.
- [63] Greg Nelson and Derek C. Oppen. 1980. Fast Decision Procedures Based on Congruence Closure. 27, 2 (1980), 356–364. doi:10.1145/322186.322198
- [64] Jens Otten. 2008. leanCoP 2.0 and leanCoP 1.2: High Performance Lean Theorem Proving in Classical and Intuitionistic Logic (System Descriptions). In *Automated Reasoning, 4th International Joint Conference, IJCAR 2008, Sydney, Australia, August 12-15, 2008, Proceedings* (2008) (*Lecture Notes in Computer Science*, Vol. 5195), Alessandro Armando, Peter Baumgartner, and Gilles Dowek (Eds.). Springer, 283–291. doi:10.1007/978-3-540-71070-7_23
- [65] Pavel Panchekha. 2021. *Binding in E-graphs*. <https://pavpanchekha.com/blog/egg-bindings.html>
- [66] Matthew Paras. 2025. *Mattwparas/Steel*. <https://github.com/mattwparas/steel>
- [67] Sadahiro Saeki. 1996. A Proof of the Existence of Infinite Product Probability Measures. 103, 8 (1996), 682–683. jstor:2974880 doi:10.2307/2974880
- [68] Thoralf Skolem. 1920. Logisch-Kombinatorische Untersuchungen Über Die Erfüllbarkeit Oder Beweisbarkeit Mathematischer Sätze Nebst Einem Theorem Über Dichte Mengen. In *Selected Works in Logic*.
- [69] Alfred Tarski. 1933. Pojęcie Prawdy w Językach Nauk Dedukcyjnych. (1933).
- [70] Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panchekha. 2021. Egg: Fast and Extensible Equality Saturation. 5, Article 23 (2021). Issue POPL. doi:10.1145/3434304
- [71] Yihong Zhang, Yisu Remy Wang, Oliver Flatt, David Cao, Philip Zucker, Eli Rosenthal, Zachary Tatlock, and Max Willsey. 2023. Better Together: Unifying Datalog and Equality Saturation. 7 (2023), 125:468–125:492. Issue PLDI. doi:10.1145/3591239

Open Science

All code, benchmarks and instructions on how to use them are available at [48].

Ethical Considerations

The research presented in this paper focuses on the development of formal verification techniques to improve the security and privacy of cryptographic protocols. As a tool for automated theorem proving, its primary purpose is to provide rigorous guarantees of security, thereby contributing to the robustness of digital infrastructure. This research does not involve human subjects, personal data, or real-world vulnerability analysis of deployed systems. Our experimental evaluation relies on established formal models of protocols and standard cryptographic assumptions.

Unlike many protocol verifiers, CRYPTOVAMPIRE2 cannot find attacks. It could nonetheless still be used by nefarious actors to build highly resilient systems for malicious purposes. We believe, however, that the benefits of lowering the barrier to building secure protocol standards and implementations for the general public far outweigh these potential harms.

Generative AI Usage

Generative AI was used for editorial purposes on this paper, to orchestrate the testing (the tests themselves were written by hand). They were used to very little success within CRYPTOVAMPIRE2, hence the code itself was overwhelmingly written by hand.

In all cases, their output was *never* blindly trusted, and we always thoroughly checked for correctness.

A Spaces, Measures and Notations

A.1 Space

We denote by $B = \{0, 1\}$ the set of symbols used by our Turing Machines; for convenience, we set it to the booleans. Then B^* is the set of words on B or *bitstrings*. Finally, $B^{\mathbb{N}}$ is the set of sequences on B ; these are Turing machine *tapes*. Notably, $B^{\mathbb{N}}$ can be given a nice uniform probability measure based on the product of B s [44, 67]. Similarly, we can build an Ω with at least countably many uniformly random independent tapes.

A.2 Notations

We use P as a probability measure, and P_X denotes the conditional measure. With w, x, y and z random elements, we use $x \prod_X y$ for the independence of x and y conditioned on X and $w \stackrel{a.e.}{=}_X z$ for equality almost everywhere over X between w and z . We will omit X when $X = \Omega$. We denote by X^c the complement $\Omega \setminus X$ of X .

Throughout this document, $\rho = (\rho_h, \rho_A)$ is a pair of uniformly and independently sampled random elements of $B^{\mathbb{N}}$. They are respectively the honest and attacker *random tapes* used to model randomized Turing Machines.

We gather here some general properties and lemmas relevant throughout the document, but to no particular section.

PROPERTY 8 (MEASURABILITY). *For all cryptographic models M , a BC term t , and security parameter η , the function $x \mapsto \llbracket t \rrbracket_{bc}^M(1^\eta, x)$ is measurable.*

PROOF. We first notice, by induction on t , that $\eta, x \mapsto \llbracket t \rrbracket_{bc}^M(1^\eta, x)$ is computable in polynomial time in η . This notably implies that there exists a polynomial p_t such that for each η the domain of $x \mapsto \llbracket t \rrbracket_{bc}^M(1^\eta, x)$ is $B^{p_t(\eta)}$ which is finite. $\square$

REMARK 5. *This result extends with the same proof to general Polynomial Time Turing Machines.*

This shows that Def. 1's (11) and Prop. 3's (23) are well defined.

In general, we refer to [50] for most of the probability results used during the proofs.

For $H \subset \mathbb{N}$ finite and $\rho_0 = (b_n)_{n \in \mathbb{N}} \in B^{\mathbb{N}}$, let $\pi_H(\rho_0) \in B^{|H|}$ be the projection of ρ_0 to $B^{|H|}$ done by selecting b_i for each $i \in H$. Importantly, it verifies Prop. 9.

PROPERTY 9. *$\pi_H(\rho_h)$ is uniform over $B^{|H|}$.*

B Complete semantics

We assume that the following functions are part of $\mathcal{F}$: If $\_$ then $\_$, else $\_$, $\ll$, $\_$, $\gg$, $\emptyset$, true, false, and fail.

DEFINITION 15 (PREVIOUS STEP). *Given $\mathbb{T}$ and a step $a \in \mathbb{S} \cup \{\text{undef}\}$, we define the previous step $\text{pred}_{\mathbb{T}}(a)$ as follows:*

$$\text{pred}_{\mathbb{T}}(a) := \begin{cases} \text{init} & \text{if } a \text{ is init} \\ \text{undef} & \text{if } a \text{ is undef} \\ \max_{\leq_S} \{b \mid b \in \mathbb{S}, b <_S a\} & \text{otherwise} \end{cases} \quad (40)$$

DEFINITION 16 (UNFOLDING). *Given a trace $\mathbb{T}$, we unfold terms of $\mathcal{S}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C})$ into $\mathcal{BC}(\mathcal{N}_{bc}, \mathcal{F}_{bc}, \mathcal{G}_{bc})$ according to Fig. 6, with*

$$\mathcal{N}_{bc} := \{n_i \mid n \in \mathcal{N}, i \in \mathbb{I}^{\llbracket \cdot \rrbracket}\} \quad (41)$$

$$\mathcal{F}_{bc} := \{f_i \mid f \in \mathcal{F}, i \in \mathbb{I}^{\llbracket \cdot \rrbracket}\} \quad (42)$$

$$\mathcal{G}_{bc} := \{\text{att}\} \quad (43)$$

C Cryptographic inference rules

C.1 Generalities

Appendix C provides proofs of the cryptography-related rules used within CRYPTOVAMPIRE2.

As stated in Section 5.2.3, subterm relations hide proofs of the existence of Turing machines. We explore this fact in Section C.5 and in cases where our rules diverge from the rest of the literature (e.g., Section C.3). We will, however, refer to [5] for the rest of the rules and simply highlight how to adapt the subterm relations.

All subterm relations follow a similar structure to that shown in Fig. 7 even though they can arbitrarily diverge from it. In the following section, we will only describe subterm relations in terms of their differences from Fig. 7.

REMARK 6. *The variants of SUBREC and SUBRECCCELL are, in practice, decided by a call to [47] instead of GOLGGE. This was illustrated by Section 5.2.3's PRFINPUT's uses of VAMPPrf_{m,p}(m, k, $\_$). Since these sections focus on correctness rather than implementation details, we will continue as though SUBREC and SUBRECCCELL were regular GOLGGE rules.*

$$\begin{aligned}
(i)^{\mathbb{T}} &:= \sigma_{\mathbb{T}}(i) & (a[\vec{i}])^{\mathbb{T}} &:= \begin{cases} a[\vec{i}^{\mathbb{T}}] & \text{if it is in } \mathcal{S} \\ \text{undef} & \text{otherwise} \end{cases} \\
(f[\vec{i}](\vec{i}))^{\mathbb{T}} &:= f_{(\vec{i})}^{\mathbb{T}}(\vec{i}^{\mathbb{T}}) & (n[\vec{i}])^{\mathbb{T}} &:= n_{(\vec{i})}^{\mathbb{T}} \\
(c[\vec{j}])^{\mathbb{T}} &:= c[\vec{j}^{\mathbb{T}}] & (\text{pred}(T))^{\mathbb{T}} &:= \overline{\text{pred}_{\mathbb{T}}}(T)^{\mathbb{T}} \\
(\text{happens}(T))^{\mathbb{T}} &:= \delta((T)^{\mathbb{T}} \neq \text{undef}) \\
(T < T')^{\mathbb{T}} &:= \delta((T)^{\mathbb{T}} <_{\mathcal{S}} (T')^{\mathbb{T}}) \\
(T = T')^{\mathbb{T}} &:= \delta((T)^{\mathbb{T}} = (T')^{\mathbb{T}}) & (i = j)^{\mathbb{T}} &:= \delta((i)^{\mathbb{T}} = (j)^{\mathbb{T}}) \\
(\text{cell}(C) @_P T)^{\mathbb{T}} &:= \overline{\text{cell}}((C)^{\mathbb{T}})_{P, \mathbb{T}}((T)^{\mathbb{T}}) \\
(M @_P T)^{\mathbb{T}} &:= \overline{M}_{P, \mathbb{T}}((T)^{\mathbb{T}}) \\
\text{With the pseudo Dirac function } \delta \text{ defined as} \\
\delta(P) &:= \begin{cases} \text{true} & \text{if } P \\ \text{false} & \text{otherwise} \end{cases} \\
(\text{Find } \vec{i}, \vec{\tau} \text{ st. } t_1 \text{ then } t_2 \text{ else } t_3)^{\mathbb{T}} &:= \\
\text{If } (t_1\{\vec{i}, \vec{\tau} \mapsto \vec{i}_1, \vec{\tau}_1\})^{\mathbb{T}} \text{ then } (t_2\{\vec{i}, \vec{\tau} \mapsto \vec{i}_1, \vec{\tau}_1\})^{\mathbb{T}} \text{ else } \dots \\
\text{If } (t_1\{\vec{i}, \vec{\tau} \mapsto \vec{i}_n, \vec{\tau}_n\})^{\mathbb{T}} \text{ then } (t_2\{\vec{i}, \vec{\tau} \mapsto \vec{i}_n, \vec{\tau}_n\})^{\mathbb{T}} \text{ else } (t_3)^{\mathbb{T}} \\
\text{where } \{\vec{i}_1, \dots, \vec{i}_n\} = \mathbb{I}^{|\vec{i}|} \text{ and } \{\vec{\tau}_1, \dots, \vec{\tau}_k\} = \mathbb{S}^{|\vec{\tau}|}. \\
\overline{\text{input}}_{P, \mathbb{T}}(\text{init}), \overline{M}_{P, \mathbb{T}}(\text{undef}) &:= \emptyset \text{ with } M \in \\
\{\text{input}, \text{frame}, \text{cell}(\_)\} & \quad \overline{\text{exec}}_{P, \mathbb{T}}(\text{undef}) := \text{false} \\
\overline{\text{exec}}_{P, \mathbb{T}}(\text{init}) &:= \text{true} \quad \overline{\text{frame}}_{P, \mathbb{T}}(\text{init}) := (\text{msg}_P(\text{init}))^{\mathbb{T}} \\
\overline{\text{input}}_{P, \mathbb{T}}(T) &:= (\text{att}(\text{frame}@_P \text{pred}(T)))^{\mathbb{T}} \\
\overline{\text{exec}}_{P, \mathbb{T}}(T) &:= (\text{cond}@_P T \wedge \text{exec}@_P \text{pred}(T))^{\mathbb{T}} \\
\overline{\text{msg}}_{P, \mathbb{T}}(T) &:= (\text{msg}_P(T))^{\mathbb{T}} \quad \overline{\text{cond}}_{P, \mathbb{T}}(T) := (\text{cond}_P(T))^{\mathbb{T}} \\
\overline{\text{cell}}(c)_{P, \mathbb{T}}(T) &:= (\text{cell}_P(c, T))^{\mathbb{T}} \\
\overline{\text{frame}}_{P, \mathbb{T}}(T) &:= \left(\begin{array}{l} \ll \text{exec}@_P T; \\ \ll \text{If } \text{exec}@_P T \text{ then } \text{msg}@_P T \text{ else fail;} \\ \text{frame}@_P \text{pred}(T) \gg \gg \end{array} \right)^{\mathbb{T}}
\end{aligned}$$

Figure 6: Unfolding of $\mathcal{S}(\mathcal{P}, \mathcal{S}, \mathcal{N}, \mathcal{F}, \mathcal{C})$

$$\begin{aligned}
& \frac{\text{SUBFALSE} \quad \vdash_P \neg h}{\triangleright_0 m|h} \quad \frac{\text{SUBFA} \quad \triangleright_0 x_1|h \quad \dots \quad \triangleright_0 x_n|h}{\triangleright_0 f(x_1, \dots, x_n)|h} \\
& \frac{\text{SUBITE} \quad \triangleright_0 c|h \quad \triangleright_0 l|h \wedge c \quad \triangleright_0 r|h \wedge \neg c}{\triangleright_0 \text{If } c \text{ then } l \text{ else } r|h} \quad \frac{\text{SUBAND} \quad \triangleright_0 a|h \quad \triangleright_0 b|h \wedge a}{\triangleright_0 a \wedge b|h} \\
& \frac{\text{SUBFIND} \quad \triangleright_0 c\{\vec{\alpha} \mapsto \vec{\alpha}\}|h \quad \triangleright_0 l\{\vec{\alpha} \mapsto \vec{\alpha}\}|h \wedge c\{\vec{\alpha} \mapsto \vec{\alpha}\} \quad \triangleright_0 r|h}{\triangleright_0 \text{Find } \vec{\alpha} \text{ st. } c \text{ then } l \text{ else } r|h} \\
& \frac{\text{SUBREC} \quad \triangleright_0 \text{msg}_P(a[\vec{i}])|h \quad \triangleright_0 \text{cond}_P(a[\vec{i}])|h \quad a[\vec{i}] \in \mathcal{S} \quad M \in \{\text{frame}, \text{exec}\} \quad \text{assuming } \triangleright_0 \text{frame}@_P T'|h \text{ when } T' \leq_S a[\vec{i}]}{\triangleright_0 M @_P T|h} \\
& \frac{\text{SUBRECCELL} \quad \triangleright_0 \text{cell}_P(c[\vec{j}_0] a[\vec{i}])|h \quad a[\vec{i}] \in \mathcal{S} \quad \text{assuming } \triangleright_0 \text{cell}(c[\vec{j}]) @_P T'|h \text{ when } T' <_S a[\vec{i}]}{\triangleright_0 \text{cell}(c[\vec{j}]) @_P \text{pred}(T)|h}
\end{aligned}$$

where $\vec{\alpha}$ and $\vec{j}_0$ are fresh variables, and $f \in \mathcal{F} \cup \{\text{att}\}$.

Figure 7: General shape of subterm relations

C.2 Nonce Freshness

NOTATIONS 2. We write $\|\mathbf{n}\|_{\eta} \subset \mathbb{N}$ for the set of indices of ρ_h used to compute $\llbracket \mathbf{n} \rrbracket_{M, \mathbb{T}}^{\eta}$. Note that it only depends on η .

For $H \subseteq \mathbb{N}$ and random tapes ρ, ρ' , let $s_H(\rho, \rho')$ be the tape obtained by replacing the bits of ρ with those of ρ' at all indices in H . Let $s_{\mathbf{n}, \eta}(\rho, \rho') := s_{\|\mathbf{n}\|_{\eta}}(\rho, \rho')$.

DEFINITION 17 (STABILITY). We say a term t is stable over $H \subseteq \mathbb{N}$ when for all random tapes ρ_A, ρ_h , and ρ' ,

$$\llbracket t \rrbracket_{M, \mathbb{T}}^{\eta}((\rho_h, \rho_A)) = \llbracket t \rrbracket_{M, \mathbb{T}}^{\eta}((s_H(\rho_h, \rho'), \rho_A)) \quad (44)$$

We write $\text{stable}_{M, \mathbb{T}}^{\eta}(t, H)$.

DEFINITION 18 (NONCE FRESHNESS). For any term m , hypothesis h , and nonce $\mathbf{n}$, $\mathbf{n} \triangleright m|h$ holds when for all $h' \equiv h$ we can find m' such that

$$m' \equiv \text{If } h' \text{ then } m \text{ else } \emptyset \quad (45)$$

$$\text{stable}_{M, \mathbb{T}}^{\eta}(h', \|\mathbf{n}\|_{\eta}) \Rightarrow \text{stable}_{M, \mathbb{T}}^{\eta}(m', \|\mathbf{n}\|_{\eta}) \quad (46)$$

We omit h when $h = \text{true}$.

In other words, $\mathbf{n} \triangleright m|h$ means that, modulo $\equiv$, if h can be computed without $\mathbf{n}$, then m can be as well.

PROPERTY 10. $\mathbf{n} \triangleright m|h$ extends Fig. 7 with NONCEINSTANCE.

$$\begin{array}{c}
\text{NONCEINSTANCE} \\
\frac{\mathbf{n} \neq \mathbf{k}}{\mathbf{n} \triangleright \mathbf{k}|h}
\end{array}$$

We will prefix the relevant rules with “NONCE” (e.g., NONCEFA points to the instantiation of SUBFA with $\mathbf{n} \triangleright m|h$).

The resulting rules hold modulo $\equiv$.

PROOF. $n \blacktriangleright m$ is trivially compatible with $\equiv$.

NONCEINSTANCE: $n \neq k$ enforces that both nonces use non-overlapping slices. Then (46) is trivial with the assumption on h .

NONCEFALSE: Take $m' := \emptyset$; it does not use ρ_h at all.

NONCEFA: Let m'_i be the m' chosen for $n \blacktriangleright x_i | h$ when using h' . Take $m' := \text{If } h' \text{ then } f(m'_1, \dots, m'_n) \text{ else } \emptyset$.

NONCEITE: We start like with **NONCEFA**. We need to prove the stability of $\text{If } h' \text{ then } (m'_c \text{ else } m'_c) \text{ else } \emptyset$ with $m'_c := \text{If } m'_c \equiv \emptyset \text{ then } \text{false} \text{ else } m'_c$. We notice that $h' \wedge m'_c \equiv h \wedge c$ and $h' \wedge \neg m'_c \equiv h \wedge \neg c$. Then case analysis on the evaluations of h' and m'_c yields the result.

NONCEAND: Consequence of **NONCEITE** and **NONCEFA**.

NONCEFIND: Same as **NONCEITE**, but we first pick a trace before picking the h 's for the assumptions.

NONCEREC: By picking a trace and using induction on the unfolding. $\square$

We can then move to properly proving Th. 1

THEOREM (1 NO GUESSING). For any nonce n and term t ,

$$n \blacktriangleright t \Rightarrow |n \bar{\neq} t| \quad (6)$$

LEMMA 2 (INDEPENDENCE). If $n \blacktriangleright m$ holds, then there exists an $m' \equiv m$ such that $\llbracket n \rrbracket_{M,T}^\eta(\rho) \coprod \llbracket m' \rrbracket_{M,T}^\eta(\rho)$.

PROOF. Take $h' := \text{true}$ and the m' given by Def. 18. Let $N := \llbracket n \rrbracket_\eta$ and ρ' an independent random tape.

We have $\pi_N(\rho_h) \coprod \rho_{\mathcal{A}}, s_N(\rho_h, \rho')$, thus

$$\llbracket n \rrbracket_{M,T}^\eta((\rho_h, \rho_{\mathcal{A}})) \coprod \llbracket m' \rrbracket_{M,T}^\eta((s_N(\rho_h, \rho'), \rho_{\mathcal{A}})). \quad (47)$$

Then (46) lets us conclude. $\square$

PROOF OF TH. 1. Fix a trace $\mathbb{T}$ and assume $n \blacktriangleright t$. Lem. 2 gives us t' such that:

$$\begin{aligned} \mathbb{T} \vdash_P t' &\equiv t & (48) \\ \mathbb{P} \left(\llbracket n \equiv t' \rrbracket_{M,T}^\eta(\rho) \right) &= \frac{1}{2^\eta} & (\text{by independence}) \end{aligned}$$

Then (48) yields an ow. sequence $(E_\eta)_{\eta \in \mathbb{N}}$ such that, if $\rho \in E_\eta$ then $\llbracket n \rrbracket_{M,T}^\eta(\rho) = \llbracket t' \rrbracket_{M,T}^\eta(\rho)$. Let $\varepsilon_\eta := \mathbb{P}(E_\eta^c) = \text{negl}(\eta)$. We have

$$\mathbb{P} \left(\llbracket n \equiv t \rrbracket_{M,T}^\eta(\rho) \right) = \frac{1}{2^\eta} + \varepsilon_\eta = \text{negl}(\eta) \quad (49)$$

$\square$

C.3 Deduce

DEFINITION (14). Given $u_1^*, u_2^*, v_1^*, v_2^*, h_1^*, h_2^* \in \mathcal{S}$, let $u_1^*, u_2^* \triangleright_\# v_1^*, v_2^* | h_1^*, h_2^*$ hold when for all h_1 and h_2 such that $h_i \equiv h_i^*$, we can find

- (1) u_1, u_2, v_1, v_2 terms
- (2) a sequence $(H_\eta)_{\eta \in \mathbb{N}}$ of subsets of $\mathbb{N}$ of size N_η ;
- (3) $\mathcal{B}$, a polynomial-time Turing machine

such that

$$v_i \equiv v_i^* \text{ and } u_i \equiv \text{If } h_i' \text{ then } u_i^* \text{ else } \emptyset \quad (33)$$

$$\text{stable}_{M,T}^\eta(h_i, H_\eta) \Rightarrow \text{stable}_{M,T}^\eta(u_i, H_\eta) \quad (34)$$

$$\mathcal{B}(1^\eta, \rho_{\mathcal{A}}, \pi_{H_\eta}(\rho_h), \llbracket u_i \rrbracket^\eta(\rho)) = \llbracket v_i \rrbracket^\eta(\rho) \text{ over } R_\eta^i \quad (35)$$

where $R_\eta^i = \{\llbracket h_i \rrbracket^\eta(\rho) = 1\}$ and N_η is a polynomial in η

PROPERTY 11 (DEDUCE AND INDISTINGUISHABILITY).

$$u_1, u_2 \triangleright_\# v_1, v_2 | \text{true}, \text{true} \wedge u_1 \sim u_2 \Rightarrow v_1 \sim v_2 \quad (50)$$

For this proof, we first need the following two lemmas.

LEMMA 3. Consider the notation of Def. 14 with $h_1 = h_2 = \text{true}$. If (34) holds, then for all polynomial-time Turing machines $\mathcal{A}$ and $\tau \in \mathbb{B}^{\llbracket H_\eta \rrbracket}$, we have

$$\pi_{H_\eta}(\rho_h) \coprod \mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \mathcal{B}(1^\eta, \rho_{\mathcal{A}}, \tau, \llbracket u_i \rrbracket^\eta(\rho))) \quad (51)$$

PROOF. Same proof as Lem. 2 starting from (34). $\square$

Then we can prove Prop. 11.

PROOF OF PROP. 11. Let $u_1, u_2, v_1, v_2 \in \mathcal{S}$ be the representatives of $u_1^*, u_2^*, v_1^*, v_2^*$ taken for $h_i := \text{true}$ following the notation of Def. 14. Noticing that $\text{stable}_{M,T}^\eta(\text{true}, H_\eta)$ holds trivially, we get:

$$u_1 \sim u_2 \quad (52)$$

$$\text{stable}_{M,T}^\eta(u_i, H_\eta) \quad (53)$$

$$\mathcal{B}(1^\eta, \rho_{\mathcal{A}}, \pi_{H_\eta}(\rho_h), \llbracket u_i \rrbracket^\eta(\rho)) = \llbracket v_i \rrbracket^\eta(\rho) \quad (54)$$

Let $\mathcal{A}$ be a polynomial-time Turing machine and $\pi_\eta := \pi_{H_\eta}(\rho_h)$.

Let us focus on:

$$\begin{aligned} \mathbb{P}(\mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \llbracket v_i \rrbracket^\eta(\rho)) = 1) \\ = \mathbb{P}(\mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \mathcal{B}(1^\eta, \rho_{\mathcal{A}}, \pi_\eta, \llbracket u_i \rrbracket^\eta(\rho))) = 1) \quad (\text{by (54)}) \end{aligned}$$

Let us write $\mathcal{A}'(1^\eta, \rho_{\mathcal{A}}, \tau, x) := \mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \mathcal{B}(1^\eta, \rho_{\mathcal{A}}, \tau, x))$

$$\begin{aligned} &= \int_\tau \mathbb{P}_{\pi_\eta = \tau}(\mathcal{A}'(1^\eta, \rho_{\mathcal{A}}, \tau, \llbracket u_i \rrbracket^\eta(\rho)) = 1) \\ &\quad \times \mathbb{P}(\pi_\eta = \tau) d\tau \\ &= \int_\tau \mathbb{P}(\mathcal{A}'(1^\eta, \rho_{\mathcal{A}}, \tau, \llbracket u_i \rrbracket^\eta(\rho)) = 1) \\ &\quad \times \mathbb{P}(\pi_\eta = \tau) d\tau \quad (\text{by Lem. 3}) \end{aligned}$$

Let τ_η be uniformly sampled random elements of $\mathbb{B}^{N_\eta}$ independent from ρ_h and $\rho_{\mathcal{A}}$.

$$\begin{aligned} &= \int_\tau \mathbb{P}(\mathcal{A}'(1^\eta, \rho_{\mathcal{A}}, \tau, \llbracket u_i \rrbracket^\eta(\rho)) = 1) \mathbb{P}(\tau_\eta = \tau) d\tau \\ &= \mathbb{P}(\mathcal{A}'(1^\eta, \rho_{\mathcal{A}}, \tau_\eta, \llbracket u_i \rrbracket^\eta(\rho)) = 1) \quad (55) \end{aligned}$$

Finally, since all the Turing machines considered are polynomial-time, they cannot read the entire random tapes. Then, since N_η is polynomial, we can build a polynomial-time Turing machine $\mathcal{A}_0$ that behaves like $\mathcal{A}'$ but fetches its own τ_η from $\rho_{\mathcal{A}}$. Thus we have:

$$\mathbb{P}(\mathcal{A}_0(1^\eta, \rho_{\mathcal{A}}, \llbracket u_i \rrbracket^\eta(\rho))) = \mathbb{P}(\mathcal{A}'(1^\eta, \rho_{\mathcal{A}}, \tau_\eta, \llbracket u_i \rrbracket^\eta(\rho))) \quad (56)$$

$\frac{\text{DEDUCE}}{u, v \triangleright_{\#} a, b \mid \text{true}, \text{true}} \quad u \sim v \vdash a \sim b$	$\frac{\text{DEDUCENONCE}}{n \triangleright u \mid h_1 \quad n \triangleright v \mid h_2} \quad u, v \triangleright_{\#} n, n \mid h_1, h_2$
$\frac{\text{DEDUCEFA}}{u, v \triangleright_{\#} x_1, y_1 \mid h_1, h_2 \quad \dots \quad u, v \triangleright_{\#} x_n, y_n \mid h_1, h_2} \quad u, v \triangleright_{\#} f(x_1, \dots, x_n), f(y_1, \dots, y_n) \mid h_1, h_2$	$\frac{\text{DEDUCEMACRO}}{\vdash_{\text{P}} h_i \Rightarrow t' \leq t \quad i = 1, 2 \quad \mathbf{M} \in \{\text{frame}, \text{exec}\}} \quad \text{frame}@_{p_1} t, \text{frame}@_{p_2} t \triangleright_{\#} \mathbf{M}@_{p_1} t', \mathbf{M}@_{p_2} t' \mid h_1, h_2$
$\frac{\text{DEDUCEOUTPUT}}{\vdash_{\text{P}} h_i \Rightarrow \text{exec}@_{p_i} t' \quad i = 1, 2 \quad \mathbf{M} \in \{\text{msg}, \text{cond}\}} \quad \text{frame}@_{p_1} t, \text{frame}@_{p_2} t \triangleright_{\#} \mathbf{M}@_{p_1} t', \mathbf{M}@_{p_2} t' \mid h_1, h_2$	$\frac{\text{DEDUCEITE}}{u, v \triangleright_{\#} c_1, c_2 \mid h_1, h_2 \quad u, v \triangleright_{\#} l_1, l_2 \mid c_1 \bar{\wedge} h_1, c_2 \bar{\wedge} h_2} \quad u, v \triangleright_{\#} r_1, r_2 \mid (\bar{\neg} c_1) \bar{\wedge} h_1, (\bar{\neg} c_2) \bar{\wedge} h_2$
$\frac{\text{DEDUCEAND}}{u, v \triangleright_{\#} a_1, a_2 \mid h_1, h_2 \quad u, v \triangleright_{\#} b_1, b_2 \mid a_1 \bar{\wedge} h_1, a_2 \bar{\wedge} h_2} \quad u, v \triangleright_{\#} a_1 \bar{\wedge} b_1, a_2 \bar{\wedge} b_2 \mid h_1, h_2$	$\frac{\text{DEDUCEFIND}}{u, v \triangleright_{\#} c_1 \{\vec{\alpha} \mapsto \vec{\alpha}'\}, c_2 \{\vec{\alpha} \mapsto \vec{\alpha}'\} \mid h_1, h_2 \quad u, v \triangleright_{\#} l_1 \{\vec{\alpha} \mapsto \vec{\alpha}'\}, l_2 \{\vec{\alpha} \mapsto \vec{\alpha}'\} \mid c_1 \{\vec{\alpha} \mapsto \vec{\alpha}'\} \bar{\wedge} h_1, c_2 \{\vec{\alpha} \mapsto \vec{\alpha}'\} \bar{\wedge} h_2} \quad u, v \triangleright_{\#} r_1, r_2 \mid (\bar{\forall} \vec{\alpha}, \bar{\neg} c_1) \bar{\wedge} h_1, (\bar{\forall} \vec{\alpha}, \bar{\neg} c_2) \bar{\wedge} h_2$
$u, v \triangleright_{\#} \text{Find } \vec{\alpha} \text{ st. } c_1 \text{ then } l_1 \text{ else } r_1, \text{ if } c_2 \text{ then } l_2 \text{ else } r_2 \mid h_1, h_2$ where $f \in \mathcal{F} \cup \{\text{att}\}$.	

Figure 8: Full DEDUCE rules

Let us simply write $\mathcal{A}_0^i := \mathcal{A}_0(1^\eta, \rho_{\mathcal{A}}, \llbracket u_i \rrbracket^\eta(\rho))$ and $\mathcal{A}^i := \mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \llbracket v_i \rrbracket^\eta(\rho))$. Then

$$\begin{aligned} & \left| \mathbb{P}(\mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \llbracket v_1 \rrbracket^\eta(\rho))) - \mathbb{P}(\mathcal{A}(1^\eta, \rho_{\mathcal{A}}, \llbracket v_2 \rrbracket^\eta(\rho))) \right| \quad (57) \\ &= \left| \mathbb{P}(\mathcal{A}^1 = 1) - \mathbb{P}(\mathcal{A}^2 = 1) \right| \quad (\text{by definition}) \\ &= \left| \mathbb{P}(\mathcal{A}_0^1 = 1) - \mathbb{P}(\mathcal{A}_0^2 = 1) \right| \quad (\text{by (55)}) \\ &= \text{negl}(\eta) \quad (\text{by (52)}) \end{aligned}$$

hence $v_1 \sim v_2$, and thus $v_1^* \sim v_2^*$. $\square$

PROPERTY 12. The rules of Fig. 8 hold modulo $\equiv$.

PROOF. $\triangleright_{\#}$ is trivially compatible with $\equiv$.

DEDUCE: This is Prop. 11.

DEDUCENONCE: By Def. 18, we take $H_\eta := \llbracket n \rrbracket_\eta$.

DEDUCEFA: For each h_i , we take $H_\eta := \cup_i H_\eta^i$, where H_η^i is given by $u, v \triangleright_{\#} x_i, y_i \mid h_1, h_2$. Since f is computable by an attacker, we conclude immediately. We then choose our u' and v' for item 1

$\frac{\text{FA}}{u \sim v \vdash \vec{t} \mathbin{++} \vec{a} \sim \vec{t}' \mathbin{++} \vec{a}} \quad u \sim v \vdash f(\vec{t}) : : \vec{a} \sim f(\vec{t}') : : \vec{b}$	$\frac{\text{FAPROJ}}{u \sim v \vdash \tau(\vec{t}) \sim \tau(\vec{t}')} \quad u \sim v \vdash \vec{t} \sim \vec{t}'$
$\frac{\text{FAFIND}}{u \sim v \vdash c\{\vec{\alpha} \mapsto \vec{\alpha}'\} : : l\{\vec{\alpha} \mapsto \vec{\alpha}'\} : : r : : \vec{a} \quad \sim c'\{\vec{\alpha}' \mapsto \vec{\alpha}'\} : : l'\{\vec{\alpha}' \mapsto \vec{\alpha}'\} : : r' : : \vec{b} \quad \ \vec{\alpha}\ = \ \vec{\alpha}'\ } \quad u \sim v \vdash \text{Find } \vec{\alpha} \text{ st. } c \text{ then } l \text{ else } r : : \vec{a} \quad \sim \text{Find } \vec{\alpha}' \text{ st. } c' \text{ then } l' \text{ else } r' : : \vec{b}$	

Using standard list notations; τ is a projection and permutation operator.

Figure 9: Rules for Function Application

from the u'_i and v'_i of the assumptions when f is not a constant.

$$u' := \text{If } \bigwedge_{i,j} u'_i \equiv u'_j \text{ then } u'_0 \text{ else} \quad (58)$$

$$v' := \text{If } \bigwedge_{i,j} v'_i \equiv v'_j \text{ then } v'_0 \text{ else} \quad (59)$$

When f is a constant, simply picking u and v is sufficient.

With this we easily get (33) and (34). Then, since f is a deterministic polynomial-time Turing machine, we quickly also get (35).

DEDUCEMACRO: When the precondition holds, we fetch $\mathbf{M}@_{p_i} t'$ directly from $\text{frame}@_{p_i} t'$, as it is known to appear in its unfolding when h_i holds.

DEDUCEOUTPUT: Same as DEDUCEMACRO, but the content of the macro is guarded by $\text{exec}@_{p_i} t'$ in the frame. Hence the extra condition.

DEDUCEITE: Same as DEDUCEFA. By case analysis on the $\llbracket c_i \rrbracket_{\mathbf{M}, \mathbf{T}}^\eta$, we notice that the tighter constraints in the assumptions are sufficient for (33).

DEDUCEAND: Consequence of DEDUCEFA and DEDUCEITE. $\square$

REMARK 7. DEDUCEOUTPUT is especially useful as it enables Section 5.2.5's technique by noticing that on Rmk. 1's publication steps, $\text{exec}@_{pT}$ is always true.

REMARK 8. No rule directly references memory cells. These are handled by equality reasoning only.

C.4 Function Application

The FA rules can be used to “forget” information about function application. To properly encode them, we use a list-like construction. We extend $\mathcal{F}$ with the standard $\_ : : \_, []$.

NOTATIONS 3 (LISTS). In Fig. 9 use vector notation to denote list, and we use $++$ as sugar for list concatenation; these are in reality dynamic rules (Section 4.3.1) that unwrap and rewrap the list.

PROPERTY 13. The rules of Fig. 9 hold.

PROOF. Lifted from [10]. $\square$

These rules are widely used across proofs of protocols in the CCSA. However, they are very adversarial to automation. Indeed, they almost always increase the number of applicable rules to a

goal, and it is almost always possible to apply the previous rules again to the goal left by an FA rule. Moreover, FAPROJ effectively has no semantic effect on a proof, but can exponentially increase the size of the search space. In Appendix 6's Table 1, we can notice that the example making use of FA spends comparatively more time in GOLGGE.

Thus, the rules of Fig. 9 are not directly implemented in CRYPTO-VAMPIRE2. Instead, we leverage Section 4.3.1's dynamic rules to apply multiple FA rules at once before handing control back to GOLGGE, thus avoiding as much back-and-forth as possible. Furthermore, to prevent the search space from exploding and to allow CRYPTO-VAMPIRE2 to terminate in reasonable times, we add the option to cap the size of the FA list. In our experiments, that cap was somewhat arbitrarily set at 4.

C.5 Pseudo Random Function

We provide a slightly improved PRF rule compared to that of SQUIRREL [5]. Since this requires a new proof, this section also serves as an example proof for a cryptographic rule in this variant of the CCSA. For subsequent rules, we will simply defer to [5] instead.

C.5.1 The PRF security game.

DEFINITION 19 (PSEUDO RANDOM FUNCTION [41]). Let $\mathcal{H} : \mathbb{B}^* \times \mathbb{B}^* \rightarrow \mathbb{B}^\eta$ be an efficient, keyed function. $\mathcal{H}$ is a Pseudo Random Function if, for all PPTMs $\mathcal{A}$ with an oracle, uniformly sampled $k \in \mathbb{B}^\eta$, and $f : \mathbb{B}^* \rightarrow \mathbb{B}^\eta$, we have:

$$\left| \mathbb{P}(\mathcal{A}^{\mathcal{H}(\_,k)}(1^\eta)) - \mathbb{P}(\mathcal{A}^f(1^\eta)) \right| = \text{negl}(\eta) \quad (60)$$

(60) is stronger than our needs as it replaces *all* messages hashed with k . Thus, we need Lem. 4.

C.5.2 Lemmas. The goal of this section is to get to Lem. 5. Informally, the idea is to apply (60) from Def. 19 twice on different instances, noting that $f(x)$ behaves like a nonce when all other parameters are fixed. This gives us Lem. 4. It is, however, limited to the very last use of $\mathcal{H}$, which is where the PRF rule in [5] stops. However, we notice that it is possible to sequentially turn calls to $\mathcal{H}$ into nonces by following the call graph and then turn the extra nonces back into calls to $\mathcal{H}$. This is Lem. 5.

Objects and Notations. In this section, we consider $t_1, \dots, t_l, \dots, t_n$ and $\mathcal{B}$ Probabilistic Polynomial-Time Turing Machines (PPTMs); $\mathcal{H} : \mathbb{B}^* \times \mathbb{B}^* \rightarrow \mathbb{B}^\eta$ a PRF-secure function as in Def. 19; and $f : \mathbb{B}^* \rightarrow \mathbb{B}^\eta$ a uniformly sampled function.

NOTATIONS 4 (ORACLES). We write $\mathcal{H}_k$ for the oracle $\mathcal{H}(\_, k)$ and $f_\star$ for the oracle $f(\_)$. For a given PPTM $x, x_{|l}^\mathcal{O}(\vec{y}) := x^\mathcal{O}(1^\eta, \vec{y})$; crucially, we omit the function application when $\vec{y}$ is empty.

Throughout these proofs, $\mathcal{O}$ will effectively be either $\mathcal{H}_k$ or $f_\star$. We will gradually turn relevant uses of $\mathcal{H}_k$ into $f_\star$ and vice-versa.

NOTATIONS 5 (VECTORIZED ARGUMENTS). As before, we write $g(\vec{x})$ for $g(x_1, \dots, x_n)$. We will also use a Haskell-inspired map operator, writing $g^\Delta(\vec{x})$ for $g(x_1), \dots, g(x_n)$.

The t_i intuitively represent the arguments to $\mathcal{H}$. We split them at index l , such that $\mathcal{H}(t_i, k)$ is the term to be replaced by n_{PRF} . Then $t_{l+1}, \dots, t_n$ are the arguments of the hashes computed *after* $\mathcal{H}(t_l, k)$.

During the proof, we will replace every instance of $\mathcal{H}(t_i, k)$ for $i \geq l$ by random bitstrings. To track the dependencies on previous calls to $\mathcal{H}(t_i, k)$, we make t_j take $j-l$ arguments and $(\dagger)$ we assume that the computation of $t_{l+1}^\mathcal{O}, \dots, t_n^\mathcal{O}, t_{l+1}^\mathcal{O}(x_l), \dots, t_n^\mathcal{O}(x_l, \dots, x_{n-1})$ only calls $\mathcal{O}$ on $t_{l+1}^\mathcal{O}, \dots, t_{n-1}^\mathcal{O}$ for any $\vec{x}$.

$$\text{Let } h := \bigwedge_{i=1}^{l-1} t_i \neq t_i$$

LEMMA 4. With k and n independently and uniformly sampled from $\mathbb{B}^\eta$ and $\vec{t} := t_1, \dots, t_{l-1}$.

$$\left| \mathbb{P}(\mathcal{B}_{|l}(\mathcal{H}_k^\Delta(\vec{t}^{\mathcal{H}_k}), \text{If } h_{|l}^{\mathcal{H}_k} \text{ then } \mathcal{H}_k(t_{l+1}^{\mathcal{H}_k}) \text{ else } \mathbf{0}) = 1) - \mathbb{P}(\mathcal{B}_{|l}(\mathcal{H}_k^\Delta(\vec{t}^{\mathcal{H}_k}), \text{If } h_{|l}^{\mathcal{H}_k} \text{ then } n \text{ else } \mathbf{0}) = 1) \right| = \text{negl}(\eta) \quad (61)$$

This lemma loosely corresponds to [5]'s PRF rule. It lets one substitute one use of $\mathcal{H}$ by a nonce under the constraint that this term is not reused in later hashes. The proof follows a similar approach to that in [28].

PROOF. Using the notation of Lem. 4 and Def. 19. We assume $\mathcal{H}$ is PRF secure; let f be a random function as in Def. 19.

The proof outline can be summarized by (62):

$$\begin{aligned} A := & \left| \mathbb{P}(\mathcal{B}_{|l}(\mathcal{H}_k^\Delta(\vec{t}^{\mathcal{H}_k}), \text{If } h_{|l}^{\mathcal{H}_k} \text{ then } \mathcal{H}_k(t_{l+1}^{\mathcal{H}_k}) \text{ else } \mathbf{0}) = 1) - \mathbb{P}(\mathcal{B}_{|l}(\mathcal{H}_k^\Delta(\vec{t}^{\mathcal{H}_k}), \text{If } h_{|l}^{\mathcal{H}_k} \text{ then } n \text{ else } \mathbf{0}) = 1) \right| \\ = & \left| \mathbb{P}(\mathcal{B}_{|l}(\mathcal{H}_k^\Delta(\vec{t}^{\mathcal{H}_k}), \text{If } h_{|l}^{\mathcal{H}_k} \text{ then } \mathcal{H}_k(t_{l+1}^{\mathcal{H}_k}) \text{ else } \mathbf{0}) = 1) - \mathbb{P}(\mathcal{B}_{|l}(f^\Delta(\vec{t}^f), \text{If } h_{|l}^f \text{ then } f(t_{l+1}^f) \text{ else } \mathbf{0}) = 1) \right. \\ & \left. + \mathbb{P}(\mathcal{B}_{|l}(f^\Delta(\vec{t}^f), \text{If } h_{|l}^f \text{ then } f(t_{l+1}^f) \text{ else } \mathbf{0}) = 1) - \mathbb{P}(\mathcal{B}_{|l}(\mathcal{H}_k^\Delta(\vec{t}^{\mathcal{H}_k}), \text{If } h_{|l}^{\mathcal{H}_k} \text{ then } n \text{ else } \mathbf{0}) = 1) \right| \end{aligned} \quad (62)$$

Focusing on the middle terms, we write $f_t^\eta := f^\Delta(\vec{t}^f)$ and $f_{t_l}^\eta := f(t_{l+1}^f)$, we have:

$$\begin{aligned} B := & \mathbb{P}(\mathcal{B}_{|l}(f_t^\eta, \text{If } h_{|l}^f \text{ then } f_{t_l}^\eta \text{ else } \mathbf{0}) = 1) \\ = & \sum_{\vec{\omega} \in (\mathbb{B}^\eta)^n} \mathbb{P}(f_t^\eta = \vec{\omega}) \mathbb{P}_{f_t^\eta = \vec{\omega}}(\mathcal{B}_{|l}(\vec{\omega}, \text{If } h_{|l}^f \text{ then } f_{t_l}^\eta \text{ else } \mathbf{0}) = 1) \\ = & \sum_{\vec{\omega} \in (\mathbb{B}^\eta)^n} \mathbb{P}(f_t^\eta = \vec{\omega}) \mathbb{P}_{f_t^\eta = \vec{\omega}}(\neg h_{|l}^f \cap \mathcal{B}_{|l}^f(\vec{\omega}, 0)) \\ & + \mathbb{P}(f_t^\eta = \vec{\omega}) \mathbb{P}_{f_t^\eta = \vec{\omega}}(h_{|l}^f \cap \mathcal{B}_{|l}^f(\vec{\omega}, f_{t_l}^\eta)) \end{aligned} \quad (63)$$

Then we notice that, since f is uniformly sampled, for any fixed $x \in \mathbb{B}^\eta \setminus \vec{\omega}$, $f(x)$ is uniformly sampled too. Moreover, h effectively no longer depends on f .

$$B = \sum_{\vec{\omega} \in (\mathbb{B}^\eta)^n} \mathbb{P}(f_t^\eta = \vec{\omega}) \mathbb{P}_{f_t^\eta = \vec{\omega}}(\neg h_{|l}^f \cap \mathcal{B}_{|l}^f(\vec{\omega}, 0)) + \mathbb{P}(f_t^\eta = \vec{\omega}) \mathbb{P}_{f_t^\eta = \vec{\omega}}(h_{|l}^f \cap \mathcal{B}_{|l}^f(\vec{\omega}, n)) \quad (64)$$

$$= \mathbb{P}(\mathcal{B}_{|l}(f_t^\eta, \text{If } h_{|l}^f \text{ then } n \text{ else } \mathbf{0}) = 1) \quad (65)$$

Then using (65) and (60) in (62) yields (61). $\square$

While Lem. 4 restricts us to hashing *before* computing m , Lem. 5 alleviates this restriction.

First, we need to introduce suitable notation and formalism for this approach. We recursively define the $w_i^{(0)}$ for $i > l$ as:

$$w_{i+1|l}^{(0)}(w) := \text{If } t_{l|l}^\mathcal{O} \neq t_{i|l}^\mathcal{O} \left(w, w_{l+1}^{(0)}, \dots, w_i^{(0)} \right) \\ \text{then } \mathcal{O} \left(t_i \left(w, w_{l+1}^{(0)}, \dots, w_i^{(0)} \right) \right) \text{ else } \mathbf{0} \quad (66)$$

These $w_i^{(0)}$ are guarded calls to the hashing oracles. Their recursive nature formalizes the fact that one call may depend on previous calls.

LEMMA 5. Let k and n be independently and uniformly sampled from $\mathbb{B}^\eta$; $\vec{t} := t_1, \dots, t_{l-1}$ and $\vec{t}' := t_1, \dots, t_{l-1}$. Let

$$y_{1|l}^\mathcal{O} := \text{If } h_{l|l}^\mathcal{O} \text{ then } \mathcal{H}_k \left(t_{l|l}^\mathcal{O} \right) \text{ else } \mathbf{0} \quad (67)$$

$$y_{2|l}^\mathcal{O} := \text{If } h_{l|l}^\mathcal{O} \text{ then } n \text{ else } \mathbf{0} \quad (68)$$

we have, with $(0)\vec{w} := w_{l+1}^{(0)}, \dots, w_n^{(0)}$,

$$\left| \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (0)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{1|l}^{\mathcal{H}_k} \right), y_{1|l}^{\mathcal{H}_k} \right) \right) \right. \\ \left. - \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (0)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{2|l}^{\mathcal{H}_k} \right), y_{2|l}^{\mathcal{H}_k} \right) \right) \right| = \text{negl}(\eta) \quad (69)$$

The lemma lets us replace $y := t_{l|l}^{\mathcal{H}_k}$ with a fresh random bitstring n despite that y may be used in later oracle calls.

PROOF. Consider for $i > l$

$$\begin{aligned} & \text{let } t := t_{l|l}^\mathcal{O} (x_{l+1}, \dots, x_i) \\ & \text{if } t = t_{l|l}^\mathcal{O} \text{ then return } \mathbf{0} \\ & \text{for } 1 \leq j < l \text{ do} \\ H_{i|l}^\mathcal{O} (x_l, \dots, x_{i-1}) := & \quad \text{if } t_{j|l}^\mathcal{O} = t \text{ then return } \mathcal{O} \left(t_{j|l}^\mathcal{O} \right) \quad (70) \\ & \text{for } l < j < i \text{ do} \\ & \quad \text{let } t' := t_{j|l}^\mathcal{O} (x_l, \dots, x_j) \\ & \quad \text{if } t' = t \text{ then return } x_j \\ & \text{return } \mathcal{O}(t) \end{aligned}$$

By case analysis, for $j \in \{1, 2\}$, we have:

$$H_{i|l}^{\mathcal{H}_k} \left(y_{j|l}^{\mathcal{H}_k}, w_{l+1}^{(0)} \left(y_{j|l}^{\mathcal{H}_k} \right), \dots, w_{i-1}^{(0)} \left(y_{j|l}^{\mathcal{H}_k} \right) \right) = w_i^{(0)} \left(y_{j|l}^{\mathcal{H}_k} \right) \quad (71)$$

The H_i explicitly define the guard to the oracle calls.

We define

$$w_{i+1|l}^{(1)}(w) := H_{i|l}^{\mathcal{H}_k} \left(w, w_{l+1}^{(1)}(w), \dots, w_{i-1}^{(1)}(w) \right) \quad (72)$$

Then

$$\left| \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (0)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{1|l}^{\mathcal{H}_k} \right), y_{1|l}^{\mathcal{H}_k} \right) \right) \right. \\ \left. - \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (0)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{2|l}^{\mathcal{H}_k} \right), y_{2|l}^{\mathcal{H}_k} \right) \right) \right| \\ = \left| \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (1)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{1|l}^{\mathcal{H}_k} \right), y_{1|l}^{\mathcal{H}_k} \right) \right) \right. \\ \left. - \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (1)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{2|l}^{\mathcal{H}_k} \right), y_{2|l}^{\mathcal{H}_k} \right) \right) \right| \quad (73)$$

Due to the position of the final return statement in (70), it is possible to apply Lem. 4 on the t_i starting from $i = n$ down to $n = l$.

Let the $n_i \in \mathbb{B}^\eta$ be mutually independent and uniformly sampled. Consider

$$\begin{aligned} & \text{let } t := t_{l|l}^\mathcal{O} (x_{l+1}, \dots, x_i) \\ & \text{if } t = t_{l|l}^\mathcal{O} \text{ then return } \mathbf{0} \\ & \text{for } 1 \leq j < l \text{ do} \\ H_{i|l}^\mathcal{O} (x_l, \dots, x_{i-1}) := & \quad \text{if } t_{j|l}^\mathcal{O} = t \text{ then return } \mathcal{O} \left(t_{j|l}^\mathcal{O} \right) \quad (74) \\ & \text{for } l < j < i \text{ do} \\ & \quad \text{let } t' := t_{j|l}^\mathcal{O} (x_l, \dots, x_j) \\ & \quad \text{if } t' = t \text{ then return } x_j \\ & \text{return } n_i \end{aligned}$$

Notice the final n_i on (74). We again define:

$$w_{i+1|l}^{(2)}(w) := H_{i|l}^{\mathcal{H}_k} \left(w, w_{l+1}^{(2)}(w), \dots, w_{i-1}^{(2)}(w) \right) \quad (75)$$

We get that

$$\left| \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (1)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{1|l}^{\mathcal{H}_k} \right), y_{1|l}^{\mathcal{H}_k} \right) \right) \right. \\ \left. - \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (2)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{2|l}^{\mathcal{H}_k} \right), y_{2|l}^{\mathcal{H}_k} \right) \right) \right| = \text{negl}(\eta) \quad (76)$$

Notice that we went all the way to using y_2 .

By the same reasoning, we apply Lem. 4 on the t_i from $i = n$ down to $i = l + 1$ in a situation where we are already using y_2 . By leveraging the independence of the n_i to put them back in the same position, we get:

$$\left| \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (1)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{2|l}^{\mathcal{H}_k} \right), y_{2|l}^{\mathcal{H}_k} \right) \right) \right. \\ \left. - \mathbb{P} \left(\mathcal{B}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), (2)\vec{w}_{l|l}^{\mathcal{H}_k} \left(y_{2|l}^{\mathcal{H}_k} \right), y_{2|l}^{\mathcal{H}_k} \right) \right) \right| = \text{negl}(\eta) \quad (77)$$

Finally, combining (73), (76), and (77) completes the proof. $\square$

C.5.3 The PRF rule. Then we can define $m, k, n_{\text{PRF}} \blacktriangleright_{\text{PRF}} t|h$ as a proof of existence of the various PPTMs of Lem. 4.

DEFINITION 20. $m, k, n_{\text{PRF}} \blacktriangleright_{\text{PRF}} t|h$ holds when for all $h' \equiv h$ we can find m', t' and all the PPTMs introduced in Section C.5.2 and $\mathcal{T}$ such that:

- (1) no Turing Machine touches the slice of ρ related to k and n_{PRF} ;
- (2) $m' \equiv m$ and $t' \equiv t$;
- (3) $\mathcal{T}_{l|l} \left(\mathcal{H}_k \left(\vec{t}_{l|l}^{\mathcal{H}_k} \right), \vec{w}_{l|l}^{\mathcal{H}_k} (y_\eta), y_\eta \right) = \llbracket t' \rrbracket_{M, T}^\eta$;
- (4) there exists an l such that $t_{l|l}^{\mathcal{H}_k} = \llbracket m' \rrbracket_{M, T}^\eta$;
- (5) overwhelmingly, we have

$$\llbracket h' \rrbracket_{M, T}^\eta \Rightarrow t_{i|l}^{\mathcal{H}_k} \neq \llbracket m' \rrbracket_{M, T}^\eta \text{ for } i < l \quad (78)$$

$$\llbracket h' \rrbracket_{M, T}^\eta \Rightarrow t_{i|l}^{\mathcal{H}_k} \left(y_\eta, w_{l+1|l}^{\mathcal{H}_k}(y_\eta), \dots, w_{i-1|l}^{\mathcal{H}_k}(y_\eta) \right) \neq \llbracket m' \rrbracket_{M, T}^\eta \text{ for } i > l \quad (79)$$

with

$$w_{i|l}^\mathcal{O}(y) := \mathcal{O} \left(t_{i|l}^\mathcal{O} \left(y, w_{l+1|l}^\mathcal{O}(y), \dots, w_{i-1|l}^\mathcal{O}(y) \right) \right) \quad (80)$$

$$\begin{array}{c}
\text{PRFFoundInstance} \\
\frac{}{m, k, n_{PRF} \blacktriangleright_{PRF} \mathcal{H}(m, k)|h} \\
\\
\text{PRFNonce} \\
\frac{\vdash_P h \Rightarrow n \neq n_{PRF} \bar{\wedge} n \neq k}{m, k, n_{PRF} \blacktriangleright_{PRF} n|h} \\
\\
\text{PRFNeqM} \\
\frac{\vdash_P h \Rightarrow m \neq m' \quad m, k, n_{PRF} \blacktriangleright_{PRF} m'|h}{m, k, n_{PRF} \blacktriangleright_{PRF} \mathcal{H}(m', k)|h} \\
\\
\text{PRFNeqK} \\
\frac{\vdash_P h \Rightarrow m \neq m' \bar{\vee} k \neq k' \quad m, k, n_{PRF} \blacktriangleright_{PRF} m'|h \quad m, k, n_{PRF} \blacktriangleright_{PRF} k'|h}{m, k, n_{PRF} \blacktriangleright_{PRF} \mathcal{H}(m', k')|h}
\end{array}$$

Figure 10: Extra rules for $m, k, n_{PRF} \blacktriangleright_{PRF} t|h$

$$y_\eta := \mathcal{H}_k \left(\llbracket m' \rrbracket_{M, T}^\eta \right) \quad (81) \quad \vec{w} := w_{l+1}, \dots, w_n \quad (83)$$

$$k := \llbracket k \rrbracket_{M, T}^\eta \quad (82) \quad \vec{t} := t_1, \dots, t_{l-1} \quad (84)$$

Where, without loss of generality, we assume that these also take their randomness from the pair of random tapes ρ .

PROPERTY 14. Extending Fig. 7 by Fig. 10 and skipping $f = \mathcal{H}_k$ in SUBFA yields a sound set of rules modulo $\equiv$ for $m, k, n_{PRF} \blacktriangleright_{PRF} t|h$.

PROOF. We will focus on PRFITE and PRFFOUNDINSTANCE as the others behave similarly.

PRFITE: The key point in PRFITE is to control the order of evaluation to make the guard c effective, this means we diverge from [28] by not fully evaluating “bottom up”.

Assume we have h' such that $h \equiv h'$ and

$$m, k, n_{PRF} \blacktriangleright_{PRF} c|h \quad (85)$$

$$m, k, n_{PRF} \blacktriangleright_{PRF} l|c \bar{\wedge} h \quad (86)$$

$$m, k, n_{PRF} \blacktriangleright_{PRF} r|\bar{c} \bar{\wedge} h \quad (87)$$

we take $h', c \bar{\wedge} h'$ and $\bar{c} \bar{\wedge} h'$ for (85), (86) and (87) respectively. We index by c (resp. l, r) all the items given by Def. 20 for (85) (resp. (86), (87)).

We employ the same trick as in the proof of Prop. 12 for our choice of m' . Otherwise, t' is chosen as expected:

$$m' := \text{If } m'_c \equiv m'_l \equiv m'_r \text{ then } m'_c \text{ else } \mathbf{0} \quad (88)$$

$$t' := \text{If } m'_c \equiv m'_l \equiv m'_r \text{ then } (\text{If } t'_c \text{ then } t'_l \text{ else } t'_r) \text{ else } \mathbf{0} \quad (89)$$

Then we construct the t_i . We take $l := l_c, n := n_c + n_l + n_r, t_1, \dots, t_{l-1} := t_{c,1}, \dots, t_{c,l-1}$. With the w_i defined as in Def. 20, we build:

$$t_{l|\eta}^\mathcal{O} := \text{If } t_{c,l_c|\eta}^\mathcal{O} = t_{l,l_l|\eta}^\mathcal{O} = t_{r,l_r|\eta}^\mathcal{O} \text{ then } t_{c,l_c|\eta}^\mathcal{O} \text{ else } \mathbf{0} \quad (90)$$

$$t_{n_c+i|\eta}^\mathcal{O}(y, w_1, \dots, w_{n_c+i-1}) := \text{If } \mathcal{T}_c(w_1, \dots, w_{n_c}, y) \text{ then } t_{l,i|\eta}^\mathcal{O} \text{ else } \mathbf{0} \text{ when } i < l \quad (91)$$

$$t_{n_c+i|\eta}^\mathcal{O}(y, w_1, \dots, w_{n_c+i-1}) := \text{If } \mathcal{T}_c(w_1, \dots, w_{n_c}, y) \text{ then } t_{l,i|\eta}^\mathcal{O}(y, w_{n_c+l}, \dots, w_{n_c+i-1}) \text{ else } \mathbf{0} \text{ when } i > l \quad (92)$$

And similarly for the t_i s of r .

$$\begin{array}{c}
\text{PRFNonce}^* \\
\frac{\vdash_P h \Rightarrow n \neq n_{PRF} \bar{\wedge} n \neq k}{m, k, n_{PRF} \blacktriangleright_{PRF}^* n|h} \\
\\
\text{PRFNeqM}^* \\
\frac{\vdash_P h \Rightarrow m \neq m' \quad m, k, n_{PRF} \blacktriangleright_{PRF}^* m'|h}{m, k, n_{PRF} \blacktriangleright_{PRF}^* \mathcal{H}(m', k)|h} \\
\\
\text{PRFNeqK}^* \\
\frac{\vdash_P h \Rightarrow m \neq m' \bar{\vee} k \neq k' \quad m, k, n_{PRF} \blacktriangleright_{PRF}^* m'|h \quad m, k, n_{PRF} \blacktriangleright_{PRF}^* k'|h}{m, k, n_{PRF} \blacktriangleright_{PRF}^* \mathcal{H}(m', k')|h}
\end{array}$$

Figure 11: Extra rules for $m, k, n_{PRF} \blacktriangleright_{PRF}^* t|h$

By Section C.5.2's item ($\dagger$), we know that t_{c,l_c} , t_{l,l_l} , and t_{r,l_r} at most call the oracle on $t_1, \dots, t_{l-1}$. Therefore, we can extract a PPTM $\mathcal{C}$ such that, with $\vec{t} := t_1, \dots, t_{l-1}$, we have

$$\mathcal{C}_{|\eta} \left(\mathcal{H}_k^{\mathcal{C}} \left(\vec{t}_{|\eta}^{\mathcal{H}_k} \right) \right) = \left(t_{c,l_c|\eta}^{\mathcal{H}_k} = t_{l,l_l|\eta}^{\mathcal{H}_k} = t_{r,l_r|\eta}^{\mathcal{H}_k} \right) \quad (93)$$

We then define $\mathcal{T}$.

$$\begin{aligned}
\mathcal{T}_{|\eta}(x_1, \dots, x_{n_c+n_l+n_r}, y) &:= \\
&\quad \text{If } \mathcal{C}_{|\eta}(x_1, \dots, x_{l-1}) \text{ then} \\
&\quad \left(\text{If } \mathcal{T}_{c|\eta}(x_1, \dots, x_{n_c}, y) \text{ then } \mathcal{T}_{l|\eta}(x_{n_c+1}, \dots, x_{n_c+n_l}, y) \right. \\
&\quad \quad \left. \text{else } \mathcal{T}_{r|\eta}(x_{n_c+n_l+1}, \dots, x_{n_c+n_l+n_r}, y) \right) \\
&\quad \text{else } \mathbf{0} \quad (94)
\end{aligned}$$

where the branches of the if are only evaluated after the condition and we skip x_l in the arguments.

Case analysis shows that all conditions of Def. 20 are fulfilled.

PRFFOUNDINSTANCE: We simply take $l = n = 1$ and

$$\mathcal{T}(y) := y \quad (95)$$

□

PROPERTY 15. We can retrieve a C for a proof of $m, k, n_{PRF} \blacktriangleright_{PRF} t|h$ made with the rules of Prop. 14 that evaluates like Def. 20's $\mathcal{T}$.

PROOF. Trivially by induction. □

We also introduce a simpler subterm search for the message itself (i.e., the argument of the hash). As the message cannot depend on itself, we forgo most of the complexity of Def. 20, notably we intuitively remove any t_i with $i > l$.

DEFINITION 21. $m, k, n_{PRF} \blacktriangleright_{PRF}^* t|h$ holds when for all $h' \equiv h$ we can find m', t' , all the PPTMs $t_1, \dots, t_{l-1}$ introduced in Section C.5.2, and $\mathcal{T}$ such that:

- (1) no Turing Machine touches the slice of ρ related to k and n_{PRF} ;
- (2) $m' \equiv m$ and $t' \equiv t$;
- (3) $\mathcal{T}_{|\eta} \left(\mathcal{H}_k^{\mathcal{C}} \left(\vec{t}_{|\eta}^{\mathcal{H}_k} \right) \right) = \llbracket t' \rrbracket_{M, T}^\eta$;
- (4) $\llbracket h' \rrbracket_{M, T}^\eta \Rightarrow t_{i|\eta}^{\mathcal{H}_k} \neq \llbracket m' \rrbracket_{M, T}^\eta$ overwhelmingly for $i < l$.

where $\vec{t} := t_1, \dots, t_{l-1}$; without loss of generality, we assume that the PPTMs also take their randomness from the pair of random tapes ρ .

Def. 21 is strictly weaker than Def. 20, it notably does not feature an equivalent for $\text{PRF\_FOUND\_INSTANCE}$.

PROPERTY 16. Extending Fig. 7 by Fig. 11 and skipping $f = \mathcal{H}_k$ in SUBFA yields a sound set of rules modulo $\equiv$ for $m, k, n_{\text{PRF}} \triangleright_{\text{PRF}}^* t|h$.

PROOF. By a similar proof as for Prop. 14. $\square$

In practice we deactivate $\text{PRF\_FOUND\_INSTANCE}$ past a PRF_REC in CRYPTOVAMPIRE2 because the content of previous steps in the trace is read-only to the tool, thus we cannot accept context holes there.

DEFINITION 22 (SHORTCUT). As a shortcut, we write

$$m, k, n_{\text{PRF}} \triangleright_{\text{PRF}} C[x], y := m, k, n_{\text{PRF}} \triangleright_{\text{PRF}} C[x] \wedge m, k, n_{\text{PRF}} \triangleright_{\text{PRF}}^* y \quad (96)$$

Then finally we have:

PROPERTY (1, PRF). When $\mathcal{H}(\_, \_)$ behaves like a Pseudo Random Function (Def. 19), the following rule holds:

$$\frac{\text{PRF} \quad u \sim v \vdash C[n_{\text{PRF}}] \sim b \quad m, k, n_{\text{PRF}} \triangleright_{\text{PRF}} C[\mathcal{H}(m, k)], m}{u \sim v \vdash C[\mathcal{H}(m, k)] \sim b}$$

where k and n_{PRF} are nonces.

PROOF. We instead use transitivity and weakening and simply prove $\vdash C[n_{\text{PRF}}] \sim C[\mathcal{H}(m, k)]$ under the assumption that $m, k, n_{\text{PRF}} \triangleright_{\text{PRF}} C[m], m$. Unfolding Def. 22 we get

$$m, k, n_{\text{PRF}} \triangleright_{\text{PRF}} C[m] \quad (97)$$

$$m, k, n_{\text{PRF}} \triangleright_{\text{PRF}}^* m \quad (98)$$

We take $h := \text{true}$ for both (97) and (98). We will index objects by (97) or (98) to denote where they are coming from. We otherwise directly reuse the notation from Def. 20 and Def. 21. Without loss of generality, we assume (97) and (98) share their t_i s for $i \neq l$. We define $\mathcal{T}_0$ as

$$\mathcal{T}_{0|_{\eta}}^{\mathcal{O}}(y) := \mathcal{T}_{(97)|_{\eta}} \left(\mathcal{O} \left(\vec{t}_{|_{\eta}}^{\mathcal{O}} \right), \vec{w}_{|_{\eta}}^{\mathcal{O}}(y), y \right) \quad (99)$$

Where $k := \llbracket k \rrbracket_{\text{M}, \text{T}}^{\eta}$ for conciseness. We notice that using Def. 20's (3) and Def. 21's (3) we have:

$$\mathcal{T}_{0|_{\eta}} \left(\mathcal{T}_{(98)|_{\eta}} \left(\mathcal{H}_k^{\mathcal{O}}(\vec{t}) \right) \right) \stackrel{\text{ow.}}{=} \llbracket C[\mathcal{H}(m, k)] \rrbracket_{\text{M}, \text{T}}^{\eta} \quad (100)$$

$$\mathcal{T}_{0|_{\eta}} \left(\llbracket n_{\text{PRF}} \rrbracket_{\text{M}, \text{T}}^{\eta} \right) \stackrel{\text{ow.}}{=} \llbracket C[n_{\text{PRF}}] \rrbracket_{\text{M}, \text{T}}^{\eta} \quad (101)$$

Let $\mathcal{B}$ be a distinguisher between $C[n_{\text{PRF}}]$ and $C[\mathcal{H}(m, k)]$.

By a simple reorganization of terms in $\mathcal{T}_0$, we can apply Lem. 5 and get

$$\left| \mathbb{P} \left(\mathcal{B} \left(\mathcal{T}_{0|_{\eta}} \left(\mathcal{T}_{(98)|_{\eta}} \left(\mathcal{H}_k^{\mathcal{O}}(\vec{t}) \right) \right) \right) \right) - \mathbb{P} \left(\mathcal{B} \left(\mathcal{T}_{0|_{\eta}} \left(\llbracket n_{\text{PRF}} \rrbracket_{\text{M}, \text{T}}^{\eta} \right) \right) \right) \right| = \text{negl}(\eta) \quad (102)$$

Thus using (100) and (101):

$$\left| \mathbb{P} \left(\mathcal{B} \left(\llbracket C[\mathcal{H}(m, k)] \rrbracket_{\text{M}, \text{T}}^{\eta} \right) \right) - \mathbb{P} \left(\mathcal{B} \left(\llbracket C[n_{\text{PRF}}] \rrbracket_{\text{M}, \text{T}}^{\eta} \right) \right) \right| = \text{negl}(\eta) \quad (103)$$

that is $C[\mathcal{H}(m, k)] \sim C[n_{\text{PRF}}]$. $\square$

$$\begin{array}{c} \text{ENCKINSTANCE} \\ \frac{}{k, k', r, m \triangleright_{\text{enc}} \{m\}_{\text{pk}(k)}^r | h} \\ \\ \text{ENCKNONCE} \\ \frac{\vdash_{\text{P}} h \Rightarrow k \neq n \wedge k' \neq n \wedge r \neq n}{k, k', r, m \triangleright_{\text{enc}} n | h} \\ \\ \text{ENCKPK} \\ \frac{k'' = k \text{ or } k'' = k'}{k, k', r, m \triangleright_{\text{enc}} \text{pk}(k'') | h} \\ \\ \text{ENCKPKNEQ} \\ \frac{\vdash_{\text{P}} h \Rightarrow k \neq n \wedge k' \neq n \wedge r \neq n \quad k, k', r, m \triangleright_{\text{enc}} n | h}{k, k', r, m \triangleright_{\text{enc}} \text{pk}(n) | h} \\ \\ \text{ENCKFASTRONG} \\ \frac{k, k', r, m \triangleright_{\text{enc}} a | h \quad k, k', r, m \triangleright_{\text{enc}}^* b | h}{k, k', r, m \triangleright_{\text{enc}}^* a : : b | h} \\ \\ \text{ENCKFAWEAK} \\ \frac{k, k' \triangleright_{\text{enc}}^{\mathcal{O}} a | h \quad r \triangleright a | h \quad k, k', r, m \triangleright_{\text{enc}}^* b | h}{k, k', r, m \triangleright_{\text{enc}}^* a : : b | h} \end{array}$$

Figure 12: Rules for $k, k', r, m \triangleright_{\text{enc}} t|h$ extending Fig. 7

$$\begin{array}{c} \text{ENCNONCE} \\ \frac{\vdash_{\text{P}} h \Rightarrow k \neq n \wedge k' \neq n}{k, k' \triangleright_{\text{enc}}^{\mathcal{O}} n | h} \\ \\ \text{ENCOPK} \\ \frac{k = k'' \text{ or } k' = k''}{k, k' \triangleright_{\text{enc}}^{\mathcal{O}} \text{pk}(k'') | h} \\ \\ \text{ENCOPKNEQ} \\ \frac{\vdash_{\text{P}} h \Rightarrow k \neq n \quad k, k', r, m \triangleright_{\text{enc}} n | h}{k, k' \triangleright_{\text{enc}}^{\mathcal{O}} \text{pk}(n) | h} \\ \\ \text{ENCODDEC} \\ \frac{k, k' \triangleright_{\text{enc}}^{\mathcal{O}} a | h \quad k, k' \triangleright_{\text{enc}}^{\mathcal{O}} \text{pk}(b) | h}{k, k' \triangleright_{\text{enc}}^{\mathcal{O}} \text{dec}(a, b) | h} \end{array}$$

Figure 13: Rules for $k, k' \triangleright_{\text{enc}}^{\mathcal{O}} m|h$

C.6 Encryption

CRYPTOVAMPIRE2 supports symmetric and asymmetric encryption through two main rules

Using the SQUIRREL [5] st notation, we have:

DEFINITION 23. $k, k', r, m \triangleright_{\text{enc}} t|h$ means that for all $h' \equiv h$, there exist m' and t' such that

- (1) $m', t' \equiv m, t$
- (2) $h \Rightarrow \bigwedge_{(x, \vec{j}, c) \in \text{st}_{\text{pk}(\cdot)}^{(r)}} \forall \vec{j}, c \Rightarrow x \neq k$
- (3) $h \Rightarrow \bigwedge_{(x, \vec{j}, c) \in \text{st}_{\text{pk}(\cdot)}^{(r)}} \forall \vec{j}, c \Rightarrow x \neq k'$
- (4) $r \not\equiv t', m'$

DEFINITION 24. $k, k' \triangleright_{\text{enc}}^{\mathcal{O}} t|h$ means that for all $h' \equiv h$, there exists t' such that

- (1) $t' \equiv t$
- (2) $h \Rightarrow \bigwedge_{(x, \vec{j}, c) \in \text{st}_{\text{pk}(\cdot), \text{dec}(\cdot)}^{(r)}} \forall \vec{j}, c \Rightarrow x \neq k$
- (3) $h \Rightarrow \bigwedge_{(x, \vec{j}, c) \in \text{st}_{\text{pk}(\cdot), \text{dec}(\cdot)}^{(r)}} \forall \vec{j}, c \Rightarrow x \neq k'$

$$\begin{array}{c}
\text{DDHINSTANCE} \\
\hline
a, b, c \triangleright_{\text{ddh}} (g^a)^b | h \\
\\
\text{DDHNONCE} \\
\hline
\vdash_{\text{P}} h \Rightarrow a \neq n \wedge b \neq n \wedge c \neq n \\
\hline
a, b, c \triangleright_{\text{ddh}} n | h
\end{array}
\quad
\begin{array}{c}
\text{DDHEXP} \\
n = a \text{ or } n = b \\
\hline
a, b, c \triangleright_{\text{ddh}} g^n | h
\end{array}$$

Figure 14: Rules for $a, b, c \triangleright_{\text{ddh}} m | h$

PROPERTY 17. Extending Fig. 7 by Fig. 12 (resp. Fig. 13) while skipping the case $f = \text{pk}$ (resp. $f = \text{pk}, \text{dec}$) yields a valid set of rules modulo $\equiv$.

PROOF. The only notable point is that the check for the decryption oracle in ENCODEC is very similar to that for the pk oracle; thus, we refer to ENCPK and ENCPKNEQ. $\square$

PROPERTY 18 (IND-CCA). If $\{\_\}^-$ is an IND-CCA-secure asymmetric encryption scheme then:

$$\begin{array}{c}
\text{IND-CCA} \\
\hline
u \sim v \vdash C \left[\left\{ 0^{\|m\|} \right\}_{\text{pk}(k)}^r \right] \sim b \\
\hline
k, k', r, m \triangleright_{\text{enc}}^* C \left[\left\{ m \right\}_{\text{pk}(k)}^r \right] \text{true} \quad k, k' \triangleright_{\text{enc}}^{\mathcal{O}} m | \text{true} \\
\hline
u \sim v \vdash C \left[\left\{ m \right\}_{\text{pk}(k)}^r \right] \sim b
\end{array}$$

PROPERTY 19 (ENC-KP). If $\{\_\}^-$ is an ENC-KP-secure asymmetric encryption scheme then:

$$\begin{array}{c}
\text{ENC-KP} \\
\hline
u \sim v \vdash C \left[\left\{ \|m\| \right\}_{\text{pk}(k')}^r \right] \sim b \\
\hline
k, k', r, m \triangleright_{\text{enc}}^* C \left[\left\{ m \right\}_{\text{pk}(k)}^r \right] \text{true} \quad k, k' \triangleright_{\text{enc}}^{\mathcal{O}} m | \text{true} \\
\hline
u \sim v \vdash C \left[\left\{ m \right\}_{\text{pk}(k)}^r \right] \sim b
\end{array}$$

PROOF OF PROP. 18 AND PROP. 19. $k, \_, r, m \triangleright_{\text{enc}}^*$ lets us split between what is computed before submitting m to the challenger (so with a decryption oracle) and what is computed after.

The rules are adapted for the corrected version of [5]. $\square$

Prop. 19 can also be extended to symmetric encryption.

C.7 Decisional Diffie-Hellman

DEFINITION 25. $a, b, c \triangleright_{\text{ddh}} t | h$ means that for all $h' \equiv h$, there exists t' such that:

- (1) $m' \equiv m$
- (2) $h \Rightarrow \bigwedge_{(x, \vec{j}, c) \in \text{st}_{\mathbb{E}_g}(t')} \forall \vec{j}, c \Rightarrow x \neq a$
- (3) $h \Rightarrow \bigwedge_{(x, \vec{j}, c) \in \text{st}_{\mathbb{E}_g}(t')} \forall \vec{j}, c \Rightarrow x \neq b$
- (4) $h \Rightarrow \bigwedge_{(x, \vec{j}, c) \in \text{st}_{\mathbb{E}_g}(t')} \forall \vec{j}, c \Rightarrow x \neq c$

PROPERTY 20. Extending Fig. 7 by Fig. 14 yields sound rules modulo $\equiv$ for $a, b, c \triangleright_{\text{ddh}} t | h$.

PROOF. Interestingly, the No-guessing theorem (Th. 1) lets us skip removing g from DDHFA. $\square$

PROPERTY 21. If g is evaluated as the generator of a group where the Decisional Diffie-Hellman problem is hard, then the following rule holds

$$\begin{array}{c}
\text{DDH} \\
u \sim v \vdash C[g^c] \sim b \quad a, b, c \triangleright_{\text{ddh}} C[(g^a)^b] | \text{true} \\
\hline
u \sim v \vdash C[(g^a)^b] \sim b
\end{array}$$

PROOF. Using transitivity on the DDH rule of [5]. $\square$

C.8 Xor

Re-using the list-like construction introduced in Section C.4, we have Prop. 22.

PROPERTY 22 (XOR). When $\oplus$ represents the XOR operator,

$$\begin{array}{c}
\text{XOR} \\
u \sim v \vdash n : : \vec{a} \sim \vec{b} \quad n \triangleright t \quad n \triangleright \vec{a} \quad \vdash_{\text{P}} \|t\| \equiv \eta \\
\hline
u \sim v \vdash (n \oplus t) : : \vec{a} \sim \vec{b}
\end{array}$$

where $\|\_\|$ evaluates to the length of the bitstring.

PROOF. Lifted from [5]. $\square$