Leveraging Cryptographic Simulator Synthesis for Formally Verifying the FOO E-Voting Protocol

David Baelde
Univ Rennes, CNRS, IRISA

Adrien Koutsos
Inria

Justine Sauvage
Inria

Abstract

Cryptographic proofs proceed in large part by reductions to cryptographic assumptions expressed as games. These reductions rely on simulators which are often tedious to write and involve a significant amount of trivial code. Thus, simulators are only sketched in pen-and-paper proofs, which is error-prone. Mechanized cryptographic proofs remove the risk of errors, but requiring users to explicitly write simulators is an unreasonable burden.

In this paper, we consider the problem of simulator synthesis in Squirrel, where cryptographic simulation is expressed as bi-deduction. Although the seminal work on bi-deduction [5] provides a proof system and a simple proof-search procedure for it, we show that it suffers from systematic failures when working with games such as IND-CCA2. We provide a significantly improved procedure, that can re-use oracle calls across recursive iterations, and generates precise invariants to justify it. We implement this procedure in Squirrel and validate it in a proof of ballot privacy for the FOO e-voting protocol, which is the first computational mechanized proof for FOO, and the most complex Squirrel proof to date.

1 Introduction

Cryptographic proofs are often structured as sequences of "game hops" [39]. Some of these "hops" are reductions to cryptographic assumptions: assuming the security of some game $\mathcal{G}$ (which is composed of two computationally indistinguishable implementations $\mathcal{G}_0$ and $\mathcal{G}_1$) we can prove that another game $\mathcal{P}$ is secure by exhibiting an adversary in $\mathcal{G}$ that can simulate $\mathcal{P}$; indeed, a distinguisher for $\mathcal{P}$ composed with that simulator would yield a distinguisher for $\mathcal{G}$. Writing simulators in detail is tedious, involving a lot of boilerplate code for a few interesting steps. As a result, they are not detailed in pen-and-paper proofs. An ideal goal for computer-aided cryptography would be to keep the user free from the burden of writing down simulators, while checking in full details that correct simulators exist. In other words, it is desirable to automatically synthesize correct cryptographic simulators.

Code-centric view of simulator synthesis. We follow a code-centric view [11] of the simulator synthesis problem. In this paradigm, the cryptographic assumption $\mathcal{G} = (\mathcal{G}_0, \mathcal{G}_1)$ is typically represented by a pair of implementations of the same oracle: said otherwise, $\mathcal{G}$ is a pair of two APIs with the same signature. The target game is a pair of programs $(\mathcal{P}_0, \mathcal{P}_1)$ which must be shown indistinguishable. As a typical example, we express indistinguishability between a pair of protocols $(\mathsf{P}_0, \mathsf{P}_1)$ by taking $\mathcal{P}_i$ to be a program representing N interactions between an adversary $\mathcal{A}$ and P_i:

```
for each k ∈ [0;N[ {
  input ← A(frame);          (* the adversary provides an input *)
  output ← P_i(input);       (* forward the input to the protocol *)
  frame ← ⟨frame, output⟩;}  (* add output to A's knowledge *)
return frame                 (* return all outputs for distinguishability test *)
```

To reduce $\mathcal{P}$ to $\mathcal{G}$, we must exhibit a single simulator $\mathcal{S}$ such that $\mathcal{S}$ computes $\mathcal{P}_0$ from $\mathcal{G}_0$ and $\mathcal{P}_1$ from $\mathcal{G}_1$. Following the structure of $\mathcal{P}$, the simulator will also feature a main loop over $[0;N[$. The simulator may call the adversary $\mathcal{A}$ as a subroutine. Thus, to simulate the body of the adversary/protocol main loop, the key difficulty lies in dealing with the protocol call $\mathsf{P}_i(\mathsf{input})$, which we must simulate using the functions provided by the API $\mathcal{G}_i$ — the differences between P_0 and P_1 must be obtained from the differences between $\mathcal{G}_0$ and $\mathcal{G}_1$ because the simulator $\mathcal{S}$ must work for both sides i. There are two difficulties there. First, we must deal with the fact that $\mathcal{G}_i$ is a *stateful* API. Thus, when we call an oracle $\mathcal{G}_i.f$, we are obtaining an output which depends on the current state of the game *and* modifying this state, which will impact future outputs. Further, $\mathcal{G}_i$ is *probabilistic*, which makes it necessary to relate the random samplings of $\mathcal{G}_i.f$ with that of P_i.

In summary, the simulator synthesis problem requires to synthesize a *looping* (or *recursive*) simulator which must exploit a *stateful* and *probabilistic* API.

CCSA and Squirrel. In this paper, we focus on the CCSA logics approach [2] implemented in the Squirrel proof assistant [26]. This approach combines a general-purpose logic

with an abstract treatment of cryptography, yielding concise proofs and easier automation. In particular, recent work [5] has shown that the cryptographic axioms used in CCSA logics since their inception [7] can be replaced by a general notion of *bi-deduction*, which is a technique allowing to establish the existence of a simulator using a cryptographic game. In [5], a proof system and proof-search procedure showed that simulator synthesis could be automated. They dealt with the three difficulties we identified as follows: the *statefulness* of the game is tracked using Hoare-style pre- and post-conditions; the *probabilistic* behaviors of the game and the target protocol are related using *probabilistic couplings* [9, 41] (technically, couplings are implicitly constructed through so-called *name constraints*); and the *looping* or *recursive* aspect is handled by an ad hoc inductive invariant inference technique. They implemented their technique as a tactic, called **crypto**, in Squirrel.

However, while the **crypto** tactic of [5] was successfully used to synthesize simulators against several games (PRF, EUF, CPA$, ...), we discovered that it lacks in flexibility and expressivity. More precisely, it fails to infer invariants that can deal with some key cryptographic assumptions, including the pervasive IND-CCA2 game. In §2, we detail an example suffering from this limitation, which allows us to identify the features missing from the simulator synthesis of [5]: the generated simulators cannot *memoize values across iterations*; further, synthesis cannot infer the *time-sensitive* invariants needed to justify the correction of the more complex simulators needed to support games such as IND-CCA2.

Contributions. In this paper, we significantly improve the bi-deduction based simulator synthesis technique of [5], obtaining a tool that can be used at scale.

i) Our main contribution is a new approach to *inductive simulator synthesis* (§5), i.e. support for recursively defined terms by the synthesis procedure. The new procedure generates simulators that memoize the output of game oracles across recursive iterations, and synthesizes times-sensitive invariants that are necessary to establish the correctness of such simulators. This solves systematic shortcomings of the earlier procedure when considering practical protocols and, e.g., CCA2. Moreover, our approach is terminating while the original procedure, relying on a fixed point computation, could in principle diverge — though we do not know of an example where this happened.

ii) As a secondary contribution (§4), we improve on the *basic synthesis* procedure of [5] and formally describe it, which is key to proving correct our inductive synthesis technique.

iii) We implement our algorithms as a richer version of Squirrel's **crypto** tactic, which we also improve on various practical aspects, including error reporting and efficiency.

iv) Building on the above contributions, we develop a formal proof of vote privacy for the FOO [29] e-voting protocol in Squirrel. As we will see (§2, §7), our novel inductive synthesis procedure is critical to deal with the complexity of the proof. Our security proof is based on the CCSA pen-and-paper proof of [6], which we generalize to an arbitrary number of voters. Our proof is the most complex Squirrel proof to date, both in terms of the diversity of the cryptographic assumptions, and in lines of code. Further, it is the first computational mechanized proof of ballot privacy for FOO.

Limitations. There are several limitations to our work that are worth discussing. First, our synthesis procedure is not complete. Like [5], we do not aim for completeness but focus on expressive synthesis while preserving predictability, so that the user can understand how to prepare the ground for **crypto**. Compared with [5], we improve on expressivity while retaining predictability. Second, we empirically argue for the expressiveness and usefulness of our approach through practical case-studies. While this paper presents some sizable early results in that direction (with several case-studies, including the large-scale proof of FOO), further work would be needed to thoroughly evaluate the usefulness of our approach across a large range of application scenarios. Third, our approach can be applied to a non-adaptive setting, or to an adaptive setting when the number of sessions is independent of the security parameter. Dealing with polynomially many sessions requires to finely control the simulator's complexity and to have precise advantage bounds (see [3]). Doing so in a fully automated approach like ours presents interesting but orthogonal (and thus out-of-scope) challenges.

We also stress that our proof of vote privacy for FOO should *only* be viewed as a large-scale validation of our techniques. It must not be taken as an encouragement to use FOO for any purpose. Indeed, we only consider a simple form of vote privacy in this paper, but stronger properties might be desirable. Further, other important properties should be carefully considered when choosing a voting protocol, such as verifiability or coercion resistance. Finally, formal proofs of protocols are only one part of a security analysis, which does not take into account implementation-level flaws and human factors.

Artifacts. The companion artifacts for this paper, including our extended version of Squirrel with our improved **crypto** tactic as well as all the proof developments, are open source and available online [37]. Further, our improved version of **crypto** has been merged into the main branch of Squirrel [26]. Details and proofs can be found in the full version of this paper, available online [38].

Outline. The rest of the paper is structured as follows: we give in §2 a motivating example which concretely shows the need for memoizing simulators and time-sensitive invariants; we present in §3 the formal setting of bi-deduction; and in §4 the (improved) basic proof-search procedure for it; we introduce in §5 our inductive proof-search technique; finally,

```
game G_b = {                    (* defines the games G_b, for b ∈ {0,1} *)
  sk ←$; ℓ ← [];                (* initializes sk and ℓ *)
  oracle pub() := { return (pub sk) }
  oracle left-right(m_0, m_1) := { r ←$;
                                   e ← enc m_b r (pub sk);
                                   ℓ ← e :: ℓ;
                                   return (if |m_0| = |m_1| then e else 0) }
  oracle decrypt(c) := { return (if c ∉ ℓ then dec c sk else 0) }}.
```

$x \leftarrow e$ assigns x to the value of e; and $x \xleftarrow{\$}$ assigns to x a randomly sampled value using a distribution based on the type of x.

Figure 1: The IND-CCA2 cryptographic game.

we present in §7 the FOO protocol and our formal proof of ballot privacy for it; and compare with related work in §8.

We use colors to help distinguish elements of distinct categories. To ensure accessibility to color-blind readers, colors only encode redundant information.

2 Motivating Example: an Abstract Mixnet

We illustrate the key difficulties encountered when synthesizing simulators for the FOO protocol. FOO uses mixnets [16, 42] to achieve vote privacy. We use a high-level abstract modeling of mixnets (in the style of [6]), which works in two phases. In the collection phase, voters sends encrypted ballots to the mixnet, using its public key (pub k). The ballots are collected by the mixnet sub-process M, which decrypts them and stores them into the ballot-box bb. Here, we consider a simple setting with a single honest voter V, but many sessions of M — dishonest voters are not explicitly modeled as we let the adversary play for them. Further, we let the adversary control V's vote, which V thus inputs from the network. After the collection phase, the mixnet sub-process P shuffles the ballot-box bb containing decrypted ballots, and publishes it.

Modeling. We define in Fig. 1 the IND-CCA2 games for our encryption scheme, where: (pub k) is the public key associated to a private key k; (enc m r pk) is the encryption of plaintext m using public key pk and randomness r; (dec c k) is the decryption of ciphertext c using private key k. The log ℓ prevents the trivial attack in which the left-right challenge is sent to the decryption oracle. The IND-CCA2 assumption states that a probabilistic polynomial-time adversary with access to the game's oracles has a negligible probability of distinguishing whether it is interacting with $\mathcal{G}_0$ or $\mathcal{G}_1$.

We describe our abstract mixnet protocol using Squirrel's process algebra in Fig. 2. We actually define *two* protocols through two process variants obtained by projecting the diff operators to their first or second component. The process obtained by projecting the diff operators to their first component corresponds to our abstract mixnet setting. We mod-

```
system MixNet =
 new r; new k;                          (* sample r and the secret key k *)
 Pk: out(c, pub k) |                    (* publish the public key pub k *)
 V: in(c,x); out(c, enc diff(x,dummy) r (pub k)) |   (* voter send its ballot *)
 M: !_i in(c,y);                        (* decrypt ballots and store them in bb *)
      bb(i) := if V < M(i) &&
                  y = enc diff(input V,dummy) r (pub k)
               then input V else dec y k) |
 P: out(c, shuffle(bb)).                (* publish stored decrypted ballots *)
```

We describe processes using the applied π-calculus [1]. The special value c denotes an arbitrary channel over which communications are taking place (e.g. the Internet). Instruction out(c,t) outputs t over c, while in(c,x) inputs a value over c and stores it in x. Then, $!_i$ P denotes the replication of the process P, where each replicated instance of P may use the session identifier i. The test V < M(i) indicates that the voter's output takes place before session i of the mixnet: this test, which cannot be implemented by a real-world process, is a modeling artifact which helps with the security analysis.

Figure 2: An abstract mixnet protocol.

ify the output of the mixnet sub-process M's using a conditional that will be instrumental when reasoning with the IND-CCA2 game, but that does not change M's behavior: indeed, (input V = dec y k) if (y = enc (input V) r (pub k)). The process obtained by taking the second projection differs in that V will encrypt a dummy message instead of its input, and M will "magically" retrieve V's input when receiving that encryption. This variant is an *idealized* version of our protocol.

Cryptographic reduction. It should be quite obvious that an adversary will not be able to distinguish whether it is interacting with the real or idealized version of the protocol, assuming that it feeds V with an input that has the same length as dummy. Showing that a protocol is indistinguishable from an idealization is a common step in cryptographic proofs: e.g., in the actual FOO protocol, it is a key step to prove that the privacy of V's vote is preserved (as we will see in §7).

To formally show that our protocol is indistinguishable from its idealized version, it suffices to exhibit an IND-CCA2 adversary that can simulate our protocols. More formally, we need a *single* IND-CCA2 adversary (the *simulator*) such that, given a trace describing the interaction of an adversary with the protocol, the simulator computes the resulting sequences of protocol outputs using the IND-CCA2 oracles: when running with $\mathcal{G}_0$ it produces the outputs of the real protocol; when running with $\mathcal{G}_1$ it produces idealized outputs.

We use a non-adaptive setting in which the adversary must interact with the protocol along a fixed trace. More precisely, we consider a finite, totally ordered set of *timestamps*, where a timestamp is an element of the data-type:

$$\mathsf{t} ::= \mathsf{init} \mid \mathsf{Pk} \mid \mathsf{V} \mid \mathsf{M(i)} \mid \mathsf{P}$$

For example, V denotes the timepoint at which the honest

```
1  frame, output, input ← [ ⊥ for t ∈ [ init; t0 ] ]  (* initialize arrays *)
2  bb ← [ ⊥ for i ∈ index ]            (* one more array initialization *)
3  ballotV ← None;                     (* to memoize V's encrypted ballot *)
4  input[init], output[init], frame[init] ← empty  (* initial timepoint*)
5  (* recursively compute input, output, frame and bb *)
6  for each t ∈ ]init; t0 ] {
7    input[t] ← att(frame[pred t]) (* the adversary computes the input *)
8    begin  (* simulate output[t] by case analysis on t *)
9      match t with
10     | Pk → output[t] ← G.pub()
11     | V →
12       ballotV ← Some (G.left-right(input[t], dummy)) (* memoize *)
13       output[t] ← Option.get(ballotV)
14     | M(i) →
15       (* in the condition below, we retrieve V's ballot from ballotV *)
16       if V < M(i) && input[t] = Option.get(ballotV)
17       then { bb[i] ← input[V] } (* bypass decryption oracle *)
18       else { bb[i] ← G.decrypt(input[t]) } (* safe decryption oracle call *)
19       output[t] ← empty (* no output there *)
20     | P → output[t] ← shuffle(bb)
21   end
22   frame[t] ← ⟨ frame[pred t], output[t] ⟩ (* add output to the frame *)
23 }
24 return (frame[t0])
```

We have Option.get (Some x) = x **and** Option.get(None) = ⊥**.**

Figure 3: Reduction to the IND-CCA2 assumption.

voter will send its encrypted ballot, and M(i) the timepoint at which the mixnet with session identifier i collects, decrypts and stores a ballot — we use the abstract type index to denote session identifiers. Timestamps are ordered by $<$, and for t $\neq$ init, we let (pred t) denote the timestamp preceding t w.r.t. $<$.

Fig. 3 shows a simulator $\mathcal{S}$ witnessing that the concrete and ideal mixnet protocols are indistinguishable up-to some timepoint t_0 by reduction to the IND-CCA2 game (noted $\mathcal{G}$). Line 1, the simulator $\mathcal{S}$ initializes a number of timestamp-indexed arrays to store intermediate values: output[t] and input[t] store, resp, the output and input of the protocol at timepoint t; frame[t] is the sequence of all outputs from init to t included. The array cell bb[i] (line 2) will store the decrypted ballot processed by M(i). Finally, ballotV is initialized to None (line 3), and will be used to memoize V's ballot.

The simulator's main loop (lines 7–23) iterates over the timestamps in the trace]init; t_0]. For each such timestamp t, the input at time t is obtained (line 7) as the result of an attacker computation att($\cdot$) — modeled as an arbitrary unspecified procedure — taking all previous outputs as arguments: input[t] = att(frame[pred(t)]). Then, the next output is simulated (lines 8–21) depending on which action is considered. For t = V, we can call the **left-right** oracle — we assume that len (input[V]) = len dummy. For t = M(i), we need to distinguish whether V has occurred before or not.

If V has not occurred before M(i), then the log ℓ is empty and we can decrypt any message computed by the simulator, including input[M(i)]: it can thus simulate output[M(i)] easily.

Otherwise V $<$ M(i), and we cannot decrypt output[V] because it has been obtained from the **left-right** oracle and is thus in the game's log ℓ. Fortunately, output[M(i)] is written in such a way that this forbidden decryption is avoided. Note, however, that the simulator needs to test whether input[M(i)] = enc diff(input[V],dummy) = output[V], thus it needs to use again the result of the call to **left-right** performed in V — calling the oracle again at this point would yield an encryption with a different random seed. Instead, our simulator exploits the fact that Option.get(ballotV) = output[V] when V $<$ M(i), i.e. the voter's encrypted has been memoized in ballotV.[1]

Finally, the sequence of all outputs up-to timepoint t is computed and stored in frame[t] (line 22).

The above analysis shows two key features (absent from [5]) needed for our simulator. First, we need simulators that **memoize the result of oracle calls** across recursive calls: the value of an oracle call performed in V must be reused later for M(i). Second, proving the correctness of such a simulator requires an invariant that tracks the value of the game's state (here, the log ℓ) after each step of the simulator's recursive process. Crucially, **time-sensitive invariants** are necessary, to express in our example that the log is empty before V but contains one element after it. Without such an invariant we would have to show that input[M(i)] $\neq$ enc diff(input[V],dummy) r k for the **else** branch, which is unnecessary.

3 Higher-Order CCSA and Bi-Deduction

We recall the key elements of the CCSA theoretical framework, sticking to syntactic aspects which are sufficient for our purposes here. Full details may be found in [4, 5].

3.1 Logic

We work within higher-order CCSA [4], using typed λ-calculus to describe (sets of) messages, propositions, etc.

Types, noted τ, are commonly seen as sets of values. In our setting, these sets vary with the security parameter η: for instance, a type of cryptographic keys might be interpreted as bitstrings of length η. Thus we interpret types as *η-indexed collections of values*. We assume an arbitrary set of base types, noted $\tau_0, \tau_0', \ldots$, featuring e.g. bitstring interpreted as $\{0,1\}^*$ for any η. A type can be tagged as finite, noted $\texttt{finite}(\tau)$, meaning that it must be interpreted as an η-indexed collection of *finite* sets of values. Further, the interpretation of a fixed type, noted $\texttt{fixed}(\tau)$, must be the same all η. E.g., bitstring is not finite, but bool and index are, and all three types are fixed.

We assume a set $\mathcal{X}$ of typed variables. Variables are introduced in *environments*, noted $\mathcal{E}, \mathcal{E}_0, \ldots$, which are collections of *declarations* and *definitions*: the former are noted $x : \tau$, while the latter are noted $x : \tau = t$. Definitions can be mutually

[1] In our simple example, we could get rid of ballotV and use output[V] instead. However, oracle calls are generally not directly used as outputs in protocols, which requires the use of memoization variables as done here.

recursive and must satisfy a well-foundedness condition [4] ensuring that x can be interpreted as t in all models.

As in all CCSA logics [2,7], we make use of special objects called *names* to represent (honest) randomness. Specifically, we assume distinguished *name symbols* indexed over a finite type, e.g. $\mathsf{n} : \mathsf{index} \to \mathsf{bitstring}$. A *name* is obtained by applying one such symbol to an index, as in $(\mathsf{n}\ i)$. Distinct names of the same type are interpreted as identical but independent (polynomial-time computable) probability distributions.

Terms are noted using the letters t, u, v and are formed from variables and names using application and λ-abstraction subject to standard typing rules. We use a conditional construct $(\mathsf{if}\ u\ \mathsf{then}\ v\ \mathsf{else}\ w)$. We write $\mathsf{type}_{\mathcal{E}}(t)$ the type of t in $\mathcal{E}$. A term of type τ is interpreted as an η-indexed collection of random variables ranging over the interpretation of τ [4]. The semantics of term constructors is the expected one, lifted to collections of random variables.

Example 1. *The mixnet condition from our motivating example can be written as the following term of type* bool*:*

$$\mathsf{V} < \mathsf{M}(i) \ \wedge\ \mathsf{input}\ (\mathsf{M}(i)) = \mathsf{enc}\ (\mathsf{input}\ \mathsf{v})\ (\mathsf{r}\ \langle\rangle)\ (\mathsf{pub}\ (\mathsf{k}\ \langle\rangle))$$

Here, k *and* r *are name symbols indexed over* unit*, and* i *is a variable of type* index*. Other symbols are variables, which we view as constant and function symbols, that would be declared in some global environment. For instance,* enc *has type* bitstring → rand → pkey → bitstring*;* input *has type* timestamp → bitstring*;* M *has type* index → timestamp*; the function symbols* $\wedge$ *and* $=$*, used in infix notation, have types, resp.,* bool → bool → bool *and* bitstring → bitstring → bool*.*

The higher-order CCSA logic features two kinds of formulas. First, *local formulas* (noted f, g) are simply terms of type bool. Local formulas make use of logical connectives and equality, with standard syntax and semantics. For example, the local formula $\forall i, j : \mathsf{index}.\ (\mathsf{n}\ i) = (\mathsf{n}\ j) \Rightarrow i = j$ states that distinct instances of n cannot be equal — we use here a symbol $\forall : (\mathsf{index} \to \mathsf{bool}) \to \mathsf{bool}$. Second, *global formulas* (noted F, G) are first-order formulas built over specific predicates. In this paper, we will only make use of the overwhelming (resp. exact) truth predicate $[f]$ (resp. $[f]_{\mathrm{e}}$) which expresses that a local formula f is true with overwhelming probability (resp. true for all η and all random samplings). For example, we have $[f]$ but not $[f]_{\mathrm{e}}$ for the local formula f of the previous example: in our models, there is a negligible but non-zero probability that $(\mathsf{n}\ i)$ and $(\mathsf{n}\ j)$ collide.

Given two terms u_0, u_1 of the same type, the *bi-term* $\#(u_0; u_1)$ represents a pair of left and right scenarios. We factorize common behavior between bi-terms. For example, $f(\#(u_0; u_1))$ and $\langle u, h(\#(v_0; v_1))\rangle$ denote, resp., $\#(f(u_0); f(u_1))$ and $\#(\langle u, h(v_0)\rangle; \langle u, h(v_1)\rangle)$. We write bi-terms in **bold** (e.g. $\boldsymbol{t}, \boldsymbol{u}, \boldsymbol{f}$) to distinguish them from standard terms. Finally, we write $[\boldsymbol{f}]$ for the (global-level) conjunction $[f_0] \wedge [f_1]$ — and similarly for $[\boldsymbol{f}]_{\mathrm{e}}$.

3.2 Bi-Deduction

We follow [5] and formally represent the existence of a simulator using the notion of cryptographic bi-deduction. This notion is parameterized by a *cryptographic game* $\mathcal{G} = (\mathcal{G}_0, \mathcal{G}_1)$, where both game variants $\mathcal{G}_i$ implement the same set of oracles, which the adversary can access. Oracles are implemented in a probabilistic imperative programming language, whose details are not relevant here. These programs may use a subset of the available function symbols of the logic, given through a particular environment $\mathcal{L}$ called the *standard library*. It is assumed that library functions are computable in deterministic polynomial time. Expressions of our programming language, which we call *program terms* are built using standard library functions (and constants) and *program variables*, which we will denote $\ell, \ell_0, \ldots$. In contrast, *logical terms*, i.e. terms of the logic, are built using function symbols from larger environments than $\mathcal{L}$ and using logical variables from $\mathcal{X}$ rather than program variables.

Roughly, the bi-deduction $\boldsymbol{u} \rhd_{\mathcal{G}} \boldsymbol{v}$ expresses that there exists an adversary $\mathcal{S}$ w.r.t. the game $\mathcal{G}$, called the simulator, such that $\mathcal{S}^{\mathcal{G}_i}(u_i) = v_i$ for each $i \in \{0, 1\}$, i.e. $\mathcal{S}$ computes v_i when ran on input u_i with oracles $\mathcal{G}_i$. In full details, a bi-deduction judgement is of the following form:

$$\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C}; (\boldsymbol{\phi}, \boldsymbol{\psi}) \vdash \boldsymbol{u} \rhd_{\mathcal{G}} \boldsymbol{v}$$

The environment $\mathcal{E}$ and set of global formulas $\boldsymbol{\Theta}$ properly describe the logical setting of the bi-deduction — bi-deducing $\boldsymbol{v}$ may crucially rely on a definition from $\mathcal{E}$ or an hypothesis from $\boldsymbol{\Theta}$. The name constraints $\boldsymbol{C}$ and assertions $\boldsymbol{\phi}$ and $\boldsymbol{\psi}$ track key aspects of the simulator's behavior and the game's memory; we describe next these two important ingredients.

Name constraints. As usual in computer-aided cryptography [9], $\mathcal{S}^{\mathcal{G}_i}$ does not directly compute logical terms such as v_i: instead, there is a probabilistic coupling between the computation's result and the interpretation of these terms [5]. Although this can largely be ignored here, one important aspect is that we must track the random samplings performed in $\mathcal{S}^{\mathcal{G}_i}$ and the corresponding names used in $\boldsymbol{v}$. Indeed, we must distinguish a name representing a random sampling performed directly by the simulator, a name representing a random sampling performed during the initialization of the game, and a random sampling performed during an oracle call. We do so using *name constraints* of the form $c = (\vec{\alpha}, \mathsf{n}, t, T, f)$ where $\vec{\alpha}$ is a sequence of variables, n is a name symbol, t is a term, f is a local formula, and T is a tag which is either $\mathtt{T}_S$ (sampling by the simulator), $\mathtt{T}^{\mathsf{loc}}_{\mathsf{G}}$ (sampling inside an oracle call) or $\mathtt{T}^{\mathsf{glob}}_{\mathsf{G},v}$ (sampling of the game's global variable v).

We use constraint systems which are multisets of name constraints. The bi-deduction judgement features one constraint system for each side of the game, i.e. $\boldsymbol{C} = \#(C_0; C_1)$. We write $\boldsymbol{C} \cdot \boldsymbol{C'}$ for the (component-wise) union of constraint systems.

Example 2. *We consider names* $\mathsf{k} : \mathsf{unit} \to \mathsf{skey}$, $\mathsf{r} : \mathsf{index} \to \mathsf{rand}$ *and* $\mathsf{n} : \mathsf{index} \to \mathsf{bitstring}$*, and we assume that latter names have the same length as* dummy*. We can bi-deduce, in the CCA2 game and with an arbitrary input* $\boldsymbol{u}$*, the bi-term* $\boldsymbol{v} := \lambda i : \mathsf{index}.\ \mathsf{enc}_{\#}(\mathsf{n}\, i; \mathsf{dummy})\ (\mathsf{r}\, i)\ (\mathsf{pub}\ (\mathsf{k}\ \langle\rangle))$ *which can be understood as a (bi-)array of encryptions for all elements of* index*. This bi-deduction is possible with the following constraint system (the same on both sides of the simulation):*

$$\boldsymbol{C} := \{(i, \mathsf{n}, i, \mathtt{T}_S, \top),\ (i, \mathsf{r}, i, \mathtt{T}^{\mathsf{loc}}_{\mathsf{G}}, \top),\ (\epsilon, \mathsf{k}, \langle\rangle, \mathtt{T}^{\mathsf{glob}}_{\mathsf{G},sk}, \top)\}$$

Assertions. In order to ensure the correctness of a simulator, and reason compositionally, we track the game's memory in the style of a Hoare logic: the assertions $\boldsymbol{\phi}$ and $\boldsymbol{\psi}$ are pre- and post-conditions on the game's memory. Taking these into account, bi-deduction states that if $\mathcal{S}^{\mathcal{G}_i}(u_i)$ is ran when the game's memory satisfies ϕ_i, it will return v_i and the final game memory will satisfy ψ_i.

Although the theory of [5] does not rely on a specific assertion language, their implementation of the **crypto** tactic uses a simple assertion language tailored for tracking the typical logs of cryptographic games. We briefly describe that language, referring the reader to the long version [38] for details. We assume a base type bitstring set for sets of bitstrings. We make use of *symbolic sets* which are sequences $(s_1, \ldots, s_k)$ where each s_i is of the form $\{t_i \mid \vec{\alpha}_i : f_i\}$ where t_i is a logical term of type bitstring set, f_i is a local formula, and $\vec{\alpha}_i$ is a sequence of (typed) variables bound in both t_i and f_i. Intuitively, a symbolic set denotes the set of all elements t_i, for arbitrary values of $\vec{\alpha}_i$ such that f_i holds. Then, an *assertion* ϕ is a finite map from program variables to symbolic sets, and a memory μ satisfies ϕ if, for each variable ℓ in the domain of ϕ, we have that $\mu(\ell)$ is included in the interpretation of $\phi(\ell)$ — the assertion *over-approximates* the sets present in memory.

Example 3. *Let* $\mathsf{pk} := \mathsf{pub}\ (\mathsf{k}\ \langle\rangle)$*. The bi-deduction of Example 2 is valid with the following pre- and post-conditions, for arbitrary bi-symbolic sets* $\boldsymbol{S}$ *and* $\boldsymbol{S'}$*:*

$$\boldsymbol{\phi} := (\ell \mapsto \boldsymbol{S})$$
$$\boldsymbol{\psi} := \left(\ell \mapsto \boldsymbol{S}, \boldsymbol{S'}, \{\mathsf{enc}_{\#}(\mathsf{n}\, i; \mathsf{dummy})\ (\mathsf{r}\, i)\ \mathsf{pk} \mid i : \mathtt{T}\}\right)$$

Indeed, the post-condition covers the elements added to the log: encryptions of $(\mathsf{n}\, i)$ *on the left, and of* dummy *on the right. It is correct to add* $\boldsymbol{S'}$ *as assertions are over-approximations.*

Example 4. *In our mixnet example, the game's memory at iteration* t *satisfies the bi-assertion* $\ell \mapsto \{\#(\mathsf{input@V}; \mathsf{dummy}) \mid \epsilon : \mathsf{V} < t\}$*. The condition* $\mathsf{V} < t$ *is not necessary for the correction of the assertion, but it makes it more precise and simplifies the verification of calls to the decryption oracle before* V*.*

3.3 Proof System

Cryptographic bi-deduction may be established by means of derivation rules provided in [5]. We briefly recall this proof system, using some selected rules shown in Fig. 4. The rules deal with bi-deduction judgments as introduced above, except that inputs and outputs are sequences of bi-terms. Moreover, we commonly use the syntactic sugar $(\boldsymbol{t} \mid \boldsymbol{f})$ for $\langle \boldsymbol{f}, \text{if } \boldsymbol{f} \text{ then } \boldsymbol{t} \text{ else dummy}\rangle$ where $\mathsf{dummy} \in \mathcal{L}$ is an arbitrary public constant of the appropriate type. Hence $\boldsymbol{u} \rhd (\boldsymbol{t} \mid \boldsymbol{f})$ means that we only have to compute $\boldsymbol{t}$ when $\boldsymbol{f}$ holds, but we also have to bi-deduce the condition $\boldsymbol{f}$.

$$\rhd.\textsc{Refl}\quad \frac{\boldsymbol{v} \in \boldsymbol{u}}{\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C}; (\boldsymbol{\phi}, \boldsymbol{\phi}) \vdash \boldsymbol{u} \rhd \boldsymbol{v}} \qquad \rhd.\textsc{Transitivity}\quad \frac{\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C}; (\boldsymbol{\phi}, \boldsymbol{\phi'}) \vdash \boldsymbol{u} \rhd \boldsymbol{v} \quad \mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C'}; (\boldsymbol{\phi'}, \boldsymbol{\psi}) \vdash \boldsymbol{u}, \boldsymbol{v} \rhd \boldsymbol{w}}{\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C} \cdot \boldsymbol{C'}; (\boldsymbol{\phi}, \boldsymbol{\psi}) \vdash \boldsymbol{u} \rhd \boldsymbol{v}, \boldsymbol{w}}$$

$$\rhd.\textsc{Oracle}_f\quad \frac{\begin{array}{c}\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C}; (\boldsymbol{\phi}, \boldsymbol{\phi'}) \vdash \boldsymbol{u} \rhd_{\mathcal{G}} (\boldsymbol{t} \mid \boldsymbol{g}), (\boldsymbol{o} \mid \boldsymbol{g}), (\boldsymbol{s} \mid \boldsymbol{g}) \\ \vec{\boldsymbol{k}} = (\mathsf{k_v}\ \boldsymbol{o_v})_{\mathsf{v} \in \mathcal{G}.\mathsf{glob}_\$} \qquad \vec{\boldsymbol{r}} = (\mathsf{r_v}\ \boldsymbol{s_v})_{\mathsf{v} \in f.\mathsf{loc}_\$} \\ \boldsymbol{C'} = \left(\Pi_{\mathsf{v} \in \mathcal{G}.\mathsf{glob}_\$}(\emptyset, \mathsf{k_v}, \boldsymbol{o_v}, \mathtt{T}^{\mathsf{glob}}_{\mathsf{G},\mathsf{v}}, \boldsymbol{g})\right) \cdot \left(\Pi_{\mathsf{v} \in f.\mathsf{loc}_\$}(\emptyset, \mathsf{r_v}, \boldsymbol{s_v}, \mathtt{T}^{\mathsf{loc}}_{\mathsf{G}}, \boldsymbol{g})\right) \\ \mathcal{E}; \boldsymbol{\Theta} \models \{\boldsymbol{\phi'}, \boldsymbol{g}\} \boldsymbol{v} \leftarrow O_f(\boldsymbol{t})[\boldsymbol{k}; \boldsymbol{r}]\{\boldsymbol{\psi}\}\end{array}}{\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C} \cdot \boldsymbol{C'}; (\boldsymbol{\phi}, \boldsymbol{\psi}) \vdash \boldsymbol{u} \rhd (\boldsymbol{v} \mid \boldsymbol{g})}$$

$$\rhd.\textsc{Rewrite-R}\quad \frac{\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C}; (\boldsymbol{\phi}, \boldsymbol{\psi}) \vdash \boldsymbol{u} \rhd \boldsymbol{v'} \qquad \mathcal{E}; \boldsymbol{\Theta} \vdash [\boldsymbol{v} = \boldsymbol{v'}]_{\mathrm{e}}}{\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C}; (\boldsymbol{\phi}, \boldsymbol{\psi}) \vdash \boldsymbol{u} \rhd \boldsymbol{v}}$$

Figure 4: Selected bi-deduction rules.

Several rules correspond to basic building blocks for simulators. For instance, the reflexivity rule $\rhd.\textsc{Refl}$ corresponds to a simulator that simply outputs one of its inputs. In the transitivity rule $\rhd.\textsc{Transitivity}$, we compose a simulator computing $\boldsymbol{v}$ from $\boldsymbol{u}$ with another one computing $\boldsymbol{w}$ from $\boldsymbol{u}$ and $\boldsymbol{v}$ to obtain a simulator computing $\boldsymbol{v}, \boldsymbol{w}$ from $\boldsymbol{u}$. Bi-deduction enjoys an induction rule, not shown here, which corresponds to iterating a simulator over all values in an enumerable type, as we did in our motivating example.

The oracle rule $\rhd.\textsc{Oracle}_f$ corresponds to a simulator that calls the game's oracle f. This rule reflects a peculiarity of the model, where adversaries choose the source of randomness used by oracles — without being able to access it, and with consistency conditions expressed through name constraints. The *oracle Hoare triple* $\{\boldsymbol{\phi}, \boldsymbol{g}\} \boldsymbol{v} \leftarrow O_f(\boldsymbol{t})[\boldsymbol{k}; \boldsymbol{r}]\{\boldsymbol{\psi}\}$ expresses that the oracle f, if called on $\boldsymbol{t}$ when both $\boldsymbol{\phi}$ and $\boldsymbol{g}$ hold, with the names $\boldsymbol{k}$ corresponding to the variables $\mathcal{G}.\mathsf{glob}_\$$ globally sampled in the game and the names $\boldsymbol{r}$ corresponding to the variables $f.\mathsf{loc}_\$$ sampled inside oracle f, will return $\boldsymbol{v}$ and leave the game's memory in a state satisfying $\boldsymbol{\psi}$.

An example of a rule that does not apply any simulator construction is $\rhd.\textsc{Rewrite-R}$: it states that if a simulator computes $\boldsymbol{v'}$, and that this bi-term is equal to $\boldsymbol{v}$, then the simulator also computes $\boldsymbol{v}$ — the equality is for each component of the bi-terms, and must be exact. Without this rule, we would be limited to derive terms following their syntactic structure.

4 Basic Simulator Synthesis

We present a basic simulator synthesis procedure based on bi-deduction, that generates simulators *without loops or recursion* — it does not use the induction rule of bi-deduction. That procedure is an improved version of the one implemented in [5], although it sticks to its general principle. However, the formal description of the procedure is new, and necessary for justifying the inductive synthesis procedure of Section 5.

Synthesis queries. The rules of [5] provide basic building blocks for deriving valid simulators but, in order to obtain a synthesis procedure, we need to determine when and how to use each rule. We first clarify what parts of a bi-deduction judgement are, respectively, *inputs* and *outputs* of the simulator synthesis procedure. To this effect, we use *synthesis queries* of the following form:

$$\mathcal{E};\boldsymbol{\Theta}_i;\boldsymbol{C}_i;\boldsymbol{\phi} \vdash \boldsymbol{u} \triangleright_{\mathcal{G}} \boldsymbol{v} \rightsquigarrow (\boldsymbol{\Theta}_o;\boldsymbol{C}_o;\boldsymbol{\psi};\boldsymbol{w})$$

The components on the left of $\rightsquigarrow$ are *inputs* of the synthesis procedure, while components on the right of $\rightsquigarrow$ are *outputs*. For the sake of clarity, we colored inputs in black and outputs in dark red in the query above. This query is valid whenever the corresponding bi-deduction judgement is valid:

$$\mathcal{E};\boldsymbol{\Theta}_i,\boldsymbol{\Theta}_o;\boldsymbol{C}_i \cdot \boldsymbol{C}_o;(\boldsymbol{\phi},\boldsymbol{\psi}) \vdash \boldsymbol{u} \triangleright_{\mathcal{G}} \boldsymbol{v},\boldsymbol{w}$$

Said otherwise, our synthesis procedure takes as inputs a set of hypotheses $\boldsymbol{\Theta}_i$ and initial constraints $\boldsymbol{C}_i$, a pre-condition $\boldsymbol{\phi}$ on the state of game $\mathcal{G}$, the simulator's input $\boldsymbol{u}$ and its target output $\boldsymbol{v}$. It attempts to synthesize a simulator $\mathcal{S}$ that computes $\boldsymbol{v}$ when given $\boldsymbol{u}$ as inputs, and returns $\mathcal{S}$'s randomness constraints $\boldsymbol{C}_o$, the resulting post-condition $\boldsymbol{\psi}$, further hypotheses $\boldsymbol{\Theta}_o$ that must hold for $\mathcal{S}$ to be correct, and additional terms $\boldsymbol{w}$ that $\mathcal{S}$ computed while computing $\boldsymbol{v}$. The extra hypotheses $\boldsymbol{\Theta}_o$ are proof obligations that will be discharged to the user at the end of the simulator synthesis, together with a formula expressing the validity of the combined constraints $\boldsymbol{C}_i \cdot \boldsymbol{C}_o$.

The additional outputs $\boldsymbol{w}$ are called *memoization hints*, and will allow to re-use the result of oracle calls across recursive iterations of our final recursive simulators (see §5). They will be added to inputs of further synthesis queries. To distinguish them from standard inputs, the inputs of synthesis queries are split into two sequences, noted $\boldsymbol{in}.\texttt{std}$ and $\boldsymbol{in}.\texttt{memo}$ for, resp., standard and memoization inputs. We may still use $\boldsymbol{in}$ as a single sequence when the distinction is irrelevant, e.g. $\boldsymbol{t} \in \boldsymbol{in}$ means that $\boldsymbol{t}$ belongs to either of the two sequences.

Synthesis query rules. We design rules for deriving synthesis queries, that provide a higher-level and more operational variant of the bi-deduction proof system. The validity of synthesis query rules can be established by combining several bi-deduction rules to derive the validity of the conclusion query from that of the premises.

We make use of a few standard automated reasoning utilities. We assume a weak head normalization function $\mathsf{whnf}^{\Theta}_{\mathcal{E}}(t)$. In practice we only normalize modulo the definitions of $\mathcal{E}$ and some basic equations of Θ, but any normalization function ensuring $\mathcal{E};\Theta \vdash [t = \mathsf{whnf}^{\Theta}_{\mathcal{E}}(t)]_{\mathrm{e}}$ is correct. Because our normalization only relies on a builtin part of Θ, we omit that component for brevity. We also use unification: if u and v are terms well-typed in $(\mathcal{E},\vec{x}:\vec{\tau})$, then $\mathsf{unify}^{\mathcal{E}}_{\vec{x}}(u = v)$ is a partial procedure that may return a substitution θ mapping a subset of $\vec{x}$ to well-typed terms in $\mathcal{E}$, such that $\mathcal{E};\Theta \vdash [u\theta = v\theta]_{\mathrm{e}}$. We actually use unification on bi-terms, with the natural specification. This loose specification may be met by various unification procedures; in practice we use a simple one which only exploits definitions in $\mathcal{E}$ and basic equations from Θ.

We describe in Fig. 5 a few representative rules, providing a more complete presentation in Appendix A. The UNREACH rule is to be used when the term to be deduced is under an infeasible condition. The negation of the condition is discharged to the user, hence it is important to use this rule only as a last resort in an automatic synthesis context: otherwise, an invalid proof obligation might render the whole synthesis useless. An opposite strategy is used in LOAD.SIMPLIFIED: there, we attempt to instantiate an input $\lambda\vec{x}.(\boldsymbol{u} \mid \boldsymbol{g})$ to obtain the desired output $(\boldsymbol{o} \mid \boldsymbol{f})$; we determine a possible value for $\vec{x}$ by unification, and we verify *automatically* that under this instantiation, $\boldsymbol{g}$ is implied by $\boldsymbol{f}$ (this is denoted $\vdash_{\mathsf{auto}}$); it then only remains to verify that the values of $\vec{x}$ can themselves be simulated. This rule requires that the implication has been verified, hence there is no risk of abusive applications as with UNREACH.

The query synthesis rule for oracle calls, ORACLE, is an effective version of $\triangleright.\mathrm{ORACLE}_f$ that relies on the following assumptions on oracle f:

- The body of f (in both sides of the game) is a sequence of random samplings of $f.\mathsf{loc}_\$$, followed by assignations and a final return statement of the form (**return** if c_f then o_f else dummy): the interesting result o_f is returned under a condition c_f; otherwise an irrelevant constant from $\mathcal{L}$ is returned.

- We further assume that the expression o_f does not contain memory locations: it may only refer to the oracle's inputs, and to local and global samplings.

We let $\vec{x}$ be the oracle's input variables. We also let $\vec{y} = \mathcal{G}.\mathsf{glob}_\$$ and $\vec{z} = f.\mathsf{loc}_\$$ for brevity. The assignation statements in f may refer to the logical variables $\vec{x},\vec{y},\vec{z}$ in addition to program variables, i.e. memory locations. Note that, although conditionals are not allowed in the oracle's body, they can be used inside the expressions of assignations. Although limited, this format is met by the CCA2 game and all cryptographic games that we have encountered so far.

Following the abstract interpretation terminology, we view assertions as abstract memories, and we define (see the long

Selected core rules.

$$\textsc{Unreach}\quad \frac{}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright (o \mid f) \rightsquigarrow ([\neg f]_e;\emptyset;\phi;\epsilon)}$$

$$\textsc{Conv}\quad \frac{in' = \mathsf{whnf}^{\Theta}_{\mathcal{E}}(in) \qquad out' = \mathsf{whnf}^{\Theta}_{\mathcal{E}}(out) \qquad \mathcal{E};\Theta;C;\phi \vdash in' \triangleright out' \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright out \rightsquigarrow (\Theta';C';\psi;w)}$$

$$\textsc{Oracle}\quad \frac{\theta = \mathsf{unify}^{\mathcal{E}}_{\vec{x},\vec{y},\vec{z}}(o = o_f) \qquad \theta(\vec{y}) = (\mathsf{k_v}\ p_v)_{v\in\mathcal{G}.\mathsf{glob}_\$} \qquad \theta(\vec{z}) = (\mathsf{r_v}\ s_v)_{v\in f.\mathsf{loc}_\$} \qquad \mathcal{E};\Theta;C;\phi \vdash in \triangleright \left(\theta(\vec{x}), (p_v)_{v\in\mathcal{G}.\mathsf{glob}_\$}, (s_v)_{v\in f.\mathsf{loc}_\$} \mid g\right) \rightsquigarrow (\Theta';C';\phi';w) \qquad C'' = \left(\Pi_{v\in\mathcal{G}.\mathsf{glob}_\$}(\emptyset,\mathsf{k_v},o_v,\top^{\mathsf{glob}}_{\mathsf{G},v},g)\right)\cdot\left(\Pi_{v\in f.\mathsf{loc}_\$}(\emptyset,\mathsf{r_v},s_v,\top^{\mathsf{loc}}_{\mathsf{G}},g)\right) \qquad g_\mu = \mathsf{b\text{-}eval}_{\phi'}(c_f\theta) \qquad \psi = \mathsf{post}^{\theta}_{f}(\phi')}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright (o \mid g) \rightsquigarrow (\Theta',[g \Rightarrow g_\mu]_e;C'\cdot C'';\psi;w)}$$

Selected destruction rules.

$$\textsc{FA.}\Rightarrow\textsc{FA.}\Rightarrow\quad \frac{\mathcal{E};\Theta;C;\phi \vdash in \triangleright (g_0 \mid f),(g_1 \mid f \wedge g_0) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright (g_0 \Rightarrow g_1 \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

$$\textsc{FA}\quad \frac{s \in \mathcal{L} \qquad \mathcal{E};\Theta;C;\phi \vdash in \triangleright (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright (s\ o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

Selected memory and memoization rules.

$$\textsc{Load.Simplified}\quad \frac{\lambda\vec{x}.(u \mid g) \in in.\mathtt{std} \qquad \theta = \mathsf{unify}^{\mathcal{E}}_{\vec{x}}(u = o) \qquad \mathsf{dom}(\theta) = \vec{x} \qquad \mathcal{E};\Theta \vdash_{\mathsf{auto}} [f \Rightarrow g\theta]_e \qquad \mathcal{E};\Theta;C;\phi \vdash in \triangleright (\vec{x}\theta \mid f) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

$$\textsc{Memoize.Store}\quad \frac{\mathcal{E};\Theta;C;\phi \vdash in \triangleright out \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright out \rightsquigarrow (\Theta';C';\psi;(w,out))}$$

$$\textsc{Memoize.Load.Simplified}\quad \frac{\lambda\vec{x}.(u \mid g) \in in.\mathtt{memo} \qquad \theta = \mathsf{unify}^{\mathcal{E}}_{\vec{x}}(u = o) \qquad \mathsf{dom}(\theta) = \vec{x} \qquad \mathcal{E};\Theta;C;\phi \vdash in \triangleright (o \mid f \wedge \neg g\theta) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \triangleright (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

Figure 5: Selected proof-search rules.

version [38] for details) an abstract interpretation function eval: given a term s of type bitstring set and an abstract memory ϕ, the symbolic set $\mathsf{eval}_\phi(s)$ over-approximates the value of s in any memory satisfying ϕ. Similarly, when f has type bool, the boolean $\mathsf{b\text{-}eval}_\phi(f)$ under-approximates f, i.e. it implies f. Given an initial abstract memory and an assignation $\ell \leftarrow e$ of a bitstring set memory location, we can abstractly evaluate e and the resulting memory, to obtain the updated abstract memory. This can be chained for all assignations in the body of f, to obtain $\mathsf{post}^{\theta}_{f}(\phi \mid g)$ where θ is a substitution of domain $\vec{x},\vec{y},\vec{z}$: this defines a valid post-condition for a call to f with the values given by θ in a context where g holds and the memory satisfies ϕ.

Our ORACLE rule proceeds as follows. First, it does not attempt to syntactically match the term to be computed with the oracle's return expression: instead, it matches it with o_f, and discharges proof obligation which ensures that c_f holds when the oracle is called. We thus unify o_f and o to determine the values of $\vec{x},\vec{y},\vec{z}$. As before, the values of $\vec{y},\vec{z}$ must be names, and the name indices as well as the arguments $\theta(\vec{x})$ must be bi-deducible. The oracle is (implicitly) called with the abstract memory ϕ' obtained after that bi-deduction, and only when g holds — which implies g_μ and then c_f. The final abstract memory is $\mathsf{post}^{\theta}_{f}(\phi' \mid g)$. Overall, our synthesis rule is justified using $\triangleright.\textsc{Oracle}_f$, relying on the Hoare triple

$$\{\phi', g \wedge g_\mu\} out \leftarrow O_f(\theta(\vec{x}))[\theta(\vec{y});\theta(\vec{z})]\{\mathsf{post}^{\theta}_{f}(\phi')\}$$

which is is valid under the assumption $[g \Rightarrow g_\mu]_e$.

The basic simulator synthesis procedure. Our basic synthesis procedure $\mathsf{synthesize}_\triangleright(\cdot)$ is a (recursive) function which takes the left part of a synthesis query (its inputs) and attempts to derive it by applying synthesis rules following a particular strategy. Upon success, it returns the right part of the synthesis query (its outputs). Our procedure is such that

$$\text{if}\quad \mathsf{synthesize}_\triangleright(\mathcal{E};\Theta;C;\phi \vdash in \triangleright out) = (\Theta';C';\psi;w)$$
$$\text{then}\quad \mathcal{E};\Theta;C;\phi \vdash in \triangleright out \rightsquigarrow (\Theta';C';\psi;w) \text{ is derivable.}$$

The procedure has four different phases, as depicted in Fig. 6 and described below. The phases are applied successively until one succeeds. A successful phase will usually generate bi-deduction premises, which are themselves resolved recursively by the procedure. Our procedure never backtracks: if a phase succeeds but generates invalid bi-deduction subgoals, the procedure will fail without trying the next phases. This makes the procedure less complete but more *efficient* and *predictable*. This latter property is crucial in a fully automated setting, as it helps the user construct an intuitive understanding of why the procedure failed and how this may be fixed.

The **reuse phase** checks if the target out can be directly obtained from standard inputs. This phase only uses the LOAD rule, attempting to use it on all terms from $in.\mathtt{std}$. Because this phase greedily uses the first relevant input, it could in principle make a wrong decision and cause the overall synthesis to fail. We never encountered this issue in practice.

The **oracle phase** tries to obtain $out = (o \mid f)$ by calling one of the oracles of the game. A high-level description of the

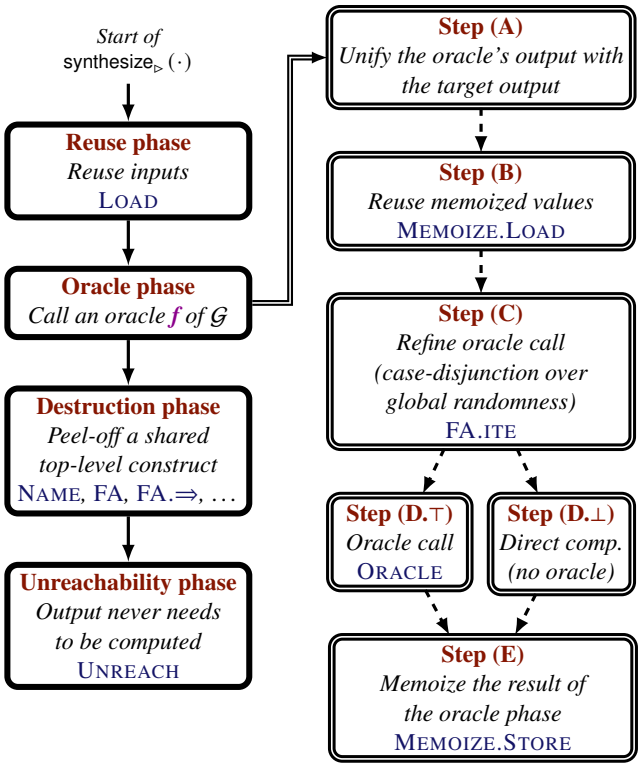

Phases are in boxes with a simple border. Steps of the oracle phase are in boxes with a double border. Phase/step names are in dark red. Each box includes the main rule they rely on, when applicable. Simple arrows indicates progression between phases (continue in case of failure). Dashed arrow indicates progression between steps of the oracle phase (continue in case of success). All phases and steps may generate new bi-deduction premises, which are resolved recursively by $\mathsf{synthesize}_{\triangleright}(\cdot)$**.**

Figure 6: Control-flow of $\mathsf{synthesize}_{\triangleright}(\cdot)$.

oracle phase is shown in the right side of Fig. 6, and detailed in the long version [38]. The oracle phase is composed of several steps applied sequentially, where the failure of any of the steps leads to a failure of the whole phase. Step **(A)** identifies an oracle f of the game whose result can be unified with the target $\boldsymbol{o}$. The next two steps change the target $(\boldsymbol{o} \mid \boldsymbol{g})$ into $(\boldsymbol{o} \mid \boldsymbol{g} \wedge \boldsymbol{g'})$ where $\boldsymbol{g'}$ describes situations where an oracle call is not necessary: this may be because a memoized value can be reused **(B)** or because random values (names) do not belong to the game, hence the target can be computed directly **(C)**. This step is crucial: calling an oracle impacts the abstract memory as well as the name constraints; it must be used sparingly to avoid unnecessary failures of the synthesis. Then, the arguments of the oracle are bi-deduced assuming $\boldsymbol{g'}$ hold **(D.⊤)**, and the output term is recursively bi-deduced assuming $\boldsymbol{g'}$ does not hold **(D.⊥)**. Finally, a memoization hint is added to mark the deduced term as reusable in future oracle phases **(E)**.

The **destruction phase** checks if the left and right components of $\boldsymbol{out}$ start with the same top-level construct, which is then computed by the simulator being synthesized. Specialized rules such as FA.⇒ are always prioritized over the generic FA rule. Finally, we use the CONV rule beforehand to put the left and right components of $\boldsymbol{out}$ in *weak-head normal form*, which helps to apply the destruction rules more often. Applying this phase too early would often lead to invalid simulators: an obvious example of this would be to use FA on an encryption when it is necessary to obtain the encryption through the corresponding oracle of the CCA game.

Lastly, the **unreachable phase** lets the simulator abandon the synthesis by asking the user to prove that $\boldsymbol{out}$ never needs to be evaluated. It relies on UNREACH and is used only as a last resort, as it may produce invalid proof obligations. Moreover, it is always possible to postpone its application, so our last resort strategy cannot hurt.

5 Inductive Simulator Synthesis

Assume a function $\boldsymbol{u}$ defined by recurrence over some type τ using the well-founded order $<$. In a beautified syntax:

$$<: \tau \to \tau \to \mathsf{bool} \qquad \textbf{let rec } \boldsymbol{u}\ (x:\tau) = \boldsymbol{u}_0$$

where $\boldsymbol{u}_0$ is the body of $\boldsymbol{u}$'s definition.[2] We let $\leq$ be the reflexive closure of $<$.

Consider $\boldsymbol{t}$ of type τ, and assume we want to bi-deduce $(\boldsymbol{u}\ \boldsymbol{t})$ from some inputs $\boldsymbol{in}$. Excluding degenerated cases, doing so will require to recursively evaluate $\boldsymbol{u}$ on points $x < \boldsymbol{t}$. We can build such simulators using the (simplified) induction rule:

$$\frac{\begin{array}{l}\mathcal{E}, x:\tau; \boxed{\mathcal{I}(<,x)} \vdash \\ \boldsymbol{in}, x, \lambda y.(\boldsymbol{u}\ y, \boxed{\boldsymbol{w}\ y} \mid y < x \wedge x \leq t) \triangleright (\boldsymbol{u}\ x, \boxed{\boldsymbol{w}\ x} \mid x \leq \boldsymbol{t}) \\ \qquad\qquad \rightsquigarrow (\boxed{\mathcal{I}(\leq,x)})\end{array}}{\mathcal{E}; \boldsymbol{\phi}_0 \vdash \boldsymbol{in} \triangleright \boldsymbol{u}\ \boldsymbol{t} \rightsquigarrow (\boxed{\mathcal{I}(\leq,\boldsymbol{t})})}$$

(For the sake of simplicity, we omit the constraints $\boldsymbol{C}$, hypotheses $\boldsymbol{\Theta}$, etc. The detailed rule is in the long version [38].

Let us unpack this definition. It states that to bi-deduce $(\boldsymbol{u}\ \boldsymbol{t})$, it is sufficient to bi-deduce $(\boldsymbol{u}\ x)$ for any $x \leq \boldsymbol{t}$, assuming by induction that we already computed $(\boldsymbol{u}\ y)$ for any $y < x$. There are two additional ingredients (in colored boxes) which we describe next.

First, $\mathcal{I}$ is a *memory invariant* describing the evolution of the game's state during the recursive evaluation of $\boldsymbol{u}$. Concretely, $\mathcal{I}(\bowtie, x)$ is an abstract memory parameterized by a relation $\bowtie$ over τ and a value x of type τ, where $\mathcal{I}(<,x)$ and $\mathcal{I}(\leq,x)$ represents the memory, resp., *before* and *after* the evaluation at iteration x. The premise of the rule assumes that the initial memory satisfies $\mathcal{I}$ strictly before x; and establishes

[2] As indicated by the bold font, $\boldsymbol{u}_0$ is a bi-term. Formally, we have two symbols s_0 and s_1 defined by recurrence over x and whose bodies are, resp., the left and right component of $\boldsymbol{u}_0$. Then, we let $\boldsymbol{u} = \#(s_0; s_1)$. For the sake of simplicity, we use a single symbol definition $\boldsymbol{u}$ using a bi-term as a body.

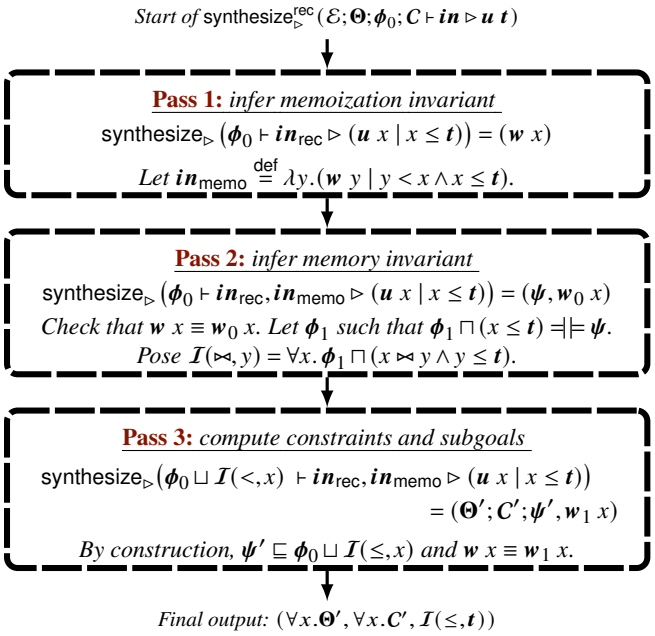

We assume $u \equiv \lambda(x:\tau).u_0$ where τ is a `fixed` and `finite` base-type. We assume a total order $<$ over τ. We let $in_{\text{rec}} \stackrel{\text{def}}{=} (in, \lambda y.(u\ y \mid y < x))$. All calls to $\text{synthesize}_{\triangleright}(\cdot)$ use the same environment $(\mathcal{E}, x:\tau)$, initial hypotheses Θ, and initial constraints C, which are thus omitted. We also omit components of $\text{synthesize}_{\triangleright}(\cdot)$'s outputs that are discarded — e.g. in the first pass, only the outputted memoization hint $(w\ x)$ is used.

Figure 7: Inductive simulator synthesis procedure $\text{synthesize}^{\text{rec}}_{\triangleright}(\cdot)$.

that after computing $(u\ x)$, the memory satisfies $\mathcal{I}$ up-to x included. The conclusion of the induction rule states that at the end of the recursive computation, the memory satisfies $\mathcal{I}$ up-to t. To initialize the induction (not shown here) we must ensure that ϕ_0 entails $\mathcal{I}(<, x_0)$ where x_0 is the smallest element w.r.t. $<$.

Second, w is a *memoization invariant*. Essentially, $(w\ y)$ represents a set of intermediate values that have been computed during the bi-deduction of $(u\ y)$ and that we decided to memoize. Thus, when bi-deducing $(u\ x)$, we may reuse past memoized values $(w\ y)$ for any $y < x$.

Invariant inference. Applying the induction rules requires to come-up with a memory invariant $\mathcal{I}(\cdot,\cdot)$ and a memoization invariant $(w\ \cdot)$. Below, we present a novel technique to do so, improving upon [5] in several ways. First, it infers time-*sensitive* invariants, as opposed to [5] which only supports time-*insensitive* invariants — which, as shown in §2, are too restrictive in practice. Second, it infers memoization invariants — [5] featured no such invariants. Third, our invariant inference technique always terminates: [5] provides no such guarantees (though their approach terminates in practice).

Concretely, we define an inductive simulator synthesis procedure $\text{synthesize}^{\text{rec}}_{\triangleright}(\cdot)$ that builds upon the basic synthesis of §4. This procedure is summarized in Fig. 7, and relies on three passes of the basic synthesis procedure $\text{synthesize}_{\triangleright}(\cdot)$, where the first pass infer a memoization invariant using the memoization hints returned by $\text{synthesize}_{\triangleright}(\cdot)$, the second pass computes a memory invariant, and the third and last pass computes the final proof-obligations Θ' to be discharged to the user and the constraints C' guaranteeing the existence of a probabilistic coupling justifying our reduction.

Crucially, the memory invariant of the second pass is an inductive invariant by construction (see Theorem 1 later). Thus, we do not need to check that this invariant is inductive.

```
1  let rec output (t : timestamp) =
2    match t with
3    | Pk ⟶ pub k
4    | V ⟶
5      let pk = pub k in
6      enc diff(input V, dummy) r pk
7    | M(i) ⟶ empty
8    | P ⟶ shuffle(bb)
9  and frame (t : timestamp) =
10   match t with
11   | init ⟶ empty
12   | _ ⟶ ⟨ frame (pred t), output t ⟩
13 and input (t : timestamp) =
14   match t with
15   | init ⟶ empty
16   | _ ⟶ att(frame(pred t))

and bb (t : timestamp) =                          17
  match t with                                    18
  | M(i) ⟶                                        19
    let pk = pub k in                             20
    let x =                                       21
      if V < M(i) &&                              22
        (input t =                                23
          enc diff(input V, dummy) r pk)          24
      then input V else dec (input t) k           25
    in                                            26
    bb[ i ⟶ x ]  (* update cell i *)              27
  | _ ⟶ bb(pred t)                                28
```

Figure 8: CCSA encoding of our abstract mixnet protocol.

Example. To understand this technical procedure, it is helpful to apply it on our abstract mixnet of §2. We give in Fig. 8 a CCSA encoding of this protocol using recursive functions. Let us reduce the equivalence of the left and right versions of frame t_0 to CCA2, where t_0 is some constant timestamp. We unroll the execution of $\text{synthesize}^{\text{rec}}_{\triangleright}(\phi_0 \vdash \emptyset \triangleright \text{frame}\ t_0)$ on the initial memory $\phi_0 = (\ell \mapsto \epsilon)$ of the CCA2 game. Let t be the inductive step variable.

Obviously, detailing all passes, phases and steps of our synthesis procedure is not realistic, and would be too verbose to be of any use. Instead, we let the reader get a intuitive understanding of how the recursive functions in Fig. 8 are simulated. Further, the simulator of Fig. 3, which we manually wrote, should help intuit what is going on.

◇ *Pass 1.* After the first pass, we have the memoization hints:

$$w\ t \stackrel{\text{def}}{=} (\text{encV} \mid t = V \wedge t \le t_0), \qquad \text{(line 6)}$$
$$\lambda i.\,(\text{encV} \mid \text{V} < \text{M(i)} \wedge t = \text{M(i)} \wedge t \le t_0) \qquad \text{(line 24)}$$
$$\lambda i.\,(\text{decM} \mid \ldots \wedge t = \text{M(i)} \wedge t \le t_0) \qquad \text{(line 25)}$$

where encV = enc diff(input V, dummy) (pub k) and decM = dec (input t) k. (We omit some details of the decM hint, as it will not be used.) We annotated each memoization hint with

the corresponding line in Fig. 8. Remark that the call to the **left-right** oracle line 24 is guarded by the test $\mathsf{V} < \mathsf{M(i)}$ (at line 22), which thus appears in the corresponding memoization hint. After some simplification, this yields the following memoized values at time t:

$$\boldsymbol{in}_{\mathsf{memo}} \stackrel{\mathsf{def}}{=} (\mathsf{encV} \mid V < t \wedge t \leq t_0),$$
$$\lambda i.\,(\mathsf{encV} \mid \mathsf{V} < \mathsf{M(i)} \wedge \mathsf{M(i)} < t \wedge t \leq t_0),\ \ldots$$

Note that the second memoized value is subsumed by the first.

◇ *Pass 2.* The second pass computes the memory footprint of our simulator. Crucially, the memoized values in $\boldsymbol{in}_{\mathsf{memo}}$ allow to reduce the number of oracle calls through re-use, which is critical to simulate our mixnet protocol. Concretely, looking at Fig. 8, we see that we may need to call the **left-right** oracle to simulate the computation at lines 6 and 24. For line 6, this adds encV to ℓ whenever $t = \mathsf{V}$. For line 24, we know that $t = \mathsf{M(i)}$ and we are operating under the condition that $\mathsf{V} < \mathsf{M(i)}$ (line 22): thus, we can re-use the memoized value in $\boldsymbol{in}_{\mathsf{memo}}$, and never need to call the oracle **left-right** oracle there. Thus, after some simplifications which we omit, we obtain the post-condition

$$\boldsymbol{\psi} \stackrel{\mathsf{def}}{=} (\ell \mapsto \{\mathsf{encV} \mid t = \mathsf{V} \wedge t \leq t_0\}),$$

which yields the memory invariant

$$\boldsymbol{\mathcal{I}}(\bowtie, t) \stackrel{\mathsf{def}}{=} (\ell \mapsto \{\mathsf{encV} \mid t_1 : t_1 = \mathsf{V} \wedge t_1 \bowtie t \wedge t \leq t_0\}),$$

or equivalently $(\ell \mapsto \{\mathsf{encV} \mid \mathsf{V} \bowtie t \wedge t \leq t_0\})$. By Theorem 1 (presented later), this is an inductive invariant of our simulator.

◇ *Pass 3.* The last pass computes the proof-obligations $\boldsymbol{\Theta}'$ and constraints $\boldsymbol{C}'$, using the memory and memoization invariants, resp, $\boldsymbol{\mathcal{I}}$ and $\boldsymbol{in}_{\mathsf{memo}}$. We do not detail them, as this would take too much space. Still, in our implementation, all generated goals are automatically discharged by Squirrel automated reasoning tactic **auto**. We refer the curious reader to the proof artifact [37] (file `proofs/motivating.sp`) for details — invariants in that file are however less readable than the ones presented here, which have been manually simplified.

Soundness. We now state the theorem ensuring the soundness of $\mathsf{synthesize}^{\mathsf{rec}}_{\triangleright}(\cdot)$. First, we must describe the class of cryptographic hardness games supported by our result. We already saw that, as in [5], we only support game whose global mutable state are logs (or sets) of bitstring values. Further, we must require that there are no complex flows between the different logs of the game. E.g., a game with two logs ℓ and ℓ', and with a update of the form $\ell \leftarrow \ell \cup \ell'$ in one of its oracle, is not supported by our procedure. Formally:

Assumption 1. *For any update $\ell \leftarrow \boldsymbol{s}$ in an oracle of $\mathcal{G}$, the only global mutable variable that $\boldsymbol{s}$ may depends upon is $\{\ell\}$.*

To our knowledge, this assumption is at no loss, as we do not know of any cryptographic game featuring such patterns. Indeed, standard games use logs to keep track of values sent to various oracles. Typically, there is one log per kind of values to be tracked, and there are no flows between distinct logs. This is exactly the class of games we support.

We can now state our main soundness theorem, whose proof can be found in the long version [38].

Theorem 1. *Let $\mathcal{G}$ satisfying Assumption 1. Let $\boldsymbol{u} \equiv \lambda(x : \tau).\boldsymbol{u}_0$ where τ is a `fixed` and `finite` base-type. Let $<$ be a total order over τ such that $\mathcal{E}; \boldsymbol{\Theta} \models \mathsf{adv}(<)$. If*

$$\mathsf{synthesize}^{\mathsf{rec}}_{\triangleright}(\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{\phi}_0; \boldsymbol{C} \vdash \boldsymbol{in} \triangleright \boldsymbol{u}\ \boldsymbol{t}) = (\boldsymbol{\Theta}', \boldsymbol{C}', \boldsymbol{\psi})$$

and $\boldsymbol{t} \in \boldsymbol{in}$, then we have:

$$\mathcal{E}; \boldsymbol{\Theta}; \boldsymbol{C}; \boldsymbol{\phi}_0 \vdash \boldsymbol{in} \triangleright \boldsymbol{u}\ \boldsymbol{t} \rightsquigarrow (\boldsymbol{\Theta}'; \boldsymbol{C}'; \boldsymbol{\psi}; \boldsymbol{w}).$$

6 Implementation and Basic Evaluation

In this section, we present the implementation of our simulator synthesis procedure in Squirrel, and evaluate its usefulness and performances through a first round of case-studies. We present our large-scale case-study, a proof of privacy for the FOO e-voting protocol, later in Section 7.

Implementation. We have implemented the simulator synthesis procedure described above in the main development line of Squirrel [26], as a new version of the **crypto** tactic. The modified tactic relies by default on the basic simulator synthesis of Section 4, but the inductive synthesis of Section 5 is only used when required by the user through a ~time_sensitive flag; otherwise, the old inductive synthesis procedure [5] is used. Thus, memoization is turned on by default but not time-sensitive invariants. While both features make the tactic strictly more powerful, enabling the former does not break any existing development, which is not the case for the latter. Indeed, the tactic generates different subgoals with time sensitive invariants. Enabling it only when explicitly required makes for a smoother transition to the new implementation, avoiding modifications in existing and ongoing developments. Moreover, **crypto** subgoals are arguably more readable without time sensitivity. Finally, our new implementation comes with several performance improvements.

Performances. We compare the performance of our new implementation with the one from [5] using a standard laptop (see specifications in Section 7) on the accompanying case studies. These are found in the `examples/crypto` directory of Squirrel's repository, and contain 13 calls to **crypto** on four different games. With the original implementation from [5], **crypto** calls take 0.1s on average (spanning from 0.06ms to 0.2s each). With our performance optimizations, the average

drops to 0.05s (individual times spanning from 2ms to 0.08s); all calls are faster except the two fastest ones. With the new default version of the tactic, featuring memoization, calls take 0.06s on average (from 4ms to 0.1s). On such simple examples, performance is not critical, but these measures show that our modifications only improve running times, while bringing more expressivity.

Benefits of memoization. We illustrate how the basic simulator synthesis of Section 4 makes our new **crypto** tactic more expressive. It is trivial to come up with artificial examples where the old **crypto** fails but the new tactic succeeds thanks to the addition of UNREACH, and practical examples of this situation will be given in the next section. More interestingly, we show in `memoization.sp` cases where the old **crypto** did not succeed while the new one does, thanks to memoization. These (artificial) examples rely on a stateless game for exclusive or. More generally, this situation arises whenever an oracle output is shared between several actions of the protocol, and the oracle does not rely on the game's internal state — otherwise time sensitivity might be needed to succeed.

Benefits of time-sensitivity. We illustrate how our full procedure, also involving the time-sensitive inductive synthesis of Section 5, enables further new successes. As explained in Section 2, both memoization and time sensitive invariants are necessary for non-degenerate reductions to the CCA2 assumption. To showcase this point, we provide in [37] a new version of the Needham–Schroeder–Lowe case study from [5]. In this example, we consider two participants, with one session each, and we seek to establish indistinguishability between the protocol and an idealized version of it, where the contents of all encryptions are replaced by zeroes. In the original version, **crypto** is used with the CCA2 game to show that the three first messages of an honest exchange are indeed equivalent to their zeroed-out variants. Then, a tedious proof shows that this equivalence carries out to the frames, essentially showing that the frames can be bi-deduced from these three messages — using vanilla bi-deduction, without any cryptographic game. Overall, the proof is 380 lines long. This approach was necessary because our simulators could not memoize oracle calls. With our new **crypto**, the proof is immediate and only 60 lines long, with all the tedious manual bi-deduction absorbed in our new tactic. We finally note that reductions to CCA2 are not the only cases where memoization and time sensitive invariants are useful: one such case involving the PRF assumption is given in `memoization.sp`, and a more realistic example involving a KDF is shown in `kdf.sp`.

7 Application: the FOO E-Voting Protocol

We now valide our technique and implementation on a large-scale case study, namely a proof of vote privacy for the FOO e-voting protocol [29]. We chose FOO because: it relies on cryptographic assumptions and primitives (CCA2, blind signatures, commitments) that were never considered in Squirrel; and there is a CCSA pen-and-paper proof of its vote privacy [6], which gave us a head start.

As in [6], we focus on proving that FOO provides vote privacy, following Benaloh's vote swapping definition [12]. While the high-level structure of our proof follows that of [6], our security analysis is significantly more complicated, for two reasons. First and foremost, we *mechanized* our proof in Squirrel, which yields a development of approximately 10 kLoC (in contrast, [6] only provides a two-page proof sketch). Second, we consider an arbitrary number of dishonest voters, where [6] only has one — thus, their protocol only has a fixed number of agents. Having an unbounded number of agents significantly complicates the security analysis, as it forces us to rely on inductive reasoning.

We describe next FOO's cryptographic primitives, the high-level structure of the protocol and the modeling of vote privacy, and finally our Squirrel proof.

Cryptographic primitives. To achieve its security goals, FOO uses mixnets, blind signatures, and commitments.

A *mixnet* [16] is a (sub-)protocol which allows a collection of agents to send messages while hiding the relations between messages and senders. Typically, the agent's messages are encrypted with the mixnet public key (or keys). Once all messages have been received, the mixnet shuffles them at random, and publishes their decryptions. To protect the users' privacy against the mixnet itself, mixnets are composed of several independent servers, where each server does one round of shuffling and decryptions, typically augmented with zero-knowledge proofs of correct shuffling.

We use the abstract modeling of shuffle presented in §2. We assume the existence of a single (fictitious) public key whose corresponding secret key is held by an honest agent representing the mixnet as a whole. Once the mixnet is done receiving messages, it decrypts them and outputs their shuffling. Following [6], we leave the shuffling function unspecified, and only require that shuffling is invariant by permutation of its inputs through the following axiom:

$$\forall f, p.\, [\mathsf{bijective}\ p \rightarrow \mathsf{shuffle}\ f = \mathsf{shuffle}(\lambda i.\, f\ (p\ i))]_{\mathrm{e}}$$

Blind signatures [17] allow a user to ask for the signature of a message m to a signer $\mathcal{V}$ *without revealing* m *to* $\mathcal{V}$. In FOO, blind signatures are used to authenticate the voters' ballots without revealing its content to the voting authority. This allows to consider a dishonest authority when analyzing vote privacy. A blind signature scheme must satisfy the blindness property, which essentially allows to swap the contents of two valid signatures even when revealing non-swapped blindings and acceptance conditions. Very roughly, this corresponds to

an indistinguishability of the form:

$$\begin{aligned} & \mathsf{b}_0, \mathsf{b}_1, \mathsf{acc}_0, \mathsf{acc}_1, \\ & \text{if } (\mathsf{acc}_0 \wedge \mathsf{acc}_1) \text{ then } (\text{unblind } \mathsf{bs}_0, \text{unblind } \mathsf{bs}_1) \text{ else } \bot \\ \sim\ & \mathsf{b}_1, \mathsf{b}_0, \mathsf{acc}_1, \mathsf{acc}_0, \\ & \text{if } (\mathsf{acc}_0 \wedge \mathsf{acc}_1) \text{ then } (\text{unblind } \mathsf{bs}_1, \text{unblind } \mathsf{bs}_0) \text{ else } \bot \end{aligned}$$

where $\mathsf{bs}_0, \mathsf{bs}_1$ are the adversarially generated blind signatures of the blinding, respectively, b_0 and b_1; acc_X states that bs_X is a valid blind signature; and (unblind bs) unblinds bs into a publicly verifiable signature ub. We refer the reader to the long version [38] for a detailed treatment of blind signatures in our security proof — which notably includes a novel presentation of the blindness cryptographic assumption that is better-suited to our needs.

Cryptographic commitments [33] allow a user to commit to a value v while hiding v until it publishes the token k_c needed to open the commit. A commitment must be *binding* (a commit of v cannot be opened into a value $v' \neq v$) and *hiding* (as long as k_c remains secret, no adversary can learn anything about v from its commit). To prove FOO's privacy, we only need the latter. Roughly, computational hiding can be captured by indistinguishabilities of the following form, assuming that the token k_c has not been leaked in m_0, m_1, u:

$$u, \text{commit } m_0\ \mathsf{k}_c \sim u, \text{commit } m_1\ \mathsf{k}_c$$

Modeling FOO. The FOO e-voting protocol, described in Fig. 9, has four phases. In the first two phases, the voters publish the commit to their vote on a public bulletin board. In the last two phases, the voters publish their commit token to the public bulletin board, which allows to open the commits to the votes and to tally the election.

In **phase 1**, the voter computes the commit c to its vote (1) using the commit token k_c. Then, blind signatures are used to let the voting authority authenticate the ballot anonymously: the voter blinds its commit (2), sends it to the voting authority (3) which sends back the blinded signature (4) that the voter verifies (5,6) and unblinds (7) unto the signature ub. Then, the voter publishes the ballot (c, ub), comprising the commit to its vote and the signature authenticating it, to the public bulletin board. Here, the voter's anonymity is preserved by interposing a mixnet between the voter and the bulletin board. Concretely, ballots are encrypted and sent to the mixnet (8,9), which decrypts and collect them (10).

Once the first voting phase is over, we move to **phase 2**, where the mixnet shuffles all ballots and sends them in bulk to the bulletin board (12) which publishes them (13) — we assume that BB is a public read-only variable, which models the fact that everybody shares the same view on the immutable bulletin board. Before sending the ballots to the bulletin board, the mixnet checks that there are no duplicates ballots (11), which is necessary for privacy (see [20]).

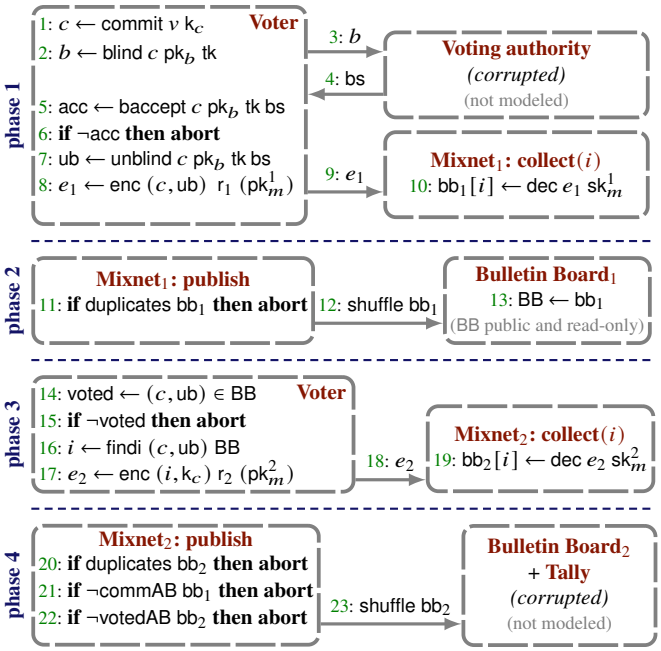

Lines are labeled by integers (1, 2, ...) for quick referencing. Integer labels only roughly indicate the ideal execution order of protocol operations: e.g., the adversary may trigger the input 9 before the voter sent its blinded ballot to be signed (output 3).

Figure 9: The FOO e-voting protocol.

In **phase 3**, the voter aborts if its ballot is not on the public bulletin board BB (14,15). Then, it retrieves the index i of its ballot on BB, and sends to the second mixnet (17,18) the pair (i, k_c) of its ballot index and its commit token, which is decrypted and collected by the mixnet (19). Finally, in **phase 4**, the mixnet publishes the shuffled commit tokens on a public bulletin board (23). As in phase 2, it first checks that there are no duplicates (20). The additional checks (21,22) are necessary to model the privacy property, and are discussed below. Finally, the public commit tokens can be used to open the commits on the first bulletin board, and to tally the election (using, e.g., vote counting or STV). As all the necessary data is available on the two public bulletin boards, anybody can tally the election. Thus, we can safely simplify the protocol model by letting the adverary run the tally.

Privacy modeling. We let the adversary control the bulletin boards, signing voting authority, and the tally. We consider two honest voters $\mathbf{Voter}_\mathsf{A}$ and $\mathbf{Voter}_\mathsf{B}$, and an arbitrary number of dishonest voters — dishonest voters are not modeled explicitly, as they are adversary-controlled, but implicitly, by letting the mixnets collect an arbitrary number of inputs. Following Benaloh's definition of privacy [12], we must show that

$$\begin{aligned} & \mathbf{P} \mid \mathbf{Voter}_\mathsf{A}(v_0) \mid \mathbf{Voter}_\mathsf{B}(v_1) \\ \sim\ & \mathbf{P} \mid \mathbf{Voter}_\mathsf{A}(v_1) \mid \mathbf{Voter}_\mathsf{B}(v_0) \end{aligned} \tag{1}$$

Group of files individual files	LoC	Time (seconds)	Calls to crypto count	 time	
Definitions and utilities	**1652**	**8**			
Reduction to Privacy_CCA	**3274**	**145.4**	**4**		
ccapk1.sp	482	12.2	2	3.8	
ccapk2.sp	525	38.2	2	11	
cca.sp	2267	95			
Deduction steps and shuffle opening	**3242**	**82**			
shuffle.sp	191	2.4			
deduction.sp	2122	64.8			
reduction.sp	929	14.8			
Cryptographic arguments	**2054**	**26.9**	**24**		
blinding.sp	384	5.9	1	0.9	(✓)
commitKeySecrecy.sp	576	11.6	4	0.3	
commitSecrecy.sp	683	8.6	8	0.2	(✓)
distinctCommits.sp	81	0.2	2	≤ 0.1	
distinctEncryptions.sp	294	0.3	8	≤ 0.1	
voteHiding.sp	36	0.3	1	≤ 0.1	
Total	**10219**	**262.3**	**28**	**33**	

All times were averaged over 10 runs and are in seconds. For calls to crypto, the time is for the *longest* call in the file, except for the total time, which is the cumulated time spent in crypto. The time spent in crypto is measured in Squirrel using the wrapper (time crypto).
We marked with a ✓ the calls to crypto that benefited from the UNREACH rule.

Table 1: Overview of the Squirrel development for FOO [37].

where v_0 and v_1 are arbitrary votes chosen by the adversary and **P** is the rest of the protocol, i.e. the mixnets and the bulletin board. As usual, we must rule-out trivial privacy attacks against the indistinguishability in Eq. (1). Indeed, an adversary can trivially breaks security by letting the first voter (**A** on the left, **B** on the right) cast its ballot, but blocking the ballot of the second voter. Further, the adversary does not cast any dishonest ballots itself. Then, inspecting the final bulletin board breaks the indistinguishability, as it only contains **A**'s vote on the left and **B**'s vote on the right. Our model rules out this trivial attack by having the final mixnet publish the ballot box only if it contains the ballots of the two honest agents, for both phases (see checks 21,22).

Proof. We define in Squirrel a pair of systems, called Privacy_real, corresponding to Eq. (1). Contrary to the informal protocol description, our system never aborts: instead, if one of the aborting conditions of Fig. 9 fails, the corresponding agent will output a dummy message. Thus it is always possible to keep executing the protocol until its very last action MOP (corresponding to 23), where the second mixnet publishes commit tokens. We consider arbitrary unbounded traces, of length independent from the security parameter η, (see [3] for a discussion of this limitation). We show that for any such trace ending with MOP, the two frames (with and without the votes swapped) are indistinguishable:

```
global theorem vote_privacy @system:Privacy_real : equiv(frame@MOP).
```

The first step in the proof is to reduce this theorem to the same one but for a modified pair of systems, called Privacy_CCA, where all encryptions are replaced by same-length encryptions of zeroes. This step is itself justified by two reductions to the CCA2 game, for sk_m^1 and sk_m^2. In the resulting systems, the messages sent to the mixnets are completely hidden from the attacker, which is necessary for several reasoning steps in the rest of the proof.

Next, we proceed by case analysis over a condition ϕ expressing that the two honest votes have successfully gone through the whole protocol. In other words, ϕ ensures that the protocol has not aborted. It notably implies votedAB. Depending on ϕ, we are going to consider the mixnets' shuffles differently, in order to apply specific cryptographic arguments.

When ϕ holds, the attacker has access to the final publication of the commit keys. Although A and B's commits can thus be linked with the vote that they contain, their voting material from the first phase is still private thanks to blind signatures. We reduce the indistinguishability on (if phi then frame@MOP)

to a simple indistinguishability

$$\begin{array}{rl} & m, c_0, c_1, \mathsf{b}_0, \mathsf{b}_1, \mathsf{acc}_0, \mathsf{acc}_1, \\ & \text{if } \mathsf{acc}_0 \wedge \mathsf{acc}_1 \text{ then } (\mathsf{ub}_0, \mathsf{ub}_1) \\ \sim & m, c_0, c_1, \mathsf{b}_1, \mathsf{b}_0, \mathsf{acc'}_1, \mathsf{acc'}_0, \\ & \text{if } \mathsf{acc}'_0 \wedge \mathsf{acc}'_1 \text{ then } (\mathsf{ub}'_0, \mathsf{ub}'_1) \end{array} \quad (2)$$

where m contains irrelevant cryptographic material; c_i, b_i, acc_i and ub_i are the commit, blinding, acceptance condition and unblinding (as in Fig. 9) of the voter who voted v_i; acc'_i and ub'_i are their variants when the voters are swapped. The acceptance conditions and unblindings depend on blinding tokens which are tied to the voter's identity: acc_0 is the verification performed by A on the left, while on the right it is acc'_1. To obtain this simple equivalence, we need to gradually decompose the original equivalence on frames, in a proof by induction which technically relies on bi-deduction. A crucial step is when we change a shuffle $\mathsf{shuffle}(\lambda i.\, f\ i)$ into its inputs $f\ i$: because of the key property of shuffles, we can change the shuffle into $\mathsf{shuffle}(\lambda i.\, f\ (p\ i))$ for any permutation p in order to obtain the inputs $f\ i$ in a convenient order. In the branch of the proof where ϕ holds, we open shuffles so as to organize the voting material by *vote* rather than by *identity*: this is how we obtain c_0, c_1 on both sides in Eq. (2), and similarly for unblindings. Finally, Eq. (2) is proved by **crypto** by reduction to the blinding game.

Some auxiliary cryptographic arguments are performed as part of the above decomposition. We notably show that, during the first phase, the commit of an honest voter cannot be confused with that of another voter (honest or not), hence the failure of the duplication check at (11) can only be caused by two dishonest ballots, which is not a concern for privacy. Finally, we show that commits and commit keys remain secret in the relevant early phases of the protocol, using the blinding and commitment hiding assumptions on truncated system where the later mixnets are ineffective.

When ϕ does not hold, commitment keys are not revealed by the second mixnet. Hence privacy simply follows from the commitment hiding property. We decompose our equivalence as before, but when opening shuffles we organize their inputs by identity, resulting in an indistinguishability of the form

$$m, c_{\mathsf{A}}, c_{\mathsf{B}} \sim m, c_{\mathsf{B}}, c_{\mathsf{A}}$$

which **crypto** reduces to the commitment's hiding game.

Proof development and performances. We provide in Table 1 an overview of how the 10 kLOC of the Squirrel development are split across the different steps of our proof. The development itself is available online [37].

Further, Table 1 contains performance evaluation numbers. Performances were evaluated on a standard laptop running Ubuntu with 32 GB of RAM and an Intel i7-12700H processor. The overall proof takes less than 5 minutes to run and contains 28 calls to **crypto**. The longest calls to **crypto** are during the initial CCA2 game-hops idealizing the content of the encrypted messages, with 11 seconds for the CCA2 call idealizing the messages sent to the second mixnet (`ccapk2.sp`), and 3.8 seconds for the idealization of the first mixnet (`ccapk1.sp`). Notice that each file `ccapkN.sp` contains two game-hops justified by **crypto**, one per side of the protocol. All other calls to **crypto** are significantly faster, taking less than a second. Overall, we see that calls to **crypto** take at most a dozen of seconds, and often less than a second. Finally, 33 out of 262 seconds (13 %) of proof verification is spent by Squirrel doing simulator synthesis. Overall, we find that the performances of **crypto** are acceptable, in view of its high degree of automation and of its intended use in interactive proof developments.

8 Related work

We discuss related work on mechanized cryptography, simulator synthesis, and formal proofs of voting protocols.

Mechanized cryptography. State-of-the-art mechanized cryptographic techniques lack general simulator synthesis.

The best system in that respect is CryptoVerif [13]: it is designed for mechanizing game hopping proofs, and can automatically recognize when a game hop is justified by a cryptographic assumption, i.e. when a simulator exists. But it does not provide a logic that may be used to discharge verification conditions justifying a game hop, limiting its applicability.

On the opposite end of the spectrum, another dominant system is EasyCrypt [8, 25], a full-featured proof assistant for a higher-logic which embeds several domain-specific languages for writing and reasoning about cryptographic code. This design makes EasyCrypt very expressive, but it is much less automated. In particular, simulators have to be explicitly given as programs by the user: there is no simulator synthesis at all. Other systems based on general-purpose proof assistants, such as SSProve [31] in Coq or CryptHOL [10] in Isabelle/HOL, suffer even more from the same trade-off.

Typing-based reduction synthesis. Another line of work [23, 24, 28] introduced the typing-based reduction synthesis (TBRS) approach, which allows to automatically establish the existence of simulators using typed interfaces and parametricity results [34]. Concretely, they provide abstract interfaces and show that, if a program is well-typed against any of these interfaces, then the concrete implementation of the underlying cryptographic primitive can be swapped for an idealized implementation (à la UC [14]).

While all these approaches are highly automated, they only support a restricted number of cryptographic primitives for which abstract interfaces and the associated parametricity meta-theorems can be designed. More precisely, there exist TBRS interfaces for games targeting encryptions, MACs, signatures, hashes, Diffie-Hellman, and RSA. Our approach

supports these games in theory. Some games (RSA, PRF-ODH) have never been used in actual Squirrel developments, but should present no difficulties.

The TBRS approach is not systematic, and designing and proving a new interface is manual and requires deep expertise (this is the kind of limitation which [5] tackled). Previous TBRS papers targeted primitives used in TLS [35] and QUIC [36] and do not support the blind signatures and commitments needed in FOO. We see no theoretical obstacles in designing TBRS interfaces and results for these games, but it remains to be done by an expert. Thus, FOO is currently out-of-scope of this approach.

Previous work using TBRS relied on the F⋆ proof-assistant [40], which lacks relational reasoning capabilities. Hence, large parts of cryptographic arguments can only be carried-out on paper. Further, the complexity analysis of the synthesized reduction is only manually checked. We do not suffer from these limitations.

There are other approaches [21, 30] to automatically building cryptographic reductions using type-based techniques, by relying on ad hoc techniques for soundness rather than parametricity. Again, these approaches only support a restricted, built-in number of cryptographic primitives.

Security of e-voting protocols. We already compared to the pen-and-paper computational proof for FOO [6]. There exist several tool-assisted proofs of security for FOO [15, 22, 32], but all of them are in the symbolic model. To our knowledge, our proof is the first mechanized *computational* cryptographic proof for FOO. Conducting a proof of FOO in another tool would be a major endeavor. Further, it would not have been possible to prove vote privacy for FOO in Squirrel without our improved **crypto** tactic. Thus, we do not compare our proof to alternative ones using other tools or approaches.

There have been a few mechanized computational cryptographic proofs of other e-voting protocols, for Helios [18], Belenios [19], and Selene [27]. All these proofs have been carried-out in EasyCrypt [25]. These protocols have a different structure than FOO, as they are not two-phased e-voting protocols. This makes it it difficult to compare our development with theirs.

9 Conclusion

We have presented a novel automated procedure for inductive simulator synthesis, which can synthesize memoizing simulators and infer the precise time-sensitive memory invariants needed to argue for the soundness of these simulators. We implemented our procedure as a Squirrel tactic, and used it in the first proof of privacy for the FOO e-voting protocol in the computational model, which is the most involved Squirrel proof to date.

Acknowledgments

This work received funding from the France 2030 program managed by the French National Research Agency under grant agreement No. ANR-22-PECY-0006.

Ethical Considerations

This work is in most part theoretical, with applications to formal proofs of cryptographic protocols, and in particular to the FOO e-voting protocol. While the theoretical development does not raise ethical concerns in itself, its applications do. A mechanized proof may have a negative outcome if it provides false confidence in the security of a protocol. This might happen if the limitations of the security model or the assumptions under which the proof is carried out are not precisely stated, a problem which we have strived to avoid. In our work, the most prominent example is the proof of vote privacy for the FOO e-voting protocol. We clearly stated in introduction that vote privacy is not the only property that one should look for in an e-voting protocol, and that formal proofs have their limitations too.

Open Science

The theoretical part of our work can be assessed from the paper, its appendices, and the cited papers which are all easily accessible.

The source code of our improved version of Squirrel's **crypto** tactic is available at [37], in the directory `squirrel`, under (mostly) the MIT licence. Further, our changes have been upstreamed in Squirrel main branch [26].

Finally, all Squirrel proof files are publicly accessible at [37], in directory `proofs`, under the MIT license — this includes the vote privacy proof for FOO, our motivating example, and an improved version of the NSL development from [5].

References

[1] Martín Abadi, Bruno Blanchet, and Cédric Fournet. The applied pi calculus: Mobile values, new names, and secure communication. *J. ACM*, 65(1):1:1–1:41, 2018.

[2] David Baelde, Stéphanie Delaune, Charlie Jacomme, Adrien Koutsos, and Solène Moreau. An interactive prover for protocol verification in the computational model. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, pages 537–554. IEEE, 2021.

[3] David Baelde, Caroline Fontaine, Adrien Koutsos, Guillaume Scerri, and Théo Vignon. A probabilistic logic

for concrete security. In *CSF*, pages 324–339. IEEE, 2024.

[4] David Baelde, Adrien Koutsos, and Joseph Lallemand. A higher-order indistinguishability logic for cryptographic reasoning. In *38th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2023, Boston, MA, USA, June 26-29, 2023*, pages 1–13. IEEE, 2023.

[5] David Baelde, Adrien Koutsos, and Justine Sauvage. Foundations for cryptographic reductions in CCSA logics. In *CCS*, pages 2814–2828. ACM, 2024.

[6] Gergei Bana, Rohit Chadha, and Ajay Kumar Eeralla. Formal analysis of vote privacy using computationally complete symbolic attacker. In *ESORICS (2)*, volume 11099 of *Lecture Notes in Computer Science*, pages 350–372. Springer, 2018.

[7] Gergei Bana and Hubert Comon-Lundh. A computationally complete symbolic attacker for equivalence properties. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 609–620. ACM, 2014.

[8] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella-Béguelin. Computer-aided security proofs for the working cryptographer. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 71–90. Springer, 2011.

[9] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella-Béguelin. Formal certification of code-based cryptographic proofs. In *POPL*, pages 90–101. ACM, 2009.

[10] David A. Basin, Andreas Lochbihler, and S. Reza Sefidgar. Crypthol: Game-based proofs in higher-order logic. *J. Cryptol.*, 33(2):494–566, 2020.

[11] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426. Springer, 2006.

[12] Josh Daniel Cohen Benaloh. *Verifiable secret-ballot elections*. Yale University, 1987.

[13] Bruno Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Trans. Dependable Secur. Comput.*, 5(4):193–207, 2008.

[14] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145. IEEE Computer Society, 2001.

[15] Rohit Chadha, Vincent Cheval, Ştefan Ciobâcă, and Steve Kremer. Automated verification of equivalence properties of cryptographic protocols. *ACM Trans. Comput. Log.*, 17(4):23, 2016.

[16] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–88, 1981.

[17] David Chaum. Blind signatures for untraceable payments. In *CRYPTO*, pages 199–203. Plenum Press, New York, 1982.

[18] Véronique Cortier, Constantin Catalin Dragan, François Dupressoir, Benedikt Schmidt, Pierre-Yves Strub, and Bogdan Warinschi. Machine-checked proofs of privacy for electronic voting protocols. In *IEEE Symposium on Security and Privacy*, pages 993–1008. IEEE Computer Society, 2017.

[19] Véronique Cortier, Constantin Catalin Dragan, François Dupressoir, and Bogdan Warinschi. Machine-checked proofs for electronic voting: Privacy and verifiability for belenios. In *CSF*, pages 298–312. IEEE Computer Society, 2018.

[20] Véronique Cortier and Ben Smyth. Attacking and fixing Helios: An analysis of ballot secrecy. *J. Comput. Secur.*, 21(1):89–148, 2013.

[21] Stéphanie Delaune, Clément Hérouard, and Joseph Lallemand. Secrecy by typing in the computational model. In *CSF*, pages 17–32. IEEE, 2025.

[22] Stéphanie Delaune, Mark Ryan, and Ben Smyth. Automatic verification of privacy properties in the applied pi calculus. In *IFIPTM*, volume 263 of *IFIP Advances in Information and Communication Technology*, pages 263–278. Springer, 2008.

[23] Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Jonathan Protzenko, Aseem Rastogi, Nikhil Swamy, Santiago Zanella-Béguelin, Karthikeyan Bhargavan, Jianyang Pan, and Jean Karim Zinzindohoue. Implementing and proving the TLS 1.3 record layer. In *IEEE Symposium on Security and Privacy*, pages 463–482. IEEE Computer Society, 2017.

[24] Antoine Delignat-Lavaud, Cédric Fournet, Bryan Parno, Jonathan Protzenko, Tahina Ramananandro, Jay Bosamiya, Joseph Lallemand, Itsaka Rakotonirina, and Yi Zhou. A security model and fully verified implementation for the IETF QUIC record layer. In *SP*, pages 1162–1178. IEEE, 2021.

[25] The Easycrypt development team. The EasyCrypt Prover repository, accessed august 2025. https://github.com/EasyCrypt/easycrypt/.

[26] The Squirrel development team. The Squirrel Prover repository, accessed august 2025. https://github.com/squirrel-prover/squirrel-prover/.

[27] Constantin Catalin Dragan, François Dupressoir, Ehsan Estaji, Kristian Gjøsteen, Thomas Haines, Peter Y. A. Ryan, Peter B. Rønne, and Morten Rotvold Solberg. Machine-checked proofs of privacy against malicious boards for selene & co. In *CSF*, pages 335–347. IEEE, 2022.

[28] Cédric Fournet, Markulf Kohlweiss, and Pierre-Yves Strub. Modular code-based cryptographic verification. In *CCS*, pages 341–350. ACM, 2011.

[29] Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A practical secret voting scheme for large scale elections. In *AUSCRYPT*, volume 718 of *Lecture Notes in Computer Science*, pages 244–251. Springer, 1992.

[30] Joshua Gancher, Sydney Gibson, Pratap Singh, Samvid Dharanikota, and Bryan Parno. Owl: Compositional verification of security protocols via an information-flow type system. In *SP*, pages 1130–1147. IEEE, 2023.

[31] Philipp G. Haselwarter, Exequiel Rivas, Antoine Van Muylder, Théo Winterhalter, Carmine Abate, Nikolaj Sidorenco, Catalin Hritcu, Kenji Maillard, and Bas Spitters. Ssprove: A foundational framework for modular cryptographic proofs in coq. *ACM Trans. Program. Lang. Syst.*, 45(3):15:1–15:61, 2023.

[32] Steve Kremer and Mark Ryan. Analysis of an electronic voting protocol in the applied pi calculus. In *ESOP*, volume 3444 of *Lecture Notes in Computer Science*, pages 186–200. Springer, 2005.

[33] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptol.*, 4(2):151–158, 1991.

[34] John C. Reynolds. Types, abstraction and parametric polymorphism. In *IFIP Congress*, pages 513–523. North-Holland/IFIP, 1983.

[35] RFC. The transport layer security (TLS) protocol version 1.3. https://datatracker.ietf.org/doc/html/rfc8446, 2018.

[36] RFC. Quic: A udp-based multiplexed and secure transport. https://datatracker.ietf.org/doc/html/rfc9000, 2021.

[37] Justine Sauvage, Adrien Koutsos, and David Baelde. Leveraging cryptographic simulator synthesis for formally verifying the foo e-voting protocol – artifacts, December 2025. https://doi.org/10.5281/zenodo.17880703.

[38] Justine Sauvage, Adrien Koutsos, and David Baelde. Leveraging cryptographic simulator synthesis for formally verifying the foo e-voting protocol – long version. Technical report, January 2026. https://inria.hal.science/hal-05453231.

[39] Victor Shoup. Sequences of games: a tool for taming complexity in security proofs. *IACR Cryptol. ePrint Arch.*, page 332, 2004.

[40] Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. Secure distributed programming with value-dependent types. *J. Funct. Program.*, 23(4):402–451, 2013.

[41] Cédric Villani. *Optimal transport – Old and new*, volume 338, pages xxii+973. 01 2008.

[42] Douglas Wikström. A universally composable mix-net. In *TCC*, volume 2951 of *Lecture Notes in Computer Science*, pages 317–335. Springer, 2004.

A Basic Simulator Synthesis

The full set of rules used to present our basic simulator synthesis procedure are described in Fig. 10, and make use of the following type annotation: if a type τ is annotated by $\mathtt{enum}_{\mathtt{poly}}(\tau)$, then τ must be a base-type in $\mathbb{B}$ and it must be the case that in any model $\mathbb{M}$, there exists a machine $\mathcal{M}$ running in polynomial time such that $\mathcal{M}(1^\eta) = [a_1, \ldots, a_n]$ (where $[a_1, \ldots, a_n]$ denotes the array of value a_1 to a_n, suitably encoded as a bit-string) and where $[\![\tau]\!]^\eta_{\mathbb{M}} = \{a_1, \ldots, a_n\}$.

Core rules. *In the* TRANS *and* CASE *rules,* $\overline{C}$ *denotes the sub-multiset of* C *where constraints with tag* $\mathtt{T}^{\mathsf{loc}}_{\mathsf{G}}$ *are removed; the remaining constraints may be duplicated without any impact on the constraints' validity.*

UNREACH

$$\frac{}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid f) \rightsquigarrow ([\neg f]_{\mathrm{e}};\emptyset;\phi;\epsilon)}$$

CONV

$$\frac{in' = \mathsf{whnf}^{\Theta}_{\mathcal{E}}(in) \qquad out' = \mathsf{whnf}^{\Theta}_{\mathcal{E}}(out) \qquad \mathcal{E};\Theta;C;\phi \vdash in' \rhd out' \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd out \rightsquigarrow (\Theta';C';\psi;w)}$$

TRANS

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o' \mid f) \rightsquigarrow (\Theta';C';\psi';w') \qquad \mathcal{E};\Theta;\overline{C}\cdot\overline{C'};\psi' \vdash in,(o' \mid f) \rhd (o'' \mid f) \rightsquigarrow (\Theta'';C'';\psi'';w'')}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o',o'' \mid f) \rightsquigarrow (\Theta',\Theta'';C'\cdot C'';\psi'';w',w'')}$$

CASE

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd (c \mid f) \rightsquigarrow (\Theta_c;C_c;\psi_c;w_c) \qquad \mathcal{E};\Theta;\overline{C}\cdot\overline{C_c};\psi_c \vdash in \rhd (o \mid f \wedge c) \rightsquigarrow (\Theta_\top;C_\top;\psi_\top;w_\top) \qquad \mathcal{E};\Theta;\overline{C}\cdot\overline{C_c}\cdot\overline{C_\top};\psi_c \vdash in \rhd (o \mid f \wedge \neg c) \rightsquigarrow (\Theta_\bot;C_\bot;\psi_\bot;w_\bot) \qquad \Theta' = \Theta_c,\Theta_\top,\Theta_\bot \qquad C' = C_c\cdot C_\top\cdot C_\bot \qquad w' = w_c,w_\top,w_\bot \qquad \psi' = (\psi_\top \sqcap c) \sqcup (\psi_\bot \sqcap \neg c)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid f) \rightsquigarrow (\Theta';C';\psi';w')}$$

ORACLE

$$\frac{\theta = \mathsf{unify}^{\mathcal{E}}_{\vec{x},\vec{y},\vec{z}}(o = o_f) \qquad \theta(\vec{y}) = (\mathsf{k_v}\ p_{\mathsf{v}})_{\mathsf{v}\in\mathcal{G}.\mathsf{glob}_\$} \qquad \theta(\vec{z}) = (\mathsf{r_v}\ s_{\mathsf{v}})_{\mathsf{v}\in f.\mathsf{loc}_\$} \qquad \mathcal{E};\Theta;C;\phi \vdash in \rhd (\theta(\vec{x}),(p_{\mathsf{v}})_{\mathsf{v}\in\mathcal{G}.\mathsf{glob}_\$},(s_{\mathsf{v}})_{\mathsf{v}\in f.\mathsf{loc}_\$} \mid g) \rightsquigarrow (\Theta';C';\phi';w) \qquad C'' = \left(\Pi_{\mathsf{v}\in\mathcal{G}.\mathsf{glob}_\$}(\emptyset,\mathsf{k_v},o_{\mathsf{v}},\mathtt{T}^{\mathsf{glob}}_{\mathsf{G},\mathsf{v}},g)\right)\cdot\left(\Pi_{\mathsf{v}\in f.\mathsf{loc}_\$}(\emptyset,\mathsf{r_v},s_{\mathsf{v}},\mathtt{T}^{\mathsf{loc}}_{\mathsf{G}},g)\right) \qquad g_\mu = \text{b-eval}_{\phi'}(c_f\theta) \qquad \psi = \mathsf{post}^{\theta}_{f}(\phi')}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid g) \rightsquigarrow (\Theta',[g \Rightarrow g_\mu]_{\mathrm{e}};C'\cdot C'';\psi;w)}$$

Destruction rules.

NAME

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (\mathsf{n}\ o \mid f) \rightsquigarrow (\Theta';C'\cdot(\emptyset,\mathsf{n},o,\mathtt{T}_S,f);\psi;w)}$$

FA.QUANT

$$\frac{Q \in \{\forall,\exists,\lambda\} \qquad \mathtt{enum}_{\mathtt{poly}}(\tau) \qquad \mathcal{E};\mathcal{E},x:\tau;\phi \vdash in \rhd (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (Q(x:\tau).\ o \mid f) \rightsquigarrow (\forall x.\Theta',\Pi_x.C';\forall x.\psi;\lambda x.w)}$$

FA.ITE

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd (g \mid f),(o_0 \mid f \wedge g),(o_1 \mid f \wedge \neg g) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (\text{if } g \text{ then } o_0 \text{ else } o_1 \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

FA.$\wedge$

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd (g_0 \mid f),(g_1 \mid f \wedge g_0) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (g_0 \wedge g_1 \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

FA.$\Rightarrow$

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd (g_0 \mid f),(g_1 \mid f \wedge g_0) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (g_0 \Rightarrow g_1 \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

FA.$\vee$

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd (g_0 \mid f),(g_1 \mid f \wedge \neg g_0) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (g_0 \vee g_1 \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

FA.MATCH

$$\frac{t \in in.\mathtt{std} \qquad \text{for any } i,\ \mathcal{E},\vec{x}_i;\Theta;C;\phi \vdash in,\vec{x}_i \rhd (u_i \mid f \wedge t = c_i\ \vec{x}_i) \rightsquigarrow (\Theta_i;C_i;\psi_i;w_i) \qquad \Theta_{\mathsf{out}} = (\forall\vec{x}_i.\Theta_i)_i \qquad C_{\mathsf{out}} = \textstyle\prod_i \forall\vec{x}_i.C_i \qquad \psi_{\mathsf{out}} = \bigsqcup_i(\forall\vec{x}_i.\psi_i \sqcap (t = c_i\ \vec{x}_i)) \qquad w_{\mathsf{out}} = (\forall\vec{x}_i.w_i)_i}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (\mathsf{match}\ t\ \mathsf{with}\ (c_i\ \vec{x}_i \mapsto u_i)_i \mid f) \rightsquigarrow (\Theta_{\mathsf{out}};C_{\mathsf{out}};\psi_{\mathsf{out}};w_{\mathsf{out}})}$$

FA

$$\frac{s \in \mathcal{L} \qquad \mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (s\ o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

Memory rule. *(For any type* τ*, we require that* $(\mathsf{witness}_\tau : \tau) \in \mathcal{L}$*. We write* witness *when the type* τ *is clear from context.)*

LOAD

$$\frac{\lambda\vec{x}.(u \mid g) \in in.\mathtt{std} \qquad \theta = \mathsf{unify}^{\mathcal{E}}_{\vec{x}}(u = o) \qquad \theta_0 = \theta[\vec{x}\backslash\mathsf{dom}(\theta) \mapsto \mathsf{witness}] \qquad \mathcal{E};\Theta \vdash_{\mathsf{auto}} [f \Rightarrow g\theta_0]_{\mathrm{e}} \qquad \mathcal{E};\Theta;C;\phi \vdash in \rhd (\vec{x}\theta_0 \mid f) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

Memoization rules.

MEMOIZE.STORE

$$\frac{\mathcal{E};\Theta;C;\phi \vdash in \rhd out \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd out \rightsquigarrow (\Theta';C';\psi;w,out)}$$

MEMOIZE.LOAD

$$\frac{\lambda\vec{x}.(u \mid g) \in in.\mathtt{memo} \qquad \theta = \mathsf{unify}^{\mathcal{E}}_{\vec{x}}(u = o) \qquad \vec{y} = \vec{x}\backslash\mathsf{dom}(\theta) \qquad \mathtt{enum}_{\mathtt{poly}}(\mathsf{type}_{\mathcal{E}}(\vec{y})) \qquad \mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid f \wedge (\forall\vec{y}.\ \neg g\theta)) \rightsquigarrow (\Theta';C';\psi;w)}{\mathcal{E};\Theta;C;\phi \vdash in \rhd (o \mid f) \rightsquigarrow (\Theta';C';\psi;w)}$$

Figure 10: Basic proof-search rules.