

Cryptographic Choreographies

Sebastian Mödersheim	Simon Lund	Alessandro Bruni	Marco Carbone	Rosario Giustolisi
<i>DTU Compute</i>	<i>DTU Compute</i>	<i>IT-University of</i>	<i>IT-University of</i>	<i>IT-University of</i>
Kgs. Lyngby, Denmark	Kgs. Lyngby, Denmark	<i>Copenhagen, Denmark</i>	<i>Copenhagen, Denmark</i>	<i>Copenhagen, Denmark</i>
0000-0002-6901-8319	0009-0005-2957-3472	0000-0003-2946-9462	0000-0001-9479-2632	0000-0002-2917-9601

Abstract—We present **CryptoChoreo**, a choreography language for the specification of cryptographic protocols. Choreographies can be regarded as an extension of Alice-and-Bob notation, providing an intuitive high-level view of the protocol as a whole (rather than specifying each protocol role in isolation). The extensions over standard Alice-and-Bob notation that we consider are nondeterministic choice, conditional branching, and mutable long-term memory. We define the semantics of **CryptoChoreo** by translation to a process calculus. This semantics entails an understanding of the protocol: it determines how agents parse and check incoming messages and how they construct outgoing messages, in the presence of an arbitrary algebraic theory and nondeterministic choices made by other agents. While this semantics entails algebraic problems that are in general undecidable, we give an implementation for a representative theory. We connect this translation to ProVerif and show on a number of case studies that the approach is practically feasible.

Acknowledgments: Partly funded by EU Horizon Europe under Grant Agreement no. 101093006 (TaRDIS).

I. INTRODUCTION

Specification languages for security protocols can be roughly divided into three classes. The most low-level one are based on multi-set rewriting rules, such as the Tamarin input language [1] and the AVISPA Intermediate Format [2], where each rule describes state transitions corresponding usually to a pair of protocol steps from the view of one honest agent: receiving a message, processing and checking it, and sending the next message. More high-level are languages based on process calculus such as ProVerif [3], where typically each role of the protocol is described like a program, often as a sequence of sending and receiving steps. The most high-level are languages based on Alice-and-Bob notation [4], [5], [6], [7], [8], [9], [10], which describe the entire protocol by an ideal run of the protocol as a sequence of $A \rightarrow B : M$ steps where role A sends message M to role B , thus describing the interplay of all roles.

Alice-and-Bob notation is very intuitive and succinct because it gives the synopsis of the protocol and leaves implicit how agents construct the messages they send, and how they parse and check the messages they receive. The latter is a non-trivial problem one has to solve when defining a formal language based on Alice-and-Bob notation, namely when giving a formal *semantics* by translation to a lower-level language. This was described by [6], [11] for models in the free term algebra, but a key question is how to deal with algebraic properties as needed, for instance, for Diffie-

Hellman. If Alice needs to construct $\exp(\exp(g, X), Y)$, given her own secret X and the public value $\exp(g, Y)$ from Bob, the semantics needs to infer that this is possible by composing $\exp(\exp(g, Y), X)$ since this is equivalent to the goal term by the algebraic properties of exponentiation. It turns out that one can define such semantics in a general, uniform, and concise way for an *arbitrary* algebraic theory as an intruder deduction problem [7], [8], [9], [10].

As a side effect of “abusing” the intruder deduction to define the behavior of honest agents, one prevents many specification errors that can easily happen in the lower-level formalisms, e.g., when the message sent by one agent are different from the messages that another agent expects, rendering the protocol unexecutable. This may lead in the worst case to a false negative (an attack of the real system is not detected because of a specification error). Many such errors are prevented by formal Alice-and-Bob approaches because the protocol would be refused as unexecutable by the compiler.

Thus, Alice-and-Bob notation is a beneficial and accessible specification language that can be used even without a deep background in formal verification. It is also striking how often scientific works that formalize a protocol using a lower-level language first summarize the protocol informally in Alice-and-Bob notation. However, existing formal Alice-and-Bob languages do not support branches in the protocol execution (by conditions or nondeterminism), unbounded repetition, or mutable long-term memory: everything is restricted to a linear session of fixed length.

While informal uses of Alice-and-Bob notation can easily be extended ad hoc, the first main contribution of this paper is a new choreography language, called *CryptoChoreo*, which extends Alice-and-Bob notation with nondeterministic choice, branching, and mutable long-term memory.¹

Such features are needed for instance if we want to formulate a protocol with a server that maintains a long-term database and that may react to a request in different ways depending on the current state of its database. Nondeterministic choice can be helpful for modelling several options in a protocol that are at a participant’s discretion, where we do not want to formalize how they make a decision. Especially this allows us to formulate an API where a user

¹Since a choreography can be executed an unbounded number of times, nondeterminism and long-term memory are sufficient to formulate repetitions and sequential composition of protocols without an explicit repetition construct.

can nondeterministically choose to send any of a number of commands to a server (who may in turn ask other servers in order to answer the request).

CryptoChoreo is, in a sense, a conservative extension of formal Alice-and-Bob languages: we give a semantics—parameterized over an arbitrary algebraic theory—that agrees with standard Alice-and-Bob languages on the subset that does not use the new constructs. A particular challenge for this semantics is to integrate the algebraic understanding of the protocol with branching: if one party has made a nondeterministic choice, other parties do not necessarily know which choice was taken. For instance, our semantics allows for the following protocol: Alice nondeterministically chooses one of two types of message that she sends encrypted for Bob over an intermediary server as an authentication service; the server is unaware in which branch the execution is, but in each case it can execute its step uniformly by checking a MAC from Alice and signing the resulting message.

The semantics is formulated as a translation to local behaviors, i.e., a process for each role of the protocol. In general, this involves algebraic problems that are not recursively computable (since, e.g., whether two terms are equal under a set of algebraic equations is in general undecidable). The second main contribution is to give a computable translation for an algebraic theory that includes standard constructors and destructors as well as exponentiation (for Diffie-Hellman).

As a third contribution, we connect our translator with ProVerif and demonstrate the effectiveness of our approach with several case studies. A particular challenge is that ProVerif’s abstraction often is not precise enough when the long-term memory induces non-monotonic behavior (e.g., when a certain action is possible only until a change of the memory state) and thus fails to verify a protocol. We have developed several heuristics to make sound encodings in ProVerif that often overcome these problems.

Note that both the translation to process calculus and the target of the translation, ProVerif, are using Dolev-Yao models, i.e., treating cryptographic operations as blackboxes. The translation formalizes how each role is *supposed to execute* the choreography, and this should not require the role to break the cryptography. Thus, it makes sense that our semantics basically requires that a Dolev-Yao intruder with the specified initial knowledge is able to correctly execute a given role. For the target of the translation one may, however, also consider computational verification like CryptoVerif. This would require specifying more details like key-sizes and required cryptographic properties of the encryption primitives, as well as different kinds of goal specifications. Similarly, the translation could also be used to generate a secure-by design implementation, in the sense that the implementation creates outgoing, and checks incoming, messages in the same way as the formal ProVerif model.

The rest of this paper is organized as follows: in Section II we define the syntax of CryptoChoreo and give an example; in Section III we define the semantics for an arbitrary algebraic theory; in Section IV, we describe the practical implementation

$$\begin{aligned}
 &A : (\nu M. \\
 &\quad A \rightarrow s : \mathit{scrypt}((B, \mathit{crypt}((\mathit{msg}, M), \mathit{ek}(B))), \mathit{shk}(A, s)). \\
 &\quad s \rightarrow B : \mathit{sign}((A, \mathit{crypt}((\mathit{msg}, M), \mathit{ek}(B))), \mathit{inv}(\mathit{sk}(s))). \\
 &\quad B \text{ ni-authenticates } A \text{ on } M) \\
 &+ (\nu K. \\
 &\quad A \rightarrow s : \mathit{scrypt}((B, \mathit{crypt}((\mathit{upd}, K), \mathit{ek}(B))), \mathit{shk}(A, s)). \\
 &\quad s \rightarrow B : \mathit{sign}((A, \mathit{crypt}((\mathit{upd}, K), \mathit{ek}(B))), \mathit{inv}(\mathit{sk}(s))). \\
 &\quad K \text{ secret between } A, B)
 \end{aligned}$$

Fig. 1. Example Choreography.

for a representative algebraic theory; in Section V we describe the connection to ProVerif and our case studies; we present related work in Section VI; and we conclude in Section VII.

II. CHOREOGRAPHY LANGUAGE

Let us start with the example choreography in Fig. 1. We discuss later the front-matter declarations (e.g., types and initial knowledge) that is needed for a full specification. Role \mathcal{A} is starting the choreography and first makes a nondeterministic choice (+) about which of the two sub-choreographies to run. In the first case, \mathcal{A} generates a new random M (representing a message to send to \mathcal{B}), pairs it with a constant tag msg (indicating that this is a message transmission), then asymmetrically encrypts it with the encryption key $\mathit{ek}(B)$, highlighted in blue. \mathcal{A} then symmetrically encrypts the blue message and the name of \mathcal{B} with a shared key $\mathit{shk}(A, s)$ with the (trusted) server s . Suppose only \mathcal{B} knows the decryption key $\mathit{inv}(\mathit{ek}(B))$, then s can only decrypt the outer symmetric encryption, but not the blue message. The next step is that s signs the blue message and the sender name A using the private signature key $\mathit{inv}(\mathit{sk}(s))$; the idea is that s is vouching that the blue message really came from A (as the symmetric encryption guarantees to s) and \mathcal{B} can verify the signature knowing the corresponding public signing key $\mathit{sk}(s)$. Finally, we have an authentication goal when \mathcal{B} receives this message: namely, that A has indeed intended to send the message M to \mathcal{B} . Non-injective (ni in ni-authenticates) here means that \mathcal{B} has no freshness guarantee (the message may be a replay). The other sub-choreography is very similar, except that here \mathcal{A} sends a different kind of message, a key update with a fresh key K with a different goal: K is secret between A and B (of course both sub-choreographies could have a secrecy and an authentication goal).

Each of the two sub-choreographies could be specified in existing formal Alice-and-Bob languages, but not the nondeterministic choice. Note that s here does not realistically know which sub-choreography was chosen by \mathcal{A} . The semantics we give below sorts this out correctly: the server shall do the same operations in both choreographies and simply handle the blue message as a black box. Note that an intruder playing role \mathcal{A} is also allowed; and this intruder may form a blue message that complies with neither sub-choreography; s will anyway accept this message if everything it can check complies with the protocol.

$C ::=$	$\mathbf{0}$	(end)
$A \rightarrow B : t.C$		(interaction)
t secret between $A_1, \dots, A_n.C$		(secrecy goal)
B authenticates A on $t.C$		(inj. auth. goal)
B ni-authenticates A on $t.C$		(non-inj. auth. goal)
$A : \mathcal{A}$		(atomic)
$\mathcal{A} ::=$	$\nu N. \mathcal{A}$	(new)
$\mathcal{A}_1 + \mathcal{A}_2$		(choice)
if $s \doteq t$ then \mathcal{A}_1 else \mathcal{A}_2		(condition)
$s := c[t].\mathcal{A}$		(memory read)
$c[t] := s.\mathcal{A}$		(memory write)
C		(choreography)

Fig. 2. Syntax of CryptoChoreo.

a) *Terms*: We build terms using an alphabet Σ of function symbols and a set of variables \mathcal{V} . We denote all function symbols with lower-case letters and all variables with upper-case letters. In the above example, e.g., *script*, *shk*, *inv*, *msg*, and *s* are function symbols (constants are function symbols with 0 arguments), while *A* and *M* are variables. Variables mean that they can take a different value in every run of the choreography. We use sans-serif font to denote *roles* of the protocol; they can be variables like *A* and *B* or constants like *s*. The latter means that there is one fixed player who cannot be the intruder—an easy way to specify a trusted third party. We will discuss below the specification of function symbols and their algebraic properties. We also use the notation (t_1, \dots, t_n) for a concatenation using a pair operator.

b) *Syntax*: The formal syntax of a choreography is shown in Fig. 2. A choreography $\mathbf{0}$ represents a terminated protocol, in which each participant has terminated. We omit trailing $\mathbf{0}$ s when this is clear from the context. An interaction $A \rightarrow B : t.C$ denotes a protocol where *A* sends a term *t* to *B* and then continues with choreography *C*. The next items represent the specification of goals (injective and non-injective authentication, and secrecy) that we discuss later in detail.

All these constructs are present in existing formal Alice-and-Bob languages. What CryptoChoreo is adding are constructs that are all local to one role. Denote this by giving a role name *A* followed by a colon and an *atomic section* \mathcal{A} of steps that *A* locally executes. Here, we have the fresh generation of a random value νN (as is standard). Next, we have the nondeterministic choice $\mathcal{A}_1 + \mathcal{A}_2$ (this is actually an internal choice for the role who runs this atomic section and an external choice for all others). Then, we have a conditional where the condition is a comparison of terms. Last, we have reading from, and writing to, long-term memory. We denote with $c[t]$ a memory cell in a family of memory cells *c*, where *c* is an identifier and *t* is an index term. On memory read and write, there will be no race conditions with other parallel sessions, because our semantics will treat each section \mathcal{A} as atomic like the name suggests.

c) *Memory Cell Example*: To illustrate memory, consider the following augmentation shown in Fig. 3 of the example in Fig. 1, where *A* may not know the encryption key of every

$A :$	$EKB := keys[B].$
	if $EKB \doteq blank$ then
	$\nu N. A \rightarrow s : (key, B, N)$
	$s \rightarrow A : sign((key, B, ek(B), N), inv(sk(s)))$
	$A : keys[B] := ek(B)$
	else ... (previous example with $ek(B) \mapsto EKB$)

Fig. 3. Extension of the example from Fig. 1.

B (but just the public signature verification key $sk(s)$ of *s*). When *A* wants to talk to *B*, she checks her memory cell $keys[B]$; if this returns the initial value *blank*, then she does not know the key of *B* and asks *s*, which we assume knows all public encryption keys via the *ek* function and can vouch for it with its signature. *A* checks the signature and stores the key (note that for *A* the term $ek(B)$ is just a blob that she cannot verify in itself; this will be shown in the semantics below). Note that we do not have a repetition operator, and this example shows why this is without loss of generality: in the case that *A* does not know the key of *B*, the run of the choreography ends with writing the key (that she received from *s*) into her memory. Thus, in any later run between the same *A* and *B*, *A* will retrieve the key from her memory and run the standard protocol with it. In other words, unbounded repetition is implicit, because a choreography can be executed any number of times (also in parallel) with arbitrary instances of the (non-constant) roles and information between different runs can be transferred using memory.

d) *Front Matter*: The definition of our choreography language is parameterized over sets Σ and \mathcal{V} (respectively, function symbols and variables) and a set of equations *E* (over Σ, \mathcal{V} -terms). The set *E* induces a congruence relation $=_E$ on terms. In the implementation of our translation, in Section IV, we instantiate *E* with a concrete choice of properties.

Some variables and functions symbols are declared as roles (set in sans-serif in this paper) and only these can be used in places where the syntax indicates sans-serif font. For each role, one must declare the *initial knowledge*: a list of terms where all occurring variables are of type role. For our example (in the augmented version where *A* does not initially know *B*'s public key), this declaration could be:

$A :$	$A, B, s, sk(s), shk(A, s)$
$s :$	$A, B, s, sk(s), inv(sk(s)), shk(A, s), ek(B)$
$B :$	$A, B, s, sk(s), ek(B), inv(ek(B))$

Note that with less knowledge the protocol would not be executable (neither would the initial version of the example where *A* cannot ask the server for *B*'s public key be).

We require that for every variable that is *not* of type role, the first occurrence is either in a new statement (like νX) or in a memory read (like $X := c[s]$). Also, in a new statement νX , we assume that *X* did not occur before in the choreography (this can be achieved by renaming). In contrast, variables in a memory read may have occurred before, e.g., $\nu X. \dots X := c[s_1]. \dots X := c[s_2]$ is legal, and it would mean

\mathcal{L}	$::=$ $\mathbf{0}$ \mid $\text{send}(r).\mathcal{L}_i$ \mid $\text{receive}(\mathcal{X}).\mathcal{L}$ \mid $\text{lock}.\mathcal{A}$	 (end) (send) (receive) (atomic local)
\mathcal{A}	$::=$ $\nu\mathcal{X}.\mathcal{A}$ \mid $\mathcal{A}_1 + \mathcal{A}_2$ \mid $\text{if } r_1 \doteq r_2 \text{ then } \mathcal{A}_1 \text{ else } \mathcal{A}_2$ \mid $\mathcal{X} := c[r_1].\mathcal{A}$ \mid $c[r_1] := r_2.\mathcal{A}$ \mid $\text{event}(r).\mathcal{A}$ \mid $\text{unlock}.\mathcal{L}$	 (new) choice (cond) (mem read) (mem write) (event) (end)

Fig. 4. Syntax of Local Behaviors.

that the value retrieved here is the same value as before—in an ideal unattacked execution of the choreography. Our semantics can tell if the respective role has the necessary knowledge to check that and insert such a check in the code, if so.

III. TRANSLATION SEMANTICS

The semantics of CryptoChoreo is now defined by a translation to a process calculus, where we define for each role of the choreography a process, representing the local behavior of this role in one execution of the choreography. We then allow arbitrary instances of all roles to run in parallel, together with an intruder who can also play any of the roles (except trusted third parties) as a normal participant (but who does not necessarily follow the protocol).

A. Local Behaviors

It is convenient for the translation and the later connection to ProVerif to define a restricted syntax and semantics for local behaviors as the target language of the translation semantics.

a) *Syntax*: The syntax of local behaviors mirrors that of CryptoChoreo from the point of view of a single role: instead of a communication step between two roles, we have now sending and receiving. We use here however a different set of symbols Σ_p which represent public functions, i.e., functions that every agent, including the intruder, can apply. This will include most of the functions from Σ like *crypt* that represent cryptographic algorithms, as well as public constants like *msg* in the example above. It will not include, however, some functions from Σ that just describe relations in the model, but do not represent actual cryptographic algorithms like *inv* (which maps public keys to the corresponding private key) or *ek* (which maps an agent to a corresponding public key). Σ_p will also include some functions that are not in Σ : observe that we have not used any functions for decryption or signature verification, because they are *destructors* or *verifiers*, i.e., functions that extract a subterm or verify the structure of a term; while the choreography is only concerned with *constructed* messages. Also, we will use a distinct set of variables called *labels*. Labels are denoted $\mathcal{X}_1, \mathcal{X}_2, \dots$ and are disjoint from Σ, Σ_p , and \mathcal{V} . The terms built from these variables and Σ_p are called *recipes* and we denote them with r, r_1, r_2 , etc.

Fig. 4 shows the syntax of local behaviors. Note that, in order to enter the atomic section, it is necessary to make a lock step and it has to end with an unlock step; the semantics of local behaviors use that as a mutual exclusion mechanism on the memory to prevent race conditions (even if one ensures by design that each memory cell belongs to a particular agent, there may be more than one run of the choreography in parallel). The other constructs are similar to choreographies. However, instead of variables like M we have now labels like \mathcal{X} , and instead of terms (over Σ and \mathcal{V}) we have recipes. Note that memory read and receiving can only “read into” a label \mathcal{X} (while on the choreography level, these can be composed terms). We require that at each receive and memory read, we use a new label (that did not occur before in the local behavior; this can be easily achieved by renaming).

b) *Frames*: To capture the knowledge of an honest agent at a state of protocol execution, we define a *frame* to be a finite mapping $F = [\mathcal{X}_1 \mapsto t_1, \dots, \mathcal{X}_n \mapsto t_n]$ where the \mathcal{X}_i are labels and the t_i are terms. We call $\{\mathcal{X}_1, \dots, \mathcal{X}_n\}$ the *domain* of F and we say F is *concrete* if the t_i contain no variables. The semantics of local behaviors will only use concrete frames. Given a recipe r , we use a frame F like a substitution and write $F(r)$ for the term that results from replacing the labels \mathcal{X}_i with the respective term t_i ; $F(r)$ is undefined if r contains labels outside the domain of F .

To each role A , we attach an initial knowledge frame F_A that is not necessarily concrete but contains only variables of type role. In the translation from CryptoChoreo to local behaviors, we take the initial knowledge of each role $A : t_1, \dots, t_n$ at the CryptoChoreo level and turn into an initial knowledge frame $F_A = [\mathcal{X}_1 \mapsto t_1, \dots, \mathcal{X}_n \mapsto t_n]$ for the local behavior.

We require that all labels in a local behavior first occur in the initial knowledge frame, in a new, in a receive, or in a memory read.

c) *Example Local Behavior*: The role A of the example of Fig. 3 will be translated by the semantics into the following local behavior:

```

 $F_A = [\mathcal{X}_1 \mapsto A, \mathcal{X}_2 \mapsto B, \mathcal{X}_3 \mapsto s, \mathcal{X}_4 \mapsto sk(s), \mathcal{X}_5 \mapsto shk(A, s)]$ 
 $\text{lock}.\mathcal{X}_K := \text{keys}[\mathcal{X}_2].$ 
 $\text{if } \mathcal{X}_K \doteq \text{blank} \text{ then}$ 
   $\nu\mathcal{X}_N.\text{unlock}.\text{send}((\text{key}, \mathcal{X}_2, \mathcal{X}_N)).$ 
   $\text{receive}(\mathcal{X}_{kc}).\text{lock}.$ 
   $\text{if } (\text{vsign}(\mathcal{X}_{kc}, \mathcal{X}_4) = \top) \text{ then}$ 
     $\text{let } (\mathcal{X}_t, \mathcal{X}_B, \mathcal{X}_{K'}, \mathcal{X}_{N'}) = \text{open}(\mathcal{X}_{kc}) \text{ in}$ 
     $\text{if } (\mathcal{X}_t \doteq \text{key} \wedge \mathcal{X}_B = \mathcal{X}_2 \wedge \mathcal{X}_{N'} = \mathcal{X}_N) \text{ then}$ 
       $\text{keys}[\mathcal{X}_2] := \mathcal{X}_{K'}.\text{unlock}.\mathbf{0}$ 
     $\text{else unlock}.\mathbf{0}$ 
   $\text{else unlock}.\mathbf{0}$ 
 $\text{else ... translation of the other part, using } \mathcal{X}_K \text{ as } ek(B)$ 

```

Here, we use some syntactic sugar: several checks can be done by one condition; and that we write *let* for parsing the content of a message, in this case expecting that it can be parsed into a quadruple. The function *vsign* is supposed to be signature verification with the property $\text{vsign}(\text{sign}(m, \text{inv}(k)), k) =_E \top$ and *open* a destructor that yields the signed message, i.e.,

$$\begin{aligned}
& (0 \uplus L, F, \mu) \rightarrow (L, F, \mu) \\
& (\text{send}(r).\mathcal{L} \uplus L, F, \mu) \rightarrow (\mathcal{L} \uplus L, F[\mathcal{X} \mapsto r], \mu) \\
& \quad \text{where } \mathcal{X} \text{ is a fresh label with regard to } F \\
& (\text{receive}(\mathcal{X}).\mathcal{L} \uplus L, F, \mu) \rightarrow (\mathcal{L}[\mathcal{X} \mapsto F(r)] \uplus L, F, \mu) \\
& \quad \text{where } r \text{ is a recipe over the domain of } F \\
& (\text{lock}.\mathcal{A} \uplus L, F, \mu) \xrightarrow{tr} (\mathcal{L} \uplus L, F, \mu') \\
& \quad \text{if } (\mathcal{A}, \mu) \xRightarrow{*} (\text{unlock}.\mathcal{L}, \mu') \\
& (L, F, \mu) \rightarrow (\sigma(F_A)(\mathcal{L}_A) \uplus L, F, \mu) \\
& \quad \text{if } \sigma \text{ is an honest instantiation of a role } F_A : \mathcal{L}_A \\
& (L, F, \mu) \rightarrow (L, F \uplus \sigma(F_A), \mu) \\
& \quad \text{if } \sigma \text{ is a dishonest instantiation of a role } F_A : \mathcal{L}_A \\
& \quad \text{and where the labels of } F_A \text{ have been freshly renamed} \\
& (\nu \mathcal{X}. \mathcal{A}, \mu) \Rightarrow (\mathcal{A}[\mathcal{X} \mapsto n], \mu) \text{ where } n \text{ is a fresh constant} \\
& (\mathcal{A}_1 + \mathcal{A}_2, \mu) \Rightarrow (\mathcal{A}_i, \mu) \text{ if } i \in \{1, 2\} \\
& (\text{if } r_1 \doteq r_2 \text{ then } \mathcal{A}_1 \text{ else } \mathcal{A}_2, \mu) \Rightarrow (\mathcal{A}_1, \mu) \text{ if } r_1 =_E r_2 \\
& (\text{if } r_1 \doteq r_2 \text{ then } \mathcal{A}_1 \text{ else } \mathcal{A}_2, \mu) \Rightarrow (\mathcal{A}_2, \mu) \text{ if } r_1 \neq_E r_2 \\
& (\mathcal{X} := c[r_1].\mathcal{A}, \mu) \Rightarrow (\mathcal{A}[\mathcal{X} \mapsto \mu(c, r_1)], \mu) \\
& (c[r_1] := r_2.\mathcal{A}, \mu) \Rightarrow (\mathcal{A}, \mu[(c, r_1) \mapsto r_2]) \\
& (\text{event}(r).\mathcal{A}, \mu) \xRightarrow{\tau} (\mathcal{A}, \mu)
\end{aligned}$$

Fig. 5. Semantics of Local Behaviors.

$\text{open}(\text{sign}(m, \text{inv}(k))) = m$. This models a signature scheme where the signed text is transmitted in plain along with a signed hash, i.e., one needs the public key only for signature verification. One can also observe that between each lock and unlock there is at most one memory read or write, so locking is in this case actually redundant as it does not prevent any race conditions.

d) Semantics of Local Behaviors: We give a simple operational semantics for a set of local behaviors $F_{A_1} : \mathcal{L}_{A_1}, \dots, F_{A_n} : \mathcal{L}_{A_n}$ (where each F_{A_i} is the initial knowledge frame of A_i). We assume a set $Ag \subseteq \Sigma \cap \Sigma_p$ of public constant of type role and that $i \in Ag$ represents a dishonest agent (“intruder”) while all other agents are honest.²

For a behavior $F_A : \mathcal{L}_A$ we call the substitution σ an *instantiation* if it maps all variables in F_A (that are by definition of type role) to elements of Ag . We say it is a *dishonest instantiation* if $\sigma(A) = i$ and an *honest instantiation* otherwise. We write $\sigma(F_A)$ for the instantiation of the initial knowledge frame and $\sigma(F_A)(\mathcal{L}_A)$ for the instantiation of the behavior itself, replacing all labels from F_A in \mathcal{L}_A by ground terms; thus all remaining labels first occur at a new, at a receive, or at a memory read.

The last ingredient for the semantics is a memory map μ that maps every memory cell $c[(t)_E]$ to a value, initially *blank*, where $(t)_E$ is the $=_E$ -equivalence class of a ground term t over Σ . As easy notation we just write $\mu(c, t)$ for this value, and we write $\mu[(c, t) \mapsto t']$ for changing the memory cell $c[(t)_E]$ to value t' .

²One may well consider more than one dishonest agent, but for simplicity we work with just one.

The semantics of local behavior is defined in Fig. 5 and consists of two transition relations \rightarrow and \Rightarrow that call each other: \rightarrow is on triples (L, F, μ) where L is a multi-set of local behaviors, F is a frame representing the intruder knowledge and μ is the memory map; the initial state is $(\emptyset, [], \mu_0)$ where μ_0 maps everything to *blank*; \Rightarrow is on tuples (\mathcal{A}, μ) where \mathcal{A} is an atomic section of a local behavior. We decorate the transition relations with a list of events that occurred upon the transitions. In this semantics, the intruder *is* the network: every message an honest agent sends gets added into the intruder knowledge, and every message an honest agent receives comes from the intruder knowledge: the intruder can choose any recipe over their knowledge, which includes encrypting and decrypting with known keys. An atomic section is handled literally atomically using the \Rightarrow relation locally at an agent until it hits the unlock; we label the transition with the trace tr of all events that the agent emitted. The last two rules regarding \rightarrow allow spawning new instances of any role A : we choose any instantiation σ of the variables in F_A with agent names; if σ is honest, i.e., $\sigma(A) \neq i$, we apply the instantiated knowledge $\sigma(F_A)$ as a substitution to the local behavior \mathcal{L}_A , leaving only labels that are introduced by new, receive, and memory read. This semantics allows running an arbitrary number of sessions in parallel and sequentially. If σ is dishonest, i.e., $\sigma(A) = i$ then this represents that the intruder plays role A under the actual name i . This models a dishonest/compromised agent. We give the intruder the initial knowledge needed to play the role, i.e., $\sigma(F_A)$ where we have to rename the labels in the frame to avoid a clash with labels in the present intruder knowledge.

B. Projection: The Semantics of CryptoChoreo

We can now give the semantics of CryptoChoreo by translation to local behaviors. We again use frames to represent the knowledge of a role at a given point in the translation, but this differs from their use in the local behavior semantics. As said, the messages in a choreography represent how messages look like in an “ideal” or unattacked run—which may differ from the shape of messages in a real run due to interference from the attacker. For instance if A is supposed to receive $\text{exp}(g, Y)$ from B for a secret Y that A does not know, there is nothing that A can check about this message. To keep track of this during translation we make an entry $[\mathcal{X}_i \mapsto \text{exp}(g, Y)]$ in the frame F of A that expresses: A has received *some* message \mathcal{X}_i and according to the choreography it is *supposed to be* $\text{exp}(g, Y)$. Given that another entry $[\mathcal{X}_j \mapsto X]$ represents the fresh value A has created for her own secret X , then the Diffie-Hellman key $\text{exp}(\text{exp}(g, X), Y)$ can be formed with the recipe $r = \text{exp}(\mathcal{X}_i, \mathcal{X}_j)$: $F(r) = \text{exp}(\text{exp}(g, Y), X) =_E \text{exp}(\text{exp}(g, X), Y)$. In this way, frames make the connection between the messages the agents have and use in their local behavior and what the messages supposedly are on the choreography level. As part of the translation semantics, we thus need to give a general definition of how an agent should generate an outgoing message from a given frame and this is the first core algorithmic problem we define on frames:

- The deduction problem: Given frame F and term t , compute a recipe r such that $F(r) =_E t$ if one exists or return fail otherwise.

Here, the failure means that there is no feasible way for the agent to create the required outgoing message. Our semantics will in this case refuse the choreography as unexecutable. This will typically happen in case of a specification error, for instance if an agent does not possess the necessary keys to participate in the protocol as suggested.

For receiving messages there is a related problem to solve. Suppose the choreography contains a step where an agent should receive a hash $h(N)$ of a random number N . Since h is one-way and N random, the agent cannot perform any checks on this message, so will have a frame entry $[\mathcal{X}_i \mapsto h(N)]$. Suppose in a later step, the agent is supposed to receive N , i.e., $[\mathcal{X}_j \mapsto N]$. Now there is something the agent can check: whether $F(\mathcal{X}_i) =_E F(h(\mathcal{X}_j))$ holds. In general, a check is a pair of recipes (here \mathcal{X}_i and $h(\mathcal{X}_j)$) that are supposed to give the same message under the present frame F . The agent will not proceed with the execution if this check fails, because at least one of the incoming messages does not comply with the protocol. In a similar way, if an agent receives a message that is supposed to be an encryption with a key for which the agent has the decryption key, then there must be a check that this decryption indeed works. It is one of the main features of CryptoChoreo that the semantics defines what checks each party must do on incoming messages, while the modeler only specifies in the choreography what the messages look like in an unattached protocol run.

However, in general there are infinitely many possible checks, for instance in the discussed situation one could also check $F(h(\mathcal{X}_i)) =_E F(h(h(\mathcal{X}_j)))$, but this check seems redundant given the other check above. To make this notion of redundant checks precise, we adapt some equational logic concepts. Define an *interpretation* \mathcal{I} as a mapping from all labels to ground terms. (A frame is similar to an interpretation but it has only a finite domain.) For a recipe r , let $\mathcal{I}(r)$ be the ground term that results from replacing every label \mathcal{X} in r by $\mathcal{I}(\mathcal{X})$. We now say that \mathcal{I} is an *E-model* of the formula $r_1 \doteq r_2$, and write $\mathcal{I} \models_E r_1 \doteq r_2$, if $\mathcal{I}(r_1) =_E \mathcal{I}(r_2)$, and we extend this to conjunctions of equations as expected. We finally say that ϕ *E-implies* ψ , and write $\phi \models_E \psi$, if every *E-model* of ϕ is also an *E-model* of ψ . For instance $\mathcal{X}_i \doteq h(\mathcal{X}_j) \models_E h(\mathcal{X}_i) \doteq h(h(\mathcal{X}_j))$.

We can now define the second core algorithmic problem: to find a finite set of checks ϕ that is complete in the sense that any other checks one could make are already implied by ϕ already:

- The complete check problem: Given a frame F , compute a finite set of checks, i.e., equations of recipes $\phi = \{r_1 \doteq r'_1, \dots, r_n \doteq r'_n\}$ such that $F(r_i) =_E F(r'_i)$ for each $1 \leq i \leq n$, that is *complete* in the sense that if for any other r_0, r'_0 we have $F(r_0) =_E F(r'_0)$, then $\phi \models_E r_0 \doteq r'_0$, or return fail if no finite set of checks satisfies that.

A failure case, i.e., E and F where every complete set

of checks is infinite, is surprisingly hard to construct, see Appendix C, so we do not consider this a problem in practice. In section IV we give an example of an algebraic theory E for which both the deduction problem and the finite complete set of checks problem are decidable, and we sketch how these algorithms work. The details and the proofs are given in Appendix B. Note that for this E , every frame has a finite complete set of checks; while the construction of recipes can of course still fail when there is insufficient knowledge in a frame to compose the outgoing message.

For arbitrary E , both these problems are in general not recursively computable (because in general even $=_E$ is undecidable). Our projection semantics is parameterized by an arbitrary equational theory, while in Section IV we provide procedures for one representative equational theory.

We first note a complete set of checks for F is not unique, however if ϕ and ψ are two complete set of checks for F , then $\phi \models_E \psi$ and $\psi \models_E \phi$, so they are equivalent and in the semantics we can leave this choice undetermined³ without making the semantics ambiguous. By abuse of notation we thus write $\phi(F)$ for a complete set of checks for F , even though it is, strictly speaking, not a function.

Second, also the deduction problem has in general many solutions, i.e., different r_1 and r_2 such that $F(r_1) =_E F(r_2) =_E t$. For instance, if $F = [\mathcal{X}_i \mapsto h(N), \mathcal{X}_j \mapsto N]$, and the agent should send $t = h(N)$, then there are two recipes to construct t : $r_1 = \mathcal{X}_i$ and $r_2 = h(\mathcal{X}_j)$. This choice of recipes could make a difference if the agent received for \mathcal{X}_i or \mathcal{X}_j terms that do not comply with the protocol. However, after the agent has successfully executed the checks $\phi(F) = \{\mathcal{X}_i \doteq h(\mathcal{X}_j)\}$, both recipes are guaranteed to produce the same term. More generally, if $F(r_1) =_E F(r_2) =_E t$, then $\phi(F) \models_E r_1 \doteq r_2$, i.e., if we have performed all the checks in $\phi(F)$, then also the choice between the two recipes r_1 and r_2 does not matter.

Third, in the semantics, we need a slight generalization of the deduction problems, namely given several frames F_1, \dots, F_n and goal terms t_1, \dots, t_n and we want a *single* recipe r that solves all *deduction* problems, i.e., $F_i(r) = t_i$ for every $1 \leq i \leq n$. Suppose we already have a set ϕ of checks that is a complete set of checks for each of the F_i , and suppose there is a solution r for all frames. Then any solution r' for one of the frames, say F_1 , must be equivalent to r , i.e., $\phi \models_E r \doteq r'$. Thus for checked frames it suffices to compute a solution for one frame and check if it works on the other frames—if not, then there is no common solution for all frames.

Nondeterminism and conditions mean that, in general, a role does not know which branch of the choreography we currently are in, and this also holds during the translation. Therefore, during the translation, the translation state contains a finite set of pairs $(F_i : \mathcal{C}_i)$ where each F_i is a frame (all F_i have the same domain) and \mathcal{C}_i is the remainder of the choreography that still needs to be translated. We note that this is to handle nondeterminism *external* to the role being translated; when

³Thus a concrete implementation is free to choose one.

an agent branches locally we handle the translation of each branch in isolation.

Definition 1. A translation state is of the form

$$(A, \phi \triangleright \psi, b, \{(F_1 : C_1), \dots, (F_n : C_n)\})$$

where A is the role we are currently translating; ϕ and ψ are sets of equations $r_1 \doteq r_2$ between recipes, where ϕ represents checks that have already been done, and ψ are checks that are pending; $b \in \{c, a\}$ is a flag indicating whether we are on the choreography level or in an atomic section; the F_i are frames with the same domain that map to terms; and the C_i are either choreographies if $b = c$ or atomic sections if $b = a$.

During translation we preserve the invariant that $\phi \cup \psi$ is covering all checks that can be made in any frame, i.e., for every frame F_i ($i \in \{1, \dots, n\}$) and any pair of recipes r_0, r'_0 with $F_i(r_0) = F_i(r'_0)$, it must hold that $\phi \cup \psi \models_E r_0 \doteq r'_0$. If $\psi \neq \emptyset$, i.e., if there are pending checks, they will be performed first before all other translation steps.

Given a choreography \mathcal{C} and a role A of that choreography, let $F_A = [\mathcal{X}_1 \mapsto t_1, \dots, \mathcal{X}_n \mapsto t_n]$ where $\{t_1, \dots, t_n\}$ is the initial knowledge of role A and the \mathcal{X}_i are distinct labels. The initial translation state for translating A in \mathcal{C} is:

$$(A, \emptyset \triangleright \emptyset, c, \{(F_A : \mathcal{C})\})$$

i.e., there is just one possibility \mathcal{C} where we are and the current knowledge is F_A .

1) *Cases of the Semantics Function:* The semantics function $\llbracket T \rrbracket$ takes a translation state T and projects the choreography to the actions of the role, yielding a local behavior for that role. We define it recursively by a case distinction on T . Since we will often require all possibilities (F_i, C_i) to start with the same kind of command, we use the following notation: $\{F_i, \nu N_i.C_i\}_{i=1}^n$ for $\{F_1, \nu N_1.C_1, \dots, F_n, \nu N_n.C_n\}$, and similar for other constructs in place of νN . For simplicity, we first present this semantics without goals.

The semantic function translating a choreography to the local behavior of a given agent, A , is defined by 14 cases, labelled a)-n). Case a) completes the projection when all the possible continuations are finished. Cases b)-d) handle communication steps: If A is supposed to send a message we must deduce a recipe to do so from the terms available in their frame. If A receives a message we must compute all the checks that may be performed on the received term and add them to the set of pending checks. Case e) handles the processing of pending checks. Case f) handles the situation where the checks have ruled out any possible continuation. Case g) handles the entering of an atomic section by the given agent, and case h) handles atomic sections of other agents. Cases i)-m) handle the different actions that may be performed in an atomic section. Each of these cases corresponds to a primitive action which is directly translated into local behavior. Finally, case n) states that the choreography is ill-specified if none of the previous cases apply.

a) $\llbracket A, \phi \triangleright \emptyset, c, \{(F_i : \mathbf{0})\}_{i=0}^n \rrbracket$ where $n > 0$: All possibilities have finished, and the translation is simply:

$\mathbf{0}$.
b) $\llbracket A, \phi \triangleright \emptyset, c, \{(F : B \rightarrow C : t.C)\} \cup \{F_i : C_i\}_{i=1}^n \rrbracket$ for $A \neq B$ and $A \neq C$: One of the possibilities is a communication step that A is not involved in and is therefore ignored. The translation is thus:

$$\llbracket A, \phi \triangleright \emptyset, c, \{(F : C)\} \cup \{F_i : C_i\}_{i=1}^n \rrbracket$$

c) $\llbracket A, \phi \triangleright \emptyset, c, \{F_i : A \rightarrow B_i : t_i.C_i\}_{i=1}^n \rrbracket$ where $n > 0$: All possibilities are send steps for A .

As explained before, we check whether there is a recipe r such that $F_i(r) = t_i$ for each $1 \leq i \leq n$. If there is no such r , then we reject the protocol as unexecutable: either there is no way for A to produce the outgoing term t_i , or the different possibilities would require different recipes, and A cannot know in which possibility they are. However, if there is such an r ,⁴ then the translation is:

$$\text{send}(r). \llbracket A, \phi \triangleright \emptyset, c, \{F_i : C_i\}_{i=1}^n \rrbracket$$

d) $\llbracket A, \phi \triangleright \emptyset, c, \{F_i : B_i \rightarrow A : t_i.C_i\}_{i=1}^n \rrbracket$ where $n > 0$: All possibilities are receive steps for A .

Let \mathcal{X} be a new recipe variable and $F'_i = F_i[\mathcal{X} \mapsto t_i]$ for every $i \in \{1, \dots, n\}$. Let ϕ_i be a complete finite set of checks for F'_i and let $\Phi = \cup_{i=1}^n \phi_i$. This represents all checks that we can do in any of the frames F_i . First we can remove from Φ all those checks that are already implied by the checks ϕ from the translation state (i.e., that have already been done before in the translation process). We can also remove from Φ any equation that is implied by the other equations. Let thus Φ' be a resulting minimal set of equations.⁵ The translation of the receive step is then obtained by adding received message to the frames and inserting the Φ' as pending checks that have to be done next:

$$\text{receive}(\mathcal{X}). \llbracket A, \phi \triangleright \Phi', c, \{F'_i : C_i\}_{i=1}^n \rrbracket$$

e) $\llbracket A, \phi \triangleright \{r_1 \doteq r_2\} \cup \psi, b, \{F_i : C_i\}_{i=1}^n \rrbracket$ where $n > 0$: There is at least one pending check $r_1 \doteq r_2$.

We partition the possibilities into those where F_i satisfies the check and those that do not:

$$\text{Let } FCs_+ = \{(F_i : C_i) \mid F_i(r_1) =_E F_i(r_2)\}$$

$$\text{and } FCs_- = \{(F_i : C_i) \mid F_i(r_1) \neq_E F_i(r_2)\}.$$

The translation is now:

$$\text{if } r_1 \doteq r_2 \text{ then } \llbracket A, \phi \cup \{r_1 \doteq r_2\} \triangleright \psi, b, FCs_+ \rrbracket$$

$$\text{else } \llbracket A, \phi \triangleright \psi, b, FCs_- \rrbracket$$

f) $\llbracket _, _ \triangleright _, _, \emptyset \rrbracket$: There are no possible continuations.

This can happen when doing a check that splits the possibilities into FCs_+ and FCs_- , and one of them is empty. It means that if we reach this branch, the agent has detected that an incoming message is not compliant with the choreography, and aborts the execution. The translation is thus simply:

$\mathbf{0}$

g) $\llbracket A, \phi \triangleright \emptyset, c, \{F_i : A : \mathcal{A}_i\}_{i=1}^n \rrbracket$ where $n > 0$: All possibilities start with an atomic section of the agent A .

⁴If there are several such recipes, the choice between them leads to equivalent translation outcomes as explained before.

⁵Again, there may be several minimal sets, e.g., if two equations imply each other; however all resulting sets from the minimization are logically equivalent.

Then the translation is simply to issue the lock and switch the atomic section flag:

lock. $\llbracket A, \phi \triangleright \emptyset, a, \{F_i : \mathcal{A}_i\}_{i=1}^n \rrbracket$

h) $\llbracket A, \phi \triangleright \emptyset, c, \{(F : B : \mathcal{A})\} \cup \{F_i : \mathcal{C}_i\}_{i=1}^n \rrbracket$ where $B \neq A$: One possibility is that another role goes into its atomic section.

Role A should ignore these steps and just extract all continuations after the atomic section, which is defined as follows:

$$\begin{aligned} cont(\mathcal{A}_1 + \mathcal{A}_2) &= cont(\mathcal{A}_1) \cup cont(\mathcal{A}_2) \\ cont(\text{if } s \doteq t \text{ then } \mathcal{A}_1 \text{ else } \mathcal{A}_2) &= cont(\mathcal{A}_1) \cup cont(\mathcal{A}_2) \\ cont(_.\mathcal{A}) &= cont(\mathcal{A}) \\ cont(\mathcal{C}) &= \{\mathcal{C}\} \end{aligned}$$

Let $\mathcal{C}_1, \dots, \mathcal{C}_m = cont(\mathcal{A})$ in the translation:

$\llbracket A, \phi \triangleright \emptyset, c, \{F : \mathcal{C}_i\}_{i=1}^m \cup \{F_i : \mathcal{C}_i\}_{i=1}^n \rrbracket$

We now come to the cases for an atomic section of the agent we translate for:

i) $\llbracket A, \phi \triangleright \emptyset, a, \{F_i : \nu N_i. \mathcal{A}_i\}_{i=1}^n \rrbracket$ for $n > 0$: All possibilities create a fresh number N_i .

We pick a fresh label \mathcal{X} and translate:

$\nu \mathcal{X}. \llbracket A, \phi \triangleright \emptyset, a, \{F_i[\mathcal{X} \mapsto N_i] : \mathcal{A}_i\}_{i=1}^n \rrbracket$

j) $\llbracket A, \phi \triangleright \emptyset, a, \{F_i : s_i := c[t_i]. \mathcal{A}_i\}_{i=1}^n \rrbracket$ where $n > 0$:

All cases start with a memory read.

Similar to the send case, we require that there is one recipe r such that $F_i(r) =_E t_i$ for each $1 \leq i \leq n$. If not, the semantics rejects the protocol as unexecutable at this point (because the agent either cannot create the proper index for the memory lookup, or there are contradicting possibilities for this index). The retrieved message s_i is treated like in the receive case: we add it to the knowledge with a new label \mathcal{X} , giving frames $F'_i = F_i[\mathcal{X} \mapsto s_i]$, and then we compute a complete set of checks ϕ_i for each F'_i , compute the union $\Phi = \cup_{i=1}^n \phi_i$, and remove redundant equations leading to a reduced Φ' . The translation is then:

$\mathcal{X} := c[r]. \llbracket A, \phi \triangleright \Phi', a, \{F'_i : \mathcal{A}_i\}_{i=1}^n \rrbracket$

k) $\llbracket A, \phi \triangleright \emptyset, a, \{F_i : c[t_i] := s_i. \mathcal{A}_i\}_{i=1}^n \rrbracket$ where $n > 0$:

All possibilities start with a write step. We require that there are recipes r_1 and r_2 such that $F_i(r_1) = t_i$ and $F_i(r_2) = s_i$ for each $i \in \{1, \dots, n\}$. (If not, this is a specification error, because it is unclear what A 's next step is.) Then the translation is:

$c[r_1] := r_2. \llbracket A, \phi \triangleright \emptyset, a, \{F_i : \mathcal{A}_i\}_{i=1}^n \rrbracket$

l) $\llbracket A, \phi \triangleright \emptyset, a, \{F_i : \text{if } s_i \doteq t_i \text{ then } \mathcal{A}_i^+ \text{ else } \mathcal{A}_i^-\}_{i=1}^n \rrbracket$

where $n > 0$: All the \mathcal{A}_i start with a condition. We require that there are recipes r_1 and r_2 such that $F_i(r_1) = t_i$ and $F_i(r_2) = s_i$ for each $i \in \{1, \dots, n\}$. (If not, this is a specification error, because it is unclear what A 's next step is.) We define

$\mathcal{T}^+ = \llbracket A, \phi \triangleright \emptyset, a, \{F_i : \mathcal{A}_i^+\}_{i=1}^n \rrbracket$ and

$\mathcal{T}^- = \llbracket A, \phi \triangleright \emptyset, a, \{F_i : \mathcal{A}_i^-\}_{i=1}^n \rrbracket$.

The translation is:

if $r_1 \doteq r_2$ then \mathcal{T}^+ else \mathcal{T}^-

m) $\llbracket A, \phi \triangleright \emptyset, a, \{F_i : \mathcal{C}_i\}_{i=1}^n \rrbracket$ where $n > 0$: Finally, if all the \mathcal{A}_i conclude the atomic section, then the translation is:

unlock. $\llbracket A, \phi \triangleright \emptyset, c, \{F_i : \mathcal{C}_i\}_{i=1}^n \rrbracket$

n) $\llbracket _ \rrbracket$ for any other translation state: this is an error, because it is unclear what A should do next.

2) *Example*: Let us continue the example choreography from Fig. 3 and let us look at the translation for the role s. The knowledge of s gives us the frame:

$F_0 = [\mathcal{X}_1 \mapsto A, \mathcal{X}_2 \mapsto B, \mathcal{X}_3 \mapsto s, \mathcal{X}_4 \mapsto sk(s),$

$\mathcal{X}_5 \mapsto inv(sk(s)), \mathcal{X}_6 \mapsto shk(A, s), \mathcal{X}_7 \mapsto ek(B)]$

and we compute $\llbracket s, \emptyset \triangleright \emptyset, c, \{F_0 : \mathcal{C}\} \rrbracket$ where \mathcal{C} is the entire choreography. The choreography begins with atomic actions of A, checking if the key of B is known, asking s if not and starting the main protocol otherwise. Thus using rule III-B1h we get a split into two possibilities $\{(F_0 : A \rightarrow s : (key, B, N) \dots), (F_0 : \mathcal{C}_0)\}$ where \mathcal{C}_0 is the initial choreography example of Fig. 1 (with $ek(B)$ replaced by the variable EKB that represents the key that A has looked up from memory). Now \mathcal{C}_0 starts with another atomic section of A (the choice to either send a message or a key update). So we apply again rule III-B1h to split that possibility into two:

$(F_0 : A \rightarrow s : (key, B, N) \dots),$

$(F_0 : A \rightarrow s : \text{scrypt}((B, \text{crypt}(\dots, EKB)), shk(A, s)) \dots)$

$(F_0 : A \rightarrow s : \text{scrypt}((B, \text{crypt}(\dots, EKB)), shk(A, s)) \dots)$

Now finally all messages are something the server can receive, so with rule III-B1d we get updated frames F_1, F_2, F_3 augmenting F_0 with a new label \mathcal{X}_8 , which is bound to the respective incoming message.

The checks that we can do for F_1 are that \mathcal{X}_8 is a triple, that the first item is constant key and the second item is \mathcal{X}_2 .⁶ For ease of notation, assume we have functions $vtuple_n$ and π_n (for all $n \in \mathbb{N}$) with the property $vtuple_n(t_1, \dots, t_n) =_E \top$ and $\pi_i(t_1, \dots, t_n) =_E t_i$. Thus the checks for F_1 are $\phi_1 = \{vtuple_3(\mathcal{X}_8) \doteq \top, \pi_1(\mathcal{X}_8) \doteq key, \pi_2(\mathcal{X}_8) \doteq \mathcal{X}_2\}$.

In F_2 and F_3 we can check that symmetric decryption of \mathcal{X}_8 with key \mathcal{X}_6 succeeds and yields a pair. The blue parts in F_2 and F_3 cannot be decrypted, so s cannot further check anything about this. In our example theory E below we have operators scrypt , vscrypt and dscrypt with the properties $\text{dscrypt}(\text{scrypt}(m, k), k) =_E m$ and $\text{vscrypt}(\text{scrypt}(m, k), k) =_E \top$. Together they model AEAD symmetric schemes, i.e., the attacker cannot modify the encrypted message m by modifying the ciphertext $\text{scrypt}(m, k)$, as this would lead to errors when decrypting; thus the decryption is an operation that fails when applied to an incorrect message or the wrong key, and we model that in the algebra by two functions, one telling us whether decryption works with the given key and one that in the positive case gives the result. We thus have in F_2 and F_3 the complete set⁷ of the checks: $\phi_2 = \phi_3 = \{\text{vscrypt}(\mathcal{X}_8, \mathcal{X}_6) \doteq \top, vtuple_2(\text{dscrypt}(\mathcal{X}_8, \mathcal{X}_6)) \doteq \top, \pi_1(\text{dscrypt}(\mathcal{X}_8, \mathcal{X}_6)) \doteq \mathcal{X}_2\}$.

⁶More realistically, s should not expect a particular name B but rather have it determined through A's request. This is why we like to model $ek(\cdot)$ as a public function (one can look up the public key of any role), but here we deliberately made this function private, so that A has to ask the server for the role. Anyway the semantics ensures that for any instantiation of the role variables with agent names, we have any number of server instances, so this comes without loss of attacks.

⁷In Section IV we give a complete set of checks procedure for an example theory E . This procedure can be used to verify that this $\phi_1 \cup \phi_2$ —plus some checks on the initial frame F_0 (like: \mathcal{X}_4 is the public key to \mathcal{X}_5) that we omitted here—is a complete set of checks.

We thus have the translation
 $\text{receive}(\mathcal{X}_8). \llbracket s, \emptyset \triangleright \phi_1 \cup \phi_2, c, \{(F_1:\dots), (F_2:\dots), (F_3:\dots)\} \rrbracket$

Applying rule III-B1e several times to process all checks, we get the possibilities partitioned (because ϕ_1 only holds in F_1 and ϕ_2 only holds in F_2 and F_3) as follows:

if ϕ_1 then $\llbracket s, \phi_1 \triangleright \emptyset, c, \{(F_1:\dots)\} \rrbracket$ else
if ϕ_2 then $\llbracket s, \phi_1 \triangleright \emptyset, c, \{(F_2:\dots), (F_3:\dots)\} \rrbracket$ else $\mathbf{0}$

Let us just look at the most interesting branch, namely the positive case under ϕ_2 . Here the next step in F_2 is $s \rightarrow B : \text{sign}((A, \text{crypt}((\text{msg}, M), \text{EKB})), \text{inv}(\text{sk}(s)))$ and in F_3 the corresponding step but with content (upd, K) . So we need to apply rule III-B1c which requires a recipe r that works in both cases. Let $r_b = \pi_2(\text{dscrypt}(\mathcal{X}_8, \mathcal{X}_6))$ which gives the “blue message part” that s cannot decrypt in either F_2 or F_3 . Now $r = \text{sign}((\mathcal{X}_1, r_b), \mathcal{X}_5)$ is the recipe that works in both cases. Thus the complete translation for the server role is:

$\text{receive}(\mathcal{X}_8). \text{if } \phi_1 \text{ then send}(\text{sign}((\text{key}, \mathcal{X}_2, \mathcal{X}_7)), \mathcal{X}_5) \text{ else}$
 $\text{if } \phi_2 \text{ then send}(\text{sign}((\mathcal{X}_1, \pi_2(\text{dscrypt}(\mathcal{X}_8, \mathcal{X}_6))), \mathcal{X}_5))$

C. Attack Semantics

In this section, we formalize our notion of security. In a sentence, we consider there to be an attack if the system can possibly develop in a way that falsifies a given query over traces.

A *security query* is built from the following grammar:

$$p ::= \text{event}(t) \mid \text{event}(s) \sqsubseteq \text{event}(t) \mid \text{intruder}(t) \mid s \doteq t \mid \perp \mid p_1 \implies p_2 \mid \forall X. p'$$

The other logical operators can be added as syntactic sugar (in particular we will use \wedge). We only consider a query well-formed if it contains no free variables.

We characterize a trace, $(L_0, F_0, \mu) \xrightarrow{t_1, \dots, t_n}^* (L_m, F_m, \mu)$, by the list of emitted events and the final knowledge of the intruder: $([t_1, \dots, t_n], F_m)$.

Following is the semantics for evaluating a query:

$$\begin{aligned} ([t_1, \dots, t_n], F) &\models \text{event}(s) \text{ iff } t_i =_E s \text{ for some } i \\ ([t_1, \dots, t_n], F) &\models \text{intruder}(s) \text{ iff } F(r) =_E s \text{ for some } r \\ ([t_1, \dots, t_n], F) &\models s \doteq t \text{ iff } s =_E t \\ ([t_1, \dots, t_n], F) &\models \perp \text{ never} \\ ([t_1, \dots, t_n], F) &\models p_1 \implies p_2 \\ &\text{iff } ([t_1, \dots, t_n], F) \not\models p_1 \\ &\text{or } ([t_1, \dots, t_n], F) \models p_2 \\ ([t_1, \dots, t_n], F) &\models \forall X. p \\ &\text{iff for all } s \text{ we have } ([t_1, \dots, t_n], F) \models p[s/X] \\ ([t_1, \dots, t_n], F) &\models \text{event}(s) \sqsubseteq \text{event}(t) \\ &\text{iff for each } t_i =_E s, \text{ there is a distinct } t_j =_E t \end{aligned}$$

We consider a configuration as secure with regard to a given query if the query is valid on all traces from the configuration.

For the rest of this section, we will show how to encode the security goals in a choreography as queries.

First, we modify the grammar of choreographies to allow the emission of events:

$$\begin{aligned} \mathcal{C} &::= \dots \\ \mathcal{A} &::= \dots \\ &\mid \text{event}(t). \mathcal{A} \end{aligned}$$

These events will be handled by the projection in the same way as sends: We check if we can find a recipe that produces t in all frames, and return a translation error if not.

A secrecy goal expresses that a certain term should be kept a secret between a given set of agents. That is, there is an attack if the intruder is not a member of the set and can produce the secret term. The following transformation checks for this by having each member of the set emit a secrecy event at the end of each protocol run, including the names of the other agents and the term in the event.

To each secrecy goal $g = \text{“}t \text{ secret between } A_1, \dots, A_n \text{”}$, we assign a unique *event name* $e_g \in \Sigma \setminus \Sigma_p$. We then replace each g with

$$A_1 : \text{event}(e_g(t, A_1, \dots, A_n)).$$

$$\vdots$$

$$A_n : \text{event}(e_g(t, A_1, \dots, A_n))$$

The goal holds iff the following query holds:

$$\forall X, A_1, \dots, A_n. \neg \left(A_1 \neq i \wedge \dots \wedge A_n \neq i \wedge \left(\text{event}(e_g(X, A_1, \dots, A_n)) \wedge \text{intruder}(X) \right) \right)$$

Inspired by the hierarchy of authentication specifications of [12], we permit the user of CryptoChoreo to specify injective and noninjective authentication goals. Here, an attack on the noninjective authentication of agent a to another agent b on term t would be if a finishes the protocol believing that b has played the same protocol, but b has either not played the protocol with a or has done so with another term than t . Injective authentication would also have an attack if a can accept the value t more often than b commits to it. We check authentication by inserting start and end events including the names of both parties and the value they authenticate on. The authenticated party commits as early as possible to the value by emitting a start event and the authenticating party emits a corresponding end event at the very end of the protocol. For noninjective authentication, we then verify that the presence of an end event in a trace implies the presence of the corresponding start event. For injective authentication, we do the same, except that there must be a distinct start event for each end event.

Our authentication goals give the guarantee to the authenticating party that the authenticated party has intended to use the term for the given protocol, but, unlike in [12], does not guarantee that they have reached the end of the protocol. One could, alternatively, place the start event right before the last message to the authenticating party, though this would give spurious attacks if the protocol ends with messages that are not really a part of the authentication mechanism (for example if an agent ends by sending an end-signal in plaintext). We

have found our version to be a pragmatic choice, that still usually has an attack if a meaningful attack exists in the stricter version. For some protocols, however, it might also be desirable to verify that the authenticated party has actually reached some given point in the protocol, which can be checked in CryptoChoreo by placing the events manually.

To a noninjective authentication goal, $g = \text{“B ni-authenticates A on } t\text{”}$, we associate unique start- and end-event names $e_{gs}, e_{ge} \in \Sigma \setminus \Sigma_p$. We then replace the goal with the end event $B : \text{event}(e_{ge}(t, A, B))$. We want to check that the occurrence of this event implies the occurrence of a corresponding start event from A. However, it is not always possible to insert that start event right at the beginning of the choreography; t might contain values that have been generated during the protocol run.

To solve this, we make the following modification to the projection semantics for A:

If we are computing $\llbracket A, \phi \triangleright \emptyset, c, \{F_i : C_i\}_{i=1}^n \rrbracket$, $B : \text{event}(e_{ge}(t, A, B))$ is in one of the C_i , and the corresponding start event has not yet been generated, try the following: If there is a recipe r such that $F_1(r) = t \wedge \dots \wedge F_n(r) = t$, return $\text{event}(e_{gs}(r, A, B)) \cdot \llbracket A, \phi \triangleright \emptyset, c, \{F_i : C_i\}_{i=1}^n \rrbracket$ while remembering that the event was generated. Otherwise, continue computing $\llbracket A, \phi \triangleright \emptyset, c, \{F_i : C_i\}_{i=1}^n \rrbracket$ as normal.

The goal is enforced by the following query:

$$\forall X, A, B. \quad A \neq i \wedge B \neq i \wedge \text{event}(e_{ge}(X, A, B)) \implies \text{event}(e_{gs}(X, A, B))$$

For each injective authentication goal, $g = \text{“B authenticates A on } t\text{”}$, we do the same procedure, except that we use the query:

$$\forall X, A, B. \quad A \neq i \wedge B \neq i \implies \text{event}(e_{ge}(X, A, B)) \sqsubseteq \text{event}(e_{gs}(X, A, B))$$

For more complicated properties, or if a custom placement of events is desired (for example if one wants an agent to emit a start event later in the protocol), CryptoChoreo permits one to manually place events and specify custom queries. Algorithm 1 demonstrates both ways of specifying goals.

IV. AUTOMATION OF THE PROJECTION

We give now an overview of algorithms for the algebraic problems underlying the projection semantics, for a representative algebraic theory. This section provides a simplified version, omitting many of details, which can be found in Appendix B. We have included brief explanations of the correspondence between the results here and the results in the appendix.

To implement the projection of Section III, one must be able to do the following:

- (*word problem*) Given two terms s and t , check if $s =_E t$.
- (*recipe composition*) Given a set of frame-term pairs $\{(F_1, t_1), \dots, (F_n, t_n)\}$, decide if there is a recipe r so $F_1(r) =_E t_1 \wedge \dots \wedge F_n(r) =_E t_n$, and return it if so.
- (*complete set of checks*) Given a frame F calculate a complete set of checks, i.e., a finite set of checks that implies all checks that can be made at all.

The projection semantics are defined to be agnostic to the precise algebraic theory, and in general these problems are not recursively computable. There are, however, practical theories of cryptographic operators for which we can solve these problems. In this section, we demonstrate this by giving procedures for a particular cryptographic model with Diffie-Hellman keys, symmetric encryption, and asymmetric encryption. We assume that for all destructors we have a corresponding verifier that can be used to check whether the destructor would succeed.

Recall that we are distinguishing terms, denoted s, t, \dots , from recipes, denoted r, r_1, r_a, \dots . On the local behavior level, we have recipes that are built over Σ_p and labels \mathcal{X}_i , where Σ_p contains only public function symbols, including destructors and verifiers. In the previous sections, terms occurred on the choreography level, and were built over the alphabet Σ and variables \mathcal{V} , where Σ did not contain any destructors or verifiers, but could contain private functions. In this section, we additionally allow terms to contain destructors, as we here make precise how a term like $F(r)$ (which may contain destructors) is equal in the E -theory to one that does not. For distinction, we call terms *constructive* if they do not contain any destructors or verifiers.

A. Algebraic Theory

As part of our algebraic theory we consider the following built-in symbols:

	Constructors	Destructors	Verifiers
Tuples	<i>pair</i>	<i>fst, snd</i>	<i>vpair</i>
Asymmetric Enc.	<i>crypt</i>	<i>ddecrypt</i>	<i>vcrypt</i>
Private Keys	<i>inv</i>	<i>pub_k</i>	<i>vinv</i>
Symmetric Enc.	<i>scrypt</i>	<i>dscrypt</i>	<i>vsencrypt</i>
Signatures	<i>sign</i>	—	<i>vsign</i>
Signatures (cont.)	—	<i>open</i>	<i>vopen</i>
Diffie-Hellman	<i>exp</i>	<i>exp⁻¹</i>	<i>vexp</i>

All the constructors are part of Σ and can occur in terms, and all except *inv* are public and thus in Σ_p and can occur in recipes. Note that we had earlier used n -tuples for simplicity, and here have only binary tuples, but this can be seen as syntactic sugar. Besides these symbols, the modeler can declare other further function symbols that can be either public (and thus both part of Σ and Σ_p), e.g., to model hash functions or public constants, or that can be private (and thus only part of Σ), e.g., to model key infrastructures or fixed secrets between agents. However, these user-defined functions cannot have any algebraic properties. We also have a public constant \top both in Σ and Σ_p . We call a recipe *constructive* if it does not contain destructors or verifiers.

Definition 2. We define our algebra $E = R \cup B$ where $R =$

$$\begin{aligned} &\{vpair(pair(x, y), \top) \doteq \top, \\ &fst(pair(x, y), \top) \doteq x, snd(pair(x, y), \top) \doteq y, \\ &vsencrypt(scrypt(x, y), y) \doteq \top, \\ &dscrypt(scrypt(x, y), y) \doteq x, \\ &vcrypt(crypt(x, y), inv(y)) \doteq \top, \\ &ddecrypt(crypt(x, y), inv(y)) \doteq x, \end{aligned}$$

$v\text{sign}(\text{sign}(x, \text{inv}(y)), y) \doteq \top$,
 $v\text{open}(\text{sign}(x, y), \top) \doteq \top$,
 $\text{open}(\text{sign}(x, y), \top) \doteq x$,
 $v\text{inv}(\text{inv}(x), \top) \doteq \top$, $\text{pub}_k(\text{inv}(x), \top) \doteq x$,
 $v\text{exp}(\text{exp}(x, y), y) \doteq \top$, $\text{exp}^{-1}(\text{exp}(x, y), y) \doteq x$,
 and $B = \{\text{exp}(\text{exp}(x, y), z) \doteq \text{exp}(\text{exp}(x, z), y)\}$.

Let $=_E$ denote the congruence induced by these equations and $=_B$ the congruence induced just by the equation in B .

The functions $v\text{pair}$, fst , snd , $v\text{inv}$, and pub_k should actually be unary functions, because they do not require a key. For uniformity, we have made them binary functions like all the other destructors and verifiers, and we use \top as a dummy value for the key-position. As all destructors can now be regarded as decryption operators we make no distinction between the concepts in the context of this algebraic theory.

The reader may be surprised to see a verifier for Diffie-Hellman exponentiation. This is because our method below requires that every destructor has a corresponding verifier. However, the verifier $v\text{exp}$ exists only pro forma: if our procedure runs into a situation where it actually employs $v\text{exp}$, it stops with an error. The only situation where it would be employed is if we have an agent who knows both x and a term (equivalent to) $\text{exp}(t, x)$, but not t , and this does not occur in standard uses of Diffie-Hellman. Thus, $v\text{exp}$ is a tool we use for our proofs, while still preserving soundness: by aborting translation we do not give this unrealistic capability to honest agents, and thus a successful translation is correct if we drop the unrealistic equation $v\text{exp}$.

A similar question may arise from the destructor and verifier for private keys. Many approaches model a public constructor that from a given private key generates a public key; we use here instead a private constructor inv to map a public key to a corresponding private key; this allows us easily model public-key infrastructures like $\text{ek}(A)$ being the public encryption of A where ek is a public function (so every agent can lookup keys). The small price to pay is that the inverse mapping from private to public key is called a destructor and that we have a verifier to check if a private key really fits with the public key.

B. Algorithms

We now sketch out how the three problems mentioned at the beginning of this section—the word problem, the recipe composition problem, and the complete set of checks problem—are all computable. The algorithms for this are close to standard protocol analysis methods, and not a main contribution of this work. Moreover, basic requirements of efficiency (e.g., avoiding repeated analysis steps and checks) make the algorithms rather involved with details. The full details and proofs are available in Appendix B.

1) *Word Problem*: The considered theory $E = R \cup B$ allows us to decide the word problem, i.e., for terms or recipes s, t , whether $s =_E t$. This is because R used as rewrite rules modulo B (i.e., $\rightarrow_{R/B}$) is convergent. Thus we only need to compare the normal forms of s and t modulo B . The

equivalence class modulo B of any term is finite and easily computable. In Appendix B, this is proven in Theorem 1

2) *Constructive Recipe Composition*: We now first solve a simplified version of the recipe composition problem, where recipes do not contain destructors or verifiers. We define a function $\text{compose}(F, t)$ that, given a frame F and a term t , obtains all *constructive* recipes r such that $F(r) =_B t$. Roughly, for every term t_0 in $[t]_B$ (the equivalence class of t modulo B) we can check if t_0 is a term in the frame, and additionally, if $t_0 = f(t_1, \dots, t_n)$ for a public f (which cannot be a destructor or verifier by construction), we recursively compute $\text{compose}(F, t_i)$ and from the results construct the solutions for $\text{compose}(F, t)$ as expected.

For example, let $F = [\mathcal{X}_1 \mapsto X, \mathcal{X}_2 \mapsto \text{exp}(g, Y)]$ (where g is a public constant) and the target is $t = \text{exp}(\text{exp}(g, X), Y)$. The equivalence class is $[t]_B = \{t, t'\}$ with $t' = \text{exp}(\text{exp}(g, Y), X)$. Neither t nor t' is directly contained in F , so we recursively check the subterms: for t we compute $\text{compose}(F, \text{exp}(g, X)) = \{\text{exp}(g, \mathcal{X}_1)\}$ and $\text{compose}(F, Y) = \emptyset$; the second call fails because we cannot obtain Y . For t' we compute $\text{compose}(F, \text{exp}(g, Y)) = \{\mathcal{X}_2\}$ and $\text{compose}(F, X) = \{\mathcal{X}_1\}$, and we thus get $\text{compose}(F, t) = \{\text{exp}(\mathcal{X}_2, \mathcal{X}_1)\}$.

In Appendix B, we show that compose returns all constructive recipes for a given term in Lemma 5.

3) *Analysis*: We further introduce the notion of an *analyzed frame*, i.e., where the frame contains every term that can be obtained using a destructor on any message in the frame, using a constructive recipe for the key term.

We define an analysis procedure that successively applies decryption steps as long as possible: for every message that potentially can be decrypted, we check if we can compose the decryption key. If so, the resulting message is added to the frame. Whenever we add an analyzed message to the frame, we also need to check again all those messages for which we previously did not have the decryption key.

For example, the frame $F = [\mathcal{X}_1 \mapsto X, \mathcal{X}_2 \mapsto \text{exp}(g, Y), \mathcal{X}_3 \mapsto \text{scrypt}((A, \text{crypt}(M, \text{pk}(B))), \text{exp}(\text{exp}(g, X), Y))]$ is not analyzed because there is a constructive recipe for the Diffie-Hellman key of the message in \mathcal{X}_3 , as seen before. We thus add to the frame $[\mathcal{X}_4 \mapsto (A, \text{crypt}(M, \text{pk}(B)))]$ and note that \mathcal{X}_4 is a shorthand for $\text{dscrypt}(\mathcal{X}_3, \text{exp}(\mathcal{X}_2, \mathcal{X}_1))$. This does not change the derivable messages of the frame of course. The resulting frame is still not analyzed as \mathcal{X}_4 is a pair, and we can add $[\mathcal{X}_5 \mapsto A, \mathcal{X}_6 \mapsto \text{crypt}(M, \text{pk}(B))]$ and note that \mathcal{X}_5 and \mathcal{X}_6 are shorthands for $\pi_1(\mathcal{X}_4, \top)$ and $\pi_2(\mathcal{X}_4, \top)$, respectively. This is now analyzed since we have no constructive recipe for $\text{pk}(B)$ (and in fact no recipe at all).

In Appendix B, we show that this procedure terminates in Theorem 2 and that produces correctly analyzed frames in Theorem 3.

The compose and analyze algorithms together solve the recipe composition problem: Suppose $F(r) =_E t$ where t does not contain destructors and verifiers, and F' is the analyzed version of F , then $\text{compose}(F', t) \neq \emptyset$. The proof is essentially that we look at any application of a destructor or verifier

in r (if there is any) that has no further destructor or verifiers as subterms, say $d(r_0, r_1)$ where r_0, r_1 are constructive. Then the analysis must have found it, i.e., we have a label \mathcal{X} in F' that produces a term that is E -equivalent to $F'(d(r_0, r_1))$. In this way we can successively replace all destructors in r by shorthands of F' until we obtain a recipe that is constructive and thus found by *compose*.

In Appendix B, Lemma 4 shows that if there is any recipe for producing a term in an analyzed frame, then *compose* will also return such a recipe.

4) *Complete Set of Checks*: Our analysis procedure also computes some checks: whenever we have $\mathcal{X} \mapsto t$ in a frame F such that t can be decomposed, say $d(t, k)$ is a redex and k can be constructed in the F , then we can also check $v(t, k) \doteq \top$ for the corresponding verifier.

For example, in the above analysis example we would derive the checks $vscript(\mathcal{X}_3, exp(\mathcal{X}_2, \mathcal{X}_1)) \doteq \top$ and $vpair(\mathcal{X}_4, \top) \doteq \top$.

In general, however, this does not yet give the complete set of checks, for instance $F = [\mathcal{X}_1 \mapsto h(N), \mathcal{X}_2 \mapsto N]$ is already analyzed, but there is still the check $\mathcal{X}_1 \doteq h(\mathcal{X}_2)$. A complete set of checks can now be found by checking for every label (like \mathcal{X}_1 here) if there is a different way to construct it. Essentially, the proof is that, given an analyzed frame and an arbitrary check $r_1 \doteq r_2$, we can reduce it to constructive recipes, and where one of the sides is a label.

In Appendix B, Theorem 4 shows how a finite complete set of checks can be derived for any frame.

5) *Extension to Multiple Frames*: Finally, the compose procedure can be extended to the recipe composition problem, i.e., given analyzed and checked F_1, \dots, F_n and goal terms t_1, \dots, t_n , find a single recipe r with $F_i(r) = t_i$ for all i , because we try to get a solution for $F_1(r) = t_1$ and if it exists, then it works in all frames, because they are checked.

In Appendix B, Lemma 9 shows that if a frame is fully checked it does not matter which recipe from *compose* one picks (in any actual run, they will all produce the same term). Theorem 5 shows how to obtain a single recipe that works for all given frame-term pairs.

V. CASE STUDIES

A. Exporting to ProVerif

In this section, we summarize how the output from the projection semantics of Section III can be further translated to ProVerif code for automatic verification. More details can be found in Appendix A. We start by unfolding the labels from the initial knowledge in each local behavior, similarly to what we do in the semantics of Figure 5. Each local behavior is almost a valid ProVerif process already, except for the use of nondeterministic choice and memory cells.

Nondeterministic choice is simple to encode: we take a message from the network (the intruder) indicating which branch to take, and then branch on the content of that message. Thus, we leave it to the intruder to pick the branch that will lead to an attack (if one exists).

Encoding memory cells is more involved, since reasoning about long-term mutable state in ProVerif is a well-known difficulty for which several extensions and front-ends have been proposed [13], [14], [15], [16], [17]. We encode each memory cell as a private channel, and ensure by construction that the channel always contains exactly one message (except when the message has been consumed in an atomic section that has not yet been left). Then, in the semantics of ProVerif, a read from the memory cell corresponds exactly to reading the term that was most recently placed in the channel.

However, when the ProVerif process is translated to Horn-clauses, certain overapproximations will mean that it is no longer guaranteed that a message on a private channel is consumed in the right order or only once. We decrease the chance of a false positive by adding a counter to each memory cell, which is incremented on every write. Adding the admissible axiom that if two values written to a memory cell are associated with the same counter value they must be equal, excludes many impossible models during the proof search.

B. Examples

We implemented a tool that automates the projection from choreographies to local behaviors, based on the definitions in Sections III and IV. It is implemented in around 5000 lines of Haskell code. The tool also supports the generation of a ProVerif file from these local behaviors, following the steps described in Section V-A.

We will in the following describe one particular example, but have made more available. We wish to highlight the following notable examples:

- `blind-forward.choreo` is similar to the example from Section II. Here a trusted third party either helps A authenticate a new encryption key with or authenticate a request to B, without knowing which branch it is in. We have verified that in the branch where B gets a new public key from A, that key is authenticated, and that in the branch where data is sent to A with the key we have secrecy of the data.
- `SSO.choreo` describes a protocol where an agent A authenticates to another agent B using a trusted third party `ttp` to establish a secure channel. We can also verify this when the behavior of `ttp` is taken from `SSO-API.choreo`, where `ttp` is implemented like an API responding to queries and saving their state using memory cells. Our tool includes an option for selecting the behavior of participants from different choreographies like this. We have verified that data sent over the channel is authenticated and secret, and that the symmetric key shared to establish the channel is secret.
- We demonstrate in `tpm-simple.choreo` and `tpm-simple-API.choreo` the security of a TPM that can either declassify a given value or delete it. In particular, this protocol is non-monotonic, as it should be transparent to the owner of the value which choice was taken, and that if one choice is made then the other cannot be made later. We have verified that if

TABLE I

File	LOC	Translation	Verification
blind-forward.choreo	28	0.06s	0.02s
SSO.choreo	32	0.10s	1.70s
tpm-simple.choreo	34	0.06s	0.02s
SSO-DH.choreo	28	0.06s	117.20s
NSLPK.choreo	23	0.05s	0.07s
tls-1.3.choreo	37	0.12s	171.80s
ASW.choreo	101	0.19s	0.13s

anyone (the intruder) comes to know the value then the “opened” event was triggered by the TPM, and that if the “refused” event was triggered by the TPM then the value cannot be known by the intruder.

- `SSO-DH.choreo` and `SSO-DH-API.choreo` demonstrates a Diffie-Hellman exchange mediated by a trusted third party. We have verified the secrecy of a value encrypted with the established Diffie-Hellman key.
- `NSLPK.choreo` contains the Needham-Schroeder public-key protocol with Lowe’s fix. We have verified injective authentication and secrecy in both directions.
- `tls-1.3.choreo` is a simplified model of TLS 1.3. We have verified injective authenticity and secrecy of the data transmitted over the established channel in both directions.
- `ASW.choreo` contains the example described in the following.

Table I shows some additional information on the examples described above. The given runtimes were obtained on an HP EliteBook 840 G10. **LOC** (Lines Of Code) is the number of lines in the choreography file, excluding comments and empty lines. **Translation** is the time in seconds CryptoChoreo used to translate the choreography file to local behaviors and then to ProVerif code. **Verification** is the time in seconds ProVerif used to verify the obtained ProVerif code.

In our experience, protocols that succeed verification usually do so in a few seconds or less, except when the protocol involves non-trivial uses of Diffie-Hellman (like `SSO-DH.choreo` and `tls-1.3.choreo`), in which case verification could take up to a few minutes. For protocols that fail verification the result are much more varied, ranging from an immediate negative result, to a negative result after a couple of minutes, to the verification procedure seemingly going into nontermination.

C. The Asokan-Shoup-Waidner Protocol

The Asokan-Shoup-Waidner (ASW) protocol [18] serves as a motivating example that demonstrates the expressiveness of our choreography language, particularly its support for explicit branching, nondeterministic choice, conditional behavior, and long-term memory access.

ASW is a fair contract signing protocol involving three participants: an originator O, a responder R, and a trusted third party TTP. The protocol mainly ensures that either both parties obtain a binding contract, or neither does. The key challenges addressed by this protocol are:

- **Timeouts and abort scenarios:** Participants may timeout, leading to different protocol continuations.
- **Stateful TTP:** The TTP must maintain memory of previous contract states to prevent inconsistent responses.
- **Conditional logic:** Protocol actions depend on checking stored memory values.

In the following, we use formats like f_1 , f_2 , f_{abort} and $f_{resolve}$. These are transparent functions that structure messages but provide no cryptographic guarantees. A formatted message $f(t_1, \dots, t_n)$ can be modeled as a tuple (c_f, t_1, \dots, t_n) where c_f is a public constant.

D. Nondeterministic Choice and Branching

Algorithm 1 presents the ASW choreography using explicit branching in the message flow. After O sends M_1 to R, the responder can either timeout and trigger an abort request (shown in red), or continue by generating NR and replying with M_2 . In the continuation, timeouts may occur again: after receiving M_2 , O can either timeout and ask the ttp to resolve, or proceed by sending NO ; symmetrically, after receiving NO , R can either timeout and resolve, or complete by sending NR . The explicit $+$ -marked alternatives capture this nondeterminism directly at the level of the choreography.

E. Conditional Behavior and Memory Access

The distinctive feature of ASW is the behavior of the trusted third party, which consults and updates long-term memory cells indexed by the contract identifier M_1 . In both ABORT and RESOLVE, the ttp first reads $ttp_mem[M_1]$ (with *blank* denoting an uninitialized cell) and then branches on the stored value.

This illustrates several language features:

- 1) **Memory read:** $blank := ttp_mem[M_1]$ reads the current state associated with M_1 and checks that it is currently uninitialized.
- 2) **Case distinction on memory:** depending on whether $ttp_mem[M_1]$ is still *blank*, has been set to *aborted*, or contains $(resolved, M_2)$, the ttp returns a consistently matching signed response.
- 3) **Memory write:** $ttp_mem[M_1] := aborted$ (in ABORT) and $ttp_mem[M_1] := (resolved, M_2)$ (in RESOLVE) record the outcome so that repeated requests cannot lead to contradictory replies.
- 4) **Non-monotonic memory:** the same cell can be read multiple times and updated across different interactions in the run.

F. Semantic Challenge: Participant Knowledge

A subtle aspect of ASW mechanization arises during projection. When O or R make decisions (e.g., “abort”), they must send signals that the ttp later interprets. However, the originating participant does not see the entire memory state of the TTP. During projection, the analyzer must determine:

- When to *accept* incoming values without immediate verification (e.g., when O receives an untagged nonce response).

Algorithm 1 ASW Choreography

```
1: O :  $\nu \text{Text}. \nu \text{NO}$ .
2: O  $\rightarrow$  R :  $\underbrace{\text{sign}(f_1(\text{O}, \text{R}, \text{ttp}, \text{Text}, h(\text{NO})), \text{inv}(pk(\text{O})))}_{=:M_1}$ .
3: R : R  $\rightarrow$  O :  $\text{timeout.ABORT}(\text{O}, M_1)$ 
4: +  $\nu \text{NR}$ .
5: R  $\rightarrow$  O :  $\underbrace{\text{sign}(f_2(M_1, h(\text{NR})), \text{inv}(pk(\text{R})))}_{=:M_2}$ .
6: O : O  $\rightarrow$  R :  $\text{timeout.RESOLVE}(\text{R}, M_1, M_2)$ 
7: + O  $\rightarrow$  R : NO.
8: R : event( $\text{obtain}(\text{R}, M_1)$ ).event( $\text{finish}(\text{R}, M_1)$ ).
9: R : R  $\rightarrow$  O :  $\text{timeout.RESOLVE}(\text{O}, M_1, M_2)$ 
10: + R  $\rightarrow$  O : NR
11: O : event( $\text{obtain}(\text{O}, M_1)$ ).event( $\text{finish}(\text{O}, M_1)$ ).verify that at most one of them can have fired in any given trace.
12: O authenticates R on  $M_2$ 
13: R ni-authenticates O on  $M_1$ 
14:
15: ABORT(A,  $M_1$ ) =
16: A  $\rightarrow$  ttp :  $\text{sign}(f_{\text{abort}}(M_1), \text{inv}(pk(\text{A})))$ 
17: ttp : blank :=  $\text{ttp\_mem}[M_1]$ .
18:  $\text{ttp\_mem}[M_1] := \text{aborted}$ .
19: ttp  $\rightarrow$  A :  $\text{sign}(f_{\text{abort}}(M_1), \text{inv}(pk(\text{ttp})))$ 
20: A : event( $\text{finish}(\text{A}, M_1)$ )
21: +  $\text{aborted} := \text{ttp\_mem}[M_1]$ .
22: ttp  $\rightarrow$  A :  $\text{sign}(f_{\text{abort}}(M_1), \text{inv}(pk(\text{ttp})))$ 
23: A : event( $\text{finish}(\text{A}, M_1)$ )
24: + ( $\text{resolved}, M_2$ ) :=  $\text{ttp\_mem}[M_1]$ .
25: ttp  $\rightarrow$  A :  $\text{sign}(f_{\text{resolve}}(M_1, M_2), \text{inv}(pk(\text{ttp})))$ 
26: A : event( $\text{obtain}(\text{A}, M_1)$ ).event( $\text{finish}(\text{A}, M_1)$ )
27:
28: RESOLVE(A,  $M_1, M_2$ ) =
29: A  $\rightarrow$  ttp :  $\text{sign}(f_{\text{resolve}}(M_1, M_2), \text{inv}(pk(\text{A})))$ 
30: ttp : blank :=  $\text{ttp\_mem}[M_1]$ .
31:  $\text{ttp\_mem}[M_1] := (\text{resolved}, M_2)$ .
32: ttp  $\rightarrow$  A :  $\text{sign}(f_{\text{resolve}}(M_1, M_2), \text{inv}(pk(\text{ttp})))$ 
33: A : event( $\text{obtain}(\text{A}, M_1)$ ).event( $\text{finish}(\text{A}, M_1)$ )
34: +  $\text{aborted} := \text{ttp\_mem}[M_1]$ .
35: ttp  $\rightarrow$  A :  $\text{sign}(f_{\text{abort}}(M_1), \text{inv}(pk(\text{ttp})))$ 
36: A : event( $\text{finish}(\text{A}, M_1)$ )
37: + ( $\text{resolved}, M_2$ ) :=  $\text{ttp\_mem}[M_1]$ .
38: ttp  $\rightarrow$  A :  $\text{sign}(f_{\text{resolve}}(M_1, M_2), \text{inv}(pk(\text{ttp})))$ 
39: A : event( $\text{obtain}(\text{A}, M_1)$ ).event( $\text{finish}(\text{A}, M_1)$ )
```

Verified Queries:

$$\forall M_1, M_2. \left(\begin{array}{l} \text{event}(\text{abort}(M_1)) \wedge \\ \text{event}(\text{resolve}(M_1, M_2)) \end{array} \right) \Rightarrow \perp$$
$$\forall M_1, A, B. \left(\begin{array}{l} A \neq \text{ttp} \wedge B \neq \text{ttp} \wedge B \neq i \wedge \\ \text{event}(\text{finish}(B, M_1)) \wedge \\ \text{event}(\text{obtain}(A, M_1)) \end{array} \right) \Rightarrow \text{event}(\text{obtain}(B, M_1))$$

- When to insert runtime *checks* that compare received values against later constraints (e.g., when verifying

$h(\text{NO}) = X$ after learning NO).

G. Verification

We have verified several queries on the protocol. Firstly, in the branch where ttp is not needed for mediation, we can show mutual authentication on the contracts. R can only noninjectively authenticate O on M_1 since O does not sign NR . We do not consider this crucial, in any case, since freshness would in practice be guaranteed by the text in the contract. The main properties of the protocol are checked by the two manual queries. Firstly, we want it to be impossible for ttp to both resolve a contract positively and abort it. We therefore insert events to record when either happens, and then verify that at most one of them can have fired in any given trace. Secondly, we want the protocol to be fair. We formalize this by adding the event $\text{obtain}(A, M_1)$ whenever an agent A has a valid contract for the text in M_1 and $\text{finish}(A, M_1)$ whenever an agent considers the protocol to be finished (either because it was aborted or because they got the contract). As long as the intruder does not block communication with ttp forever, it is easy to see that each agent will be able to reach the finish event. We then verify that if one agent has finished and the other has a valid contract, then the first must have a valid contract too.

VI. RELATED WORK

Alice-and-Bob notation. Several lines of work have given Alice-and-Bob notation a precise semantics by compiling them to lower-level role-based specifications, and by rejecting notations that are not executable because e.g. senders cannot construct a message or receivers cannot check it. Early semantics were developed for the free term algebra and later extended to equational theories, often by reducing executability questions to intruder deduction and unification problems [6], [11], [7]. Our semantics follows this tradition, but lifts it from linear notations to *choreographies* with nondeterministic choice, branching, and mutable long-term memory. This combination is essential for modeling modern protocol interactions with stateful services and APIs, while maintaining the compact global view that motivates Alice-Bob notations.

Protocol models and tool support. At the other end of the spectrum, multi-set rewriting languages (e.g., as used by Tamarin and the AVISPA family) provide an explicit account of state and message flows and are well-suited for reasoning with rich adversary models [1], [2], [19]. Process-calculus based tools such as ProVerif offer a more program-like view of each role and highly automated verification, typically by a sound over-approximation [20], [3]. Our contribution is complementary: CryptoChoreo aims to be a high-level *specification* language that can be translated into such backends. In this respect, our approach is aligned with lines of work that provide source-to-source translations or front-ends for existing verification tools, such as SAPIC/SAPIC+ [21], [22], but with a focus on preserving the readability and single global story of a choreography.

Choreographies and global types. Choreographic programming and the theory of multiparty session types/global types study how a global description of multiparty interaction can be projected to local behaviors, with correctness guarantees such as deadlock freedom [23], [24], [25]. Our work is inspired by the same global-to-local methodology, but targets the Dolev-Yao setting with an active adversary, cryptographic constructors/destructors, equational theories, and explicit attacker knowledge.

State and APIs in protocol models. Stateful extensions and encodings have been studied both at the specification level and in tool-oriented front-ends [21]. Reasoning about long-lived global state is in particular a known difficulty for ProVerif’s Horn-clause abstraction, which has motivated several targeted extensions and front-ends, including set-membership abstractions and AIF [13], Set-Pi [14], StatVerif [15], GSVerif [16], and the PSPSP framework [17]. CryptoChoreo makes state explicit at the choreography level via memory cells, and our ProVerif export takes inspiration from these works to mitigate the same overapproximation issues at the translation level.

Equational reasoning and automation. Reasoning about message construction, parsing, and checks in the presence of equational theories is a classic challenge in symbolic protocol analysis [26]. We give algorithms for a representative theory combining standard constructors/destructors with Diffie-Hellman exponentiation, and yields an effective projection procedure [27]. This is intentionally backend-agnostic at the level of the core semantics, while our ProVerif export demonstrates one concrete and practical target.

VII. CONCLUSIONS

We introduced CryptoChoreo, a choreography language for cryptographic protocols that extends Alice-and-Bob notation with nondeterministic choice, conditional branching, and mutable long-term memory, together with a projection-based semantics. The semantics is defined for an arbitrary algebraic theory E and computes how honest agents execute the protocol, namely how they compose outgoing messages and decompose and check incoming messages. This offers the modelers a great tool to write specifications, because they just specify how the messages look like in an unattacked protocol run, and the translation semantics figures out which steps the actual implementation has to do. We can thereby avoid many specification mistakes (where a part of a message that should be checked is forgotten in the implementation). Thus, this is a basis for generating secure-by-design implementations where, simply, functions such as encryption and decryption need to be connected to real cryptographic functions etc., and where the implementation is in a one-to-one relationship with the formal model fed into verification tools. In fact, we plan to investigate if it can also be used with cryptographic verification tools like CryptoVerif.

While for a general theory E , the translation is not computable, we give the procedure for an example theory and the backend ProVerif. The example theory covers the usual

constructor/destructor theories as well as Diffie-Hellman; here we use a verifier that could not be implemented in reality, but that allows to integrate exponentiation uniformly with the rest of the method based on analyzing all terms as far as possible.

We initially considered translation to SAPIC+ [22] because it is similarly on the level of process calculus and has translators to both ProVerif and Tamarin. However, we ran into difficulties with long-term mutable state (which is always a challenge for infinite state verification) and since handling mutable state is an essential feature of choreographies, we opted for translation directly to ProVerif. Indeed our aim is similar to SAPIC+ to give modelers a way to address several methods without deep technical knowledge of those, thus we will investigate as part of future work how to connect more methods, possibly using SAPIC+. Further future work includes widening the class of supported equational theories, and improving automation for non-monotonic state encodings.

REFERENCES

- [1] S. Meier, B. Schmidt, C. Cremers, and D. Basin, “The Tamarin prover for the symbolic analysis of security protocols,” in *Computer Aided Verification (CAV)*, 2013.
- [2] A. Armando, D. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuéllar, P. H. Drielsma, P. Héam, Olga Kouchnarenko, J. Mantovani, S. Mödersheim, David von Oheimb, M. Rusinowitch, Judson Santiago, M. Turuani, L. Viganò, and L. Vigneron, “The AVISPA tool for the automated validation of internet security protocols and applications,” in *Computer Aided Verification (CAV)*, 2005.
- [3] B. Blanchet, *ProVerif: Automatic Cryptographic Protocol Verifier, User Manual and Tutorial*, 2016.
- [4] J. K. Millen, “CAPSL: common authentication protocol specification language,” in *Workshop on New Security Paradigms*. ACM, 1996, p. 132.
- [5] G. Lowe, “Casper: A compiler for the analysis of security protocols,” *J. Comput. Secur.*, vol. 6, no. 1-2, pp. 53–84, 1998.
- [6] F. Jacquemard, M. Rusinowitch, and L. Vigneron, “Compiling and verifying security protocols,” in *Logic for Programming and Automated Reasoning (LPAR)*, 2000.
- [7] S. Mödersheim, “Algebraic properties in alice and bob notation,” in *International Conference on Availability, Reliability and Security (ARES)*, 2009.
- [8] Y. Chevalier and M. Rusinowitch, “Compiling and securing cryptographic protocols,” *Inf. Process. Lett.*, vol. 110, no. 3, pp. 116–122, 2010.
- [9] O. Almousa, S. Mödersheim, and L. Viganò, “Alice and Bob: Reconciling formal models and implementation,” in *Festschrift in honor of Pierpaolo Degano*, 2015.
- [10] D. A. Basin, M. Keller, S. Radomirovic, and R. Sasse, “Alice and bob meet equational theories,” in *Festschrift in honor of José Meseguer*, 2015.
- [11] C. Caleiro, L. Viganò, and David Basin, “On the semantics of alice & bob specifications of security protocols,” *Theoretical Computer Science*, 2006.
- [12] G. Lowe, “A hierarchy of authentication specifications,” in *Proceedings of the 10th IEEE Workshop on Computer Security Foundations*. IEEE Computer Society, 1997.
- [13] S. Mödersheim, “Abstraction by set-membership: Verifying security protocols and web services with databases,” in *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS)*. ACM, 2010, pp. 351–360.
- [14] A. Bruni, S. Mödersheim, F. Nielson, and H. R. Nielson, “Set-pi: Set membership π -calculus,” in *28th IEEE Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, 2015, pp. 185–198.
- [15] M. Arapinis, E. Ritter, and M. D. Ryan, “StatVerif: Verification of stateful processes,” in *24th IEEE Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, 2011, pp. 33–47.

- [16] V. Cheval, V. Cortier, and M. Turuani, “A little more conversation, a little less action, a lot more satisfaction: Global states in proverif,” in *31st IEEE Computer Security Foundations Symposium*. IEEE Computer Society, 2018, pp. 344–358.
- [17] A. V. Hess, S. Mödersheim, A. D. Brucker, and A. Schlichtkrull, “Performing security proofs of stateful protocols,” in *34th IEEE Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, 2021, pp. 1–16.
- [18] N. Asokan, V. Shoup, and M. Waidner, “Asynchronous protocols for optimistic fair exchange,” in *1998 IEEE Symposium on Security and Privacy*. IEEE Computer Society, 1998, pp. 86–99.
- [19] D. Basin, S. Mödersheim, and L. Viganò, “OFMC: A symbolic model checker for security protocols,” in *International Journal of Information Security*, 2005.
- [20] B. Blanchet, “An efficient cryptographic protocol verifier based on prolog rules,” in *Computer Security Foundations Workshop (CSFW)*, 2001.
- [21] S. Kremer and R. Künnemann, “Automated analysis of security protocols with global state,” *Journal of Computer Security*, 2016.
- [22] V. Cheval, C. Jacomme, S. Kremer, and R. Künnemann, “SAPIC+: Protocol verifiers of the world, unite!” in *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, 2022.
- [23] M. Carbone, K. Honda, and N. Yoshida, “Structured communication-centered programming for web services,” *ACM Trans. Program. Lang. Syst.*, vol. 34, no. 2, pp. 8:1–8:78, 2012. [Online]. Available: <https://doi.org/10.1145/2220365.2220367>
- [24] M. Carbone and F. Montesi, “Deadlock-freedom-by-design: multiparty asynchronous global programming,” in *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, R. Giacobazzi and R. Cousot, Eds. ACM, 2013, pp. 263–274. [Online]. Available: <https://doi.org/10.1145/2429069.2429101>
- [25] K. Honda, N. Yoshida, and M. Carbone, “Multiparty asynchronous session types,” *J. ACM*, vol. 63, no. 1, pp. 9:1–9:67, 2016. [Online]. Available: <https://doi.org/10.1145/2827695>
- [26] Y. Chevalier and L. Vigneron, “Automated unbounded verification of security protocols,” in *Computer Aided Verification (CAV)*, 2002.
- [27] B. Schmidt, S. Meier, C. Cremers, and D. Basin, “Automated analysis of diffie-hellman protocols and advanced security properties,” in *IEEE Computer Security Foundations Symposium (CSF)*, 2012.
- [28] R. Küsters and T. Truderung, “Using proverif to analyze protocols with diffie-hellman exponentiation,” in *Proceedings of the 22nd IEEE Computer Security Foundations Symposium*. IEEE Computer Society, 2009, pp. 157–171.
- [29] S. Mödersheim, “Diffie-hellman without difficulty,” in *Formal Aspects of Security and Trust - 8th International Workshop*, ser. Lecture Notes in Computer Science, G. Barthe, A. Datta, and S. Etalle, Eds., vol. 7140. Springer, 2011, pp. 214–229.

APPENDIX

A. Exporting to ProVerif (extended)

a) *ProVerif encoding*: In this section, we will describe how to use the projection semantics from Section III to automatically convert a choreography into a ProVerif process. For the rest of this section, we assume the following:

- We have a choreography \mathcal{C}
- In \mathcal{C} there are roles A_1, \dots, A_{n_A}
- The initial knowledges are given in $F_{A_1}, \dots, F_{A_{n_A}}$
- We have projected the choreography for each role: $\mathcal{L}_{A_1} = F_{A_1}(\llbracket F_{A_1} : \mathcal{C} \rrbracket_{A_1}), \dots, \mathcal{L}_{A_{n_A}} = F_{A_{n_A}}(\llbracket F_{A_{n_A}} : \mathcal{C} \rrbracket_{A_{n_A}})$
- Of all the roles, U_1, \dots, U_{n_U} are untrusted.
- For each role A_i , $U_{A_i 1}, \dots, U_{A_i n_{U_{A_i}}}$ are the untrusted roles occurring in \mathcal{L}_{A_i} and F_{A_i} (i.e. the parameters to the role)
- In \mathcal{C} we have the cell families c_1, \dots, c_{n_c}

We start by encoding the algebra. For the most part, we encode the algebra E as is, though a simplification must be

made with regard to exp . Directly inserting the equation in B in ProVerif will lead to nontermination. However, as shown in [28] and [29], it is sound to use the following encoding for standard Diffie-Hellman:

```

fun exp(bitstring, bitstring): bitstring.
const g: bitstring.
equation forall x: bitstring, y: bitstring;
  exp(exp(g, y), x) = exp(exp(g, x), y).

```

We will omit most of the definition of the algebra here, but include the encoding of symmetric encryption:

```

type skkey.
fun sscript(bitstring, skkey): bitstring.
reduc forall x: bitstring, k: skkey;
  dscript(sscript(x, k), k) = x.
fun vscript(bitstring, skkey): bool
reduc forall x: bitstring, k: skkey;
  vscript(sscript(x, k), k) = true
otherwise forall x: bitstring, k: skkey;
  vscript(x, k) = false.

```

As another preliminary step, we include the type of agents and the name of the intruder:

```

type agent.
free i: agent.

```

ProVerif processes are written in the π -calculus, and look much like our local behaviors, except for network communication, nondeterministic choice, and memory cells. We write $\llbracket \mathcal{L} \rrbracket_{\text{pv}}$ for the ProVerif process obtained from the local behavior \mathcal{L} , and in the following only describe the cases where the local behaviors differ from the resulting ProVerif processes.

In our local behaviors, we simply use $\text{send}(X)$ and $\text{receive}(X)$ to send the content of variable X to the public network or receive a value from the network and assign it to X . In ProVerif, the corresponding statements are $\text{out}(c, X)$ and $\text{in}(c, X:\text{bitstring})$, where c is a public channel. To make a clear conceptual divide, we declare two public channels: c is the public channel the communication declared in the choreography happens over, while we use ic when information must be given to or taken from the intruder for the purpose of modeling.

```

free ic: channel.
free c: channel.

```

We then define $\llbracket \text{send}(t).\mathcal{L} \rrbracket_{\text{pv}} = \text{out}(c, t); \llbracket \mathcal{L} \rrbracket_{\text{pv}}$ and $\llbracket \text{receive}(X).\mathcal{L} \rrbracket_{\text{pv}} = \text{in}(c, X:\text{bitstring}); \llbracket \mathcal{L} \rrbracket_{\text{pv}}$.

When a process makes a nondeterministic choice, our encoding simply has the intruder decide:

```

 $\llbracket \mathcal{L}_1 + \mathcal{L}_2 + \dots + \mathcal{L}_n \rrbracket_{\text{pv}} =$ 
in(ic, Branch:nat);
if Branch = 0 then
   $\llbracket \mathcal{L}_1 \rrbracket_{\text{pv}}$ 
else if Branch = 1 then
   $\llbracket \mathcal{L}_2 \rrbracket_{\text{pv}}$ 
   $\vdots$ 
else if Branch = n then

```

$\llbracket \mathcal{L}_n \rrbracket_{pv}$

With regard to memory cells, we restrict ourselves to choreographies with the following properties:

- In every atomic section there is at most one read and one write to each cell family on every branch, and if a branch contains both they must use the same address.
- All writes happen at the end of the atomic section.
- All addresses are public.⁸

For each cell family, c_i , we declare a function:

fun cell_ c_i (bitstring): channel [*private*].

We can use this function to create a private channel for each cell in the family. Each such channel will contain at most one message at any given time, representing the current value stored in the cell.

Information sent on a private channel is of course not revealed to the intruder. Furthermore, communication over private channels is synchronous, and we can use this to enforce atomicity when there is a read and write to the same cell in one atomic section.

The same trick can be used to create atomic sections in general. We define

free atomic_lock: channel [*private*].

const atomic_baton: bitstring.

To enter an atomic section, a process must obtain the baton, and they should send it back when they leave. We have not found it beneficial to enforce atomicity of all atomic sections from local behaviors, but use it strategically for some parts (for example the memory-initializer processes below).

We ensure that when an agent wants to read from a cell there is always a message available by including initializer processes for each cell family. ProVerif supports *tables*, which are set-like structure we can use to keep track of which cells have been initialized, thus making sure that we only initialize a cell once and enforce the property that each cell channel contains at most one message. For each c_i , we define:

table c_i _initializer_table(bitstring).

let c_i _initializer()
in(atomic_lock, = atomic_baton);

in(ic, Addr: bitstring);

get c_i _initializer_table(= Addr) **in** 0

else

insert c_i _initializer_table(Addr);

(**out**(cell_ c_i (Addr), blank) |

out(atomic_lock, atomic_baton)).

The outputs at the end must be put in parallel, so that one does not block the other.

We can now have the intruder initialize our memory cells:

$\llbracket X := c[t].\mathcal{L} \rrbracket_{pv} = \mathbf{out}(ic, t); \mathbf{in}(cell_c(t), X: \text{bitstring}); \llbracket \mathcal{L} \rrbracket_{pv}$

Also, writing should not block the rest of the process:

$\llbracket c[t] := s.\mathcal{L} \rrbracket_{pv} = (\mathbf{out}(cell_c(t), s) | \llbracket \mathcal{L} \rrbracket_{pv})$.

⁸This is not technically required for the translation to be sound. However, our translation will reveal addresses to the intruder, so you will get false attacks if this is not the case.

There are two remaining problems: If on a branch in an atomic section there is a read from a cell but no write, the value in the cell will be consumed and no other process will be able to read from it. Similarly, if there is a write but no read, the cell will contain two values, and either is a possible value when the next process reads from the cell (we also lose the atomicity mentioned above). We solve this by a simple transformation: If we are done translating a branch of an atomic section where we have processed $X := c[t]$ but no corresponding write and need to return the rest of the translation $\llbracket \mathcal{L} \rrbracket_{pv}$, return instead $(\mathbf{out}(cell_c(t), s) | \llbracket \mathcal{L} \rrbracket_{pv})$. Similarly, if we want to translate $\llbracket c_1[t_1] := s_1 \dots c_n[t_n] := s_n.\mathcal{L} \rrbracket_{pv}$ (recall that writes must be at the end of an atomic section) and have not read from c_i , we return $\mathbf{out}(ic, t_i); \mathbf{in}(cell_c_i(t_i), _ : \text{bitstring}); \llbracket c_1[t_1] := s_1 \dots c_n[t_n] := s_n.\mathcal{L} \rrbracket_{pv}$.

We now have a way to translate local behaviors to ProVerif processes. However, we still need to compose all these into a single process that can be verified by ProVerif. Furthermore, this process must handle the multiple sessions and instantiations supported by the semantics in Figure 5.

For each role A_i with parameters $U_{A_i1}, \dots, U_{A_i n_{U_{A_i}}}$ we define the *process spawner* for that agent:

let process $A_i(U_{A_i1}: \text{agent}, \dots, U_{A_i n_{U_{A_i}}}: \text{agent}) = \llbracket \mathcal{L}_{A_i} \rrbracket_{pv}$.

let spawn $A_i() =$

in(ic, ($U_{A_i1}: \text{agent}, \dots, U_{A_i n_{U_{A_i}}}: \text{agent}$));

process $A_i(U_{A_i1}, \dots, U_{A_i n_{U_{A_i}}})$.

For each untrusted U_i we must also give the intruder the associated initial knowledge. Let $U_{U_i1}, \dots, U_{U_i n}$ be all the parameters of the role U_i except the role U_i itself and t_1, \dots, t_m be the terms in the knowledge of U_i (i.e. $\mathcal{F}_{U_i} = \{\mathcal{X}_1 \mapsto t_1, \dots, \mathcal{X}_m \mapsto t_m\}$). We define:

let knowledge $U_i() =$

in(ic, ($U_{U_i1}: \text{agent}, \dots, U_{U_i n}: \text{agent}$));

out(ic, ($t_1[i/U_i], \dots, t_m[i/U_i]$)).

Finally, we can define the main ProVerif process:

process

(**out**(atomic_lock, atomic_baton) |

!(**new** a: agent; **out**(ic, a) |

!(spawn $A_1()$ |

⋮

!(spawn $A_{n_A}()$ |

!(knowledge $U_1()$ |

⋮

!(knowledge $U_{n_U}()$ |

! $c_{1_initializer}()$

⋮

! $c_{n_c_initializer}()$

b) *Improved Encoding of Memory Cells*: When verifying ProVerif code produced by the encoding above, we run into the

issue that the overapproximations associated with the abstractions of ProVerif make many secure protocols unverifiable. By default, ProVerif will consider the values in a cell as a set, and when you read you will not necessarily get the last value that was written. This means that for example the TPM example included with our submission does not verify, since ProVerif considers runs where *deleted* is written to the cell and then *classified* is later read from it.

This kind of overapproximation around mutable state is well-known and has been addressed by several extensions and front-ends to ProVerif and related tools [13], [14], [15], [16], [17]. We improve the accuracy of our encoding by a trick inspired by the “precise” option in ProVerif, originally presented in [16]. When writing t to a cell c , we instead write (i_c, t) where i_c is a natural number indicating that this was the i_c th value written to c . Furthermore, we trigger an event registering what was written, every time a value is written to a call. The new encoding of cell c_i becomes:

```

fun cell_ $c_i$ (bitstring): channel [private].
event write_ $c_i$ (bitstring, nat, bitstring).
table  $c_i$ _initializer_table(bitstring).
let  $c_i$ _initializer() =
  in(atomic_lock, = atomic_baton);
  in(ic, Addr: bitstring);
  get  $c_i$ _initializer_table(= Addr) in 0
  else
    event write_ $c_i$ (Addr, 0, blank);
    insert  $c_i$ _initializer_table(Addr);
    (out(cell_ $c_i$ (Addr), (0, blank)) |
     out(atomic_lock, atomic_baton)).

```

Furthermore, we redefine the process translation with the following:

```

[[ $X := c[t].\mathcal{L}$ ]]pV = out(ic, t);
                   in(cell_ $c$ (t), (Counter $_c$ : nat, X: bitstring));
                   [[ $\mathcal{L}$ ]]pV
[[ $c[t] := s.\mathcal{L}$ ]]pV = event write_ $c$ (t, Counter $_c$ , s);
                   (out(cell_ $c$ (t), (Counter $_c$  + 1, s)) | [[ $\mathcal{L}$ ]]pV)

```

When inserting writes on branches that only have reads, we can omit incrementing the counter, as we will necessarily just write back the value that was already there. In our experiments, this omission increases the efficiency of ProVerif in many cases.

Our encoding ensures that every time a new value is put in the cell, it is associated with a new counter value. Thus, it is sound to add the following axiom:

```

axiom  $t$ : bitstring,  $C$ : nat,  $s_1$ : bitstring,  $s_2$ : bitstring;
       event(write_ $c_i$ ( $t$ ,  $C$ ,  $s_1$ ))  $\wedge$  event(write_ $c_i$ ( $t$ ,  $C$ ,  $s_2$ ))
        $\implies s_1 = s_2$ .

```

B. Automation of the Projection in Detail

We will here show in more detail how to solve the *word problem*, *recipe-composition problem*, and *complete-set-of-checks problem* described in Section IV.

In addition to the Σ and Σ_p from Section IV, we use Σ_d to denote the set of destructors (+ pair projections, Diffie-Hellman inverse, etc.) and Σ_v for the set of verifiers.

Unlike in the main matter of the paper, we will here permit terms s, t, \dots to contain all function symbols, as we here make precise how these terms are then reduced to ones built only from Σ .

For the benefit of the following proofs, we define the syntactic size of terms and checks: $|\mathcal{X}| = 1$, $|f(t_1, \dots, t_n)| = 1 + |t_1| + \dots + |t_n|$, and $|t \doteq s| = |t| + |s|$.

Additionally, we need the following concepts:

Definition 3. We define $[t]_{=B} = \{t' \mid t =_B t'\}$. We write \rightarrow_R for the rewriting system obtain by applying the equations in R from left to right. We can then define the rewriting system modulo B -equivalence classes: $s \rightarrow_{R/B} t$ iff $\exists s' t'. s =_B s' \wedge s' \rightarrow_R t' \wedge t' =_B t$. We use $t \downarrow_{R/B}$ to denote the normal form of t .

Lemma 1. If $s =_B t$ then $|s| = |t|$ and if $s \rightarrow_{R/B} t$ then $|s| > |t|$.

Proof. The first part can be shown by induction on the derivation of $s =_B t$, and the second part follows by case analysis on $s' \rightarrow_R t'$ after unfolding the definition of $\rightarrow_{R/B}$. \square

Lemma 2. $s =_E t$ if and only if either:

- $s =_B t$,
- or there is s' so $s \rightarrow_{R/B} s'$ and $s' =_E t$,
- or there is t' so $t \rightarrow_{R/B} t'$ and $s =_E t'$

Proof. We do induction on the derivation of $s =_E t$ in the equational logic induced by E . \square

We can now show that our rewriting system is well-behaved:

Lemma 3. (A) $[t]_{=B}$ is finite for all t . (B) If $s \rightarrow_{R/B} t$ then $s =_E t$. (C) $\rightarrow_{R/B}$ is convergent, modulo B . (D) $t \downarrow_{R/B}$ is defined and unique for all t , modulo B .

Proof. (A) is easy to see (the equation in B is a kind of commutativity). (B) follows from transitivity of $=_E$ and the definitions of \rightarrow_R and $\rightarrow_{R/B}$. (D) follows (C) by definition.

To show (C), we combine (B) with the following: (1) $\rightarrow_{R/B}$ has no infinite chains. (2) If $s =_E t$ and $\rightarrow_{R/B}$ applies to neither s nor t , then $s =_B t$.

(1) follows from the fact that if $t_1 \rightarrow_{R/B} t_2 \rightarrow_{R/B} \dots$ then $|t_1| > |t_2| > \dots$.

(2) follows from Lemma 2. \square

With this, we can solve the *word problem* from the start of the section:

Theorem 1. $s =_E t$ if and only if $s \downarrow_{R/B=B} t \downarrow_{R/B}$.

Proof. That $s \downarrow_{R/B=B} t \downarrow_{R/B}$ implies $s =_E t$ follows from Lemma 3 (B).

If $s =_E t$ then $s \downarrow_{R/B=E} t \downarrow_{R/B}$ by Lemma 3 (B), and since $\rightarrow_{R/B}$ applies to neither $s \downarrow_{R/B}$ nor $t \downarrow_{R/B}$ we have $s \downarrow_{R/B=B} t \downarrow_{R/B}$ by Lemma 2. \square

In the following, we will analyze a set of frames by alternating between applying verifiers and destructors to obtain new terms. We extend frames with additional information, including which checks have been performed, and how the content in a label was derived.

Definition 4 (Enhanced Frames). *Frames are built over the following grammar:*

$$\begin{array}{l}
F ::= 0 \\
\quad | F_0.\mathcal{X} \mapsto t \\
\quad | F_0.\mathcal{X} \leftarrow r \mapsto t \\
\quad | F_0.r_1 \doteq r_2
\end{array}$$

We let e range over the entries of the form $\mathcal{X} \mapsto t$, $\mathcal{X} \leftarrow r \mapsto t$, and $r_1 \doteq r_2$.

- The domain of a frame is defined as follows. $\text{dom}(0) = \emptyset$, $\text{dom}(F_0.\mathcal{X} \mapsto t) = \text{dom}(F_0) \cup \{\mathcal{X}\}$, $\text{dom}(F_0.\mathcal{X} \leftarrow r \mapsto t) = \text{dom}(F_0) \cup \{\mathcal{X}\}$, and $\text{dom}(F_0.r_1 \doteq r_2) = \text{dom}(F_0)$.
- Given a recipe r for frame F , we define $F(r)$ as expected, ignoring $r_1 \doteq r_2$ and $\leftarrow r$.

In the following, we only consider well-formed frames, which satisfy these properties:

- In a frame $F_0.\mathcal{X} \mapsto t$, $\mathcal{X} \notin \text{dom}(F_0)$ and t is constructive.
- In a frame $F_0.\mathcal{X} \leftarrow r \mapsto t$, $\mathcal{X} \notin \text{dom}(F_0)$, $\text{fv}(r) \subseteq \text{dom}(F_0)$, $F_0(r) = t$, and there exists constructive t' such that $t =_E t'$.
- In a frame $F_0.r_1 \doteq r_2$, $\text{fv}(r_1) \cup \text{fv}(r_2) \subseteq \text{dom}(F_0)$, and $F_0(r_1) =_E F_0(r_2)$.
- In a frame $F_0.e$, F_0 is well-formed.

It is easy to see that the projection semantics above and the following procedures preserve well-formedness.

The checks contained in a frame F , written $\text{checks}(F)$, is the set of all the equations $r_1 \doteq r_2$ contained in the frame together with the equations $l \doteq r$ for each entry $l \leftarrow r \mapsto t$ in the frame.

In the projection semantics, we consider sets of frames. There is a notion that these frames are identical with regard to all operations we have performed so far, which we formalize in the following:

Definition 5 (Compatible Frames). *We define two frames being compatible, written $F \simeq F'$, as the least relation satisfying the following rules:*

$$\begin{array}{c}
\frac{0 \simeq 0}{F_0 \simeq F'_0} \\
\frac{F_0.\mathcal{X} \mapsto t \simeq F'_0.\mathcal{X} \mapsto t'}{F_0 \simeq F'_0} \\
\frac{F_0.\mathcal{X} \leftarrow r \mapsto t \simeq F'_0.\mathcal{X} \leftarrow r \mapsto t'}{F_0 \simeq F'_0} \\
\frac{F_0.r_1 \doteq r_2 \simeq F'_0.r_1 \doteq r_2}{}
\end{array}$$

Thus, frames are compatible if they have passed the same checks and performed the same operations to construct new terms, in the same order.

Our analysis should extend a frame to contain all derivable subterms. This goal is captured in the following:

Definition 6. *A label $\mathcal{X} \in \text{dom}(F)$ is analyzed (in F) if for any constructive key-recipe r_k and destructor $d \in \Sigma_d$, if d applies to $(F(\mathcal{X}), F(r_k))$, there exists an entry $\mathcal{X}' \leftarrow d(\mathcal{X}, r'_k) \mapsto t$ where $F(r_k) =_E F(r'_k)$. Furthermore, for the associated verifier v there must be an entry $v(\mathcal{X}, r'_k) \doteq \top$. A frame F is fully analyzed if all labels in $\text{dom}(F)$ are analyzed.*

Note that since F is well-formed, we must have $t =_E d(F(\mathcal{X}), F(r_k))$.

When all the subterms have been added to a frame, we can obtain any obtainable term by *composition*, i.e. a constructive recipe. The following recursive function returns all constructive recipes for a given term:

Definition 7. *We define $\text{compose}(F, t) = \bigcup \{ \text{compose}_{lc}(F, s) \mid s \in [t]_{=B} \}$ where $\text{compose}_{lc}(F, s) = \text{compose}_l(F, s) \cup \text{compose}_c(F, s)$, $\text{compose}_l(F, s) = \{ \mathcal{X} \mid \mathcal{X} \in \text{dom}(F) \wedge F(\mathcal{X}) = s \}$, and $\text{compose}_c(F, s) = \{ f(s'_1, \dots, s'_n) \mid s'_1 \in \text{compose}(F, s_1) \wedge \dots \wedge s'_n \in \text{compose}(F, s_n) \}$ if $s = f(s_1, \dots, s_n)$ for $f \in \Sigma_c$ and $\text{compose}_c(F, s) = \{ \}$ otherwise.*

That $\text{compose}(F, t)$ is finite can be seen by induction on $|t|$ and that $[t]_{=B}$ is finite.

The following shows that for fully analyzed frames, compose can create any obtainable term:

Lemma 4. *Assume that F is a fully analyzed frame.*

If there is a recipe r and a constructive term t such that $F(r) =_E t$, then we have that

- (1) $\text{compose}(F, t)$ is non-empty, and
- (2) for every $r_c \in \text{compose}(F, t)$ we have $F(r_c) =_E t$.

Proof. (1) We show that if $F(r) \rightarrow_{R/B}^* t$ then there is $r_c \in \text{compose}(F, t)$ such that $F(r_c) \rightarrow_{R/B}^* t$.

We do induction on $|r|$. If $r = \mathcal{X}$, we are done. If $r = f(r_1, \dots, r_n)$ for $f \in \Sigma_c$, we apply the induction hypothesis and are done. If $r = d(r_1, \dots, r_n)$ for $f \in \Sigma_d$, then either the recipe reduces and we are done by the induction hypothesis, or $r = d(\mathcal{X}, r_{key})$ and there is another label \mathcal{X}' and (by the induction hypothesis) a constructive recipe r'_{key} such that $F(\mathcal{X}') = F(d(\mathcal{X}, r'_{key}))$.

(2) It is easy to prove that for every $r_c \in \text{compose}_{lc}(F, t)$ we have $F(r_c) =_E t$ by induction on $|t|$. Furthermore, it is by definition the case that for every $s \in [t]_{=B}$ we have $s =_E t$. \square

The following shows completeness of compose :

Lemma 5. *For any frame F (even if not fully analyzed), if r is a constructive recipe such that $F(r) =_E t$ and t is constructive, then $r \in \text{compose}(F, t)$.*

Proof. Follows by induction on $|r|$. \square

In the following, we define an analysis procedure that can be triggered during the projection semantics of Section III any time a new term is added the frame.

For the analysis of frame F , we require a *marking* of the type $M : \text{dom}(F) \rightarrow \{\star, +, \checkmark\}$. All labels are initially marked \star , and whenever a label is added to the frame it is also marked \star .

Definition 8. A marking M of frame F is accurate if any label $\mathcal{X} \in \text{dom}(F)$ where $M(\mathcal{X}) = \checkmark$ is analyzed in F .

Definition 9 (Analysis Procedure). We want to calculate the analyzed extensions of a set of frames, $\{F_1, \dots, F_n\}$, and marking, M , where the marking is accurate for each frame and the frames are pairwise compatible. The analysis-extension function, $\text{Ana}(M, \{F_1, \dots, F_n\})$, is defined by three cases:

If the given set is empty ($n = 0$), return $\{\}$.

If there is a label \mathcal{X} marked \star , try the following in order:

- $F_1(\mathcal{X}), \dots, F_n(\mathcal{X})$ are all terms for which no verifier exists, then return $\text{Ana}(M[\mathcal{X} \mapsto \checkmark], \{F_1, \dots, F_n\})$.
- (decomposition) We have a frame F_i , a verifier $v \in \Sigma_v$, and a constructive term t such that v applies to $(F_i(\mathcal{X}), t)$ (for any d and $F_i(\mathcal{X})$ at most one such t exists and it is easy to find). Furthermore, $\text{compose}(F_i, t)$ is non-empty, containing at least r_{key} .

Then, let d_1, \dots, d_k be all destructors associated with v and $\mathcal{X}_1, \dots, \mathcal{X}_k$ fresh labels, and set $r_1 = d_1(\mathcal{X}, r_{\text{key}}), \dots, r_k = d_k(\mathcal{X}, r_{\text{key}})$.

Let M' be the marking that for any \mathcal{X}' returns \checkmark if $\mathcal{X} = \mathcal{X}'$, \star if $\mathcal{X}' \in \{\mathcal{X}_1, \dots, \mathcal{X}_k\}$ or $M(\mathcal{X}') = +$, and $M(\mathcal{X}')$ otherwise.

Finally, let

$$FS^+ = \{F.v(\mathcal{X}, r_{\text{key}}) \doteq \top.X_1 \leftarrow r_1 \mapsto F(r_1). \dots .X_k \leftarrow r_k \mapsto F(r_k) \mid F \in \{F_1, \dots, F_n\} \wedge F(v(\mathcal{X}, r_{\text{key}})) =_E \top\}$$

and

$$FS^- = \{F \mid F \in \{F_1, \dots, F_n\} \wedge F(v(\mathcal{X}, r_{\text{key}})) \neq_E \top\}.$$

We then return $\text{Ana}(M', FS^+) \cup \text{Ana}(M, FS^-)$.

- Return $\text{Ana}(M[\mathcal{X} \mapsto +], \{F_1, \dots, F_n\})$.

If no labels are marked \star , return $\{(M, \{F_1, \dots, F_n\})\}$.

Note that the choice of r_{key} does not matter, since if multiple recipes are supposed to produce the same term but actually do not, we are not in a run of the protocol that agrees with the given frame. This will then be discovered when we later do the checks.

Theorem 2. The analysis procedure terminates.

Proof. Let $\text{raw}(F)$ be all the entries of the form $\mathcal{X} \mapsto t$ (i.e. the label assignments that are not the product of analysis). Let $|F| = \sum \{ |t| \mid \mathcal{X} \mapsto t \in \text{raw}(F) \}$ and $|\{F_1, \dots, F_n\}| = \sum_{i=1}^n |F_i|$. Let $|M|^\checkmark$ be the number of labels marked \checkmark in M and $|M|^+$ be the number of labels marked $+$. Lexicographically measure $(|\{F_1, \dots, F_n\}| - |M|^\checkmark, |\{F_1, \dots, F_n\}| - |M|^+)$.

The number of labels in a marking will never exceed the number of subterms of the raw terms of the frames the marking

is based on, so $|\{F_1, \dots, F_n\}| - |M|^\checkmark$ and $|\{F_1, \dots, F_n\}| - |M|^+$ will always be at least 0.

It is easy to see that no step in the analysis makes $|\{F_1, \dots, F_n\}| - |M|^\checkmark$ increase, and that each step will either decrease $|\{F_1, \dots, F_n\}| - |M|^\checkmark$ or $|\{F_1, \dots, F_n\}| - |M|^+$. \square

Theorem 3. If F_1, \dots, F_n are compatible frames, M is an accurate marking, and $\text{Ana}(M, \{F_1, \dots, F_n\}) = \{(M_1, FS_1), \dots, (M_m, FS_m)\}$, then all frames in each FS_i are fully analyzed, pairwise compatible, and M_1, \dots, M_m are accurate.

Proof. It is easy to see that the procedure preserves compatibility and accuracy of markings.

When the procedure is done, no label will be marked \star . If one of F_{ij} is then not fully analyzed, there must be a label \mathcal{X} , a constructive term t , and (by Lemma 5) a recipe $r_k \in \text{compose}(F, t)$, such that the (decomposition) case is applicable for \mathcal{X} (which must mean, by accuracy, that \mathcal{X} is marked $+$). However, this cannot be the case, as if this r_k existed when the marking of \mathcal{X} was changed from \star to $+$ then the wrong case was applied, and if it did not exist then the frame must have changed since this, which would remark \mathcal{X} as \star . \square

After distinguishing and extending a set of frames with this procedure, defining procedures for solving the two remaining problems stated at the beginning of this section becomes easier.

We start by showing that in an analyzed frame, all checks can be done in a constructive way:

Lemma 6. If F is a fully analyzed frame, then for any check $r_a \doteq r_b$ where $F(r_a) =_E F(r_b)$ there exists a set of constructive checks ϕ_c such that $\bigwedge \text{checks}(F) \wedge \bigwedge \phi_c \models_E r_a \doteq r_b$.

If r_a is of the form $f(r'_a, r_k)$ where $f \in \Sigma_v \cup \Sigma_d$ and r'_a, r_k and r_b are constructive, then there is such a ϕ_c where for each check $r_1 \doteq r_2 \in \phi_c$ we have that $|r_1 \doteq r_2| \leq \max(2 * |r_k|, |r'_a| + |r_b|)$.

Proof. We prove the first part by induction on $|r_a \doteq r_b|$.

The second part of the lemma can be proven by considering the cases of r'_a . First, we note that the proof is immediate if it is not the case that f is a destructor or verifier that applies to $(F(r'_a), F(r_k))$, and we therefore only consider this case.

If $r'_a = \mathcal{X}$ where \mathcal{X} is a label, then we know that for the applicable verifier v and destructors d_1, \dots, d_n we have a constructive recipe r'_k and labels $\mathcal{X}_1, \dots, \mathcal{X}_n$ such that $\{v(\mathcal{X}, r'_k) \doteq \top, d_1(\mathcal{X}, r'_k) \doteq \mathcal{X}_1 \dots d_n(\mathcal{X}, r'_k) \doteq \mathcal{X}_n\} \subseteq \text{checks}(F)$. If $f = v$ we set $\phi_c = \{r_k \doteq r'_k, r_b \doteq \top\}$ and are done. If $f = d_i$ we set $\phi_c = \{r_k \doteq r'_k, r_b \doteq \mathcal{X}_i\}$ and are done.

If r'_a is a composed recipe (i.e. of the form $g(r'_{a1}, \dots, r'_{an})$) then it must be the case that $f(r'_a, r_k) \rightarrow_R r''_a$ for some constructive recipe r''_a . We then set $\phi_c = \{r''_a \doteq r_b\}$ and are done. \square

Definition 10. For a term F we define the set of all checks that hold in the frame:

$\phi(F) = \{r_1 \doteq r_2 \mid r_1 \in \mathcal{T}(\Sigma_p, \text{dom}(F)) \wedge r_2 \in \mathcal{T}(\Sigma_p, \text{dom}(F)) \wedge F(r_1) =_E F(r_2)\}$

We also define the set of all checks between labels and recipes we can compose:

$\phi_{lc}(F) = \{\mathcal{X} \doteq r \mid \mathcal{X} \in \text{dom}(F) \wedge r \in \text{compose}(F, F(\mathcal{X}))\}$

Since a frame is finite, and for any t we have that $\text{compose}(F, t)$ is finite, it must be that $\phi_{lc}(F)$ is finite.

We can now solve the *complete set of checks* problem:

Theorem 4. *If F is fully analyzed, then for any $r_1 \doteq r_2 \in \phi(F)$ we have $\bigwedge \text{checks}(F) \wedge \bigwedge \phi_{lc}(F) \models_E r_1 \doteq r_2$*

Proof. Because of Lemma 6 we can assume that r_1 and r_2 are constructive. We do induction on $|r_1 \doteq r_2|$.

If either r_1 or r_2 is a label, we are done.

If $r_1 = f(r_{11}, \dots, r_{1n})$ and $r_2 = g(r_{21}, \dots, r_{2n})$, we must have $f = g$. Furthermore, if $r_{11} \doteq r_{21}, \dots, r_{1n} \doteq r_{2n}$ are valid checks in F , and we apply the induction hypothesis to get checks ϕ_1, \dots, ϕ_n . We have $\bigwedge \text{checks}(F) \wedge \bigwedge \phi_1 \wedge \dots \wedge \bigwedge \phi_n \models_E r_1 \doteq r_2$ and are done.

The only hard part is if $r_1 = \text{exp}(r_{1a}, r_{1b})$ and $r_2 = \text{exp}(r_{2a}, r_{2b})$ when $F(r_{1a}) \neq_E F(r_{2a})$ or $F(r_{1b}) \neq_E F(r_{2b})$. We then have either that $F(r_{1a}) =_E \text{exp}(s, t)$ and $t =_E F(r_{2b})$ or $F(r_{2a}) =_E \text{exp}(s, t)$ and $t =_E F(r_{1b})$, for some constructive s and t . We consider the former and the proof for the latter is symmetric.

We get $\text{exp}^{-1}(F(r_{1a}), F(r_{2b})) = s$ and $\text{vexp}(F(r_{1a}), F(r_{2b})) = \top$. From $\text{exp}^{-1}(F(r_{1a}), F(r_{2b})) = s$ and Lemma 4, we get a constructive recipe r' such that $F(r') = s$. Observe that $|r'| < |r_{1a}|$.

By Lemma 6 we get that there must be constructive checks $\phi^{r'}$ such that $\bigwedge \text{checks}(F) \wedge \bigwedge \phi^{r'} \models_E \text{exp}^{-1}(r_{1a}, r_{2b}) \doteq r' \wedge \text{vexp}(r_{1a}, r_{2b}) = \top$ and for all $r_\phi \doteq r'_\phi \in \phi^{r'}$ we have $|r_\phi \doteq r'_\phi| \leq \max(2 * |r_{2b}|, |r_{1a}| + |r'|)$. Since $|r_{2b}| < |r_{1a}|$ we further have that for all $r_\phi \doteq r'_\phi \in \phi^{r'}$ we have $|r_\phi \doteq r'_\phi| < |r_1 \doteq r_2|$. The induction hypothesis then gives a set of checks $\phi_{lc}^{r'} \subseteq \phi_{lc}(F)$ such that $\bigwedge \text{checks}(F) \wedge \bigwedge \phi_{lc}^{r'} \models_E \bigwedge \phi^{r'}$.

Furthermore, we have that $F(\text{exp}(r', r_{1b})) =_E F(r_{2a})$, which together with the induction hypothesis gives a set of checks $\phi_{lc}^{r_{2a}} \subseteq \phi_{lc}(F)$ such that $\bigwedge \text{checks}(F) \wedge \bigwedge \phi_{lc}^{r_{2a}} \models_E \text{exp}(r', r_{1b}) \doteq r_{2a}$.

Combining the above with

$\text{exp}^{-1}(r_{1a}, r_{2b}) \doteq r' \wedge \text{vexp}(r_{1a}, r_{2b}) = \top \models_E r_{1a} \doteq \text{exp}(r', r_{2b})$

and

$r_{1a} \doteq \text{exp}(r', r_{2b}) \wedge \text{exp}(r', r_{1b}) \doteq r_{2a} \models_E \text{exp}(r_{1a}, r_{1b}) \doteq \text{exp}(r_{2a}, r_{2b})$

yields

$\bigwedge \text{checks}(F) \wedge \bigwedge \phi_{lc}^{r'} \wedge \bigwedge \phi_{lc}^{r_{2a}} \models_E r_1 \doteq r_2$

and we are done. \square

It still remains to solve *recipe composition*. We have already shown that *compose* will compose a recipe for a term in a single fully analyzed frame if it exists, but it remains to be

shown that we can use it to compose a recipe that will work for all given frames. To prove this, we will leverage the complete set of checks.

Definition 11. *A frame F is fully checked if for any $r_1 \doteq r_2 \in \phi(F)$ we have $\text{checks}(F) \models_E r_1 \doteq r_2$.*

The following checking procedure simply mimics the one from Section III:

Definition 12 (Checking Procedure). *We want to calculate the checked extensions of a set of frames, $\{F_1, \dots, F_n\}$ where the frames are pairwise compatible. This is given by $\text{Check}(\phi_{lc}(F_1) \cup \dots \cup \phi_{lc}(F_n), \{F_1, \dots, F_n\})$ where Check is the check-extension function. $\text{Check}(\phi, \{F_1, \dots, F_n\})$ is defined by three cases:*

If $\phi = \{\}$, return $\{\{F_1, \dots, F_n\}\}$.

If $n = 0$, return $\{\}$.

If $r_1 \doteq r_2 \in \phi$, let $\phi' = \phi \setminus r_1 \doteq r_2$, $FS^+ = \{F.r_1 \doteq r_2 \mid F \in \{F_1, \dots, F_n\} \wedge F(r_1) =_E F(r_2)\}$

and $FS^- = \{F \mid F \in \{F_1, \dots, F_n\} \wedge F(r_1) \neq_E F(r_2)\}$,

and then return $\text{Check}(\phi', FS^+) \cup \text{Check}(\phi', FS^-)$

This procedure is sound:

Lemma 7. *If $\{F_1, \dots, F_n\}$ is a set of compatible and analyzed frames, then every $FS \in \text{Check}(\phi_{lc}(F_1) \cup \dots \cup \phi_{lc}(F_n), \{F_1, \dots, F_n\})$ will contain compatible, analyzed and checked frames.*

Proof. It is easy to see that compatibility is preserved by the procedure, and that they are still analyzed (no terms are added to the frames).

If $F \in FS$ then $\phi_{lc}(F) \subseteq \text{checks}(F)$, and we are done by Theorem 4. \square

As an immediate result, we get that there is no way to distinguish the frames that are grouped together by Check :

Lemma 8. *If F_1, \dots, F_n are compatible and fully checked frames, for any F_i and any check $r_1 \doteq r_2$ we have that $F_i(r_1) =_E F_i(r_2)$ if and only if $F_1(r_1) =_E F_1(r_2) \wedge \dots \wedge F_n(r_1) =_E F_n(r_2)$.*

Proof. We have $F_i(r_1) =_E F_i(r_2)$ if and only if $\text{checks}(F_i) \models_E r_1 \doteq r_2$. Furthermore, by compatibility, we have $\text{checks}(F_1) = \dots = \text{checks}(F_n)$. \square

We can also show a version where the frames are not necessarily compatible, but one frame fulfills all the checks in the complete set of checks of another. This implies that when we have done all the checks induced by the frame during the translation procedure, it does not matter which recipe is picked for constructing a term during actual execution (they will all refer to the same term anyways).

Lemma 9. *If ϕ is a complete set of checks for F and F' fulfills all the checks in ϕ , then for any two recipes r_1 and r_2 such that $F(r_1) =_E F(r_2)$ we will have $F'(r_1) =_E F'(r_2)$.*

Proof. Since $r_1 \doteq r_2$ is a valid check for F it must be logically implied by ϕ , which, since F' satisfies ϕ , means that it is also a valid check in F' . \square

For a set of compatible, fully analyzed, and fully checked frames, the *recipe composition* problem is simply solved by *compose*:

Theorem 5. *If F_1, \dots, F_n are compatible, fully checked, and fully analyzed frames, t_1, \dots, t_n are constructive terms, and there exists an r such that $F_1(r) =_E t_1 \wedge \dots \wedge F_n(r) =_E t_n$, then $\text{compose}(F_1, t_1)$ is non-empty and for any $r' \in \text{compose}(F_1, t_1)$ we have $F_1(r') =_E t_1 \wedge \dots \wedge F_n(r') =_E t_n$.*

Proof. We get the non-emptiness of $\text{compose}(F_1, t_1)$ and that for each r' we have $F_1(r') =_E t_1$ by Lemma 4. Since $F_1(r) =_E F_1(r')$, we have by Lemma 8 that $F_2(r) =_E F_2(r') \wedge \dots \wedge F_n(r) =_E F_n(r')$, and are done. \square

C. A Theory with no Finite Set of Checks

Let E be the following set of equations:

$$\{ \text{ord}(0) \doteq \top, \text{ord}(s(x)) \doteq \text{ord}(x), \text{ord}(\omega) \doteq \top, \\ \text{lt}(0, x) \doteq \text{ord}(x), \text{lt}(s(x), s(y)) \doteq \text{lt}(x, y), \\ \text{lt}(s(x), \omega) \doteq \text{lt}(x, \omega), g(x, h(y)) \doteq \text{lt}(x, y) \}.$$

where $g(\cdot, \cdot), 0, s(\cdot)$, and \top are public, and the others are private (note that ω is also private). Let $F = [\mathcal{X}_1 \mapsto h(\omega)]$. Note that the corresponding term rewriting system (i.e., with $\text{ord}(0) \rightarrow \top$ etc.) is convergent as can be shown with the critical pair methods (there are no critical pairs).

Theorem 6. *No finite set of equations is a complete set of checks for F .*

Proof. We assume a finite set ϕ of checks that are correct for F (i.e., $F(r) =_E F(r')$ for every $r \doteq r' \in \phi$). Without loss of generality, assume all checks are normalized terms.

Let n_0 be the maximum n such that $s^n(0)$ occurs in ϕ , or 0 if $s^n(0)$ does not occur in ϕ . (Such an n_0 must exist since ϕ is finite.)

Let \mathcal{I} be any interpretation such that $\mathcal{I}(X_1) = h(s^{n_0+1}(0))$. Now it suffices to show that $\mathcal{I} \models \phi$, because $\mathcal{I} \not\models g(s^{n_0+1}(0), \mathcal{X}_1) \doteq \top$, which, since $F(g(s^{n_0+1}(0), \mathcal{X}_1)) =_E \top$, shows that ϕ is not complete.

Let $(r, r') \in \phi$ and we can rule out $r = r'$, and note that $F(r) =_E F(r')$. We give a case distinction on r :

- $r = X_1$ and $r = 0$ is excluded because that can only be a sound check if $r = r'$.
- $r = s(r_0)$ then necessarily $r' = s(r'_0)$ with $F(r_0) =_E F(r'_0)$; thus we can just reduce this to the case of $r_0 \doteq r'_0$.
- $r = g(r_n, r_m)$. Note that r_m cannot be of the form $h(\cdot)$ because $h(\cdot)$ is private. Now case distinction on r' :
 - $r' = \top$. Thus $F(r) =_E \top$, thus $F(r_m) =_E h(m)$ for some ground term m and also $\text{lt}(F(r_n), m) =_E \top$. This forces $r_m = \mathcal{X}_1$ and $\text{ord}(F(r_n)) =_E \top$. That means r_n can only be of the form $s^n(0)$ for some $n \leq n_0$. Thus, $\mathcal{I} \models_E r \doteq r'$.
 - $r' = g(r'_n, r'_m)$. If both $F(r_n) =_E F(r'_n)$ and $F(r_m) =_E F(r'_m)$, then we can again reduce to the simpler checks $r_n \doteq r'_n$ and $r_m \doteq r'_m$.

Otherwise, both $F(r) =_E \text{lt}(x, y)$ and $F(r') =_E \text{lt}(x', y')$ for some ground terms x, x', y, y' and $\text{lt}(x, y) =_E \text{lt}(x', y')$. In this case both $F(r_m) =_E h(m)$ and $F(r'_m) =_E h(m')$ for some terms m and m' , thus $r_m = r'_m = \mathcal{X}_1$. Since $F(r_n) \neq_E F(r'_n)$ we have $\text{ord}(F(r_n)) =_E \top$ and $\text{ord}(F(r'_n)) =_E \top$. Thus, both $r_n = s^n(0)$ and $r'_n = s^n(0)$ for some $n, m \leq n_0$. Therefore, $\mathcal{I} \models_E r \doteq r'$.

- r' cannot be of any other form while $F(r) =_E F(r')$.
- $r = \top$. Then reduce to the case $r' \doteq r$. \square