

Secure Synthesis of Distributed Cryptographic Applications

Coşku Acay
Cornell University
Ithaca, NY, USA
cacay@cs.cornell.edu

Joshua Gancher
Carnegie Mellon University
Pittsburgh, PA, USA
jgancher@andrew.cmu.edu

Rolph Recto
Cornell University
Ithaca, NY, USA
rr729@cornell.edu

Andrew C. Myers
Cornell University
Ithaca, NY, USA
andru@cs.cornell.edu

Abstract—Developing secure distributed systems is difficult, and even harder when advanced cryptography must be used to achieve security goals. Following prior work, we advocate using *secure program partitioning* to synthesize cryptographic applications: instead of implementing a system of communicating processes, the programmer implements a *centralized, sequential* program, which is automatically compiled into a secure distributed version that uses cryptography.

While this approach is promising, formal results for the security of such compilers are limited in scope. In particular, no security proof yet simultaneously addresses subtleties essential for robust, efficient applications: multiple cryptographic mechanisms, malicious corruption, and asynchronous communication.

In this work, we develop a compiler security proof that handles these subtleties. Our proof relies on a novel unification of simulation-based security, information-flow control, choreographic programming, and sequentialization techniques for concurrent programs. While our proof targets hybrid protocols, which abstract cryptographic mechanisms as idealized functionalities, our approach offers a clear path toward leveraging Universal Composability to obtain end-to-end, modular security results with fully instantiated cryptographic mechanisms.

Finally, following prior observations about simulation-based security, we prove that our result guarantees *robust hyperproperty preservation*, an important criterion for compiler correctness that preserves all source-level security properties in target programs.

I. INTRODUCTION

Ensuring security for modern distributed applications remains a difficult challenge, as such systems can cross administrative boundaries and involve parties that do not fully trust each other. To defend their security policies, some applications employ sophisticated mechanisms such as complex distributed protocols [1, 2], trusted hardware [3, 4, 5], and advanced uses of cryptography. These technologies add significant complexity to software development and require expertise to use successfully [6, 7, 8].

To ease the development of secure distributed applications, prior work leverages compilers that translate high-level programs into distributed protocols that employ advanced security mechanisms. Unfortunately, most compilers only target a single mechanism—such as multiparty computation [9, 10, 11, 12], zero-knowledge proofs [13, 14, 15, 16], or homomorphic encryption [17, 18, 19]—and thus do not support secure combinations of mechanisms. On the other hand, compilers that perform *secure program partitioning* [20, 21, 22, 23, 24, 25]

do combine mechanisms, but come with limited or informal correctness guarantees.

In this work, we give the first formal security result for program partitioning that targets multiple cryptographic mechanisms, arbitrary corruption, and adversarially controlled scheduling. Our work proves the correctness of a reasonably faithful model of the compilation process used in the Viaduct compiler [25]. We formalize our result in the *simulation-based* security framework, which establishes a modular foundation for cryptographic protocol security [26]. Our security proof is primarily concerned with program partitioning itself, and thus does not reason about the concrete instantiation of cryptographic mechanisms; however, we discuss how to extend our results to reason about concrete mechanisms.

Our security proof incorporates multiple techniques for simulation-based security: information-flow type systems [27] to define the security policy and to guide partitioning, choreographies [28] to define global programs for distributed executions, and a novel information-flow guided technique for concurrent program sequentialization [29].

- We formalize a variant of Simplified Universal Composability (SUC) [30] enriched with information flow, allowing us to capture distributed protocols in the presence of adversarial scheduling and corruption.
- We show how to model secure program partitioning using security-typed choreographies [28]. The input to program partitioning is a sequential program representing an idealized execution on a single, trusted security domain, while the output is a distributed protocol with message-passing concurrency between mutually distrusting agents.
- We prove simulation-based security for our model of program partitioning. Informally, we show that any adversary interacting with the compiled distributed program is no more powerful than a corresponding adversary (a *simulator*) interacting with the source program.
- We show that, in our setting, simulation implies *robust hyperproperty preservation* [31], a strong criterion for compiler correctness that ensures security conditions of source programs are preserved in target programs.

II. OVERVIEW

We illustrate compilation via the classic Millionaires’ Problem [32], expressed as the source program in fig. 2a. Here,




Figure 1. Overview of compilation and the correctness proof. Left-to-right arrows are compilation steps; \geq are proof steps. Term w is a choreography, $\llbracket \cdot \rrbracket$ is endpoint projection, $\text{src}(\cdot)$ is the inverse of protocol synthesis, and $\langle \cdot \rangle$ models corruption. Dashed components represent proof sketches.

```

let a : A = input Alice;
let b : B = input Bob;
let x = declassify(endorse a < endorse b,
                  A  $\wedge$  B  $\rightarrow$  A  $\sqcap$  B);
output x to Alice
output x to Bob

```

(a) Source program with information-flow labels.

```

let Alice.a = input;
Alice.a  $\rightsquigarrow$  MPC(Alice, Bob).a';
let Bob.b = input;
Bob.b  $\rightsquigarrow$  MPC(Alice, Bob).b';
let MPC(Alice, Bob).x =
  declassify(endorse a' < endorse b');
MPC(Alice, Bob).x  $\rightsquigarrow$  Alice.x1;
MPC(Alice, Bob).x  $\rightsquigarrow$  Bob.x2;
output x1 to Alice
Alice.0  $\rightsquigarrow$  Bob._; // Sync outputs
output x2 to Bob

```

(b) Choreography with explicit communication and synchronization.

<pre> // Alice let a = input; send a to MPC(Alice, Bob) let x₁ = receive MPC(...); output x₁ send 0 to Bob // Sync </pre>	<pre> // Bob let b = input; send b to MPC(Alice, Bob) let x₂ = receive MPC(...); let _ = receive Alice; // Sync output x₂ </pre>
<pre> // MPC(Alice, Bob) let a' = receive Alice; let b' = receive Bob; let x = declassify(endorse a' < endorse b'); send x to Alice send x to Bob </pre>	

(c) Hybrid distributed program derived by projecting choreography.

<pre> // Alice let a = input; // Calls to MPC library output x₁ send 0 to Bob // Sync </pre>	<pre> // Bob let b = input; // Calls to MPC library let _ = receive Alice; // Sync output x₂ </pre>
---	--

(d) Concrete distributed program derived by instantiating idealized hosts.

Figure 2. Compiling the Millionaires' Problem

Alice and Bob learn who is richer without revealing their net worth to each other. To do so, the program collects inputs from Alice and Bob representing their net worth (lines 1 and 2); compares these (line 3), and outputs the result to Alice and Bob (lines 4 and 5).

A. Information Flow Control

Source programs prevent insecure information flows using a security type system [27], which assigns a *label* to every variable. Labels track the *confidentiality* and *integrity* of data. Our security type system follows prior work [33, 34] in using *downgrading mechanisms*—`declassify` and `endorse` expressions—to selectively allow information flows that would otherwise be deemed insecure. As in prior work [34], the type system constrains these downgrading mechanisms to prevent improper usage. These constraints turn out to be crucial.

In fig. 2a, the `declassify` expression explicitly allows revealing the result of the comparison $a < b$ to Alice and Bob, which is by default disallowed since the computation reads secrets from both parties. Dually, the `endorse` expressions allow untrusted data coming from Alice and Bob to influence the output from the comparison, which must be trusted since the value is output to both parties.

Downgrading require explicit *source* and *target* labels. Figure 2a suppresses these labels for the `endorse` expressions, but shows the `declassify` expression that moves from $A \wedge B$ to $A \sqcap B$. Label $A \wedge B$ is too secret for either Alice or Bob to see the value; label $A \sqcap B$ allows *both* parties to see it.

B. Compilation

Source programs serve as *specifications* of intended behavior, and correspond to *ideal functionalities* from Universal Composability [26]. Source programs act as trusted third parties, perfectly and securely executing the program on behalf of the involved hosts. The source language is high-level by design, and its simple, sequential semantics facilitate reasoning. Our compiler generates a concurrent distributed program that correctly implements the same behavior.

Figure 1 gives an overview of the compilation pipeline. First, *protocol synthesis* compiles the source program into a *choreography*, a single, centralized program that represents a distributed computation between many hosts. In addition to *parties* such as Alice and Bob, choreographies may mention *idealized hosts* such as `MPC(Alice, Bob)`, which represents a maliciously secure multiparty computation protocol between Alice and Bob. Idealized hosts can perform computations that require high confidentiality or integrity. Next, *endpoint projection*, a standard procedure in choreographic programming [35, 36], partitions the choreography into a distributed program, where hosts run in parallel and interact via message passing. The distributed program still contains idealized hosts, so it corresponds to a *hybrid program* in UC. Finally, *cryptographic instantiation* replaces idealized hosts with concrete cryptographic algorithms.

Figure 2b shows a choreography where *Alice* and *Bob* perform their respective **input** and **output** statements, while $\text{MPC}(\text{Alice}, \text{Bob})$ does the comparison. Explicit communication statements move data between hosts. Choreographies have *concurrent* semantics, so statements at different hosts may step out of program order. The penultimate line has *synchronization* between *Alice* and *Bob*: *Bob* must wait on an input from *Alice* before delivering output. This synchronization step is necessary for the distributed program to match the sequential source program, in which *Bob*’s output happens after *Alice*’s. Figure 2c shows the distributed program obtained by projecting the choreography in fig. 2b. Endpoint projection converts communication statements to **send/receive** pairs, and projects local computations to their corresponding hosts. Finally, fig. 2d shows the result of cryptographic instantiation. Each **send/receive** statement that interacts with $\text{MPC}(\text{Alice}, \text{Bob})$ is replaced with a call to a cryptographic library.

C. Defining Correctness

Our main contribution is a proof that compilation is correct. A correct compiler preserves properties of source programs in generated target programs. For generality, we demand that the compiler preserve all *hyperproperties* [37] guaranteed by the source program. Hyperproperties capture many common notions of security, including secure information flow.

Formally, preserving all hyperproperties is defined via *robust hyperproperty preservation* (RHP) [31]. Following prior observations [38, 39], we do not prove RHP directly, but instead prove *simulation*, which we show implies RHP in our framework. Simulation requires every attack by an adversary against the target program to be possible against the source program. Ideally, the source program is “obviously secure,” meaning it has straightforward, sufficiently abstract semantics and a narrow attack surface. In contrast, the target program faithfully models real code and has a realistic attack surface.

As with UC, our framework abstracts concrete cryptographic mechanisms as idealized hosts, yielding an *extensible* proof that is generic over the set of supported cryptographic mechanisms. Indeed, we sketch how our framework may be embedded into UC to leverage the UC composition theorem [26]; using the composition theorem, we can instantiate idealized hosts with cryptographic mechanisms proven secure separately.

a) Threat Model: Simulation demands we carefully define the capabilities of adversaries for each language in the pipeline. Adversaries are characterized by two sets of labels \mathcal{P} and \mathcal{U} representing *public* and *untrusted* labels. These label sets induce a partitioning of hosts into three sets: *honest* (secret and trusted), *semi-honest* (public and trusted), and *malicious* (public and untrusted). Intuitively, malicious hosts are fully controlled by the adversary, and semi-honest hosts follow the protocol but leak all their data to the adversary [40]. We say a host is *dishonest* if it is semi-honest or malicious, and *nonmalicious* if it is honest or semi-honest.

Source programs are fully sequential, so adversaries do not control scheduling. Moreover, adversaries cannot read or change intermediate data within a source program. However,

adversaries can read messages from **input/output** expressions involving dishonest hosts, read the results of **declassify** expressions with public target labels, and change the results of **endorse** expressions with untrusted source labels.

Choreographies and hybrid distributed programs have the same semantics, so their adversaries have equal power. These programs are concurrent and the adversary controls all scheduling. The adversary also fully controls malicious hosts and can read messages involving at least one semi-honest host. However, the adversary cannot carry out computational attacks, since cryptographic mechanisms are modeled as idealized hosts.

The adversary can view all message *headers* (source and destination), but not necessarily message *content*. The adversary may not drop, duplicate, or modify messages. This abstraction of secure channels can be realized by standard techniques, such as TLS [41]. In our model as in most models of cryptographic protocols [26, 42], the adversary can exploit *timing* and *progress* channels since it controls scheduling. These channels make secret control flow insecure: any discrepancy in timing or progress behavior between different control-flow paths can be detected by the adversary. Therefore, we only prove security for programs that make control flow decisions based on public, trusted data.

D. Roadmap of Correctness Proof

To define and prove our compilation pipeline secure, we make use of simulation (§III), which defines a relation \leq between semantic configurations. Intuitively, $W_1 \leq W_2$ means that any adversary against configuration W_1 is no more powerful than an equivalent adversary against W_2 . Crucially, the attacker often has more choices in W_1 (e.g., low-level scheduling decisions); the role of simulation is to show that these extra choices are benign, and thus W_1 is as secure as W_2 .

Our proof exploits the transitivity of simulation, using multiple intermediate simulations depicted in Figure 1. (To follow the flow of compilation, the figure uses \geq , which is defined as expected.)

a) Correctness of Protocol Synthesis: We first prove in §VI that protocol synthesis is correct: sequential source programs (e.g., fig. 2a) are simulated by their choreographies (e.g., fig. 2b).

There is a wide semantic gap between source programs and choreographies: while source programs are sequential and use **declassify/endorse** expressions to interface with the adversary, choreographies are concurrent and allow the adversary to read and corrupt data controlled by dishonest hosts. To bridge this gap, we break the protocol synthesis proof itself into three steps (fig. 12). These steps allow us to reason separately about the semantic features of choreographies.

Aside from employing transitivity, each proof step $W_1 \leq W_2$ requires us to define an appropriate *simulator* $\mathcal{S}(\cdot)$ such that for any adversary \mathcal{A} , W_1 running alongside \mathcal{A} is identical in behavior to $\mathcal{S}(\mathcal{A})$ running alongside W_2 . Thus, after defining the simulator as a labeled transition system, we show that the two resulting configurations are bisimilar using standard

Hosts	h	$\in \mathbb{H}$
Endpoints	c	$\in \mathbb{C} = \{\text{Adv}, \text{Env}\} \cup \mathbb{H}$
Values	v	$\in \mathbb{V} = \{0, \dots\}$
Messages	$m \in \mathbb{M}$	$::= c_1 c_2 v$
Actions	$a \in \mathbb{A}$	$::= ?m \mid !m$

$$\boxed{\text{actor}(a) = c} \quad \text{actor}(?c_1 c_2 v) = c_2 \quad \text{actor}(!c_1 c_2 v) = c_1$$

Figure 3. Syntax of messages and actions.

Processes	w
Configurations	$W ::= w_1 \parallel \dots \parallel w_n$

$$\boxed{W \xrightarrow{a} W'}$$

$$\frac{W\text{-INPUT} \quad \forall i. w_i \xrightarrow{?m} w'_i}{w_1 \parallel \dots \parallel w_n \xrightarrow{?m} w'_1 \parallel \dots \parallel w'_n}$$

$$\frac{W\text{-OUTPUT} \quad \begin{array}{l} w_i \xrightarrow{!m} w'_i \quad \forall j \neq i. w_j \xrightarrow{?m} w'_j \\ w_1 \parallel \dots \parallel w_n \xrightarrow{!m} w'_1 \parallel \dots \parallel w'_n \end{array}}$$

Figure 4. Syntax and semantics of configurations.

information-flow arguments (e.g., by defining an appropriate notion of *low-equivalence* between the configurations).

b) Correctness of Endpoint Projection: Second, we prove in §VII-C that choreographies are simulated by their corresponding distributed programs (e.g., fig. 2d). Our proof (theorem VII.1) largely follows the choreographic programming literature [28, 35, 36, 43, 44], but deals with extra complications arising from an adversary who may corrupt hosts.

c) Cryptographic Instantiation: Finally, we sketch in §VIII how hybrid distributed programs are simulated by concrete distributed programs, which make use of actual cryptographic mechanisms. In particular, we show how to embed our framework in the SUC [30] framework, which in turn embeds into the full UC framework. We can then leverage the UC composition theorem to instantiate idealized hosts one at a time, appealing to existing correctness proofs for cryptographic mechanisms.

III. SEMANTIC FRAMEWORK

We capture the semantics of programs using labeled transition systems (LTSs), where labels are *actions* a drawn from the grammar in fig. 3. An action a is either the input $?m$ or the output $!m$ of a message m . A message m specifies the endpoints c_1 and c_2 of communication and carries a value v . An endpoint is either a host h , the adversary Adv , or the external environment Env . Values are drawn from an arbitrary set \mathbb{V} , which we assume contains at least 0. We define $\text{actor}(a)$ as the host performing a : the sender performs output actions and the receiver performs input actions. Internal steps are represented as self-communication $!hh0$, which identifies a host h making progress without requiring a new syntactic form.

a) Configurations and Parallel Composition: A configuration, W , is the parallel composition of a finite set of processes w_i , which are arbitrary LTSs. Following prior work [45, 46], processes must be *input-total*: for every state w and input message $?m$, there exists a state w' such that $w \xrightarrow{?m} w'$. Figure 4 gives the semantics of configurations. A configuration W steps with an input $?m$ if all processes in W do, and steps with an output $!m$ if one of the processes outputs m , and the rest input m .

b) Adversaries: As with processes, an adversary \mathcal{A} or \mathcal{S} is an arbitrary LTS. The rules for running an adversary in parallel with a configuration are the same as in fig. 4. In contrast to processes, adversaries are *not* input-total, which enables adversaries to control scheduling: to schedule an endpoint c_1 , \mathcal{A} refuses to step with actions of the form $?c'_1 c_2 m$ where $c'_1 \neq c_1$, but accepts actions $?c_1 c_2 m$.

Due to the definition of parallel composition, a copy of every message from the configuration and the environment is delivered to the adversary; and any output of the adversary is delivered to the configuration and the environment. However, the adversary can only read a message if at least one endpoint is dishonest, and can only forge messages from malicious hosts.

Definition III.1 (Adversary Interface). For all \mathcal{A} :

- 1) If c_1 and c_2 are honest, then $\mathcal{A} \xrightarrow{?c_1 c_2 v_1} \mathcal{A}'$ if and only if $\mathcal{A} \xrightarrow{?c_1 c_2 v_2} \mathcal{A}'$ for all v_1 and v_2 .
- 2) If $\mathcal{A} \xrightarrow{!c_1 c_2 v}$, then either $c_1 = \text{Adv}$ or c_1 is malicious.

c) Determinism: To match UC, the adversary must resolve all nondeterminism, so that $\mathcal{A} \parallel W$ is deterministic. We ensure determinism with the following restrictions.

- Configurations and adversaries are *internally deterministic*: if $w \xrightarrow{a} w_1$ and $w \xrightarrow{a} w_2$, then $w_1 = w_2$.
- Adversaries are *output deterministic*: if $\mathcal{A} \xrightarrow{!m_1}$ and $\mathcal{A} \xrightarrow{!m_2}$, then $m_1 = m_2$.
- Configurations are *output deterministic per channel*: if $w \xrightarrow{!m_1}$, $w \xrightarrow{!m_2}$, and $\text{actor}(!m_1) = \text{actor}(!m_2)$, then $m_1 = m_2$.
- Adversaries are *channel selective*: if $\mathcal{A} \xrightarrow{?m_1}$ and $\mathcal{A} \xrightarrow{?m_2}$, then $\text{actor}(?m_1) = \text{actor}(?m_2)$.

d) Simulation: Simulation determines when a configuration W_2 securely realizes configuration W_1 : that is, if every adversary \mathcal{A} interacting with W_2 can be simulated by another adversary \mathcal{S} (with the same interface) running against W_1 [26]. The latter adversary is called a *simulator*.

Definition III.2 (Simulation). W_1 is simulated by W_2 , written $W_1 \geq W_2$, when the two systems are indistinguishable to *any* external environment:

$$\forall \mathcal{A}. \exists \mathcal{S}. \mathbb{T}_{\text{Env}}(\mathcal{S} \parallel W_1) = \mathbb{T}_{\text{Env}}(\mathcal{A} \parallel W_2)$$

Here, $\mathbb{T}_{\text{Env}}(W)$ is the set of *traces* of W but containing only the actions that communicate with the environment. Given trace $tr = a_1, \dots, a_n$, we have $\mathbb{T}_{\text{Env}}(W) = \{tr|_{\text{Env}} \mid W \xrightarrow{tr}\}$, where restriction $tr|_{\text{Env}}$ removes all actions in tr where neither the source nor the destination is Env .

Our definition of simulation guarantees *perfect* (i.e., information-theoretic) security. In §VIII, we discuss how to transfer our results to the SUC framework, which is based on computational security.

IV. SPECIFYING SECURITY POLICIES

To succinctly capture both security policies and the adversary’s power, we use a label model that can describe confidentiality and integrity simultaneously [47, 48, 49].

A security label $\ell \in \mathbb{L}$ is a pair of the form $\langle p, q \rangle$ where p and q are elements of an arbitrary bounded distributive lattice \mathbb{P} . Here, p describes confidentiality and q describes integrity. Elements of \mathbb{P} are called *principals*. Principals can be thought of as negation-free boolean formulas over a set $\{A, B, C, \dots\}$ of *atomic principals*.

The *acts-for* relation (\Rightarrow) orders principals by authority, and coincides with logical implication: for example, $p \wedge q \Rightarrow p$ and $q \Rightarrow p \vee q$. The most powerful principal is $\mathbf{0}$ and the least powerful, $\mathbf{1}$, so we have $\mathbf{0} \Rightarrow p \Rightarrow \mathbf{1}$ for any principal p .

We lift \wedge , \vee , and \Rightarrow to labels in the obvious pointwise manner. Whenever appropriate, we write p for the security label $\langle p, p \rangle$. Confidentiality and integrity projections ℓ^{\rightarrow} and ℓ^{\leftarrow} completely weaken the other component of a label: $\langle p, q \rangle^{\rightarrow} = \langle p, \mathbf{1} \rangle$ and $\langle p, q \rangle^{\leftarrow} = \langle \mathbf{1}, q \rangle$.

As in FLAM [49], the authority ordering on principals defines secure information flow. Flow policies become more restrictive as they become *more* secret and *less* trusted:

$$\begin{aligned} \langle p_1, q_1 \rangle \sqsubseteq \langle p_2, q_2 \rangle &\iff p_2 \Rightarrow p_1 \text{ and } q_1 \Rightarrow q_2 \\ \langle p_1, q_1 \rangle \sqcup \langle p_2, q_2 \rangle &= \langle p_1 \wedge p_2, q_1 \vee q_2 \rangle \\ \langle p_1, q_1 \rangle \sqcap \langle p_2, q_2 \rangle &= \langle p_1 \vee p_2, q_1 \wedge q_2 \rangle \end{aligned}$$

The least restrictive information flow policy is $\mathbf{0}^{\leftarrow}$ (“public trusted”), describing information that can be used anywhere, while the most restrictive is $\mathbf{0}^{\rightarrow}$ (“secret untrusted”).

A. Authority of Hosts

Protocol synthesis places computations on hosts that have enough authority to securely execute them. The authority of each host h is captured with a label $\mathbb{L}(h)$ [20]. For our example, we take $\mathbb{L}(\text{Alice}) = \mathbf{A}$ and $\mathbb{L}(\text{Bob}) = \mathbf{B}$. Following an insight from Viaduct [25], idealized hosts like $\text{MPC}(\text{Alice}, \text{Bob})$ have a derived label that conservatively approximates the security guarantees of the cryptographic mechanism. For maliciously secure MPC, a reasonable label is $\mathbb{L}(\text{MPC}(\text{Alice}, \text{Bob})) = \mathbb{L}(\text{Alice}) \wedge \mathbb{L}(\text{Bob}) = \mathbf{A} \wedge \mathbf{B}$, meaning that $\text{MPC}(\text{Alice}, \text{Bob})$ may view secrets of *Alice* and *Bob*, and also has enough integrity to compute values for them.

B. Capturing Attacks with Labels

The power of the adversary is determined by partitioning labels \mathbb{L} across the two axes: public/secret and trusted/untrusted; we denote these sets as \mathcal{P}/\mathcal{S} and \mathcal{T}/\mathcal{U} , respectively. We only consider sets that form *valid attacks* [34]. Intuitively, a valid attack requires that all untrusted labels are public, so that the adversary cannot modify secret data; we define valid

	Variables $x \in \mathbb{X}$	Labels $\ell \in \mathbb{L}$	Operators $f \in \mathbb{F}$
Atomic Expr.	t	$::= v \mid x$	
Expressions	e	$::= f(t_1, \dots, t_n)$ declassify $(t, \ell_f \rightarrow \ell_t)$ endorse $(t, \ell_f \rightarrow \ell_t)$ input output t	
Statements	s	$::= \text{let } h.x = e; s$ if $(h.t, s_1, s_2)$ skip	
Buffers	B	$\in \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{V}^*$	
Processes	w	$::= \langle \mathbb{H}, B, s \rangle$	

Figure 5. Syntax of the source language.

attacks formally in the technical report [50]. The rest of our development is parameterized over a valid partitioning of labels.

Recalling the threat model from §II-C, an honest host has a secret, trusted label ($\mathbb{L}(h) \in \mathcal{S} \cap \mathcal{T}$); a semi-honest host has a public, trusted label ($\mathbb{L}(h) \in \mathcal{P} \cap \mathcal{T}$); and a malicious host has a public, untrusted label ($\mathbb{L}(h) \notin \mathcal{T}$). A host with a secret, untrusted label does not make any sense: an untrusted host is fully controlled by the adversary, so it cannot hide information from the adversary. We rule out such corruptions by requiring all host labels to be *uncompromised* [51]. A valid partitioning never classifies an uncompromised label as secret and untrusted.

Definition IV.1 (Uncompromised Label). Label $\ell = \langle p, q \rangle$ is uncompromised, written $\blacktriangledown \ell$, if it is at least as trusted as it is secret: $q \Rightarrow p$.

Theorem IV.2. *Under a valid attack, if $\blacktriangledown \ell$, then we have $\ell \notin \mathcal{S} \cap \mathcal{U}$.*

V. PROTOCOL SYNTHESIS

The first program transformation, protocol synthesis, takes a sequential source program to a choreography.

A. Source Language

Figure 5 gives the syntax of source programs. The language supports an abstract set of operators f over values. We distinguish pure, atomic expressions t from expressions e that may have side effects. The **declassify** expression marks locations where private data is explicitly allowed to flow to public data. Similarly, **endorse** marks where untrustworthy data may influence trusted data. The **input/output** expressions allow programs to interact with the external environment [52, 53].

Statement **let** $h.x = e; s$ performs the local computation e at host h , binds the result to variable x , and continues as s . In source programs, h is only relevant for **input** and **output** expressions: we write **let** $h.x = \text{input}$ and **let** $h.x = \text{output } t$ in contrast to our example in fig. 2a, where we write **let** $x = \text{input } h$ and **let** $x = \text{output } t \text{ to } h$. For all other expressions, h is $*$, a single fully trusted host. Representing source programs with host annotations allows smoothly extending the language later on.

Expressions	$e ::= \dots \mid \mathbf{receive} \ h \mid \mathbf{send} \ t \ \mathbf{to} \ h$
Statements	$s ::= \dots \mid h_1.t \rightsquigarrow h_2.x; s \mid h_1[v] \rightsquigarrow h_2; s$
Processes	$w ::= \langle H \subseteq \mathbb{H}, B, s \rangle$

Figure 6. Syntax of choreographies as an extension to source syntax (fig. 5). The **send/receive** expressions are only relevant for the security proof.

$$\boxed{h.e \xrightarrow{a}_i v} \quad \frac{\ell_f \notin \mathcal{P} \quad \ell_t \in \mathcal{P}}{h.\mathbf{declassify}(v, \ell_f \rightarrow \ell_t) \xrightarrow{!hAdvv}_i v}$$

$$\frac{\ell_f \notin \mathcal{T} \quad \ell_t \in \mathcal{T}}{h.\mathbf{endorse}(v, \ell_f \rightarrow \ell_t) \xrightarrow{?Advhv'}_i v'}$$

$$\frac{\mathbb{L}(h) \in \mathcal{T}}{h.\mathbf{input} \xrightarrow{?Envhv}_i v} \quad \frac{\mathbb{L}(h) \in \mathcal{T}}{h.\mathbf{output} \ v \xrightarrow{!hEnvv}_i 0}$$

$$\boxed{s \xrightarrow{a}_r s'}$$

$$h_1.v \rightsquigarrow h_2.x; s \xrightarrow{!h_1h_2v}_r s[v/x] \quad h_1[v] \rightsquigarrow h_2; s \xrightarrow{!h_1h_2v}_r s$$

$$\boxed{s \xrightarrow{a}_\alpha^c s'} \quad \frac{s \xrightarrow{a}_\alpha^c s' \quad \mathbf{actor}(a) \notin \mathbf{hosts}(E)}{\mathbf{let} \ h.x = e; s \xrightarrow{a}_\alpha^c \mathbf{let} \ h.x = e; s'}$$

Figure 7. Select ideal, real, and concurrent stepping rules.

A source-program configuration is a logically centralized process $\langle \mathbb{H}, B, s \rangle$. The component \mathbb{H} indicates the process acts for all hosts. The second component, B , is a *buffer* mapping pairs of endpoints to first-in-first-out queues of values. Processes buffer input so that output on other processes (the adversary and the environment) is nonblocking.

B. Choreography Language

Choreographies are centralized representations of distributed computations. Unlike source programs, they make explicit the location of computations, storage, and communication.

Figure 6 gives the syntax of choreographies, which we present as an extension of the source syntax. Host annotations on **let** and **if** statements are no longer restricted to $*$ and can be any host $h \in \mathbb{H}$. The *global communication* statement $h_1.t \rightsquigarrow h_2.x; s$ represents host h_1 sending the value of t to h_2 , which stores it in variable x . The *selection* statement $h_1[v] \rightsquigarrow h_2; s$ communicates control flow decisions, and is used to establish *knowledge of choice* [28, 35]. We extend expressions with **send** and **receive**, however, the compiler never generates choreographies with these expressions; they are only used in proofs to model malicious corruption (§V-E).

As for source programs, configurations are single processes $\langle H, B, s \rangle$, but H may be a strict subset of \mathbb{H} and does not contain the ideal process $*$.

C. Operational Semantics of Choreographies

Following §III, we give operational semantics to programs using labeled transition systems. Since choreographies strictly extend the syntax of source programs, it suffices to define a semantics for choreographies.

Following fig. 1, we define two transition relations: *ideal* stepping \rightarrow_i gives meaning to source programs and to idealized choreographies (an intermediate language for our simulation proof), and *real* stepping \rightarrow_r gives meaning to choreographies and distributed programs. Additionally, we lift ideal and real stepping to concurrent versions, written \rightarrow_i^c and \rightarrow_r^c , to capture the concurrent semantics of choreographies. Figure 7 gives a selection of key rules; we defer full definitions to the technical report [50].

1) *Ideal Semantics*: We write $h.e \xrightarrow{a}_i v$ to mean expression e evaluates to value v at host h with action a . We assume operators are total; partial operators (like division) can be made total using defaults. Formally, we give meaning to operator application assuming a denotation function $\mathbf{eval} : \mathbb{F} \times \mathbb{V}^* \rightarrow \mathbb{V}$. We model **declassify** and **endorse** expressions as *interactions* with the adversary endpoint Adv. When a value is declassified from a secret label to a public one, the program *outputs* the value to Adv. Dually, when a value is endorsed from an untrusted label to a trusted one, the program takes *input* from Adv, and uses that value instead. When the confidentiality/integrity of the value does not change, these expressions act as the identity function and take internal steps. The **input/output** expressions communicate with the environment endpoint Env, except on malicious hosts; there, they step internally and always return 0. In source programs and idealized choreographies, **receive/send** expressions only communicate with malicious hosts; we suppress them (they take internal steps and always return 0) to give less power to the adversary in the ideal setting.

For statements, we write $s \xrightarrow{a}_i s'$ to mean statement s steps to s' with action a . Statement stepping rules are as expected: **let** statements step using substitution, **if** statements pick a branch based on their conditional, and communication and selection statements step internally, naming the “sending host” as the host performing the action.

2) *Real Semantics*: Real stepping rules modify ideal stepping rules. The **declassify/endorse** expressions always step internally instead of communicating with Adv. The **receive/send** expressions communicate a value with the specified host. Finally, communication and selection statements step with a visible action instead of internally.

3) *Concurrent Lifting for Choreographies*: Concurrent stepping rules, written $s \xrightarrow{a}_\alpha^c s'$, lift an underlying statement-stepping judgment (\rightarrow_i or \rightarrow_r). Concurrent stepping allows choreographies to step statements at different hosts out of program order as long as there are no dependencies between the hosts, and is the standard way choreographies model the behavior of a distributed system [28]. The key rule allows skipping over **let** statements to step a statement in the middle of a program. This rule requires the actor of the performed action to be different from the hosts of the statements being

$$\boxed{\text{src}(s) = s'}$$

$$\text{src}(\mathbf{let} \ h.x = e; s) = \begin{cases} \mathbf{let} \ h.x = e; \text{src}(s) & \text{I/O}(e) \\ \mathbf{let} \ *.x = e; \text{src}(s) & \neg \text{I/O}(e) \end{cases}$$

$$\text{src}(h_1.t \rightsquigarrow h_2.x; s) = \text{src}(s)[t/x]$$

$$\text{src}(h_1[v] \rightsquigarrow h_2; s) = \text{src}(s)$$

$$\text{src}(\mathbf{if}(h.t, s_1, s_2)) = \mathbf{if}(*.t, \text{src}(s_1), \text{src}(s_2))$$

$$\text{src}(\mathbf{skip}) = \mathbf{skip}$$

$$\text{I/O}(e) = (e = \mathbf{input}) \vee (\exists t. e = \mathbf{output} \ t)$$

Figure 8. Canonical source program from a choreography.

skipped over, matching the behavior of target programs where code running on a single host is single-threaded.

a) *Synchronous vs. Asynchronous Choreographies*: When skipping over **let** statements, requiring only the *actor* to be missing from the context leads to an *asynchronous* semantics [54]. In a *synchronous* setting, the side condition would require both endpoints to be missing: $\text{hosts}(a) \cap \text{hosts}(E) = \emptyset$. Consider the following program:

```
let Alice.x1 = input;
Bob.0  $\rightsquigarrow$  Alice.x2;
```

Alice is waiting on an input, so is not ready to receive from **Bob**. In a synchronous setting, these statements must execute in program order since **Bob** can only send if **Alice** is ready to receive. In an asynchronous setting, sends are nonblocking, so the second statement can execute first.

4) *Processes*: A buffer B behaves as a FIFO queue for each channel: it can input a message by appending the received value at the end of the corresponding queue, and can output the value at the beginning of any queue. Buffers guarantee in-order delivery within a single channel c_1c_2 , but messages across different channels may be reordered. A process w forwards its input to its buffer if the message is addressed to a relevant host; otherwise, w discards the message. A process takes an internal step when its buffer delivers a message to its statement, and an output step when its statement outputs.

D. Compiling to Choreographies

Instead of committing to a specific algorithm, we give validity criteria for the output of protocol synthesis, which generalizes our results to and beyond prior work [21, 25, 55]. Because a source program can be realized as many different choreographies, protocol synthesis cannot be modeled as a function from source programs to choreographies. Instead, we capture a valid protocol synthesis as a mapping from choreographies to source programs.

Definition V.1 (Valid Protocol Synthesis). Choreography s' is a valid result of protocol synthesis on source program s if $\text{src}(s') = s$, $\epsilon \vdash s'$, and $\Delta \Vdash s'$ for some Δ .

Figure 8 defines the function $\text{src}(\cdot)$, which maps a choreography to its canonical source program by removing communication and selection statements and replacing all host annotations

with $*$ (except those associated with **input** and **output**). The judgment $\Gamma \vdash s$ denotes that choreography s has secure information flows, and $\Delta \Vdash s$ denotes s is well-synchronized. We define these judgments next.

1) *Information-Flow Type System*: First, we give a type system for choreographies based on information-flow control [20, 21, 22, 25] which validates that hosts have enough authority to execute their assigned statements.

Figure 9 gives the typing rules. A label context Γ maps a variable to its host and label, as a set of bindings $x : h.l$. The judgment $\Gamma \vdash e : h.l$, means that e at host h has label l in the context Γ . Rule ℓ -VARIABLE ensures hosts only use variables they own. Rules ℓ -DECLASSIFY and ℓ -ENDORSE enforce nonmalleable information flow control (NMIFC) [34] by requiring source and target labels to be uncompromised [51, 34]. NMIFC requires declassified data to be trusted, enforcing *robust declassification*, and endorsed data to be public, enforcing *transparent endorsement*. These restrictions prevent the adversary from exploiting downgrades. Enforcing NMIFC is crucial for our simulation result, which we discuss in §VI-C.

In choreographies, **receive/send** expressions model communication with malicious hosts. Choreographies exclude code for malicious hosts, which exhibit arbitrary behavior; thus, labels for **receive/send** expressions must be approximated. Rule ℓ -RECEIVE ensures data coming from malicious hosts is considered untrusted; it treats the data as fully public since we do not care about preserving the confidentiality of malicious hosts. Rule ℓ -SEND ensures secret data is not sent to malicious hosts; it ignores integrity since malicious hosts are untrusted.

Statement checking rules have the form $\Gamma \vdash s$; they are largely standard [27], but do not track program counter labels since we require programs to only branch on public, trusted values. Rules ℓ -LET and ℓ -COMMUNICATE check that the host storing a variable has enough authority to do so. This is the key condition governing secure host selection and prevents, for example, **Bob**'s secret data being placed on **Alice**, or high-integrity data being placed on an untrusted host. Rule ℓ -SELECT ensures that if host h_1 informs h_2 of a branch being taken, then h_1 has at least as much integrity as h_2 . So malicious hosts cannot influence control flow on nonmalicious hosts. Finally, rule ℓ -IF requires control flow to be public and trusted.

2) *Synchronization Checking*: Next, we define a novel *synchronization* judgment, $\Delta \Vdash s$, which guarantees that all external actions in s happen in sequential program order. For example, any **endorse** statement that happens after a **declassify** must logically *depend* on the **declassify**. Since these statements may be run on different hosts, the **declassify** could happen before the **endorse**, violating program order. To prevent this, we require that the host running the **endorse** *synchronizes* with the host running the **declassify**.

Synchronization becomes more complex with *corruption*. For example, if **Alice** and **Bob** synchronize through another host h (**Alice** $\rightsquigarrow h \rightsquigarrow$ **Bob**) and h is malicious, h might give **Bob** the go-ahead before confirming with **Alice**. We use integrity labels to ensure synchronization even under corruption.

$\Gamma \vdash t : h.\ell$	$\Gamma \vdash e : h.\ell$				
ℓ -VALUE	ℓ -VARIABLE	ℓ -OPERATOR	ℓ -DECLASSIFY		
$\frac{}{\Gamma \vdash v : h.\ell}$	$\frac{\ell' \sqsubseteq \ell}{\Gamma, x : h.\ell' \vdash x : h.\ell}$	$\frac{\forall i. \Gamma \vdash t_i : h.\ell}{\Gamma \vdash f(t_1, \dots, t_n) : h.\ell}$	$\frac{\Gamma \vdash t : h.\ell_f \quad \ell_f^{\leftarrow} = \ell_t^{\leftarrow} \quad \nabla \ell_f \quad \nabla \ell_t \quad \ell_t \sqsubseteq \ell}{\Gamma \vdash \mathbf{declassify}(t, \ell_f \rightarrow \ell_t) : h.\ell}$		
ℓ -ENDORSE	ℓ -INPUT	ℓ -OUTPUT	ℓ -RECEIVE	ℓ -SEND	
$\frac{\Gamma \vdash t : h.\ell_f \quad \ell_f^{\rightarrow} = \ell_t^{\rightarrow} \quad \nabla \ell_f \quad \nabla \ell_t \quad \ell_t \sqsubseteq \ell}{\Gamma \vdash \mathbf{endorse}(t, \ell_f \rightarrow \ell_t) : h.\ell}$	$\frac{\mathbb{L}(h) \sqsubseteq \ell}{\Gamma \vdash \mathbf{input} : h.\ell}$	$\frac{\Gamma \vdash t : h.\mathbb{L}(h)}{\Gamma \vdash \mathbf{output} t : h.\ell}$	$\frac{\mathbb{L}(h')^{\leftarrow} \sqsubseteq \ell}{\Gamma \vdash \mathbf{receive} h' : h.\ell}$	$\frac{\Gamma \vdash t : h.\mathbb{L}(h')^{\rightarrow}}{\Gamma \vdash \mathbf{send} t \text{ to } h' : h.\ell}$	
$\Gamma \vdash s$					
ℓ -LET	ℓ -COMMUNICATE	ℓ -SELECT	ℓ -IF	ℓ -SKIP	
$\frac{\Gamma \vdash e : h.\ell \quad \mathbb{L}(h) \Rightarrow \ell \quad \Gamma, x : h.\ell \vdash s}{\Gamma \vdash \mathbf{let} h.x = e; s}$	$\frac{\Gamma \vdash t : h_1.\ell \quad \mathbb{L}(h_2) \Rightarrow \ell \quad \Gamma, x : h_2.\ell \vdash s}{\Gamma \vdash h_1.t \rightsquigarrow h_2.x; s}$	$\frac{\mathbb{L}(h_1)^{\leftarrow} \sqsubseteq \mathbb{L}(h_2)^{\leftarrow} \quad \Gamma \vdash s}{\Gamma \vdash h_1[v] \rightsquigarrow h_2; s}$	$\frac{\Gamma \vdash t : h.\mathbf{0}^{\leftarrow} \quad \Gamma \vdash s_1 \quad \Gamma \vdash s_2}{\Gamma \vdash \mathbf{if}(h.t, s_1, s_2)}$	$\frac{\Gamma \vdash s}{\Gamma \vdash \mathbf{skip}}$	

Figure 9. Information-flow typing rules for expressions and statements in choreographies.

$\Delta \Vdash s$	SYNC-EXTERNAL	SYNC-INTERNAL	
	$\frac{\text{external}(e) \quad \text{reset}(\Delta, h) \Vdash s \quad \text{outputting}(e) \Rightarrow \text{synched}(\Delta, h)}{\Delta \Vdash \mathbf{let} h.x = e; s}$	$\frac{\text{internal}(e) \quad \Delta \Vdash s}{\Delta \Vdash \mathbf{let} h.x = e; s}$	
	SYNC-COMMUNICATE	SYNC-SELECT	SYNC-IF
	$\frac{\text{sync}(\Delta, h_1 \rightsquigarrow h_2) \Vdash s}{\Delta \Vdash h_1.t \rightsquigarrow h_2.x; s}$	$\frac{\text{sync}(\Delta, h_1 \rightsquigarrow h_2) \Vdash s}{\Delta \Vdash h_1[v] \rightsquigarrow h_2; s}$	$\frac{\Delta \Vdash s_1 \quad \Delta \Vdash s_2}{\Delta \Vdash \mathbf{if}(h.t, s_1, s_2)}$
			SYNC-SKIP
			$\frac{}{\Delta \Vdash \mathbf{skip}}$

$\text{synched}(\Delta, h)$	$\text{reset}(\Delta, h) = \Delta'$	$\text{sync}(\Delta, h_1 \rightsquigarrow h_2) = \Delta'$
$\text{synched}(\Delta, h) = \forall h'. \Delta(h', h) \sqsubseteq \mathbb{L}(h') \vee \mathbb{L}(h)$	$\text{reset}(\Delta, h) = \Delta[h, * := \mathbf{1}][h, h := \mathbb{L}(h)]$	
$\text{sync}(\Delta, h_1 \rightsquigarrow h_2) = \Delta[* , h_2 := \Delta(*, h_2) \wedge (\Delta(*, h_1) \vee \mathbb{L}(h_2))]$		

Figure 10. Checking that a concurrent choreography has sequential behavior.

Figure 10 defines the synchronization-checking judgment $\Delta \Vdash s$. Intuitively, a choreography is well-synchronized when for any *external* (input or output) expression e , a high-integrity communication path exists from e to all *output* expressions following e in the program order.¹ Integrity of a communication path $h_1 \rightsquigarrow \dots \rightsquigarrow h_n$ is determined by the hosts in the path:

$$\mathbb{L}(h_1 \rightsquigarrow \dots \rightsquigarrow h_n) = \mathbb{L}(h_1)^{\leftarrow} \vee \dots \vee \mathbb{L}(h_n)^{\leftarrow}$$

Hosts can be malicious, so each host on the path weakens integrity, which is captured by disjunction (\vee). Multiple paths between the same hosts increase integrity, which we capture by taking the conjunction (\wedge) of path labels:

$$\mathbb{L}(\text{paths}(h_1, h_2)) = \bigwedge_{\text{path} \in \text{paths}(h_1, h_2)} \mathbb{L}(\text{path})$$

We track the integrity of paths using the context Δ , which maps pairs of hosts $\Delta(h_1, h_2)$ to the integrity label $\mathbb{L}(\text{paths}(h_1, h_2))$.

¹Input expressions are **input** and **endorse**; output expressions are **output** and **declassify**.

Rule SYNC-EXTERNAL checks a **let** statement that executes an external expression e on h . The continuation is checked under a context where the label of all paths *from* h to any other host are set to $\mathbf{1}$ (this corresponds to removing the paths), since these hosts now need to synchronize with h . Also, if e is an output expression, h must be synchronized with all hosts through the following condition, which ensures that if neither h_1 nor h_2 is malicious, a communication path exists from h_1 to h_2 that could not have been influenced by the adversary:

$$\mathbb{L}(\text{paths}(h_1, h_2)) \sqsubseteq \mathbb{L}(h_1) \vee \mathbb{L}(h_2) \quad (1)$$

Rules SYNC-COMMUNICATE and SYNC-SELECT update Δ using $\text{sync}(\Delta, h_1 \rightsquigarrow h_2)$. The function captures that a path from some h to h_1 implies there is a path from h to h_2 that goes through h_1 . Further, all existing paths are still valid.

E. Modeling Malicious Corruption

Malicious hosts are fully controlled by the adversary. Following UC [26], we entirely remove processes that correspond to malicious hosts in hybrid distributed programs, and allow

$$\langle\!\langle s \rangle\!\rangle = s'$$

$$\begin{aligned} \langle\!\langle \text{let } h.x = e; s \rangle\!\rangle &= \begin{cases} \text{let } h.x = e; \langle\!\langle s \rangle\!\rangle & h \in \mathcal{T} \\ \langle\!\langle s \rangle\!\rangle & \text{o/w} \end{cases} \\ \langle\!\langle h_1.t \rightsquigarrow h_2.x; s \rangle\!\rangle &= \begin{cases} h_1.t \rightsquigarrow h_2.x; \langle\!\langle s \rangle\!\rangle & h_1, h_2 \in \mathcal{T} \\ \text{let } h_1._ = \text{send } t \text{ to } h_2; \langle\!\langle s \rangle\!\rangle & h_1 \in \mathcal{T} \\ \text{let } h_2.x = \text{receive } h_1; \langle\!\langle s \rangle\!\rangle & h_2 \in \mathcal{T} \\ \langle\!\langle s \rangle\!\rangle & \text{o/w} \end{cases} \\ \langle\!\langle h_1[v] \rightsquigarrow h_2; s \rangle\!\rangle &= \begin{cases} h_1[v] \rightsquigarrow h_2; \langle\!\langle s \rangle\!\rangle & h_1, h_2 \in \mathcal{T} \\ \text{let } h_1._ = \text{send } v \text{ to } h_2; \langle\!\langle s \rangle\!\rangle & h_1 \in \mathcal{T} \\ \langle\!\langle s \rangle\!\rangle & \text{o/w} \end{cases} \\ \langle\!\langle \text{if}(h.t, s_1, s_2) \rangle\!\rangle &= \begin{cases} \text{if}(h.t, \langle\!\langle s_1 \rangle\!\rangle, \langle\!\langle s_2 \rangle\!\rangle) & h \in \mathcal{T} \\ \perp & \text{o/w} \end{cases} \\ \langle\!\langle \text{skip} \rangle\!\rangle &= \text{skip} \end{aligned}$$

$$\langle\!\langle w \rangle\!\rangle = w' \quad \langle\!\langle H, B, s \rangle\!\rangle = \langle\!\langle \{h \in H \mid h \in \mathcal{T}\}, B, \langle\!\langle s \rangle\!\rangle \rangle\!\rangle$$

Figure 11. Modeling malicious corruption. We write $h \in \mathcal{T}$ for $\mathbb{L}(h) \in \mathcal{T}$.

the adversary to forge arbitrary messages in their stead. This is reflected in choreographies by rewriting them to elide statements that involve malicious hosts.

Figure 11 defines the *corruption* $\langle\!\langle s \rangle\!\rangle$ of a choreography s . The operation considers each statement in turn. If *all* hosts involved in a statement are nonmalicious, the statement stays as is. If *all* hosts involved in a statement are malicious, the statement is removed entirely. Otherwise, only *some* involved hosts are malicious, and we rewrite the statement. Communication statements become either a **send** or a **receive**, depending on whether the sending or the receiving host is nonmalicious. Selection statements are similar, except we cannot have a malicious sending host and a nonmalicious receiving host (rule ℓ -SELECT). Similarly, **if** statements cannot be at a malicious host since we require trusted control flow (rule ℓ -IF).

VI. CORRECTNESS OF PROTOCOL SYNTHESIS

We prove the correctness of protocol synthesis by demonstrating a simulation between source programs and their corresponding choreographies. For $w = \langle H, B, s \rangle$, we write $\Gamma \vdash w$ and $\Delta \Vdash w$ if $\Gamma \vdash s$ and $\Delta \Vdash s$, respectively, and define $\text{src}(w) = \langle H, B, \text{src}(s) \rangle$.

Theorem VI.1. *If $\epsilon \vdash w$, and $\Delta \Vdash w$ for some Δ , then $\langle\!\langle \text{src}(w), \rightarrow_i \rangle\!\rangle \geq \langle\!\langle w, \rightarrow_r \rangle\!\rangle$.*

We prove theorem VI.1 through a series of intermediate simulations, following fig. 12 from left to right. First, in §VI-A, we show idealized, sequential choreographies simulate their canonical source programs. Then, we show in §VI-B that our synchronization judgment ensures all externally visible actions happen in program order. Finally, in §VI-C, we move from the ideal semantics \rightarrow_i^c to the real semantics \rightarrow_r^c .

For each simulation, we define a simulator that emulates the adversary “in its head”. We ensure that the emulated adversary’s view is the same as the real adversary even though the simulator only has access to public information. Concretely, we establish

a (weak) bisimulation relation [56, 57] between the ideal world (simulator running against ideal configuration) and the real world (adversary running against real configuration).

A. Correctness of Host Selection

First, we show that the original source program is simulated by the sequential choreography:

Theorem VI.2. *If $\epsilon \vdash w$, then $\langle\!\langle \text{src}(w), \rightarrow_i \rangle\!\rangle \geq \langle\!\langle w, \rightarrow_i \rangle\!\rangle$.*

This is shown via two simpler simulations. First, we add host annotations and explicit communication to the source program:

Lemma VI.3. *If $\epsilon \vdash w$, then $\langle\!\langle \text{src}(w), \rightarrow_i \rangle\!\rangle \geq \langle\!\langle w, \rightarrow_i \rangle\!\rangle$.*

Proof. Statements removed by $\text{src}(\cdot)$ only produce internal actions, which the simulator can recreate. Host annotations affect program behavior only by changing the source and destination of internal actions and actions generated by **declassify/endorse**; the simulator must recover the original host names before forwarding messages from/to the adversary.

The simulator maintains a public view of w and runs the adversary against this view. When the adversary steps w , the simulator steps $\text{src}(w)$ only if the statement is preserved by $\text{src}(\cdot)$; it does nothing otherwise. To handle **declassify**, whenever the simulator receives a message of the form $*\text{Adv}v$, the simulator inspects its copy of w to determine the sending host h , and sends $h\text{Adv}v$ to the adversary instead. Similarly, to handle **endorse**, the simulator replaces h with $*$ in messages $\text{Adv}hv$ from the adversary. \square

Second, we add corruption, obtaining a choreography which steps sequentially:

Lemma VI.4. *If $\epsilon \vdash w$, then $\langle\!\langle w, \rightarrow_i \rangle\!\rangle \geq \langle\!\langle w, \rightarrow_i \rangle\!\rangle$.*

Proof. Corruption only removes statements at malicious hosts, however, these statements only generate internal actions: **input/output** expressions always step internally using ideal rules, and typing ensures **declassify/endorse** expressions step internally. This means $\langle\!\langle w \rangle\!\rangle$ and w have the same external behavior, except w takes extra internal steps. The simulator follows the control flow and acts like the adversary, but whenever the adversary schedules $\langle\!\langle w \rangle\!\rangle$, the simulator schedules w multiple times until the head statement is at a nonmalicious host; then, it schedules w again.

A small caveat: in $\langle\!\langle w \rangle\!\rangle$, all data from malicious hosts is explicitly replaced with 0, whereas malicious hosts may store arbitrary data in w . Since data from malicious hosts is untrusted, our type system ensures this data does not influence trusted data, which includes all output messages. Formally, we only require and maintain that $\langle\!\langle w \rangle\!\rangle$ and w agree on *trusted* values. \square

B. Correctness of Sequentialization

Next, we show that if the choreography is well-synchronized, then the choreography stepping sequentially is simulated by itself stepping concurrently:

Theorem VI.5. *If $\epsilon \vdash w$, and $\Delta \Vdash w$ for some Δ , then $\langle\!\langle w, \rightarrow_i \rangle\!\rangle \geq \langle\!\langle w, \rightarrow_i^c \rangle\!\rangle$.*




Figure 12. Intermediate simulation steps for proving the correctness of protocol synthesis.

The adversary, interacting with the concurrent version of the choreography, can schedule a statement that is not next in program order. If the statement produces an externally visible action, the simulator must schedule the same statement. Since the simulator interacts with the sequential version, it must “unwind” the choreography by scheduling every statement leading up to the desired statement. Synchronization ensures unwinding does not fail due to a statement blocked on input (**input** or **endorse**), or a statement that performs a different visible action (**output** or **declassify**).

The concurrent and sequential choreographies necessarily fall out of sync during simulation: the adversary may schedule steps for the concurrent choreography that the simulator cannot immediately match, and the simulator might schedule steps for the sequential choreography while unwinding, steps the adversary did not schedule. Nevertheless, the two choreographies remain *joinable*: they can reach a common choreography via only internal actions. We prove choreographies are *confluent* [58, 59, 60], which ensures joinable processes remain joinable throughout the simulation.

Proof sketch for theorem VI.5. The simulator maintains a public view of the concurrent process, and runs the adversary against this view. When the adversary schedules an output, the simulator schedules the sequential process until it performs the same output; the simulator does nothing for input and internal actions. Well-synchronization guarantees the sequential program can perform the output. The primary invariant, that the concurrent and sequential processes remain joinable, is ensured by confluence. See the technical report [50] for details. \square

C. Correctness of Ideal Execution

Finally, we move from a concurrent choreography stepping with ideal rules to a concurrent choreography stepping with real rules:

Theorem VI.6. *If $\epsilon \vdash w$, then $\langle \langle w \rangle, \rightarrow_i^c \rangle \geq \langle \langle w \rangle, \rightarrow_r^c \rangle$.*

The main difference between the two semantics is the interface with the adversary. In the real semantics, dishonest hosts actively leak data to the adversary (through **send** expressions and communication statements), and the adversary controls all data coming from malicious hosts (through **receive**). In contrast, the ideal semantics interacts with the adversary only via **declassify** and **endorse**. In effect, the ideal semantics causes leakage and corruption to become coarse-grained. Additionally, by eliminating all blocking **receive** expressions (which communicate with the adversary), the ideal semantics

can make progress in a manner independent of the adversary; this aids the sequentialization proof in §VI-B.

To bridge the gap between the real and ideal semantics, we show that the simulator can use **declassify** expressions to recreate all data no longer leaked through communication, and **endorse** expressions to corrupt all data no longer corruptible through **receive** expressions. This is possible because our type system ensures secrets are not directly sent to dishonest hosts (rules ℓ -SEND and ℓ -COMMUNICATE), and data from malicious hosts cannot directly influence trusted data (rule ℓ -RECEIVE).

Proof sketch for theorem VI.6. The simulator maintains a public view of the real process, and runs the adversary against this view. The simulator flips the input/output behavior of **declassify** and **endorse** expressions: when the ideal process *outputs* data through a **declassify** expression, the simulator *inputs* this data instead. Similarly, the simulator *outputs* data with **endorse** expressions, which it sends to the ideal process. The key invariant is that the simulator’s version of the process matches the real one on *public values*, and the ideal process matches the real one on *trusted values*.

This invariant is strong enough to witness simulatability. Since the simulator’s version of the process matches the real one on public values, the adversary in the real configuration has a view identical to the adversary running inside of the simulator (the adversary only sees public data). Similarly, since the real process matches the ideal one on trusted values, the environment has the same view in both (the environment is only sent trusted data).

Next, we argue that our simulator can preserve the invariant. For the simulator to have an accurate view of public values in the real process, the ideal process must output values through **declassify** expressions that match the values in the real process. The information-flow type system provides necessary restrictions. *Robust declassification* [34] guarantees only trusted values are declassified. Since the ideal process matches the real one on trusted values, data coming from **declassify** expressions is accurate. Dually, when the simulator sends data to the ideal process through **endorse** expressions, the values must match the ones in the real process. *Transparent endorsement* [34] guarantees only public values are endorsed. Thus, the result follows from the simulator holding accurate public values. See the technical report [50] for details. \square

VII. ENDPOINT PROJECTION

The second stage of compilation, endpoint projection, transforms a choreography into a distributed program.

Expressions	$e ::= \dots \mid \mathbf{receive} \ h \mid \mathbf{send} \ t \ \mathbf{to} \ h$
Statements	$s ::= \dots \mid \mathbf{case} \ (h_1 \rightsquigarrow h_2) \ \{v \mapsto s_v\}_{v \in V}$
Processes	$w ::= \langle \{h \in \mathbb{H}\}, B, s \rangle$

Figure 13. Hybrid distributed syntax as an extension to source (fig. 5).

A. Hybrid Distributed Language

As for choreographies, fig. 13 specifies the syntax of hybrid distributed programs by extending the syntax of source programs. Hybrid distributed programs are configurations containing multiple processes, with each process acting on behalf of a single host $h \in \mathbb{H}$. These processes communicate data via **send/receive** and agree on control flow via **send** and **case**. The **case** statement, on receiving a value on the expected channel, steps to the specified branch.

Operational Semantics: Each process transitions using the real stepping rules (\rightarrow_r), and the distributed configuration steps using the parallel composition rules in fig. 4.

B. Compiling to Distributed Programs

Given a choreography s and a host h , the endpoint projection $\llbracket s \rrbracket_h$ defines the local program that h should run. The distributed system $\llbracket s \rrbracket$ is derived by independently projecting onto each host in the choreography.

Our notion of endpoint projection is entirely standard, so we defer the formal definition to the technical report [50]. Our proof is agnostic to how endpoint projection is defined, and only relies on its soundness and completeness, properties extensively studied in prior work [28, 35, 36, 43, 44].

C. Correctness of Endpoint Projection

Let $\langle W \rangle$ remove from W all processes for malicious hosts.

Theorem VII.1. *If $\epsilon \vdash w$, then $\langle \langle w \rangle, \rightarrow_r^c \rangle \geq \langle \langle \llbracket w \rrbracket \rangle, \rightarrow_r \rangle$.*

A choreography and its endpoint projection match each other action-for-action; once we prove this fact, showing simulation is trivial since we can pick $\mathcal{S} = \mathcal{A}$. This perfect correspondence between a choreography and its projection is studied extensively in the literature [28, 35, 36, 43, 44], and formalized as *soundness* and *completeness* of endpoint projection. However, the standard methods of proving soundness and completeness must be modified to handle malicious corruption and asynchronous communication. Existing work relates w to $\llbracket w \rrbracket$, which we must extend to relate $\langle w \rangle$ to $\langle \llbracket w \rrbracket \rangle$. This is trivial since $\langle \llbracket w \rrbracket \rangle = \langle \langle w \rangle \rangle$.

The presence of *asynchrony* breaks the perfect correspondence between the projected program and the choreography: a **send/receive** pair reduces in two steps in a projected program, but the corresponding communication statement reduces in only one. We follow prior work [54, 61] and add syntactic forms to choreographies for partially reduced **send/receive** pairs: messages that have been sent and buffered but not yet received. These run-time terms exist only to restore the correspondence, and are never generated by the compiler. An additional, simple simulation then shows that a choreography with these run-time

terms simulates one without, removing the need to reason about run-time terms in other proof steps. For details, see the technical report [50].

VIII. CRYPTOGRAPHIC INSTANTIATION

Our simulation result is a necessary and novel first step toward constructing a verified, secure compiler for distributed protocols that use cryptography. We have abstracted all cryptographic mechanisms into idealized hosts (e.g., $\mathbf{MPC}(\mathbf{Alice}, \mathbf{Bob})$); thus, to achieve a full end-to-end security proof, these idealized hosts must be securely instantiated with cryptographic subprotocols (e.g., BGW [62] for multiparty computation). Such an instantiation would imply UC security for all compiled programs, in contrast to existing formalization efforts for individual protocols [63, 64, 65].

To this end, we show how the distributed protocols arising from our compilation correspond to hybrid protocols in the Simplified Universal Composability (SUC) framework [30]. Then, we show how to take advantage of the *composition* theorem in SUC to obtain secure, concrete instantiations of cryptographic protocols.

a) *Simplified UC:* Let s be a choreography with partitioning $\llbracket s \rrbracket$. We construct a corresponding SUC protocol $\llbracket s \rrbracket^{\text{SUC}}$ which behaves identically to the partitioned choreography, with minor differences due to the differing computational models. Each host in s is either a *local* host (e.g., \mathbf{Alice}), or an *idealized* host standing in for cryptography, such as $\mathbf{MPC}(\mathbf{Alice}, \mathbf{Bob})$. Local hosts map onto SUC parties, while idealized hosts map onto *ideal functionalities* in SUC.

Protocol execution in SUC happens through *activations* scheduled by the adversary: a party runs for some steps, delivers messages to a central *router*, and cedes execution to the adversary. To faithfully capture the behavior of host h in $\llbracket s \rrbracket$, the party/functionality for h in $\llbracket s \rrbracket^{\text{SUC}}$ is essentially a *wrapper* around the projected host $\llbracket s \rrbracket_h$, who steps $\llbracket s \rrbracket_h$ accordingly and forwards correct messages to the router.

Each wrapper needs to explicitly model *corruption*, which our framework captures by labels: if host h is semi-honest ($\mathbb{L}(h) \in \mathcal{P} \cap \mathcal{T}$), the wrapper for h allows the adversary to query h for its current message transcript so far. Similarly, if h is malicious ($\mathbb{L}(h) \notin \mathcal{T}$), the wrapper for h should enable the adversary to take complete control over h . By using labels to model corruption, we model *static* security in SUC.

b) *Communication Model:* In SUC, all messages between local hosts are fully public, while messages between hosts and functionalities contain *public headers* (e.g., the source/destination addresses) and *private content* (the message payload). In our system, we do not stratify message privacy along the party/functionality axis, but rather along the information flow lattice: the adversary can read the messages intended for semi-honest hosts, and can forge messages from malicious hosts. Indeed, information flow policies allow more flexible security policies for communication.

However, we can encode our communication model into SUC with the aid of additional functionalities. To do so, we make use of a secure channel functionality \mathcal{R}_{sec} , which guarantees

in-order message delivery and enables secret communication between honest hosts. We can realize \mathcal{R}_{sec} in SUC via a standard subprotocol using a public key infrastructure.

For ideal functionalities in $\llbracket s \rrbracket^{\text{SUC}}$, we need to ensure that they only communicate with local hosts, and not with other ideal functionalities. This property is preserved by compilation, so we only need to ensure that host selection produces a choreography s that has this property. Indeed, our synchronization judgment $\Delta \Vdash s$ makes it possible for choreographies to stay well-synchronized, even when the ideal hosts do not communicate with each other.

c) Adversaries and Environments: In our framework, we prove perfect security against non-probabilistic adversaries. However, allowing the adversary to use probability (as in SUC) does not weaken our simulation result.² Additionally, in UC/SUC, the environment is given by a concurrently running process that outputs a *decision bit*, whereas our model uses a trace semantics to model the environment. Security for the latter easily implies the former, since our simulation result proves equality of environment views between the two worlds.

Finally, UC (and SUC) require all parties—the adversary, environment, simulator, and hosts—to run in polynomial time since cryptographic schemes can be broken given enough time. Since our simulators query the adversary and emulate source and target programs in a straightforward manner, all simulators we define run in time bounded by a polynomial in the run times of the adversary and the source and target programs.

A. Secure Instantiation of Cryptography

To securely instantiate cryptographic mechanisms, we appeal to the *composition* theorem in SUC, which says that ideal SUC-functionalities \mathcal{F} may be substituted for SUC protocols that securely realize them. Concrete cryptographic protocols are obtained by applying this theorem iteratively to each ideal host.

Ideal hosts in our model correspond closely to the broad class of *reactive, deterministic straight-line* functionalities in SUC, including MPC [30, 66] and Zero-Knowledge Proofs (ZKP) [67]. The main difference is that our model allows the adversary to corrupt ideal functionalities (both semi-honestly and maliciously), while SUC functionalities are incorruptible. However, we guarantee that the adversary does not gain more power in our model by restricting the possible corruption models via authority labels for ideal hosts.

For example, we have $\text{MPC}(\text{Alice}, \text{Bob})$ has label $A \wedge B$, meaning that $\text{MPC}(\text{Alice}, \text{Bob})$ is semi-honest (resp. malicious) only if *both* Alice and Bob are semi-honest (resp. malicious). Thus, any power the adversary gains in corrupting $\text{MPC}(\text{Alice}, \text{Bob})$ can be instead achieved using Alice and Bob alone. Similar security concerns for label-based host selection have been discussed for Viaduct [25]. We can formalize this intuition via a simulation of the form $\langle W, \rightarrow_i \rangle \leq \langle W, \rightarrow'_i \rangle$, where W uses $\text{MPC}(\text{Alice}, \text{Bob})$, and \rightarrow'_i is modified so that corruption of $\text{MPC}(\text{Alice}, \text{Bob})$ is impossible.

²The dummy adversary theorem [26] implies that security against non-probabilistic adversaries guarantees security against probabilistic adversaries.

IX. SECURITY PRESERVATION

We use simulation to define the correctness of compilation, and show that it corresponds to a well-studied correctness criterion, *robust hyperproperty preservation* (RHP) [31]. RHP states that hyperproperties [37] satisfied by source programs under any context are also satisfied by target programs under any context. RHP is important because common notions of information-flow security such as termination-insensitive noninterference, observational determinism [68], and nonmalleable information flow control [34] are hyperproperties. With a compiler that satisfies RHP, one only needs to prove security of source programs; security of target programs immediately follows.

Definition IX.1 (Robust Hyperproperty Preservation (RHP)). Let \downarrow be a compiler from a source program to a target program, \bowtie be an operator that composes a program with its context, and \mathbb{B} be a behavior function that returns the set of possible traces generated from a whole program (i.e., a program composed with a context). Then \downarrow satisfies RHP over source program set \mathbb{P} , source context set \mathbb{C}_S , and target context set \mathbb{C}_T when, given program $P \in \mathbb{P}$, for all $C_T \in \mathbb{C}_T$ there exists $C_S \in \mathbb{C}_S$ such that $\mathbb{B}(C_S \bowtie P) = \mathbb{B}(C_T \bowtie P \downarrow)$.³

Patrignani et al. [38, 39] previously observed a correspondence between UC simulation and robust hyperproperty preservation; it also holds for our notion of simulation.

Theorem IX.2 (Simulation Implies RHP). *Define $\mathcal{A} \bowtie W = \mathcal{A} \parallel W$. Then given behavior function $\mathbb{B}(\cdot) = \mathbb{T}_{\text{Env}}(\cdot)$ and an operator \downarrow between configurations such that $W \downarrow \leq W$ for any configuration $W \in \mathbb{W}$, we have that \downarrow satisfies RHP over source program set \mathbb{W} and source and target context set $\{\mathcal{A}\}$.*

Corollary IX.3 (Partitioning Satisfies RHP). *The function $\lambda w. \langle \llbracket w \rrbracket, \rightarrow_r \rangle$ satisfies RHP over source and target context set $\{\mathcal{A}\}$ and source program set*

$$\{\langle \text{src}(w), \rightarrow_i \rangle \mid \epsilon \vdash w, \Delta \Vdash w\}$$

Proof. Theorem IX.2 follows from definitions III.2 and IX.1. Corollary IX.3 follows from theorems VI.1 and IX.2. \square

X. RELATED WORK

Secure Program Partitioning and Compilers for Secure Computation: Prior work on secure program partitioning [69, 20, 21, 25] focuses largely on the engineering effort of compiling security-typed source programs to distributed code with the aid of cryptography. Liu et al. [70] does give an informal UC simulation proof of a compiler, but it is limited to two party semi-honest MPC and oblivious RAM, and they do not consider integrity.

There is also a large literature on compilers for secure computation that target specific cryptographic mechanisms such as MPC [9, 11, 71, 12, 72], ZKP [73, 13, 14], and homomorphic encryption [18, 17, 19, 74, 75]. Again, the

³This is the “property-free” definition of RHP as given by Abate et al. [31]. An equivalent but more direct definition is that \downarrow satisfies RHP given that if for some hyperproperty HP it is the case that $\mathbb{B}(C_S \bowtie P) \in \text{HP}$ for any $C_S \in \mathbb{C}_S$, then $\mathbb{B}(C_T \bowtie P \downarrow) \in \text{HP}$ for any $C_T \in \mathbb{C}_T$.

focus of this literature is efficient implementations, not formal guarantees.

A long line of work [23, 24, 76, 77, 78, 79] focuses on enforcing computational noninterference for information-flow typed programs by using standard cryptographic mechanisms, such as encryption. However, computational noninterference guarantees little in the presence of downgrading. In contrast, our compiler enjoys *simulation-based security*, which guarantees preservation of *all* hyperproperties, even for programs using declassification and endorsement.

Our compilation model is closest to that of Viaduct [25], which like this work also approximates security guarantees of cryptographic mechanisms with information-flow labels. However, this work differs from Viaduct, and from most of the literature on secure program partitioning and compilers for secure computation, in that our goal is to provide a *model* of a secure end-to-end compilation process, not to provide an implementation of an actual compiler. Our model currently does not analyze a source language with functions, loops, or mutable arrays, which Viaduct supports. Additionally, our model does not capture some minor subtleties of Viaduct, such as allowing secret-dependent conditionals whenever they may be eliminated via multiplexing the two branches. An interesting research direction would be to close the gap between our model and existing compilers by providing a verified compiler implementation akin to CompCert [80].

Simulation-based Security: Simulation-based cryptographic frameworks, such as Universal Composability [26], Reactive Simulatability [42], and Constructive Cryptography [81], allow modular proofs of distributed cryptographic protocols, and Liao et al. [82] give a core language for formalizing UC protocols. We abstract away concrete cryptography, so we do not explicitly model some subtleties of these systems: probability, computational complexity, and cryptographic hardness assumptions. But our approach should be compatible with these frameworks.

Prior verification efforts [63, 83, 64] show simulation-based security for concrete cryptographic mechanisms. Our work is orthogonal: simulation-based security for compiler correctness, rather than proofs for individual protocols.

Secure Compilation: Standard notions of compiler correctness are derived from full abstraction and hyperproperty preservation [31]. Patrignani et al. [38, 39] argue that robust hyperproperty preservation and Universal Composability are directly analogous. We affirm this hypothesis by proving that our simulation-based security result guarantees RHP. To our knowledge, we are the first to make this connection formally.

Choreographies: The use of choreographies is central to our compilation process and to the proof of its correctness. The primary concern in the extensive literature on choreographies [28, 35, 36, 43, 44], is proving deadlock freedom; very little prior work considers security [84]. Our extension of choreographies with an information-flow type system, modeling semi-honest and malicious corruption, is novel.

XI. CONCLUSION AND FUTURE WORK

This work presents a novel simulation-based security result for a compiler from sequential source programs to distributed programs, using idealizations of cryptographic mechanisms. Our simulation result guarantees that the security properties of source programs are preserved in the compiled protocols.

We believe this work opens up many opportunities for future research; in particular, it gives a clear path toward building a fully verified cryptographic compiler. Aside from adding more language features (e.g., subroutines, loops, and mutable arrays) to our source language to better match real-world cryptographic compilers, the remaining verification effort largely consists of verifiably instantiating all abstract components we assume. For example, we model protocol selection via abstract judgments on choreographies; in turn, outputs of a concrete protocol selection algorithm must satisfy these judgments. Additionally, we assume that parties communicate over secure channels, and use ideal functionalities for performing cryptography; these assumptions can be eliminated using cryptographic instantiations verified in the Simplified UC framework.

Our security result holds for a strong attacker that knows when communication occurs between hosts and controls its scheduling, but secret control flow is therefore precluded. To improve expressive power, weaker attacker models should be explored.

We have focused on confidentiality and integrity properties; however, some protocols also explicitly address availability [85], which would be another interesting direction for exploration.

XII. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under NSF grant 1704788. Any opinions, findings, conclusions or recommendations are those of the authors and do not necessarily reflect the views of any of these funders. We thank the reviewers for their thoughtful suggestions. We also thank Silei Ren, Ethan Cecchetti, Drew Zagieboylo, and Suraaj Kanniwadi for discussions and feedback.

REFERENCES

- [1] L. Lamport, “The part-time parliament,” *ACM Trans. on Computer Systems*, vol. 16, no. 2, pp. 133–169, May 1998. Available: <http://doi.acm.org/10.1145/279227.279229>
- [2] M. Castro and B. Liskov, “Practical Byzantine fault tolerance and proactive recovery,” *ACM Trans. on Computer Systems*, vol. 20, p. 2002, 2002.
- [3] F. McKeen, I. Alexandrovich, A. Berenzon, C. Rozas, H. Shafi, V. Shanbhogue, and U. Savagaonkar, “Innovative instructions and software model for isolated execution,” in *Workshop on Hardware and Architectural Support for Security and Privacy*, R. B. Lee and W. Shi, Eds. ACM, 2013, p. 10.
- [4] A. Gollamudi, S. Chong, and O. Arden, “Information flow control for distributed trusted execution environments,” in *32nd IEEE Computer Security Foundations Symp. (CSF)*. IEEE, 2019, pp. 304–318. Available: <https://doi.org/10.1109/CSF.2019.00028>
- [5] V. Costan, I. Lebedev, and S. Devadas, “Sanctum: Minimal hardware extensions for strong software isolation,” in *USENIX Security Symposium*, 2016.
- [6] M. Egele, D. Brumley, Y. Fratantonio, and C. Kruegel, “An empirical study of cryptographic misuse in Android applications,”

- in *ACM Conf. on Computer and Communications Security (CCS)*, 2013, pp. 73–84. Available: <http://doi.acm.org/10.1145/2508859.2516693>
- [7] M. Georgiev, S. Iyengar, S. Jana, R. Anubhai, D. Boneh, and V. Shmatikov, “The most dangerous code in the world: validating SSL certificates in non-browser software,” in *19th ACM Conf. on Computer and Communications Security (CCS)*, T. Yu, G. Danezis, and V. D. Gligor, Eds. ACM, 2012, pp. 38–49. Available: <https://doi.org/10.1145/2382196.2382204>
- [8] C. Decker and R. Wattenhofer, “Bitcoin transaction malleability and MtGox,” in *19th European Symposium on Research in Computer Security*. Springer, 2014, pp. 313–326.
- [9] D. Malkhi, N. Nisan, B. Pinkas, and Y. Sella, “Fairplay—a secure two-party computation system,” in *13th Usenix Security Symposium*, Aug. 2004, pp. 287–302. Available: <http://www.usenix.org/publications/library/proceedings/sec04/tech/malkhi.html>
- [10] C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi, “ObliVM: A programming framework for secure computation,” in *25th ACM Symp. on Operating System Principles (SOSP)*. IEEE, 2015, pp. 359–376. Available: <https://doi.org/10.1109/SP.2015.29>
- [11] A. Aly, D. Cozzo, M. Keller, E. Orsini, D. Rotaru, P. Scholl, N. P. Smart, and T. Wood, *SCALE-MAMBA v1.6 : Documentation*, 2019. Available: <https://homes.esat.kuleuven.be/~nsmart/SCALE>
- [12] A. Rastogi, M. A. Hammer, and M. Hicks, “Wysteria: A programming language for generic, mixed-mode multiparty computations,” in *IEEE Symp. on Security and Privacy*, May 2014, pp. 655–670.
- [13] B. Parno, J. Howell, C. Gentry, and M. Raykova, “Pinocchio: Nearly practical verifiable computation,” in *IEEE Symp. on Security and Privacy*. IEEE, 2013, pp. 238–252.
- [14] C. Costello, C. Fournet, J. Howell, M. Kohlweiss, B. Kreuter, M. Naehrig, B. Parno, and S. Zahur, “Geppetto: Versatile verifiable computation,” in *IEEE Symp. on Security and Privacy*. IEEE, 2015, pp. 253–270.
- [15] R. S. Wahby, S. Setty, Z. Ren, A. J. Blumberg, and M. Walfish, “Efficient RAM and control flow in verifiable outsourced computation,” in *Network and Distributed System Security Symp.* The Internet Society, 2015. Available: <https://www.ndss-symposium.org/ndss2015/efficient-ram-and-control-flow-verifiable-outsourced-computation>
- [16] A. Kosba, C. Papamanthou, and E. Shi, “xJsnark: A framework for efficient verifiable computation,” in *IEEE Symp. on Security and Privacy*. IEEE, 2018, pp. 944–961.
- [17] R. Dathathri, O. Saarikivi, H. Chen, K. Laine, K. E. Lauter, S. Maleki, M. Musuvathi, and T. Mytkowicz, “CHET: an optimizing compiler for fully-homomorphic neural-network inferencing,” in *40th ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)*, K. S. McKinley and K. Fisher, Eds. ACM, 2019, pp. 142–156.
- [18] R. Dathathri, B. Kostova, O. Saarikivi, W. Dai, K. Laine, and M. Musuvathi, “EVA: an encrypted vector arithmetic language and compiler for efficient homomorphic computation,” in *41st ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)*, A. F. Donaldson and E. Torlak, Eds. ACM, 2020, pp. 546–561.
- [19] M. Cowan, D. Dangwal, A. Alaghi, C. Trippel, V. T. Lee, and B. Reagen, “Porcupine: A synthesizing compiler for vectorized homomorphic encryption,” in *42nd ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)*, 2021, pp. 375–389.
- [20] S. Zdancewic, L. Zheng, N. Nystrom, and A. C. Myers, “Secure program partitioning,” *ACM Trans. on Computer Systems*, vol. 20, no. 3, pp. 283–328, Aug. 2002. Available: <http://www.cs.cornell.edu/andru/papers/sosp01/spp-tr.pdf>
- [21] L. Zheng, S. Chong, A. C. Myers, and S. Zdancewic, “Using replication and partitioning to build secure distributed systems,” in *IEEE Symp. on Security and Privacy*, May 2003, pp. 236–250. Available: <http://www.cs.cornell.edu/andru/papers/sp03.pdf>
- [22] S. Chong, J. Liu, A. C. Myers, X. Qi, K. Vikram, L. Zheng, and X. Zheng, “Secure web applications via automatic partitioning,” in *21st ACM Symp. on Operating System Principles (SOSP)*, Oct. 2007, pp. 31–44. Available: <http://www.cs.cornell.edu/andru/papers/swift-sosp07.pdf>
- [23] C. Fournet and T. Rezk, “Cryptographically sound implementations for typed information-flow security,” in *35th ACM Symp. on Principles of Programming Languages (POPL)*, Jan. 2008, pp. 323–335. Available: <https://doi.org/10.1145/1328438.1328478>
- [24] C. Fournet, G. le Guernic, and T. Rezk, “A security-preserving compiler for distributed programs: From information-flow policies to cryptographic mechanisms,” in *16th ACM Conf. on Computer and Communications Security (CCS)*, Nov. 2009, pp. 432–441. Available: <https://doi.org/10.1145/1653662.1653715>
- [25] C. Acay, R. Recto, J. Gancher, A. Myers, and E. Shi, “Viaduct: An extensible, optimizing compiler for secure distributed programs,” in *42nd ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)*. ACM, Jun. 2021, pp. 740–755.
- [26] R. Canetti, “Universally composable security: A new paradigm for cryptographic protocols,” in *42nd Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society, 2001, pp. 136–145.
- [27] A. Sabelfeld and A. C. Myers, “Language-based information-flow security,” *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 1, pp. 5–19, Jan. 2003. Available: <http://www.cs.cornell.edu/andru/papers/jsac/sm-jsac03.pdf>
- [28] F. Montesi, *Introduction to Choreographies*. Cambridge: Cambridge University Press, 2023.
- [29] A. Bakst, K. von Gleissenthall, R. G. Kıcı, and R. Jhala, “Verifying distributed programs via canonical sequentialization,” *Proc. ACM Program. Lang.*, vol. 1, no. OOPSLA, Oct. 2017.
- [30] R. Canetti, A. Cohen, and Y. Lindell, “A simpler variant of universally composable security for standard multiparty computation,” *Cryptology ePrint Archive*, Report 2014/553, 2014. Available: <https://eprint.iacr.org/2014/553>
- [31] C. Abate, R. Blanco, D. Garg, C. Hritcu, M. Patrignani, and J. Thibault, “Journey beyond full abstraction: Exploring robust property preservation for secure compilation,” in *32nd IEEE Computer Security Foundations Symp. (CSF)*. IEEE Computer Society, 2019, pp. 256–271. Available: <https://doi.org/10.1109/CSF.2019.00025>
- [32] A. C. Yao, “Protocols for secure computations,” in *23rd annual IEEE Symposium on Foundations of Computer Science*, 1982, pp. 160–164. Available: <https://doi.org/10.1109/SFCS.1982.38>
- [33] A. C. Myers, A. Sabelfeld, and S. Zdancewic, “Enforcing robust declassification,” in *17th IEEE Computer Security Foundations Workshop (CSFW)*, Jun. 2004, pp. 172–186. Available: <http://www.cs.cornell.edu/andru/papers/csfw04.pdf>
- [34] E. Cecchetti, A. C. Myers, and O. Arden, “Nonmalleable information flow control,” in *24th ACM Conf. on Computer and Communications Security (CCS)*. ACM, Oct. 2017, pp. 1875–1891. Available: <http://www.cs.cornell.edu/andru/papers/nmifc>
- [35] F. Montesi, “Choreographic programming,” Ph.D. dissertation, IT University of Copenhagen, 2013. Available: <https://www.fabriziomontesi.com/files/choreographic-programming.pdf>
- [36] L. Cruz-Filipe and F. Montesi, “A core model for choreographic programming,” *Theoretical Computer Science*, vol. 802, pp. 38–66, 2020.
- [37] M. R. Clarkson and F. B. Schneider, “Hyperproperties,” in *21st IEEE Computer Security Foundations Symp. (CSF)*, Jun. 2008, pp. 51–65.
- [38] M. Patrignani, R. S. Wahby, and R. Künneman, “Universal composability is secure compilation,” arXiv ePrint 1910.08634,

- 2019.
- [39] M. Patrignani, R. Künnemann, and R. S. Wahby, “Universal composability is robust compilation,” arXiv ePrint 1910.08634, 2022.
- [40] Y. Lindell, “How to simulate it — A tutorial on the simulation proof technique,” in *Tutorials on the Foundations of Cryptography*. Springer International Publishing, 2017, pp. 277–346.
- [41] T. Dierks and E. Rescorla, “The transport layer security (TLS) protocol version 1.1,” Internet RFC-4346, Apr. 2006.
- [42] M. Backes, B. Pfizmann, and M. Waidner, “The reactive simulatability (RSIM) framework for asynchronous systems,” *Information and Computation*, vol. 205, no. 12, pp. 1685–1720, 2007.
- [43] L. Cruz-Filipe, F. Montesi, and M. Peressotti, “A formal theory of choreographic programming,” *CoRR*, vol. abs/2209.01886, 2022.
- [44] A. K. Hirsch and D. Garg, “Pirouette: Higher-order typed functional choreographies,” *Proc. ACM Program. Lang.*, vol. 6, no. POPL, pp. 1–27, Jan. 2022.
- [45] N. A. Lynch and M. R. Tuttle, “Hierarchical correctness proofs for distributed algorithms,” in *6th ACM Symp. on Principles of Distributed Computing*, F. B. Schneider, Ed. ACM, 1987, pp. 137–151.
- [46] W. Rafnsson and A. Sabelfeld, “Compositional information-flow security for interactive systems,” in *27th IEEE Computer Security Foundations Symp. (CSF)*. IEEE Computer Society, 2014, pp. 277–292. Available: <https://doi.org/10.1109/CSF.2014.27>
- [47] A. C. Myers and B. Liskov, “Protecting privacy using the decentralized label model,” *ACM Transactions on Software Engineering and Methodology*, vol. 9, no. 4, pp. 410–442, Oct. 2000. Available: <http://www.cs.cornell.edu/andru/papers/iflow-tosem.pdf>
- [48] D. Stefan, A. Russo, D. Mazières, and J. C. Mitchell, “Disjunction category labels,” in *Proceedings of the 16th Nordic conference on Information Security Technology for Applications*, ser. Lecture Notes in Computer Science, P. Laud, Ed., vol. 7161. Springer, 2011, pp. 223–239.
- [49] O. Arden, J. Liu, and A. C. Myers, “Flow-limited authorization,” in *28th IEEE Computer Security Foundations Symp. (CSF)*, Jul. 2015, pp. 569–583. Available: <http://www.cs.cornell.edu/andru/papers/fflam>
- [50] C. Acay, J. Gancher, R. Recto, and A. C. Myers, “Secure synthesis of distributed cryptographic applications (technical report),” Jan. 2024. Available: <https://arxiv.org/abs/2401.04131>
- [51] D. Zagieboylo, G. E. Suh, and A. C. Myers, “Using information flow to design an ISA that controls timing channels,” in *32nd IEEE Computer Security Foundations Symp. (CSF)*, Jun. 2019. Available: <https://www.cs.cornell.edu/andru/papers/hyperisa>
- [52] K. R. O’Neill, M. R. Clarkson, and S. Chong, “Information-flow security for interactive programs,” in *19th IEEE Computer Security Foundations Workshop (CSFW)*. IEEE Computer Society, 2006, pp. 190–201. Available: <https://doi.org/10.1109/CSFW.2006.16>
- [53] D. Clark and S. Hunt, “Non-interference for deterministic interactive programs,” in *5th Workshop on Formal Aspects in Security and Trust (FAST)*, ser. Lecture Notes in Computer Science, vol. 5491. Springer, 2008, pp. 50–66. Available: https://doi.org/10.1007/978-3-642-01465-9_4
- [54] L. Cruz-Filipe and F. Montesi, “On asynchrony and choreographies,” in *ICE@DisCoTec*, ser. EPTCS, M. Bartoletti, L. Bocchi, L. Henrio, and S. Knight, Eds., vol. 261, 2017, pp. 76–90.
- [55] M. Ishaq, A. Milanova, and V. Zikas, “Efficient MPC via program analysis: A framework for efficient optimal mixing,” in *26th ACM Conf. on Computer and Communications Security (CCS)*. ACM, 2019, pp. 1539–1556.
- [56] S. Nain and M. Y. Vardi, “Branching vs. linear time: Semantical perspective,” in *ATVA*, ser. Lecture Notes in Computer Science, K. S. Namjoshi, T. Yoneda, T. Higashino, and Y. Okamura, Eds., vol. 4762. Springer, 2007, pp. 19–34.
- [57] M. Hennessy and R. Milner, “Algebraic laws for nondeterminism and concurrency,” *J. ACM*, vol. 32, no. 1, pp. 137–161, 1985.
- [58] M. H. A. Newman, “On theories with a combinatorial definition of “equivalence”,” *Annals of Mathematics*, vol. 43, no. 2, pp. 223–243, 1942.
- [59] A. Church and J. B. Rosser, “Some properties of conversion,” *Transactions of the American Mathematical Society*, vol. 39, no. 3, pp. 472–482, 1936.
- [60] F. Baader and T. Nipkow, *Term rewriting and all that*. Cambridge University Press, 1998.
- [61] J. Å. Pohjola, A. Gómez-Londoño, J. Shaker, and M. Norrish, “Kalas: A verified, end-to-end compiler for a choreographic language,” in *ITP*, ser. LIPIcs, J. Andronick and L. de Moura, Eds., vol. 237. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 27:1–27:18.
- [62] M. Ben-Or, S. Goldwasser, and A. Wigderson, “Completeness theorems for non-cryptographic fault-tolerant distributed computation,” in *20th Annual ACM Symposium on Theory of Computing*, J. Simon, Ed., 1988, pp. 1–10. Available: <https://doi.org/10.1145/62212.62213>
- [63] R. Canetti, A. Stoughton, and M. Varia, “EasyUC: Using EasyCrypt to mechanize proofs of universally composable security,” in *32nd IEEE Computer Security Foundations Symp. (CSF)*. IEEE, 2019, pp. 167–183.
- [64] M. Barbosa, G. Barthe, B. Grégoire, A. Koutsos, and P. Strub, “Mechanized proofs of adversarial complexity and application to universal composability,” in *CCS ’21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*, Y. Kim, J. Kim, G. Vigna, and E. Shi, Eds. ACM, 2021, pp. 2541–2563. Available: <https://doi.org/10.1145/3460120.3484548>
- [65] J. Gancher, K. Sojakova, X. Fan, E. Shi, and G. Morrisett, “A core calculus for equational proofs of cryptographic protocols,” *Proc. ACM Program. Lang.*, vol. 7, no. POPL, pp. 866–892, 2023.
- [66] M. Hastings, B. Hemenway, D. Noble, and S. Zdancewic, “SoK: general purpose compilers for secure multi-party computation,” in *IEEE Symp. on Security and Privacy*, 2019, pp. 1220–1237. Available: <https://doi.org/10.1109/SP.2019.00028>
- [67] Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai, “Zero-knowledge from secure multiparty computation,” in *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, ser. STOC ’07. New York, NY, USA: ACM, 2007, pp. 21–30. Available: <http://doi.acm.org/10.1145/1250790.1250794>
- [68] S. Zdancewic and A. C. Myers, “Observational determinism for concurrent program security,” in *16th IEEE Computer Security Foundations Workshop (CSFW)*, Jun. 2003, pp. 29–43. Available: <http://www.cs.cornell.edu/andru/papers/csfw03.pdf>
- [69] E. Tilevich and Y. Smaragdakis, “J-Orchestra: Enhancing Java programs with distribution capabilities,” *ACM Trans. Softw. Eng. Methodol.*, vol. 19, no. 1, pp. 1:1–1:40, Aug. 2009. Available: <http://doi.acm.org/10.1145/1555392.1555394>
- [70] C. Liu, Y. Huang, E. Shi, J. Katz, and M. Hicks, “Automating efficient RAM-model secure computation,” in *IEEE Symp. on Security and Privacy*. IEEE, 2014, pp. 623–638.
- [71] N. Büscher, D. Demmler, S. Katzenbeisser, D. Kretzmer, and T. Schneider, “HyCC: Compilation of hybrid protocols for practical secure computation,” in *25th ACM Conf. on Computer and Communications Security (CCS)*. New York, NY, USA: ACM, 2018, p. 847–861.
- [72] E. Chen, J. Zhu, A. Ozdemir, R. S. Wahby, F. Brown, and W. Zheng, “Silph: A framework for scalable and accurate generation of hybrid MPC protocols,” *Cryptology ePrint Archive*, 2023.

- [73] A. Ozdemir, F. Brown, and R. S. Wahby, “CirC: Compiler infrastructure for proof systems, software verification, and more,” in *IEEE Symp. on Security and Privacy*, 2022, pp. 2248–2266.
- [74] A. Viand, P. Jattke, M. Haller, and A. Hithnawi, “HECO: automatic code optimizations for efficient fully homomorphic encryption,” *CoRR*, vol. abs/2202.01649, 2022. Available: <https://arxiv.org/abs/2202.01649>
- [75] R. Malik, K. Sheth, and M. Kulkarni, “Coyote: A compiler for vectorizing encrypted arithmetic circuits,” in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, 2023, pp. 118–133.
- [76] A. Askarov, D. Hedin, and A. Sabelfeld, “Cryptographically-masked flows,” *Theor. Comput. Sci.*, vol. 402, no. 2-3, pp. 82–101, 2008.
- [77] P. Laud, “On the computational soundness of cryptographically masked flows,” in *35th ACM Symp. on Principles of Programming Languages (POPL)*, G. C. Necula and P. Wadler, Eds. ACM, 2008, pp. 337–348.
- [78] R. Küsters, T. Truderung, and J. Graf, “A framework for the cryptographic verification of Java-like programs,” in *25th IEEE Computer Security Foundations Symp. (CSF)*, S. Chong, Ed. IEEE Computer Society, 2012, pp. 198–212.
- [79] R. Küsters, T. Truderung, B. Beckert, D. Bruns, M. Kirsten, and M. Mohr, “A hybrid approach for proving noninterference of Java programs,” in *28th IEEE Computer Security Foundations Symp. (CSF)*, C. Fournet, M. W. Hicks, and L. Viganò, Eds. IEEE Computer Society, 2015, pp. 305–319.
- [80] X. Leroy, “Formal verification of a realistic compiler,” *Commun. ACM*, vol. 52, pp. 107–115, 2009.
- [81] U. Maurer, “Constructive cryptography – A new paradigm for security definitions and proofs,” in *Theory of Security and Applications*, ser. Lecture Notes in Computer Science, S. Mödersheim and C. Palamidessi, Eds., vol. 6993. Springer, 2011, pp. 33–56.
- [82] K. Liao, M. A. Hammer, and A. Miller, “ILC: a calculus for composable, computational cryptography,” in *40th ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)*, Jun. 2019, pp. 640–654.
- [83] A. Lochbihler, S. R. Sefidgar, D. Basin, and U. Maurer, “Formalizing constructive cryptography using CryptHOL,” in *32nd IEEE Computer Security Foundations Symp. (CSF)*. IEEE, 2019, pp. 152–166.
- [84] A. Lluch-Lafuente, F. Nielson, and H. R. Nielson, “Discretionary information flow control for interaction-oriented specifications,” in *Logic, Rewriting, and Concurrency*, ser. Lecture Notes in Computer Science, N. Martí-Oliet, P. C. Ölveczky, and C. L. Talcott, Eds., vol. 9200. Springer, 2015, pp. 427–450.
- [85] L. Zheng and A. C. Myers, “A language-based approach to secure quorum replication,” in *9th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security (PLAS)*, Aug. 2014. Available: <http://www.cs.cornell.edu/andru/papers/plas14>