ILC: A Calculus for Composable, Computational Cryptography

Kevin Liao
University of Illinois
Urbana-Champaign, USA
kliao6@illinois.edu

Matthew A. Hammer
DFINITY, USA
matthew@dfinity.org

Andrew Miller
University of Illinois
Urbana-Champaign, USA
soc1024@illinois.edu

Abstract

The universal composability (UC) framework is the established standard for analyzing cryptographic protocols in a modular way, such that security is preserved under concurrent composition with arbitrary other protocols. However, although UC is widely used for on-paper proofs, prior attempts at systemizing it have fallen short, either by using a symbolic model (thereby ruling out computational reduction proofs), or by limiting its expressiveness.

In this paper, we lay the groundwork for building a concrete, executable implementation of the UC framework. Our main contribution is a process calculus, dubbed the Interactive Lambda Calculus (ILC). ILC faithfully captures the computational model underlying UC—interactive Turing machines (ITMs)—by adapting ITMs to a subset of the π-calculus through an affine typing discipline. In other words, *well-typed ILC programs are expressible as ITMs.* In turn, ILC's strong confluence property enables reasoning about cryptographic security reductions. We use ILC to develop a simplified implementation of UC called SaUCy.

CCS Concepts • **Security and privacy → Formal security models**;

Keywords Provable security, universal composability, process calculus, type systems

ACM Reference Format:
Kevin Liao, Matthew A. Hammer, and Andrew Miller. 2019. ILC: A Calculus for Composable, Computational Cryptography. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '19), June 22–26, 2019, Phoenix, AZ, USA.* ACM, New York, NY, USA, 31 pages. https://doi.org/10.1145/3314221.3314607

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
PLDI '19, June 22–26, 2019, Phoenix, AZ, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6712-7/19/06…$15.00
https://doi.org/10.1145/3314221.3314607

1 Introduction

In cryptography, a proof of security in the simulation-based universal composability (UC) framework is considered the gold standard for demonstrating that a protocol "does its job securely" [16]. In particular, a UC-secure protocol enjoys the strongest notion of compositionality—it maintains all security properties even when run *concurrently* with arbitrary other protocol instances. This is in contrast with weaker property-based notions that only guarantee security in a standalone setting [37] or under sequential composition [26]. Thus, the benefit of using UC is *modularity*—it supports analyzing complex protocols by composing simpler building blocks. However, the cost of using UC is that security proofs tend to be quite complicated. We believe that applying a PL-style of systemization to UC can help simplify its use, bring new clarity, and provide useful tooling. We envision a future where modularity of cryptographic protocol composition translates to modular implementation as well.

Reviewing prior efforts of applying PL techniques to cryptography, we find they run up against challenges when importing the existing body of UC theory. Either they do not support computational reasoning (which considers issues of probability and computational complexity) [10], do not support message-passing concurrency for distributed protocols [4], or are too expressive (allow for expressing nondeterminism with no computational interpretation) [2].

Our observation is that these approaches diverge from UC at a low level: UC is defined atop the underlying (concurrent) computational model of *interactive Turing machines* (ITMs). The significance of ITMs is that they have a clear computational interpretation, so it is straightforward to relate execution traces to a probabilistic polynomial time computation, as is necessary for cryptographic reduction proofs. The presence of (non-probabilistic) nondeterminism in alternative models of concurrency would frustrate such reduction proofs. ITMs sidestep this issue by having a deterministic (modulo random coin tosses), "single-threaded" execution semantics. That is, processes pass control from one to another each time a message is sent so that *exactly one process is active at any given time*, and, moreover, *the order of activations is fully determined.*

In this paper, we take up the challenge of faithfully capturing these idioms by designing a new process calculus called the Interactive Lambda Calculus (ILC), which adapts

ITMs to a subset of the π-calculus [43] through an affine typing discipline. In other words, *well-typed ILC programs are expressible as ITMs.* We then use ILC to build a concrete, executable implementation of a simplified UC framework, dubbed SaUCy.

1.1 Interactive Lambda Calculus

Why do we need another process calculus in the first place? Where do existing ones fall short? On the one hand, process calculi such as the π-calculus [43] and its cryptography-oriented variants [1, 2, 35] are not a good fit to ITMs, since they permit non-confluent reductions by design (i.e., non-probabilistic nondeterminism). On the other hand, various other calculi that do enjoy confluence are overly restrictive, only allowing for fixed or two-party communications [10, 24, 30].

ILC fills this gap by adapting ITMs to a subset of the π-calculus through an affine typing discipline. To maintain that only one process is active (can write) at any given time, processes implicitly pass around an affine "write token" by virtue of where they perform read and write effects: When process A writes to process B, process A "spends" the write token and process B "earns" the write token. Moreover, to maintain that the order of activations is fully determined, the read endpoints of channels are (non-duplicable) affine resources, and so each write operation corresponds to a single, unique read operation. Together, these give ILC its central metatheoretic property of *confluence*.

The importance of confluence is that the only nondeterminism in an ILC program is due to random coin tosses taken by processes, which have a well-defined distribution. Additionally, any apparent concurrency hazards, such as adversarial scheduling of messages in an asynchronous network, are due to an explicit adversary process rather than uncertainty built into the model itself. This eliminates non-probabilistic nondeterminism, and so ILC programs are amenable to the reasoning patterns necessary for establishing computational security guarantees.

1.2 Contributions

To summarize, our main contributions are these:

- We design a foundational calculus for the purpose of systemizing UC called the Interactive Lambda Calculus, which exhibits confluence and is a faithful abstraction of ITMs.
- We use ILC to build a concrete, executable implementation of a simplified UC framework called SaUCy.
- We then use SaUCy to port over a sampling of theory from UC literature, including a composition theorem, an instantiation proof of UC commitments [19], and an examination of a subtle definitional issue involving reentrant concurrency [14].

2 Overview

We first provide background on the universal composability framework and then give a tour of ILC.

2.1 Background on Universal Composability

Security proofs in the UC framework follow the real/ideal paradigm [26]. To carry out some cryptographic task in the real world, we define a distributed protocol that achieves the task across *many untrusted processes*. Then, to show that it is secure, we compare it with an idealized protocol in which processes simply rely on a *single trusted process* to carry out the task for them (and so security is satisfied trivially).

The program for this single trusted process is called an *ideal functionality* as it provides a uniform way to describe all the security properties we want from the protocol. Roughly speaking, we say a protocol π *realizes* an ideal functionality $\mathcal{F}$ (i.e., it meets its specification) if every adversarial behavior in the real world can also be exhibited in the ideal world.

Once we have defined π and $\mathcal{F}$, proving realization formally follows a standard rhythm:

1. The first step is a construction: We must provide a *simulator* $\mathcal{S}$ that translates any attack $\mathcal{A}$ on the protocol π into an attack on $\mathcal{F}$.
2. The second step is a relational analysis: We must show that running π under attack by any adversary $\mathcal{A}$ (the real world) is *indistinguishable* from running $\mathcal{F}$ under attack by $\mathcal{S}$ (the ideal world) to any distinguisher $\mathcal{Z}$ called the *environment*.

In particular, $\mathcal{Z}$ is an adaptive distinguisher: It interacts with both the real world and the ideal world, and the simulation is sound if no $\mathcal{Z}$ can distinguish between the two.

As mentioned, the primary goal of this framework is *compositionality*. Suppose a protocol π is a protocol module that realizes a functionality $\mathcal{F}$ (a specification of the module), and suppose a protocol ρ, which relies on $\mathcal{F}$ as a subroutine, in turn realizes an application specification functionality $\mathcal{G}$. Then, the composed protocol $\rho \circ \pi$, in which calls to $\mathcal{F}$ are replaced by calls to π, also realizes $\mathcal{G}$. Instead of analyzing the composite protocol consisting of ρ and π, it suffices to analyze the security of ρ itself in the simpler world with $\mathcal{F}$, the idealized version of π.

Finally, the UC framework is defined atop the underlying computational model of interactive Turing machines (ITMs). In the ITM model, processes pass control from one to another each time a message is sent so that *exactly one process is active at any given time*, and, moreover, *the order of activations is fully determined*. This gives ITMs a clear computational interpretation, which is necessary for the above proofs (in particular, cryptographic reductions) to go through.

2.2 ILC by Example

We make the above more concrete by running through an example of *commitment*, an essential building block in many

$\mathcal{F}_{\text{COM}}$ proceeds as follows, running with committer P and receiver Q.

1. Upon receiving a message (Commit, b) from P, where $b \in \{0, 1\}$, record the value b and send the message (Receipt) to Q. Ignore any subsequent Commit messages.
2. Upon receiving a message (Open) from P, proceed as follows: If some value b was previously recorded, then send the message (Open, b) to Q and halt. Otherwise, halt.

```
fCom :: Wr Msg ⟶ Rd Msg ⊸ 𝟙
let fCom toQ frP =
  let (!(Commit b), frP) = rd frP in
    wr Receipt ⟶ toQ ;
    let (!Open, frP) = rd frP in
      wr (Opened b) ⟶ toQ
```

Figure 1. An ideal functionality for a one-time commitment scheme in prose (left) and in ILC (right).

cryptographic protocols [11]. The idea behind commitment is simple: A *committer* provides a *receiver* with the digital equivalent of a "sealed envelope" containing some value that can later be revealed. The commitment scheme must be *hiding* in the sense that the commitment itself reveals no information about the committed value, and *binding* in the sense that the committer can only open the commitment to a single value. For security under composition, an additional *non-malleability* property is required, which roughly prevents an attacker from using one commitment to derive another related one.

All of these properties are captured at once using an ideal functionality. In Figure 1 (left), we show a simplified ideal functionality for one-time bit commitment, $\mathcal{F}_{\text{COM}}$, as it would appear in the cryptography literature [19]. The functionality simply waits for the committer P to commit to some bit b, notifies the receiver Q that it has taken place, and reveals b to Q upon request by P. Notice that Q never actually sees a commitment to b (only the (Receipt) message), so the three properties hold trivially.

In Figure 1 (right), we implement a simplified version of $\mathcal{F}_{\text{COM}}$ in ILC to highlight some key features of the language. The function fCom takes two channel endpoints as arguments. The first is a write endpoint toQ : Wr Msg (for sending messages of type Msg to Q), and the second is a read endpoint frP : Rd Msg (for receiving messages of type Msg from P). At a high level, it should be clear how the communication pattern in fCom follows that in $\mathcal{F}_{\text{COM}}$, but there are a few details that require further explanation. These details are better explained in the context of ILC's type system, which we give a quick tour of next.

2.3 ILC Type System

ILC terms have either an unrestricted type, meaning they can be freely copied, or an affine type, meaning they can be used at most once. Affine typing serves a special purpose, namely, to ensure that ILC processes have a determined sequence of activations, as is required in ITMs. This is achieved through the following invariants:

- *Only one process is active at any given time.* Processes implicitly pass around an affine "write token" ⓦ by virtue of where they perform read and write effects. In order for process A to write to process B, process A must first own the write token. Because the write token is unique, at most one process owns the write token ("is active" or "can write") at any given time. When process B reads the message from A, process B earns the write token, thereby conserving its uniqueness and now allowing process B to write to some other process.
- *The order of activations is deterministic.* Each channel (or "tape" in ITM parlance) has a read endpoint and a write endpoint. The read endpoint is an affine resource, and so it is owned by at most one process. This ensures that each write operation corresponds to a single, unique read operation.

Intuitively, the first invariant rules out the possibility of write nondeterminism. Consider the case in which two processes are trying to execute writes in parallel, which would lead to a race condition. This does not typecheck, since the affine write token belongs to at most one process. One might justifiably wonder why write endpoints are unrestricted and read endpoints are affine. Note that if two processes are trying to write in parallel, the two write endpoints need not be the same, so making write endpoints affine would not help our case in eliminating write nondeterminism.

Dually, the second invariant rules out the possibility of read nondeterminism. Consider the case in which two processes A and B are listening on the same read endpoint. If a process C writes on the corresponding write endpoint, which of A or B (or both) gets activation? If only one of them is activated, then we have a source of nondeterminism. If both are activated, now A and B both own write tokens, violating its affinity. In any case, this does not typecheck since read endpoints are affine resources, making it impossible for two processes A and B to listen on the same read endpoint. Together, these invariants ensure that processes have a determined sequence of activations as desired.

A Selection of Typing Rules. To see these invariants in action, we walk through the typing rules for fork, write, and read expressions. We read the typing judgement $\Delta; \Gamma \vdash e : U$ as "under affine context Δ and unrestricted context Γ, expression e has type U." The metavariables U and V range over all types (both unrestricted and affine).

The fork expression $e_1 \mid\triangleright\ e_2$ spawns a child process e_1 and continues as e_2.

$$\frac{\Delta_1; \Gamma \vdash e_1 : U \qquad \Delta_2; \Gamma \vdash e_2 : V}{\Delta_1, \Delta_2; \Gamma \vdash e_1 \mid\triangleright\ e_2 : V}\ \text{fork}$$

Its typing rule says that if we can partition the affine context as Δ_1, Δ_2 such that e_1 has type U under contexts $\Delta_1; \Gamma$ and e_2 has type V under contexts $\Delta_2; \Gamma$, then the expression has type V. Notice that affine resources (e.g., read endpoints and the write token) must be split between the child process and the parent process, thereby preventing their duplication.

The write expression $\mathsf{wr}(e_1, e_2)$ sends the value that e_1 evaluates to on the write endpoint that e_2 evaluates to. One thing to mention is that only values of a sendable type (ranged over by S) can be sent over channels (more on this later).

$$\frac{\Delta_1; \Gamma \vdash e_1 : S \qquad \Delta_2; \Gamma \vdash e_2 : \mathsf{Wr}\, S}{\Delta_1, \Delta_2, ⓦ; \Gamma \vdash \mathsf{wr}(e_1, e_2) : \mathbb{1}}\ \text{wr}$$

Its typing rule says that if we own the write token and we can partition the affine context as Δ_1, Δ_2 such that e_1 has type S under contexts $\Delta_1; \Gamma$ and e_2 evaluates to a write endpoint (of type $\mathsf{Wr}\, S$) under contexts $\Delta_2; \Gamma$, then the expression has type $\mathbb{1}$ (unit). Notice that typing a write expression spends the write token, and so it cannot execute another write until it gets "reactivated" by reading from some other process.

The read expression $\mathsf{rd}(e_1, x.e_2)$ reads a value on the read endpoint that e_1 evaluates to and binds the value-endpoint pair as x in the affine context of e_2. Rebinding the read endpoint allows it to be reused.

$$\frac{ⓦ \notin \Delta_2 \qquad \Delta_1; \Gamma \vdash e_1 : \mathsf{Rd}\, S \qquad \Delta_2, ⓦ, x : !\, S \otimes \mathsf{Rd}\, S; \Gamma \vdash e_2 : U}{\Delta_1, \Delta_2; \Gamma \vdash \mathsf{rd}(e_1, x.e_2) : U}\ \text{rd}$$

Its typing rule says that if we can partition the affine context as Δ_1, Δ_2 such that e_1 evaluates to a read endpoint (of type $\mathsf{Rd}\, S$) under contexts $\Delta_1; \Gamma$, and e_2 has type U under contexts $\Delta_2, ⓦ, x : !\, S \otimes \mathsf{Rd}\, S; \Gamma$, then the expression has type U.

There are a few things to unpack here. First, we explain the affine product type $!\, S \otimes \mathsf{Rd}\, S$. Since sendable values are unrestricted and read endpoints are affine, the value read on the channel is wrapped in a ! operator (pronounced "bang") so that it can be placed in an affine pair. Next, observe that ⓦ is available in the body e_2 of the read expression (i.e., it is conserved), but only under the condition that it is not already in the affine context Δ_2 (otherwise, a process could arbitrarily mint write tokens, violating its affinity).

Revisiting fCom. Having gone through several typing rules, we now revisit fCom from Figure 1 (right). In particular, we should convince ourselves that fCom respects the invariants of the type system.

The type signature tells us that toQ : Wr Msg is unrestricted (→ is the type connective for unrestricted arrows) and frP : Rd Msg is affine (⊸ is the type connective for affine arrows). As we mentioned, write endpoints *are not* affine, since this restriction does not help in preventing write nondeterminism; read endpoints *are* affine, which does prevent read nondeterminism. To see that frP is being used affinely, notice that it is rebound when deconstructing the value-endpoint pair from each read operation, so it can be used again.

To see that the write token is being passed around appropriately, notice that the read and write effects are interleaved. Before each read operation, fCom does not own the write token: In the first read operation, only frP : RdMsg is present in the affine context; in the second read operation, the first write operation has already spent the write token. Before each write operation, fCom does own the write token: Each is preceded by a read operation.

3 Interactive Lambda Calculus

We now present the Interactive Lambda Calculus in full, formalizing its syntax, static semantics, and dynamic semantics.

3.1 Syntax

The syntax of ILC is given in Figure 2. Types (written U, V) are bifurcated into unrestricted types (written A, B) and affine types (written X, Y).

A subset of the unrestricted types are sendable types (written S, T), i.e., the types of values that can be sent over channels. This restriction ensures that channels model network channels, which send only data. The sendable types include unit ($\mathbb{1}$), products ($S \times T$), and sums ($S + T$).

The unrestricted types include the sendable types, write endpoint types ($\mathsf{Wr}\, S$), products ($A \times B$), sums ($A + B$), arrows ($A \to_\infty U$ or simply $A \to U$), and write arrows ($A \to_\mathsf{w} U$). Write arrows specify unrestricted abstractions for which the write token can be moved into the affine context of the abstraction body during β-reduction.

The affine types include bang types ($!\, A$), read endpoint types ($\mathsf{Rd}\, S$), products ($X \otimes Y$), sums ($X \oplus Y$), and arrows ($X \to_1 U$ or simply $X \multimap U$). Notice that the write token ⓦ lives in the affine context, though it cannot be bound to any variable. Instead, it flows around implicitly by virtue of where read and write effects are performed.

For concision, certain syntactic forms are parameterized by a multiplicity π to distinguish between the unrestricted (∞) and affine (1) counterparts; other syntactic forms are parameterized by a syntax label ℓ, which includes the multiplicity labels and the write label w (related to write effects). On introduction and elimination forms for functions (abstraction, application, and fixed points), the label w denotes variants that move around the write token as explained above. On introduction and elimination forms for products and sums, the label w denotes the sendable variants.

All types	$U, V ::= A \mid X$	Syntax labels	$\ell ::= \pi \mid \mathsf{w}$
Sendable types	$S, T ::= \mathbb{1} \mid S \times T \mid S + T$	Multiplicity labels	$\pi ::= 1 \mid \infty$
Unrestricted types	$A, B ::= S \mid \mathsf{Wr}\, S \mid A \times B \mid A + B \mid A \to_{\infty \mid \mathsf{w}} U$	Unrestricted typings	$\Gamma ::= \cdot \mid \Gamma, x : A$
Affine types	$X, Y ::= !\, A \mid \mathsf{Rd}\, S \mid X \otimes Y \mid X \oplus Y \mid X \to_1 U$	Affine typings	$\Delta ::= \cdot \mid \Delta, x : X \mid \Delta, \textcircled{w}$

Values	$v ::= () \mid (v_1, v_2)_\ell \mid \mathsf{inj}^1_\ell(v) \mid \mathsf{inj}^2_\ell(v) \mid \lambda_\ell\, x.\, e \mid c \mid !\, v$
Channel endpoints	$c ::= \mathsf{Read}(d) \mid \mathsf{Write}(d)$
Channel names	$d ::= \cdots$
Expressions	$e ::= x \mid () \mid (e_1, e_2)_\ell \mid \mathsf{inj}^i_\ell(e) \mid \mathsf{split}_\ell(e_1, x_1.x_2.e_2) \mid \mathsf{case}_\ell(e, x_1.e_1, x_2.e_2)$ $\mid \lambda_\ell\, x.\, e \mid (e_1\, e_2)_\ell \mid \mathsf{fix}_\ell(x.e) \mid \mathsf{let}_\pi(e_1, x.e_2) \mid !\, e \mid ¡\, e$ $\mid \nu(x_1, x_2).\, e \mid \mathsf{wr}(e_1, e_2) \mid \mathsf{rd}(e_1, x.e_2) \mid \mathsf{ch}(e_1, x_1.e_3, e_2, x_2.e_4) \mid e_1 \mid\!\triangleright\, e_2$

Figure 2. ILC Syntax.

Values in ILC (written v) include unit, pairs, sums, lambda expressions, channel endpoints (written c), and banged values. We distinguish between the names of channel endpoints—$\mathsf{Read}(d)$ and $\mathsf{Write}(d)$—and the channel d itself that binds them. ILC supports a fairly standard feature set of expressions. Bang-typed values have introduction form $!\, e$ and elimination form $¡\, e$. The more interesting expressions are those related to communication and concurrency:

- *Restriction:* $\nu(x_1, x_2).\, e$ binds a read endpoint x_1 and a corresponding write endpoint x_2 in e.
- *Write:* $\mathsf{wr}(e_1, e_2)$ sends the value that e_1 evaluates to on the write endpoint that e_2 evaluates to.
- *Read:* $\mathsf{rd}(e_1, x.e_2)$ reads a value from the read endpoint that e_1 evaluates to and binds the value-endpoint pair as x in e_2.
- *Choice:* $\mathsf{ch}(e_1, x_1.e_3, e_2, x_2.e_4)$ allows a process to continue as either e_3 or e_4 based on some initial read event on one of the read endpoints that e_1 and e_2 evaluate to. The value read over the channel and the two read endpoints are rebound in a 3-tuple as x_1 in e_3 or x_2 in e_4. Here, we show only binary choice, but it can be generalized to the n-ary case.
- *Fork:* $e_1 \mid\!\triangleright\, e_2$ spawns a child process e_1 and continues as e_2.

3.2 Static Semantics

The typing rules of ILC are given in Figure 3. An algorithmic version of the rules appears in the appendix.

To recap, the typing rules maintain that only one process is active at any given time (unique ownership of the write token), and the order of activations is deterministic (unique ownership of read endpoints). We read the typing judgement $\Delta; \Gamma \vdash e : U$ as "under affine context Δ and unrestricted context Γ, expression e has type U." In full detail, the typing judgement also includes a typing context Ψ, which maps channel names d to sendable types S. However, it is only used in two special rules for typing channel endpoints that do not arise for source level programs, but will be needed to typecheck a running program that has performed channel allocation:

$$\frac{\Psi(d) = S}{\Psi; \Delta; \Gamma \vdash \mathsf{Read}(d) : \mathsf{Rd}\, S}\ \text{rdend}$$

$$\frac{\Psi(d) = S}{\Psi; \Delta; \Gamma \vdash \mathsf{Write}(d) : \mathsf{Wr}\, S}\ \text{wrend}$$

This pair of rules establish the canonical forms for the types of channel endpoints, $\mathsf{Rd}\, S$ and $\mathsf{Wr}\, S$. We use the metavariable c to range over these two canonical forms.

The typing rules for the functional fragment of ILC are fairly standard, except that they now have unrestricted and affine variants (and for some, sendable variants).

The rule for unrestricted abstraction (uabs) extends the unrestricted context Γ with $x : A$ before checking the body e of the abstraction. Notice that because unrestricted abstractions can be duplicated, the body must be affinely closed (cannot contain free affine variables).

The rule for write abstraction (wabs) is similar to uabs. The only difference is that wabs extends the affine context with the write token before checking the body e of the abstraction. Dually, the write application rule (wapp) stipulates that a process must own the write token in order to apply a write abstraction.

The rule for affine abstraction (aabs) is analagous to uabs, but notice that the body *need not* be affinely closed, since affine abstractions cannot be duplicated. It turns out that most affine functions we write *are* affinely closed, and so such a function $f : X \multimap U$ can be made into an unrestricted function $g : A \to X \multimap U$ by adding a leading unrestricted argument.

The bang rule turns an unrestrictedly typed expression $e : A$ into an affinely typed expression $e : !\, A$. Dually, the gnab rule turns an affinely typed expression $e : !\, A$ into an unrestrictedly typed expression $e : A$.

The typing rules for fork, write, and read were covered in Section 2.3, so this leaves channel restriction (nu) and external choice (choice) as the remaining typing rules related to communication and concurrency.

$\boxed{\Delta;\Gamma \vdash e : U}$ Under affine context Δ and unrestricted context Γ, expression e has type U.

$$\frac{\Gamma(x) = A}{\Delta;\Gamma \vdash x : A}\ \text{uvar} \qquad \frac{\Delta(x) = X}{\Delta;\Gamma \vdash x : X}\ \text{avar} \qquad \frac{}{\Delta;\Gamma \vdash () : \mathbb{1}}\ \text{unit} \qquad \frac{\Delta_1;\Gamma \vdash e_1 : A_1 \qquad \Delta_2;\Gamma \vdash e_2 : A_2}{\Delta_1,\Delta_2;\Gamma \vdash (e_1,e_2)_\infty : A_1 \times A_2}\ \text{upair}$$

$$\frac{\Delta_1;\Gamma \vdash e_1 : S_1 \qquad \Delta_2;\Gamma \vdash e_2 : S_2}{\Delta_1,\Delta_2;\Gamma \vdash (e_1,e_2)_\mathsf{w} : S_1 \times S_2}\ \text{spair} \qquad \frac{\Delta_1;\Gamma \vdash e_1 : X_1 \qquad \Delta_2;\Gamma \vdash e_2 : X_2}{\Delta_1,\Delta_2;\Gamma \vdash (e_1,e_2)_1 : X_1 \otimes X_2}\ \text{apair} \qquad \frac{i \in \{1,2\} \qquad \Delta;\Gamma \vdash e : A_i}{\Delta;\Gamma \vdash \mathsf{inj}^i_\infty(e) : A_1 + A_2}\ \text{uinj}$$

$$\frac{i \in \{1,2\} \qquad \Delta;\Gamma \vdash e : S_i}{\Delta;\Gamma \vdash \mathsf{inj}^i_\mathsf{w}(e) : S_1 + S_2}\ \text{sinj} \qquad \frac{i \in \{1,2\} \qquad \Delta;\Gamma \vdash e : X_i}{\Delta;\Gamma \vdash \mathsf{inj}^i_1(e) : X_1 \oplus X_2}\ \text{ainj} \qquad \frac{\begin{array}{c}\Delta_1;\Gamma \vdash e_1 : A_1 \times A_2 \\ \Delta_2;\Gamma, x_1 : A_1, x_2 : A_2 \vdash e : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{split}_\infty(e_1, x_1.x_2.e_2) : U}\ \text{usplit}$$

$$\frac{\begin{array}{c}\Delta_1;\Gamma \vdash e_1 : S_1 \times S_2 \\ \Delta_2;\Gamma, x_1 : S_1, x_2 : S_2 \vdash e : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{split}_\mathsf{w}(e_1, x_1.x_2.e_2) : U}\ \text{ssplit} \qquad \frac{\begin{array}{c}\Delta_1;\Gamma \vdash e_1 : X_1 \otimes X_2 \\ \Delta_2, x_1 : X_1, x_2 : X_2;\Gamma \vdash e : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{split}_1(e_1, x_1.x_2.e_2) : U}\ \text{asplit}$$

$$\frac{\begin{array}{c}\Delta_1;\Gamma \vdash e : A_1 + A_2 \\ \Delta_2;\Gamma, x_1 : A_1 \vdash e_1 : U \qquad \Delta_2;\Gamma, x_2 : A_2 \vdash e_2 : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{case}_\infty(e, x_1.e_1, x_2.e_2) : U}\ \text{ucase} \qquad \frac{\begin{array}{c}\Delta_1;\Gamma \vdash e : S_1 + S_2 \\ \Delta_2;\Gamma, x_1 : S_1 \vdash e_1 : U \qquad \Delta_2;\Gamma, x_2 : S_2 \vdash e_2 : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{case}_\mathsf{w}(e, x_1.e_1, x_2.e_2) : U}\ \text{scase}$$

$$\frac{\begin{array}{c}\Delta_1;\Gamma \vdash e : X_1 \oplus X_2 \\ \Delta_2, x_1 : X_1;\Gamma \vdash e_1 : U \qquad \Delta_2, x_2 : X_2;\Gamma \vdash e_2 : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{case}_1(e, x_1.e_1, x_2.e_2) : U}\ \text{acase} \qquad \frac{\cdot;\Gamma, x : A \vdash e : U}{\Delta;\Gamma \vdash \lambda_\infty\, x.\, e : A \to_\infty U}\ \text{uabs} \qquad \frac{\textcircled{w};\Gamma, x : A \vdash e : U}{\Delta;\Gamma \vdash \lambda_\mathsf{w}\, x.\, e : A \to_\mathsf{w} U}\ \text{wabs}$$

$$\frac{\Delta, x : X;\Gamma \vdash e : U}{\Delta;\Gamma \vdash \lambda_1\, x.\, e : X \to_1 U}\ \text{aabs} \qquad \frac{\Delta_1;\Gamma \vdash e_2 : A \qquad \Delta_2;\Gamma \vdash e_1 : A \to_\infty U}{\Delta_1,\Delta_2;\Gamma \vdash (e_1\, e_2)_\infty : U}\ \text{uapp}$$

$$\frac{\Delta_1;\Gamma \vdash e_2 : A \qquad \Delta_2;\Gamma \vdash e_1 : A \to_\mathsf{w} U}{\Delta_1,\Delta_2,\textcircled{w};\Gamma \vdash (e_1\, e_2)_\mathsf{w} : U}\ \text{wapp} \qquad \frac{\Delta_1;\Gamma \vdash e_2 : X \qquad \Delta_2;\Gamma \vdash e_1 : X \to_1 U}{\Delta_1,\Delta_2;\Gamma \vdash (e_1\, e_2)_1 : U}\ \text{aapp}$$

$$\frac{\cdot;\Gamma, x : A \to_\infty U \vdash e : A \to_\infty U}{\Delta;\Gamma \vdash \mathsf{fix}_\infty(x.e) : A \to_\infty U}\ \text{ufix} \qquad \frac{\cdot;\Gamma, x : A \to_\mathsf{w} U \vdash e : A \to_\mathsf{w} U}{\Delta;\Gamma \vdash \mathsf{fix}_\mathsf{w}(x.e) : A \to_\mathsf{w} U}\ \text{wfix} \qquad \frac{x : X \to_1 U;\Gamma \vdash e : X \to_1 U}{\Delta;\Gamma \vdash \mathsf{fix}_1(x.e) : X \to_1 U}\ \text{afix}$$

$$\frac{\begin{array}{c}\Delta_1;\Gamma \vdash e_1 : A \\ \Delta_2;\Gamma, x : A \vdash e_2 : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{let}_\infty(e_1, x.e_2) : U}\ \text{ulet} \qquad \frac{\begin{array}{c}\Delta_1;\Gamma \vdash e_1 : X \\ \Delta_2, x : X;\Gamma \vdash e_2 : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{let}_1(e_1, x.e_2) : U}\ \text{alet} \qquad \frac{\Delta;\Gamma \vdash e : A}{\Delta;\Gamma \vdash\, !\, e : !\, A}\ \text{bang} \qquad \frac{\Delta;\Gamma \vdash e : !\, A}{\Delta;\Gamma \vdash\, ¡\, e : A}\ \text{gnab}$$

$$\frac{\Delta, x_1 : \mathsf{Rd}\, S;\Gamma, x_2 : \mathsf{Wr}\, S \vdash e : U}{\Delta;\Gamma \vdash \nu(x_1, x_2).\, e : U}\ \text{nu} \qquad \frac{\begin{array}{c}\Delta_1;\Gamma \vdash e_1 : S \\ \Delta_2;\Gamma \vdash e_2 : \mathsf{Wr}\, S\end{array}}{\Delta_1,\Delta_2,\textcircled{w};\Gamma \vdash \mathsf{wr}(e_1, e_2) : \mathbb{1}}\ \text{wr} \qquad \frac{\begin{array}{c}\textcircled{w} \notin \Delta_2 \qquad \Delta_1;\Gamma \vdash e_1 : \mathsf{Rd}\, S \\ \Delta_2, \textcircled{w}, x : !\, S \otimes \mathsf{Rd}\, S;\Gamma \vdash e_2 : U\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash \mathsf{rd}(e_1, x.e_2) : U}\ \text{rd}$$

$$\frac{\begin{array}{c}\textcircled{w} \notin \Delta_3 \qquad \Delta_1;\Gamma \vdash e_1 : \mathsf{Rd}\, S \qquad \Delta_2;\Gamma \vdash e_2 : \mathsf{Rd}\, T \\ \Delta_3, \textcircled{w}, x_1 : !\, S \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T;\Gamma \vdash e_3 : U \\ \Delta_3, \textcircled{w}, x_2 : !\, T \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T;\Gamma \vdash e_4 : U\end{array}}{\Delta_1,\Delta_2,\Delta_3;\Gamma \vdash \mathsf{ch}(e_1, x_1.e_3, e_2, x_2.e_4) : U}\ \text{choice} \qquad \frac{\begin{array}{c}\Delta_1;\Gamma \vdash e_1 : U \\ \Delta_2;\Gamma \vdash e_2 : V\end{array}}{\Delta_1,\Delta_2;\Gamma \vdash e_1 \mid\!\triangleright\, e_2 : V}\ \text{fork}$$

Figure 3. ILC typing rules.

Process names	$p, q ::= \cdots$
Name sets	$\Sigma ::= \varepsilon \mid \Sigma, d \mid \Sigma, p$
Process pools	$\pi ::= \varepsilon \mid \pi, p : e$
Configurations	$C ::= \langle \Sigma; \pi \rangle$

Evaluation contexts
$$E ::= \bullet \mid (E, e)_\ell \mid (v, E)_\ell \mid \mathsf{inj}^i_\ell(E) \mid \mathsf{split}_\ell(E, x_1.x_2.e) \mid \mathsf{case}_\ell(E, x_1.e_1, x_2.e_2) \mid (E\,e)_\ell \mid (v\,E)_\ell \mid \mathsf{let}_\pi(E, x.e) \mid !\,E \mid ¡\,E \mid \mathsf{wr}(E, e) \mid \mathsf{wr}(v, E) \mid \mathsf{rd}(E, x.e) \mid \mathsf{ch}(E, x_1.e_3, e_2, x_2.e_4) \mid \mathsf{ch}(c, x_1.e_3, E, x_2.e_4)$$

Figure 4. ILC dynamic syntax.

$\boxed{C_1 \equiv C_2}$ Configurations C_1 and C_2 are equivalent.

$$\frac{\pi_1 \equiv_{\mathsf{perm}} \pi_2}{\langle \Sigma; \pi_1 \rangle \equiv \langle \Sigma; \pi_2 \rangle}\ \text{permProcs}$$

$\boxed{c_1 \rightsquigarrow c_2}$ Write endpoint c_1 connects to read endpoint c_2.

$$\frac{}{\mathsf{Write}(d) \rightsquigarrow \mathsf{Read}(d)}\ \text{bind}$$

$\boxed{C_1 \longrightarrow C_2}$ Configuration C_1 reduces to C_2.

$$\frac{e_1 \longrightarrow e_2}{\langle \Sigma; \pi, p : E[e_1] \rangle \longrightarrow \langle \Sigma; \pi, p : E[e_2] \rangle}\ \text{local} \qquad \frac{q \notin \Sigma}{\langle \Sigma; \pi, p : E[e_1 \mid\!\triangleright\ e_2] \rangle \longrightarrow \langle \Sigma, q; \pi, q : e_1, p : E[e_2] \rangle}\ \text{fork}$$

$$\frac{C_1 \equiv C_1' \qquad C_1' \longrightarrow C_2' \qquad C_2' \equiv C_2}{C_1 \longrightarrow C_2}\ \text{congr} \qquad \frac{d \notin \Sigma}{\langle \Sigma; \pi, p : E[\nu(x_1, x_2).\, e] \rangle \longrightarrow \langle \Sigma, d; \pi, p : E[[\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e] \rangle}\ \text{nu}$$

$$\frac{c_2 \rightsquigarrow c_1}{\langle \Sigma; \pi, p : E_1[\mathsf{rd}(c_1, x.e)], q : E_2[\mathsf{wr}(v, c_2)] \rangle \longrightarrow \langle \Sigma; \pi, p : E_1[[(!v, c_1)_1/x]e], q : E_2[()] \rangle}\ \text{rw}$$

$$\frac{c \rightsquigarrow c_i \qquad i \in \{1, 2\}}{\langle \Sigma; \pi, p : E_1[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)], q : E_2[\mathsf{wr}(v, c)] \rangle \longrightarrow \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1/x_i]e_i], q : E_2[()] \rangle}\ \text{cw}$$

$\boxed{e_1 \longrightarrow e_2}$ Expression e_1 reduces to e_2.

$$\frac{}{\mathsf{let}_\pi(v, x.e) \longrightarrow [v/x]e}\ \text{let} \qquad \frac{}{((\lambda_\ell\, x.\, e)\, v)_\ell \longrightarrow [v/x]e}\ \text{app} \qquad \frac{}{\mathsf{split}_\ell((v_1, v_2)_\ell, x_1.x_2.e) \longrightarrow [v_1/x_1][v_2/x_2]e}\ \text{split}$$

$$\frac{}{\mathsf{case}_\ell(\mathsf{inj}^i_\ell(v), x_1.e_1, x_2.e_2) \longrightarrow [v/x_i]e_i}\ \text{case} \qquad \frac{}{\mathsf{fix}_\ell(x.e) \longrightarrow [\mathsf{fix}_\ell(x.e)/x]e}\ \text{fix} \qquad \frac{}{¡\,(!\,v) \longrightarrow v}\ \text{gnab}$$

Figure 5. ILC reduction rules.

The nu rule extends the affine context Δ with a read endpoint $x_1 : \mathsf{Rd}\, S$ and the unrestricted context Γ with a corresponding write endpoint $x_2 : \mathsf{Wr}\, S$ before typing the body e.

The choice rule partitions the affine context as $\Delta_1, \Delta_2, \Delta_3$. The first two affine contexts are used to type $e_1 : \mathsf{Rd}\, S$ and $e_2 : \mathsf{Rd}\, T$, respectively. The third affine context Δ_3 is extended with the affine write token and a variable x_1 (or x_2) binding an affine 3-tuple containing the read value and the two read endpoints before checking the continuation e_3 (or e_4). While somewhat cumbersome, the generality of this rule allows both read endpoints to be used in either continuation.

3.3 Dynamic Semantics

Figures 4 and 5 define the dynamic syntax and semantics of ILC, respectively. We define a *configuration* C as a tuple of dynamic channel and process names Σ, and a pool of running and terminated processes π.

We read the configuration reduction judgment $C_1 \longrightarrow C_2$ as "configuration C_1 steps to configuration C_2," and the local stepping judgment $e_1 \longrightarrow e_2$ for a single process e as "expression e_1 steps to expression e_2." The rules of local stepping follow a standard call-by-value semantics, where we streamline the definition with an evaluation context E.

Configuration stepping consists of six rules. These include a congruence rule congr that permits some of the other rules to be simpler, by making the order of the pool unimportant. The relation $\pi_1 \equiv_{\mathsf{perm}} \pi_2$ holds when π_2 is a permutation of π_1. The other five rules consist of local stepping (via local), creating new processes (via fork), creating new channels (via nu), read-write interactions (via rw), and choice-write interactions (via cw). To avoid allocating the same name

twice, the name set Σ records names of allocated channels and processes. We define the relation $c_1 \rightsquigarrow c_2$ to hold when c_1 is the write endpoint of a corresponding read endpoint c_2.

4 ILC Metatheory

Intuitively, ILC's type system design enforces that a configuration's reduction consists of a unique (*deterministic*) sequence of reader-writer process pairings, and is *confluent* with any other reduction choice that exchanges the order of *other* (non-interactive) reduction steps. As explained in Section 3, ILC's type system does so by restricting the write effects (via an affine write token) and read effects (via affine read endpoints) of processes. The proofs of type soundness, whose statements we discuss next, establish the validity of these invariants. These language-level invariants support confluence theorems, also stated below. These theorems include *full confluence*: Any two full reductions of a configuration yield a pair of equivalent configurations (isomorphic, up to a renaming of nondeterministic name choices).

4.1 Type Soundness

We prove type soundness of ILC via mostly-standard notions of progress and preservation. To state these theorems, we follow the usual recipe, except that we give a special definition of program termination that permits deadlocks. (Recall that ILC is concerned with enforcing *confluence* as its central metatheoretic property, *not* deadlock freedom.) Informally, C **term** holds when either:

1. C is fully normal: Every process in C is normalized (consists of a value), or
2. C is (at least partially) deadlocked: Some (possibly empty) portion of C is normal, and there exists one or more reading processes in C, or there exists one or more writing processes in C, however, no reader-writer process pair exists for a common channel.

We also extend the type system given in Section 3.2 with typing rules for configurations, including process pool typings Φ from process names p to types U. These details, along with the proofs of progress and preservation, can be found in the appendix.

Theorem 4.1 (Progress). *If $\Psi \vdash C : \Phi$, then either C **term** or there exists C' such that $C \longrightarrow C'$.*

Theorem 4.2 (Preservation). *If $\Psi \vdash C : \Phi$ and $C \longrightarrow C'$, then there exists $\Psi' \supseteq \Psi$ and $\Phi' \supseteq \Phi$ such that $\Psi' \vdash C' : \Phi'$.*

4.2 Confluence

Confluence implies, among other things, that the order of reduction steps is inconsequential, and that no process scheduling choices will affect the final outcome. ILC's type system enforces confluence up to nondeterministic naming choices in rules nu and fork (Figure 5). To account for different choices of dynamically-named channels and processes, respectively, we state and prove confluence with respect to a renaming function f, which consistently renames these choices in a related configuration:

Theorem 4.3 (Single-step confluence). *For all well-typed configurations C, if $C \longrightarrow C_1$ and $C \longrightarrow C_2$, then there exists a renaming function f such that either:*

1. *$C_1 = f(C_2)$, or*
2. *there exists C_3 such that $C_1 \longrightarrow C_3$ and $f(C_2) \longrightarrow C_3$.*

Intuitively, the sister configuration C_2 is either different because of a name choice (case 1), or a different process scheduling choice (case 2). In either case, there exists a renaming of any choice made to reach C_2, captured by function f. By composing multiple uses of this theorem, and the renaming functions that they construct, we prove a multi-step notion of confluence that reduces a single configuration C to two equivalent terminal configurations, C_1 and C_2:

Theorem 4.4 (Full confluence). *For all well-typed configurations C, if $C \longrightarrow^* C_1$ and $C \longrightarrow^* C_2$ and C_1 **term** and C_2 **term**, then there exists renaming function f such that $C_1 = f(C_2)$.*

The proofs of these statements can be found in the appendix.

5 Implementation

Using this on-paper design as a guide, we have implemented an ILC interpreter in Haskell, which at present consists of 2.3K source lines of code. The implementation of ILC and our concrete implementation of the UC framework called SaUCy (Section 6) are publicly available. Access to the latest developments can be found here:

https://github.com/initc3/SaUCy.

6 SaUCy

Using ILC, we build a concrete, executable implementation of a simplified UC framework, dubbed SaUCy. Then, we demonstrate the versatility of SaUCy in three ways:

1. We define a protocol composition operator and prove its associated composition theorem.
2. We walk through an instantiation of UC commitments.
3. We use ILC's type system to reason about "reentrancy," a subtle definitional issue in UC that has only recently been studied.

6.1 Probabilistic Polynomial Time in ILC

The goal of cryptography reduction is to relate every bad event in a protocol to a *probabilistic polynomial time computation* that solves a hard problem. The ILC typing rules do not guarantee termination, let alone polynomial time normalization, so we must tackle this in metatheory. Also, since ILC is effectively deterministic (confluent), we will need to express random choices some other way. To meet these needs,

we define a judgment about ILC terms that take a security parameter and a stream of random bits.

Definition 6.1 (Polynomial time normalization)**.** The judgment that e is polynomial time normalizable, written PPT e, is defined as follows:

$$\frac{\begin{array}{c}\cdot;\cdot \vdash e : \mathsf{Nat} \to [\mathsf{Bit}] \to \mathsf{Bit} \\ \forall\, k \in \mathsf{Nat}.\ \forall\, r \in [\mathsf{Bit}]^{\mathrm{poly}(k)}.\ e\ k\ r \to^{\mathrm{poly}(k)} v\end{array}}{\mathsf{PPT}\ e}\ \text{ppt}$$

This says that if for all security parameters k and all bitstrings r (of length polynomial in k) the term $e\ k\ r$ normalizes to a value v in poly(k) steps, then PPT e.

Here, we have chosen a simple definition of polynomial time defined only for closed terms (i.e., an entire system of ITMs), and that requires polynomial time normalization for every choice of random bits, not just in expectation or with high probability.

We note that most UC variants use a more nuanced definition in which the individual ITM entities, such as the environment or protocol, can be judged polynomial time independently of their surrounding context [3, 16, 27]. Looking ahead to Section 6.3, this choice will constrain our definition of secure protocol emulation. Hofheinz et al. [28] give a detailed discussion of subtle issues arising with various polynomial time definitions and their consequences for defining UC security. Regardless, the present notion suffices for our examples. We consider this issue complementary to the design of ILC itself, and adapting other notions of polynomial time to ILC as important future work. As an example, the polynomial time notion used in IITMs [3] relies on a distinction between "invited" and "uninvited" messages, which could be captured through refinement types à la the RCF calculus [12].

Definition 6.2 (Value Distribution)**.** Because processes are confluent, we know that if $e\ k\ r \to^* v$, then the value v is unique. We can therefore define the probability distribution ensemble $D(e) = \{D_{e,k}\}_k$ based on a uniform distribution U_k over k-bit strings r, so the distribution $D_{e,k}$ is given as

$$D_{e,k}(v) = \sum_{r \in R} U_k(r), \quad \text{for } R = \{r \mid e\ k\ r \to^* v\}.$$

Definition 6.3 (Indistinguishability)**.** What remains is to define a notion of indistinguishability for value distributions. However, we need to clarify when polynomial time normalization is an assumption or a proof obligation. To simplify things later, we define a partial order $e_1 \leq e_2$, which captures that e_2 must be PPT if e_1 is PPT, and if so, that their value distributions are statistically similar.

$$\frac{\mathsf{PPT}\ e_1 \implies (\mathsf{PPT}\ e_2 \text{ and } D(e_1) \sim D(e_2))}{e_1 \leq e_2}\ \text{indist}$$

6.2 SaUCy Execution Model

The implementation of SaUCy is centered around a definition of the UC execution model in ILC. For space and readability, we elide endpoint allocation/distribution with ellipses. We also abbreviate the type signature (e.g., A_z is the type of z). More details can be found in the appendix.

```
execUC :: ∀ ... . A_z →_w A_p × A_q → A_f → A_a →
                  Crupt → Nat → [Bit] → Bit
let execUC z (p,q) f a crupt k r =
  ν ... . let (rf,ra,rp,rq,rz) = splitBits r in
            f k rf crupt ...
          |▷ a k ra crupt ...
          |▷ corruptOrNot p k rp (crupt == CruptP) ...
          |▷ corruptOrNot q k rq (crupt == CruptQ) ...
          |▷ z k rz ...
```

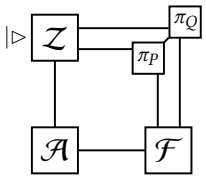

Figure 6. execUC.

The function execUC takes as arguments an environment z, a pair of protocol processes (p, q), a functionality f, an adversary a, a corruption model crupt, a security parameter k, and a random bitstring r. At a high level (ignoring details related to corruptions for now), it runs each of the processes (allocating random bits to each of them) and connects channels as illustrated in Figure 6. (The protocol processes p and q correspond to π_P and π_Q, respectively.) The execution is centered on the environment z in the sense that z first gets the write token (notably, it has type $\mathsf{Nat} \to_w \cdots \to \mathsf{Bit}$), and the experiment concludes when z returns a single bit value.

Next, we explain some of our main modeling choices and the consequences they have for the ILC implementation. To start with, we make several simplifications to standard UC, for example, focusing on the special case of two-party protocols (à la Simplified UC [18]). We also only aim to show the case of *static* corruptions, in which the corrupt parties are determined at the onset. This is achieved by parameterizing the entire experiment by a value crupt : Crupt denoting which parties are corrupt (if any). The data type Crupt is defined as follows.

```
data Crupt = CruptP | CruptQ | CruptNone
```

For a more general model with adaptive corruptions, execUC would need to accept requests from the environment to add to the crupt list as the execution proceeds.

Our corruption model is Byzantine, meaning the adversary gets to exert complete control over the corrupted parties. For each party, depending on the value of crupt, either we run a copy of the honest party, or connect the channels to the adversary. This is implemented in the function corruptOrNot.

```
fwd :: ∀ a b . Wr a → Rd a ⊸ b
letrec fwd toR frS =
```

```
  let (!msg, frS) = rd frS in wr msg → toR ; fwd toR frS

corruptOrNot :: ∀ ... . A_p → Nat → [Bit] → Bool → ···
let corruptOrNot p k bits iscrupt toZ toF toA toQ
                                    frZ frF frA frQ =
    if iscrupt then
        let _ = rd frZ in error "Z can't wr to corrupt"
      |▷ fwd toA frF
      |▷ fwd toA frQ
      |▷ fwd toF frA
    else
      p k bits toZ toF toQ frZ frF frQ
```

The fwd function simply forwards messages received on the read endpoint frS to the write endpoint toR. In corruptOrNot, if a party is corrupted, messages from the functionality and the other protocol party are forwarded to the adversary; messages from the adversary are forwarded to the functionality. Otherwise, the party is run as normal.

We also model a strong form of communication channels between the parties: P and Q are connected by a pair of raw ILC channels. Communication over these channels happens immediately, without activating the adversary or leaking even the existence of the message. In a more realistic model, the parties would only be able to communicate over a network channel modeled as a functionality, $\mathcal{F}_{\text{SMT}}$ or $\mathcal{F}_{\text{SYN}}$ [16]. Consequently our $\mathcal{F}_{\text{COM}}$ functionality would need to be weakened by leaking some (model-specific) information about the message to the adversary.

6.3 Defining UC Security in ILC

The central security definition in UC is protocol emulation. The guiding principle is that π emulates ϕ if the environment cannot distinguish between the two protocols. Our first attempt is the following, where $\mathcal{S}$ is the simulator that translates every attack in the real world into an attack expressed in the ideal world:

$$\frac{\forall\, \mathcal{Z}.\ \text{execUC}\ \mathcal{Z}\ \pi\ \mathcal{F}_1\ \mathbb{1}_{\mathcal{A}} \leq \text{execUC}\ \mathcal{Z}\ \phi\ \mathcal{F}_2\ \mathcal{S}}{\mathcal{S} \vdash (\pi, \mathcal{F}_1) \approx (\phi, \mathcal{F}_2)}\ \text{~~emulate~~}$$

To remark on a few notational choices: We make the functionality explicit, so emulation is a relationship between protocol-functionality pairs. Here, $\mathbb{1}_{\mathcal{A}}$ is the dummy adversary, which just relays messages between the environment and the parties/functionality. We elide the standard dummy lemma that shows this is without loss of generality; the intuition is that whatever an adversary can do, the environment can achieve using $\mathbb{1}_{\mathcal{A}}$.

Unfortunately this simple definition turns out to be vacuous: a degenerate protocol π can emulate anything simply failing to be PPT, e.g., by diverging. To put it another way, the problem is the definition imposes a proof obligation on the simulator $\mathcal{S}$ but not on π. What we want to say is that the real world protocol $(\pi, \mathcal{F}_1)$ must be well behaved whenever the ideal world $(\phi, \mathcal{F}_2)$ is. However, even a reasonable protocol can result in non-PPT executions if paired with a divergent environment. To solve this problem, we define protocol emulation by requiring a simulation in both directions, so every behavior in the ideal world must correspond to a behavior in the real world and vice versa.

Definition 6.4 (Protocol Emulation)**.** The judgment that one protocol-functionality pair $(\pi, \mathcal{F}_1)$ securely emulates another $(\phi, \mathcal{F}_2)$ (as proven by the simulators $\mathcal{S}_{\mathcal{R}}, \mathcal{S}_{\mathcal{I}}$) is defined as

$$\frac{\begin{array}{c}\forall\, \mathcal{Z}.\ \text{execUC}\ \mathcal{Z}\ \phi\ \mathcal{F}_2\ \mathbb{1}_{\mathcal{A}} \leq \text{execUC}\ \mathcal{Z}\ \pi\ \mathcal{F}_1\ \mathcal{S}_{\mathcal{R}} \\ \text{execUC}\ \mathcal{Z}\ \pi\ \mathcal{F}_1\ \mathbb{1}_{\mathcal{A}} \leq \text{execUC}\ \mathcal{Z}\ \phi\ \mathcal{F}_2\ \mathcal{S}_{\mathcal{I}}\end{array}}{\mathcal{S}_{\mathcal{R}}, \mathcal{S}_{\mathcal{I}} \vdash (\pi, \mathcal{F}_1) \approx (\phi, \mathcal{F}_2)}\ \text{emulate}$$

We remark this definition goes against the UC convention of requiring simulation in one direction only. One direction is preferable intuitively because it should be fine if the protocol is even more secure than its specification. This does not pose any problem for our commitment example; however, a protocol that leaks even less information than its ideal functionality requires would be impossible to prove secure under this definition. In any case, the benefit is this simplifies the polynomial time notion: vacuous protocols are clearly ruled out by the top condition, and both simulations are only required to be PPT when the environment $\mathcal{Z}$ is well-behaved.

6.4 A Composition Theorem in SaUCy

As a first demonstration of SaUCy, we work through the development of a composition operator, and give a theorem explaining its use.

Definition 6.5 (UC realizes)**.** To set out, we introduce the notation of "realizes," which views a protocol as a way of instantiating a specification functionality $\mathcal{F}_2$ from a setup assumption functionality $\mathcal{F}_1$,

$$\frac{(\pi,\ \mathcal{F}_1) \approx (\text{id}_\pi, \mathcal{F}_2)}{\mathcal{F}_1 \xrightarrow{\pi} \mathcal{F}_2}\ \text{realizes}$$

where id_π is the *dummy protocol*, which simply relays messages between the environment and the functionality. This notation is convenient because it suggests a categorical approach to composition.

Theorem 6.1 (Composition Theorem)**.**

$$\frac{\mathcal{F}_1 \xrightarrow{\pi} \mathcal{F}_2 \qquad \mathcal{F}_2 \xrightarrow{\rho} \mathcal{F}_3}{\mathcal{F}_1 \xrightarrow{\rho \circ \pi} \mathcal{F}_3}$$

The idea is that the $\rho \circ \pi$ can be defined in a natural way, where the ideal functionality channel of ρ is connected to the environment channel of π, as illustrated and defined in Figure 7.

Proof. To prove the theorem we construct the simulators $S_{\mathcal{R},\rho} \circ S_{\mathcal{R},\pi}$ (respectively $S_{\mathcal{I},\rho} \circ S_{\mathcal{I},\pi}$) in the natural way as

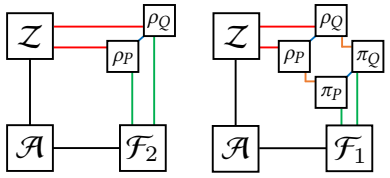


```
let (∘) (ρP, ρQ) (πP, πQ)
                wρP2Z wρQ2Z wρP2F wρQ2F
                wρP2ρQ wρQ2ρP rZ2ρP rZ2ρQ
                rF2ρP rF2ρQ rρQ2ρP rρP2ρQ =
  ν ... . πP wπP2ρP wρP2F wπP2πQ rρP2πP rF2ρP rπQ2πP
    |▷ πQ wπQ2ρQ wρQ2F wπQ2πP rρQ2πQ rF2ρQ rπP2πQ
    |▷ ρP wρP2Z wρP2πP wρP2ρQ rZ2ρP rπP2ρP rρQ2ρP
    |▷ ρQ wρQ2Z wρQ2πQ wρQ2ρP rZ2ρQ rπQ2ρQ rρP2ρQ
```

Figure 7. Protocol composition operator.

well (see the appendix). Our proof obligation is to introduce an arbitrary environment $\mathcal{Z}$ and conclude

$$\text{execUC}\ \mathcal{Z}\ (\rho \circ \pi)\ \mathcal{F}_1\ \mathbb{1}_{\mathcal{A}} \leq \text{execUC}\ \mathcal{Z}\ \mathbb{1}_{\pi}\ \mathcal{F}_3\ (\mathcal{S}_{I,\rho} \circ \mathcal{S}_{I,\pi}).$$

The main idea is to notice that that we can bring ρ from the composed protocol into the environment as $(\mathcal{Z} \circ \rho)$, reflecting the fact that the environment is meant to represent arbitrary outer protocols. This transformation results in an equivalent term, given that ILC configurations are invariant to channel renaming and reordering of processes in a configuration (as in Section 4). The following derivation completes the proof:

$$\begin{aligned}
&\text{execUC}\ \mathcal{Z}\ (\rho \circ \pi)\ \mathcal{F}_1\ \mathbb{1}_{\mathcal{A}} \\
\equiv\ &\text{execUC}\ (\mathcal{Z} \circ \rho)\ \pi\ \mathcal{F}_1\ \mathbb{1}_{\mathcal{A}} && \text{(By equivalence)} \\
\leq\ &\text{execUC}\ (\mathcal{Z} \circ \rho)\ \text{id}_{\pi}\ \mathcal{F}_2\ \mathcal{S}_{I,\pi} && \text{(From } \mathcal{F}_1 \xrightarrow{\pi} \mathcal{F}_2\text{)} \\
\equiv\ &\text{execUC}\ (\mathcal{S}_{I,\pi} \circ \mathcal{Z})\ \rho\ \mathcal{F}_2\ \mathbb{1}_{\mathcal{A}} && \text{(By equivalence)} \\
\leq\ &\text{execUC}\ (\mathcal{S}_{I,\pi} \circ \mathcal{Z})\ \text{id}_{\pi}\ \mathcal{F}_3\ \mathcal{S}_{I,\rho} && \text{(From } \mathcal{F}_2 \xrightarrow{\rho} \mathcal{F}_3\text{)} \\
\equiv\ &\text{execUC}\ \mathcal{Z}\ \text{id}_{\pi}\ \mathcal{F}_3\ (\mathcal{S}_{I,\pi} \circ \mathcal{S}_{I,\rho}) && \text{(By equivalence)}
\end{aligned}$$

The remaining case for $\mathcal{S}_{\mathcal{R},\rho} \circ \mathcal{S}_{\mathcal{R},\pi}$ is symmetric. □

Other notions of composition. Our composition operator above is just a starting point. The "universal composition" [16] operator essentially multiplexes sessions identified by unique tags (*session ids*), while a joint state composition theorem collapses multiple subroutines into one [20]. Despite its name, development in UC often involves defining additional composition operators. For example, interesting composition often happens "in the functionality" through higher order "wrapper" functionalities [29, 31] which we would express through abstraction. Some security properties require a generalized notion of ideal functionality that the environment can interact with directly. All the above motivate the development of the ILC core calculus as a flexible foundation; developing them in ILC is important future work.

6.5 Instantiating UC Commitments

We next walk through an instantiation of UC commitments (à la Canetti and Fischlin [19]). Instantiation proofs in SaUCy follow a standard rhythm. We start with a security definition as an ideal functionality (such as $\mathcal{F}_{\text{COM}}$), give the protocol, construct a simulator, and finally complete the relational analysis on paper.

While commitments are one of the simplest UC primitives, as a case study, this serves two main purposes. First, the proof demonstrates several representative UC techniques [36], in particular the simulator makes use of a "trusted setup" and extracts inputs from a corrupt sender. Second, the protocol makes use of computational primitives and thus requires a reduction step in the proof, which can go through because of ILC's confluent design.

Extending ILC with cryptographic primitives. The UC commitment protocol makes use of a cryptographic primitive, namely a trapdoor pseudorandom generator. This is provided by extending ILC with new syntactic forms, along with their static and dynamic semantics (given in the appendix). While in a symbolic setting we would instantiate these with algebraic data, in ILC we give the stepping rule in terms of an arbitrary pseudorandom function family, i.e., the actual computational definition. This can be instantiated concretely for execution (e.g., with an RSA-based function) or treated abstractly in the metatheory when we get to the reduction step of the proof.

The commitment protocol also relies on a "trusted setup," or common reference string (CRS), which is essentially public parameters generated ahead of time. The common reference string is modeled as an ideal functionality $\mathcal{F}_{\text{CRS}}$ (implemented in ILC as fCrs in the appendix).

Commitment Protocol. We implement the commitment protocol by Canetti and Fischlin [19] in ILC as follows:

```
committer :: ∀ ... . Nat → [Bit] → ··· ⊸ 𝟙
let committer k bits crupt toZ toF toQ frZ frF frQ =
  let (!(Commit b), frZ)= rd frZ in
    wr GetCRS → toF ;
    let (!(PublicStrings σ pk0pk1), frF) = rd frF in
    let r = take k bits in
    let x = if b == 0 then prg pk0 r
            else xors (prg pk1 r) σ in
      wr Commit' x → toQ ;
      let (!Open, frZ)= rd frZ in
        wr (Open' b r) → toQ

receiver :: ∀ ... . Nat → [Bit] → ··· ⊸ 𝟙
let receiver k bits crupt toZ toF toP frZ frF frP =
  let (!(Commit' x), frP) = rd frP in
    wr GetCRS → toF ;
    let (!(PublicStrings σ pk0pk1), frF) = rd frF in
```

```
    wr Receipt ⟶ toZ ;
    let (!(Open' b r), frP) = rd frP in
      if (b == 0 && x == prg pk0 r) ||
         (b == 1 && x == xors (prg pk1 r) σ)
      then wr (Opened b) ⟶ toZ
      else error "Cannot occur in honest case."
```

To briefly summarize what is going on: The setup CRS functionality fCom samples a random string σ and two trapdoor pseudorandom generator (PRG) keys pk_0 and pk_1. To commit to b, the committer produces a string y that is the result of applying one or the other of the PRGs, and if $b = 1$ additionally applying xor with σ. The intuitive explanation why this is hiding is that without the trapdoor, it is difficult to tell whether a random $4k$-bit string is in the range of either PRG. To open the commitment, the committer simply reveals the preimage and the receiver checks which of the two cases applies. The intuitive explanation why this is binding is that it is difficult to find a pair $y, y \oplus \sigma$ that are respectively in the range of both PRGs.

Defining the simulator. The SaUCy proof consists of two simulators, one for the ideal world and one for the real world. The ideal world simulator is ported directly from the UC literature [19]. The nonstandard real world simulator, given in the appendix, is trivial, but necessary because our protocol emulation definition requires simulation in both directions.

The ideal world simulator generates its own "fake" CRS for which it stores the trapdoors. The string σ is not truly random, but instead is the result of combining two evaluations of the PRGs. In Figure 8, we show the case that the committer P is corrupt (the other case is in the appendix). The simulator is activated when $\mathcal{Z}$ sends a message (Commit' y); in the real world, this is relayed by the dummy adversary to Q, who outputs Receipt back to the environment. Hence to achieve the same effect in the ideal word, the simulator must send (Commit b) to $\mathcal{F}_{\text{COM}}$. To extract b from y, the simulator makes use of the PRG trapdoor check which one has y in its range. It is necessary to argue by cryptographic reduction that this simulation is sound, which we do next.

Relational argument. The goal of the relational analysis is to show that an environment's output in the real world is indistinguishable from its output in the ideal world. The proof follows the one in Canetti and Fischlin [19].

Proof Sketch. Consider the following ensembles:

$$D_{\mathcal{R}} = D(\text{execUC } \mathcal{Z} \text{ (committer, receiver) fCrs dummyA})$$
$$D'_{\mathcal{R}} = D(\text{execUC } \mathcal{Z} \text{ (committer, receiver) bCrs dummyA})$$
$$D_{\mathcal{I}} = D(\text{execUC } \mathcal{Z} \text{ (dummyP, dummyQ) fCom simI})$$

The ensemble $D_{\mathcal{R}}$ is over the output of $\mathcal{Z}$ in a real world execution. The ensemble $D'_{\mathcal{R}}$ is similar, except $\mathcal{Z}$ runs with a bad functionality bCrs (see appendix) that computes fake public strings in the same way that the simulator does. The

```
let simI k bits crupt toZ toF toP toQ frZ frF frP frQ =
  let (pk0,td0) = kgen k in
  let (pk1,td1) = kgen k in
  let (r0, bits) = sample k bits in
  let (r1, bits) = sample k bits in
  let σ = xors (prg pk0 r0) (prg pk1 r1) in
    match crupt with
    | CruptP ⇒
      let (!GetCRS, frZ) = rd frZ in
        wr (X2Z (PublicStrings σ pk0 pk1)) ⟶toZ ;
        let (!(A2P (Commit' y)), frZ) = rd frZ in
          if check td0 pk0 y then
            wr (Commit 0) ⟶ toP
          else
            if check td1 pk1 (xors y σ) then
              wr (Commit 1) ⟶ toP
            else error "Fail" ;
          let (!(A2P (Open' b r)), frZ) = rd frZ in
            if b == 0 && y == prg pk0 r ||
               b == 1 && y == xors (prg pk1 r) σ
            then wr Open ⟶ toP
            else error "Fail"
    | ...
```

Figure 8. Ideal world simulator (excerpt) for UC commitment (full version in appendix).

ensemble $D_{\mathcal{I}}$ is over the output of $\mathcal{Z}$ in an ideal world execution. The goal is to show that $D_{\mathcal{R}} \sim D_{\mathcal{I}}$. The proof proceeds by first showing that breaking the pseudorandomness of the PRG reduces to distinguishing between $D_{\mathcal{R}}$ and $D'_{\mathcal{R}}$ (hence, $D_{\mathcal{R}} \sim D'_{\mathcal{R}}$), and then by showing that breaking the pseudorandomness of the PRG also reduces to distinguishing between $D'_{\mathcal{R}}$ and $D_{\mathcal{I}}$ (hence, $D'_{\mathcal{R}} \sim D_{\mathcal{I}}$). By the transitivity of indistinguishability, we have that $D_{\mathcal{R}} \sim D_{\mathcal{I}}$. □

Here, ILC's confluence property plays a critical role: It is necessary for defining the probability ensembles $D_{\mathcal{R}}$, $D'_{\mathcal{R}}$, and $D_{\mathcal{I}}$, without which we would not be able to obtain a reduction from some computationally hard problem to distinguishing the real world and ideal world ensembles.

6.6 Reentrancy in SaUCy

Camenisch et al. [14] recently identified subtleties in defining UC ideal functionalities (related to reentrancy and the scheduling of concurrent code) such that several functionalities in the literature are ambiguous as ITMs. Although concerning, these issues have no cryptographic flavor, and so they are better addressed from a PL standpoint. To illustrate, consider the following (untypeable) ILC process reentrantF, which

allows an adversary $\mathcal{A}$ to control the delivery schedule of messages from P to Q (i.e., an asynchronous channel):

```
loop :: ∀ a b . (a ⟶ b) ⟶ Rd a ⊸ b
letrec loop f frS = let (!v, frS) = rd frS in f v; loop f frS
let reentrantF ... frP frA =
  loop (λ msg . (let (!Ok, frA) = rd frA in wr msg ⟶ toQ)
                |▷ wr msg ⟶ toA) frP
```

After receiving input from party P, it notifies the adversary, then forks a background thread to wait for Ok before delivering the message. This introduces a race condition: Suppose input message m_1 is sent by P, but then $\mathcal{A}$, before sending Ok, instead returns control to $\mathcal{Z}$, which passes P a second input m_2. Now there are two queued messages. Which one gets delivered when the adversary sends Ok?

To resolve this issue, notice that reentrantF is untypeable in ILC. The race condition occurs because the read endpoint frA is duplicated (appears free in an unrestricted function). Camenisch et al. [14] identified several strategies for resolving this problem in UC, which in turn are expressible ILC. One approach is to make the process explicitly sequential, such that the arrival of a second message before the first is delivered causes execution to get stuck:

```
letrec sequentialF ... frP frA =
  let (!msg, frP) = rd frP in
    wr msg ⟶ toA ;
    let (!Ok, frA) = rd frA in
      wr msg ⟶ toQ ;
      sequentialF ... frP frA
```

Alternatively, we may discard such messages arriving out of order, returning them to sender; we express this in ILC using the external choice operator:

```
letrec discardingF ... frP frA =
  let (!msg, frP) = rd frP in
    wr msg ⟶ toA ;
    letrec iloop () frP frA =
      choice
      | (_,frP,frA)@(rd frP) ⇒ wr Discard ⟶ toP ;
                               iloop () frP frA
      | (_,frP,frA)@(rd frA) ⇒ wr msg ⟶ toQ ;
                               discardingF ... frP frA
      in iloop () frP frA
```

Ultimately, Camenisch et al. propose a different strategy, which is to restrict how the environment/adversary respond to certain "urgent" messages that are used to exchange meta-information (modeling related messages). That is, upon receiving an urgent message from process P, the environment (or adversary) must return control back to P immediately. Modeling this solution is left as future work, but ILC provides an ideal starting point—restrictions on the environment/adversary could be expressed by behavior refinements: upon receiving an urgent message from P, the environment (or adversary) must not send a message on its other channels before sending a message to P.

7 Related Work

7.1 Process Calculi

Process calculi have a long and rich history. ILC occupies a point in this space that is particularly suited to faithfully capturing interactive Turing machines (and hence, computational cryptography), but plenty of existing calculi are also cryptographically-flavored and/or enjoy similar properties to ILC. We survey some of them here.

With symbolic semantics. Two early adaptations of process calculi for reasoning about cryptographic protocols were the spi calculus [2] and the applied π-calculus [1], both of which extend the π-calculus with cryptographic operations [43]. Symbolic UC [10] is a simulation-based security framework in this setting. However, protocols proven secure in the symbolic setting may not be realizable with any cryptographic primitives based on hardness assumptions.

With computational semantics. Naturally, ensuing work has turned to bridging the gap between this PL-style of formalization and the computational model of cryptography by outfitting these calculi with a computational semantics. Lincoln et al. [35] give a computational semantics to a variant of the π-calculus, which allows one to define communicating probabilistic polynomial-time processes; Mateus et al. [38] adapts their calculus to explore (sequential) compositionality properties in protocols. A drawback of these protocols is that they embed probabilistic choices directly into the definition—essentially when faced with nondeterminism, each path has equal probability. Laud [33] gives a computational semantics to the spi calculus, which additionally includes a type system for ensuring well-typed protocols preserve the secrecy of messages given to it by users.

With confluence. There are a number of other process calculi that enjoy confluence. Berger et al. [6] describe a type system for capturing deterministic (sequential) computation in the π-calculus. The type system uses affineness and stateless replication to achieve deterministic computation. Fowler et al. [24] present a core linear lambda calculus with (binary) session-typed channels and exception handling that enjoys confluence and termination. The calculus only considers two-party protocols, so for our multiparty setting, ILC requires a sophisticated type system to achieve confluence.

7.2 Tools for Cryptographic Analysis

Computer-aided tools for cryptographic analysis operate in either the symbolic model or the computational model. The survey by Blanchet [8] highlights some of their differences.

Symbolic tools include the NRL protocol analyzer [41], Maude-NPA [23], and Proverif [9]. In the symbolic setting,

cryptographic operations are abstracted as term algebras (a variant of the applied π-calculus in the case of Proverif), and adversary capabilities are nondeterministic applications of deduction rules over these terms. Here, nondeterminism allows the adversary to find attack traces (if there are any), whereas the presence of nondeterminism in the computational setting would frustrate cryptographic reduction proofs.

Computational tools include CertiCrypt [5], EasyCrypt [4], CryptoVerif [7], Cryptol [34], and F* [45]. These tools focus on game-based security, which, in contrast to simulation-based definitions (such as UC), only guarantee security in a standalone setting (no composition guarantees). While they are not specifically purposed for simulation-based proofs, it would be interesting to embed ILC into EasyCrypt or F* to use their tooling.

7.3 Variations of Universal Composability

A number of models for universal composability have been proposed in the literature [3, 10, 13–18, 20, 27, 39, 40, 44]. We highlight a few that have similar goals to ours.

In contrast with UC, which uses ITMs as its computational model, the reactive simulatability framework (RSIM) [3] uses probabilistic IO automata, which are amenable to automated reasoning. In contrast with RSIM, ILC is intended to be the basis for a convenient and flexible programming language to which we can easily port existing UC pseudocode.

Models based on inexhaustible interactive Turing machines (IITMs) [15, 32] aim to address drawbacks of UC models for which polynomial time ITMs can be "exhausted" (by having other machines send useless messages, forcing them to halt). In turn, models with exhaustible ITMs are less expressive. Because IITMs maintain the "single-threaded" execution semantics of ITMs, ILC can be used to build a concrete programming model for IITM-based frameworks as well.

The abstract cryptography framework [40] advocates a top-down approach: developing theory at an abstract level (ignoring low level details such as computational models and complexity notions) to simplify definitions. While we stick to a bottom-up approach, we aim to simplify UC via PL formalisms.

Simplified universal composability (SUC) [18] gives a simpler and restricted variant of the UC framework. The main difference from vanilla UC [16] is that the set of parties is fixed, which greatly simplifies polynomial time reasoning and protocol composition while maintaining the same strong properties. We follow this in our execUC implementation.

8 Conclusion and Future Work

The universal composability (UC) framework is widely used in cryptography for proofs. SaUCy takes a step towards mechanizing UC as a programming framework for constructing and analyzing large systems. We envision using SaUCy to tackle, for example, applications involving blockchains and smart contracts [21, 22, 42], which comprise an array of cryptography and distributed computing components and suffer from increasingly unwieldy formalisms.

We can view ILC typechecking of simulators in SaUCy as a partial mechanization of UC proofs, though the indistinguishability analysis is still on paper. Even partial mechanization is useful for catching bugs; we imagine using SaUCy to systematically implement functionalities and protocols from the literature and fuzz test them. Future work would be to embed ILC within a mechanized proof system, such as F* or EasyCrypt.

Acknowledgements

We thank our shepherd, Amal Ahmed, and the anonymous reviewers for their valuable feedback. This material is based on work supported by the National Science Foundation under Grant No. 1801321 and a Graduate Research Fellowhip.

References

[1] Martín Abadi and Cédric Fournet. 2001. Mobile values, new names, and secure communication. In *ACM Sigplan Notices*, Vol. 36. ACM, 104–115.

[2] Martın Abadi and Andrew D Gordon. 1999. A calculus for cryptographic protocols: The spi calculus. *Information and computation* 148, 1 (1999), 1–70.

[3] Michael Backes, Birgit Pfitzmann, and Michael Waidner. 2007. The reactive simulatability (RSIM) framework for asynchronous systems. *Information and Computation* 205, 12 (2007), 1685–1720.

[4] G. Barthe, B. Grégoire, S. Heraud, and S. Béguelin. 2011. Computer-aided security proofs for the working cryptographer. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*.

[5] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. 2009. Formal certification of code-based cryptographic proofs. *ACM SIGPLAN Notices* 44, 1 (2009), 90–101.

[6] Martin Berger, Kohei Honda, and Nobuko Yoshida. 2001. Sequentiality and the π-calculus. In *International Conference on Typed Lambda Calculi and Applications*. Springer, 29–45.

[7] Bruno Blanchet. 2007. CryptoVerif: Computationally sound mechanized prover for cryptographic protocols. In *Dagstuhl seminar "Formal Protocol Verification Applied*. 117.

[8] Bruno Blanchet. 2012. Security protocol verification: Symbolic and computational models. In *Proceedings of the First international conference on Principles of Security and Trust*. Springer-Verlag, 3–29.

[9] Bruno Blanchet, V Cheval, X Allamigeon, and B Smyth. 2010. Proverif: Cryptographic protocol verifier in the formal model. *URL http://prosecco. gforge. inria. fr/personal/bblanche/proverif* (2010).

[10] Florian Böhl and Dominique Unruh. 2016. Symbolic universal composability. *Journal of Computer Security* 24, 1 (2016), 1–38.

[11] Gilles Brassard, David Chaum, and Claude Crépeau. 1988. Minimum disclosure proofs of knowledge. *J. Comput. System Sci.* 37, 2 (1988), 156–189.

[12] Michele Bugliesi, Stefano Calzavara, Fabienne Eigner, and Matteo Maffei. 2015. Affine refinement types for secure distributed programming. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 37, 4 (2015), 11.

[13] Jan Camenisch, Manu Drijvers, and Björn Tackmann. [n. d.]. Multi-Protocol UC and its Use for Building Modular and Efficient Protocols.

([n. d.]).

[14] Jan Camenisch, Robert R Enderlein, Stephan Krenn, Ralf Küsters, and Daniel Rausch. 2016. Universal composition with responsive environments. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 807–840.

[15] Jan Camenisch, Stephan Krenn, Ralf Küsters, and Daniel Rausch. [n. d.]. iUC: Flexible Universal Composability Made Simple (Full Version). ([n. d.]).

[16] R. Canetti. 2001. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*.

[17] Ran Canetti, Ling Cheung, Dilsun Kaynar, Moses Liskov, Nancy Lynch, Olivier Pereira, and Roberto Segala. 2008. Analyzing security protocols using time-bounded task-PIOAs. *Discrete Event Dynamic Systems* 18, 1 (2008), 111–159.

[18] Ran Canetti, Asaf Cohen, and Yehuda Lindell. 2015. A simpler variant of universally composable security for standard multiparty computation. In *Annual Cryptology Conference*. Springer, 3–22.

[19] Ran Canetti and Marc Fischlin. 2001. Universally composable commitments. In *Annual International Cryptology Conference*. Springer, 19–40.

[20] Ran Canetti and Tal Rabin. 2003. Universal composition with joint state. In *Annual International Cryptology Conference*. Springer, 265–281.

[21] Stefan Dziembowski, Lisa Eckey, Sebastian Faust, and Daniel Malinowski. [n. d.]. *Perun: Virtual payment channels over cryptographic currencies*. Technical Report.

[22] Stefan Dziembowski, Sebastian Faust, and Kristina Hostáková. 2018. General State Channel Networks. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 949–966.

[23] Santiago Escobar, Catherine Meadows, and José Meseguer. 2009. Maude-NPA: Cryptographic protocol analysis modulo equational properties. In *Foundations of Security Analysis and Design V*. Springer, 1–50.

[24] Simon Fowler, Sam Lindley, J Garrett Morris, and Sára Decova. 2018. Session Types without Tiers. (2018).

[25] Simon J Gay and Vasco T Vasconcelos. 2010. Linear type theory for asynchronous session types. *Journal of Functional Programming* 20, 1 (2010), 19–50.

[26] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to play any mental game. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*. ACM, 218–229.

[27] Dennis Hofheinz and Victor Shoup. 2015. GNUC: A new universal composability framework. *Journal of Cryptology* 28, 3 (2015), 423–508.

[28] Dennis Hofheinz, Dominique Unruh, and Jörn Müller-Quade. 2013. Polynomial runtime and composability. *Journal of Cryptology* 26, 3 (2013), 375–441.

[29] Jonathan Katz. 2007. Universally composable multi-party computation using tamper-proof hardware. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 115–128.

[30] Naoki Kobayashi, Benjamin C Pierce, and David N Turner. 1999. Linearity and the pi-calculus. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 21, 5 (1999), 914–947.

[31] Ahmed Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. 2016. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *2016 IEEE symposium on security and privacy (SP)*. IEEE, 839–858.

[32] Ralf Kusters. 2006. Simulation-based security with inexhaustible interactive turing machines. In *Computer Security Foundations Workshop, 2006. 19th IEEE*. IEEE, 12–pp.

[33] Peeter Laud. 2005. Secrecy types for a simulatable cryptographic library. In *Proceedings of the 12th ACM conference on Computer and communications security*. ACM, 26–35.

[34] Jeffrey R Lewis and Brad Martin. 2003. Cryptol: High assurance, retargetable crypto development and validation. In *Military Communications Conference, 2003. MILCOM'03. 2003 IEEE*, Vol. 2. IEEE, 820–825.

[35] Patrick Lincoln, John Mitchell, Mark Mitchell, and Andre Scedrov. 1998. A probabilistic poly-time framework for protocol analysis. In *Proceedings of the 5th ACM conference on Computer and communications security*. ACM, 112–121.

[36] Yehuda Lindell. 2017. How to simulate it–a tutorial on the simulation proof technique. In *Tutorials on the Foundations of Cryptography*. Springer, 277–346.

[37] Yehuda Lindell and Jonathan Katz. 2014. *Introduction to modern cryptography*. Chapman and Hall/CRC.

[38] Paulo Mateus, J Mitchell, and Andre Scedrov. 2003. Composition of cryptographic protocols in a probabilistic polynomial-time process calculus. In *International Conference on Concurrency Theory*. Springer, 327–349.

[39] Ueli Maurer. 2011. Constructive cryptography–a new paradigm for security definitions and proofs. In *Theory of Security and Applications*. Springer, 33–56.

[40] Ueli Maurer and Renato Renner. 2011. Abstract cryptography. In *In Innovations in Computer Science*. Citeseer.

[41] Catherine Meadows. 1996. The NRL protocol analyzer: An overview. *The Journal of Logic Programming* 26, 2 (1996), 113–131.

[42] Andrew Miller, Iddo Bentov, Ranjit Kumaresan, and Patrick McCorry. 2017. Sprites: Payment channels that go faster than lightning. *CoRR abs/1702.05812* (2017).

[43] Robin Milner. 1999. *Communicating and mobile systems: the pi calculus*. Cambridge university press.

[44] Birgit Pfitzmann and Michael Waidner. 2001. A model for asynchronous reactive systems and its application to secure message transmission. In *Security and Privacy, 2001. S&P 2001. Proceedings. 2001 IEEE Symposium on*. IEEE, 184–200.

[45] N. Swamy, C. Hriţcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, et al. 2016. Dependent types and multi-monadic effects in F*. In *Proceedings of the Symposium on Principles of Programming Languages (POPL)*.

A Algorithmic Typing Rules

$\boxed{\Delta_{in}; \Gamma \vdash e : U \dashv \Delta_{out}}$ Under input contexts Δ_{in} and Γ, expression e has type U and output context Δ_{out}.

$$\frac{}{\Delta; \Gamma, x : A \vdash x : A \dashv \Delta}\ \text{uvar} \qquad \frac{}{\Delta, x : X; \Gamma \vdash x : X \dashv \Delta}\ \text{avar} \qquad \frac{}{\Delta; \Gamma \vdash () : \mathbb{1} \dashv \Delta}\ \text{unit}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : A_1 \dashv \Delta_2 \quad \Delta_2; \Gamma \vdash e_2 : A_2 \dashv \Delta_3}{\Delta_1; \Gamma \vdash (e_1, e_2)_\infty : A_1 \times A_2 \dashv \Delta_3}\ \text{upair} \qquad \frac{\Delta_1; \Gamma \vdash e_1 : S_1 \dashv \Delta_2 \quad \Delta_2; \Gamma \vdash e_2 : S_2 \dashv \Delta_3}{\Delta_1; \Gamma \vdash (e_1, e_2)_\mathsf{w} : S_1 \times S_2 \dashv \Delta_3}\ \text{spair}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : X_1 \dashv \Delta_2 \quad \Delta_2; \Gamma \vdash e_2 : X_2 \dashv \Delta_3}{\Delta_1; \Gamma \vdash (e_1, e_2)_1 : X_1 \otimes X_2 \dashv \Delta_3}\ \text{apair}$$

$$\frac{i \in \{1, 2\} \quad \Delta_1; \Gamma \vdash e : A_i \dashv \Delta_2}{\Delta_1; \Gamma \vdash \mathsf{inj}^i_\infty(e) : A_1 + A_2 \dashv \Delta_2}\ \text{uinj} \qquad \frac{i \in \{1, 2\} \quad \Delta_1; \Gamma \vdash e : S_i \dashv \Delta_2}{\Delta_1; \Gamma \vdash \mathsf{inj}^i_\mathsf{w}(e) : S_1 + S_2 \dashv \Delta_2}\ \text{sinj}$$

$$\frac{i \in \{1, 2\} \quad \Delta_1; \Gamma \vdash e : X_i \dashv \Delta_2}{\Delta_1; \Gamma \vdash \mathsf{inj}^i_1(e) : X_1 \oplus X_2 \dashv \Delta_2}\ \text{ainj}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : A_1 \times A_2 \dashv \Delta_2 \quad \Delta_2; \Gamma, x_1 : A_1, x_2 : A_2 \vdash e : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{split}_\infty(e_1, x_1.x_2.e_2) : U \dashv \Delta_3}\ \text{usplit}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : S_1 \times S_2 \dashv \Delta_2 \quad \Delta_2; \Gamma, x_1 : S_1, x_2 : S_2 \vdash e : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{split}_\mathsf{w}(e_1, x_1.x_2.e_2) : U \dashv \Delta_3}\ \text{ssplit}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : X_1 \otimes X_2 \dashv \Delta_2 \quad \Delta_2, x_1 : X_1, x_2 : X_2; \Gamma \vdash e : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{split}_1(e_1, x_1.x_2.e_2) : U \dashv \Delta_3 \div (x_1 : X_1, x_2 : X_2)}\ \text{asplit}$$

$$\frac{\Delta_1; \Gamma \vdash e : A_1 + A_2 \dashv \Delta_2 \quad \Delta_2; \Gamma, x_1 : A_1 \vdash e_1 : U \dashv \Delta_3 \quad \Delta_2; \Gamma, x_2 : A_2 \vdash e_2 : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{case}_\infty(e, x_1.e_1, x_2.e_2) : U \dashv \Delta_3}\ \text{ucase}$$

$$\frac{\Delta_1; \Gamma \vdash e : S_1 + S_2 \dashv \Delta_2 \quad \Delta_2; \Gamma, x_1 : S_1 \vdash e_1 : U \dashv \Delta_3 \quad \Delta_2; \Gamma, x_2 : S_2 \vdash e_2 : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{case}_\mathsf{w}(e, x_1.e_1, x_2.e_2) : U \dashv \Delta_3}\ \text{scase}$$

$$\frac{\Delta_1; \Gamma \vdash e : X_1 \oplus X_2 \dashv \Delta_2 \quad \Delta_2, x_1 : X_1; \Gamma \vdash e_1 : U \dashv \Delta_3 \qquad \Delta_2, x_2 : X_2; \Gamma \vdash e_2 : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{case}_1(e, x_1.e_1, x_2.e_2) : U \dashv \Delta_3 \div (x_1 : X_1, x_2 : X_2)}\ \text{acase}$$

$$\frac{\cdot; \Gamma, x : A \vdash e : U \dashv \cdot}{\Delta; \Gamma \vdash \lambda_\infty x.\, e : A \to_\infty U \dashv \Delta}\ \text{uabs} \qquad \frac{\text{ⓦ}; \Gamma, x : A \vdash e : U \dashv \cdot}{\Delta; \Gamma \vdash \lambda_\mathsf{w} x.\, e : A \to_\mathsf{w} U \dashv \Delta}\ \text{wabs} \qquad \frac{\Delta_1, x : X; \Gamma \vdash e : U \dashv \Delta_2}{\Delta_1; \Gamma \vdash \lambda_1 x.\, e : X \to_1 U \dashv \Delta_2 \div (x : X)}\ \text{aabs}$$

$$\frac{\Delta_1; \Gamma \vdash e_2 : A \dashv \Delta_2 \quad \Delta_2; \Gamma \vdash e_1 : A \to_\infty U \dashv \Delta_3}{\Delta_1; \Gamma \vdash (e_1\ e_2)_\infty : U \dashv \Delta_3}\ \text{uapp} \qquad \frac{\Delta_1; \Gamma \vdash e_2 : A \dashv \Delta_2 \quad \Delta_2; \Gamma \vdash e_1 : A \to_\mathsf{w} U \dashv \Delta_3}{\Delta_1, \text{ⓦ}; \Gamma \vdash (e_1\ e_2)_\mathsf{w} : U \dashv \Delta_3}\ \text{wapp}$$

$$\frac{\Delta_1; \Gamma \vdash e_2 : X \dashv \Delta_2 \quad \Delta_2; \Gamma \vdash e_1 : X \to_1 U \dashv \Delta_3}{\Delta_1; \Gamma \vdash (e_1\ e_2)_1 : U \dashv \Delta_3}\ \text{aapp}$$

$$\frac{\cdot; \Gamma, x : A \to_\infty U \vdash e : A \to_\infty U \dashv \cdot}{\Delta; \Gamma \vdash \mathsf{fix}_\infty(x.e) : A \to_\infty U \dashv \Delta}\ \text{ufix} \qquad \frac{\cdot; \Gamma, x : A \to_\mathsf{w} U \vdash e : A \to_\mathsf{w} U \dashv \cdot}{\Delta; \Gamma \vdash \mathsf{fix}_\mathsf{w}(x.e) : A \to_\mathsf{w} U \dashv \Delta}\ \text{wfix} \qquad \frac{x : X \to_1 U; \Gamma \vdash e : X \to_1 U \dashv \cdot}{\Delta; \Gamma \vdash \mathsf{fix}_1(x.e) : X \to_1 U \dashv \Delta}\ \text{afix}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : A \dashv \Delta_2 \quad \Delta_2; \Gamma, x : A \vdash e_2 : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{let}_\infty(e_1, x.e_2) : U \dashv \Delta_3}\ \text{ulet} \qquad \frac{\Delta_1; \Gamma \vdash e_1 : X \dashv \Delta_2 \quad \Delta_2, x : X; \Gamma \vdash e_2 : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{let}_1(e_1, x.e_2) : U \dashv \Delta_3 \div (x : X)}\ \text{alet} \qquad \frac{\Delta_1; \Gamma \vdash e : A \dashv \Delta_2}{\Delta_1; \Gamma \vdash\ !\, e : !\, A \dashv \Delta_2}\ \text{bang}$$

$$\frac{\Delta_1; \Gamma \vdash e : !\, A \dashv \Delta_2}{\Delta_1; \Gamma \vdash\ ¡\, e : A \dashv \Delta_2}\ \text{gnab} \qquad \frac{\Delta_1, x_1 : \mathsf{Rd}\, S; \Gamma, x_2 : \mathsf{Wr}\, S \vdash e : U \dashv \Delta_2}{\Delta_1; \Gamma \vdash \nu(x_1, x_2).\, e : U \dashv \Delta_2 \div (x_1 : \mathsf{Rd}\, S)}\ \text{nu} \qquad \frac{\Delta_1; \Gamma \vdash e_1 : S \dashv \Delta_2 \quad \Delta_2; \Gamma \vdash e_2 : \mathsf{Wr}\, S \dashv \Delta_3}{\Delta_1, \text{ⓦ}; \Gamma \vdash \mathsf{wr}(e_1, e_2) : \mathbb{1} \dashv \Delta_3}\ \text{wr}$$

$$\frac{\textcircled{w} \notin \Delta_2 \qquad \Delta_1; \Gamma \vdash e_1 : \mathsf{Rd}\, S \dashv \Delta_2 \qquad \Delta_2, \textcircled{w}, x : !\, S \otimes \mathsf{Rd}\, S; \Gamma \vdash e_2 : U \dashv \Delta_3}{\Delta_1; \Gamma \vdash \mathsf{rd}(e_1, x.e_2) : U \dashv \Delta_3 \div (\textcircled{w}, x : !\, S \otimes \mathsf{Rd}\, S)}\ \text{rd}$$

$$\frac{\begin{array}{c}\textcircled{w} \notin \Delta_3 \qquad \Delta_1; \Gamma \vdash e_1 : \mathsf{Rd}\, S \dashv \Delta_2 \qquad \Delta_2; \Gamma \vdash e_2 : \mathsf{Rd}\, T \dashv \Delta_3 \\ \Delta_3, \textcircled{w}, x_1 : !\, S \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T; \Gamma \vdash e_3 : U \dashv \Delta_4 \\ \Delta_3, \textcircled{w}, x_2 : !\, T \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T; \Gamma \vdash e_4 : U \dashv \Delta_4\end{array}}{\begin{array}{c}\Delta_1; \Gamma \vdash \mathsf{ch}(e_1, x_1.e_3, e_2, x_2.e_4) : U \dashv \Delta_4 \div \\ (\textcircled{w}, x_1 : !\, S \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T, x_2 : !\, T \otimes \mathsf{Rd}\, T \otimes \mathsf{Rd}\, S)\end{array}}\ \text{choice}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : U \dashv \Delta_2 \qquad \Delta_2; \Gamma \vdash e_2 : V \dashv \Delta_3}{\Delta_1; \Gamma \vdash e_1 \mid\!\triangleright\, e_2 : V \dashv \Delta_3}\ \text{fork}$$

Figure 9. Algorithmic typing rules.

B Type Soundness

We first define syntax for process and channel typings, which each map a kind of identifier (process name or channel name) to its associated type:

$$\begin{array}{ll}\text{Process pool typings} & \Phi ::= \cdot \mid \Phi, p : U \\ \text{Channel typings} & \Psi ::= \cdot \mid \Psi, d : S\end{array}$$

Using the syntax above, we define configuration typing as a straightforward extension of single-process typing, given in Section 3.2:

$\boxed{\Psi \vdash C : \Phi}$ Configuration C is well-typed.

$$\frac{}{\Psi \vdash \langle \Sigma; \varepsilon \rangle : \cdot}\ \text{empty} \qquad \frac{\Psi \vdash e : U \qquad \Psi \vdash \langle \Sigma; \pi \rangle : \Phi}{\Psi \vdash \langle \Sigma; \pi, p : e \rangle : \Phi, (p : U)}\ \text{cons}$$

B.1 Progress

Progress for the functional fragment of ILC (local progress) is fairly standard. We follow the usual recipe, except that we give a special definition of local process termination:

$\boxed{e\ \mathbf{lterm}}$ Expression e is locally terminated.

$$\frac{}{v\ \mathbf{lterm}}\ \text{val} \qquad \frac{}{E[\mathsf{rd}(c, x.e)]\ \mathbf{lterm}}\ \text{rdterm} \qquad \frac{}{E[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)]\ \mathbf{lterm}}\ \text{chterm} \qquad \frac{}{E[\mathsf{wr}(v, c)]\ \mathbf{lterm}}\ \text{wrterm}$$

In other words, e **lterm** holds when e is a value, is reading (either as a standalone read or an external choice), or is writing.

Lemma B.1 (Local Progress). *If $\Psi \vdash e : U$, then either e **lterm** or there exists e' such that $e \longrightarrow e'$.*

Proof. By structural induction on the derivation of $\Psi \vdash e : U$. □

To state progress on configurations, we give a special definition of "program termination" that permits deadlocks:

$\boxed{C\ \mathbf{term}}$ Configuration C is terminated.

$$\frac{\begin{array}{c}\forall (p : e) \in \pi.\ e\ \mathbf{lterm} \\ \mathrm{RdChans}(\pi) = \Sigma_1 \qquad \mathrm{WrChans}(\pi) = \Sigma_2 \\ \{(c_1, c_2) \mid c_1 \in \Sigma_1, c_2 \in \Sigma_2, c_2 \rightsquigarrow c_1\} = \varnothing\end{array}}{\langle \Sigma; \pi \rangle\ \mathbf{term}}\ \text{Cterm}$$

$$\begin{array}{rl}\mathrm{RdChans}(\varepsilon) = \cdot & \mathrm{WrChans}(\varepsilon) = \cdot \\ \mathrm{RdChans}(\pi, p : E[\mathsf{rd}(c, x.e)]) = \mathrm{RdChans}(\pi), c & \mathrm{WrChans}(\pi, p : E[\mathsf{rd}(c, x.e)]) = \mathrm{WrChans}(\pi) \\ \mathrm{RdChans}(\pi, p : E[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)]) = \mathrm{RdChans}(\pi), c_1, c_2 & \mathrm{WrChans}(\pi, p : E[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)]) = \mathrm{WrChans}(\pi) \\ \mathrm{RdChans}(\pi, p : E[\mathsf{wr}(v, c)]) = \mathrm{RdChans}(\pi) & \mathrm{WrChans}(\pi, p : E[\mathsf{wr}(v, c)]) = \mathrm{WrChans}(\pi), c \\ \mathrm{RdChans}(\pi, p : v) = \mathrm{RdChans}(\pi) & \mathrm{WrChans}(\pi, p : v) = \mathrm{WrChans}(\pi)\end{array}$$

In other words, C **term** holds when either:

1. C is fully normal: Every process in C is normalized (consists of a value), or
2. C is (at least partially) deadlocked: Some (possibly empty) portion of C is normal, and there exists one or more reading processes in C, or there exists one or more writing processes in C, however, no reader-writer process pair exists for a common channel.

Theorem B.2 (Progress). *If $\Psi \vdash C : \Phi$, then either C **term** or there exists C' such that $C \longrightarrow C'$.*

Proof. By structural induction on the derivation of $\Psi \vdash C : \Phi$.

Case

$$\frac{}{\Psi \vdash \langle \Sigma; \varepsilon \rangle : \cdot}\ \text{empty}$$

$\forall (p : e) \in \varepsilon.\ e$ **lterm**	Vacuous
$\Sigma_1 = \text{RdChans}(\varepsilon) = \cdot$	By definition of RdChans
$\Sigma_2 = \text{WrChans}(\varepsilon) = \cdot$	By definition of WrChans
$\{(c_1, c_2) \mid c_1 \in \Sigma_1, c_2 \in \Sigma_2, c_2 \rightsquigarrow c_1\} = \varnothing$	
$\langle \Sigma; \varepsilon \rangle$ **term**	By rule Cterm

Case

$$\frac{\Psi \vdash e : U \qquad \Psi \vdash \langle \Sigma; \pi \rangle : \Phi}{\Psi \vdash \langle \Sigma; \pi, p : e \rangle : \Phi, (p : U)}\ \text{cons}$$

e **lterm** or $\exists\ e'$ s.t. $e \longrightarrow e'$	By i.h.
$\langle \Sigma; \pi \rangle$ **term** or $\exists \langle \Sigma'; \pi' \rangle$ s.t. $\langle \Sigma; \pi \rangle \longrightarrow \langle \Sigma'; \pi' \rangle$	By i.h.
Subcase $\exists\ e'$ s.t. $e \longrightarrow e'$	
Subsubcase local	
$e = E[e_1]$ and $e' = E[e_2]$	Suppose
$\langle \Sigma; \pi, p : E[e_1] \rangle \longrightarrow \langle \Sigma; \pi, p : E[e_2] \rangle$	By rule local
Subsubcase fork	
$e = E[e_1 \mid\!\triangleright\ e_2],\ e' = E[e_2],$ and $q \notin \Sigma$	Suppose
$\langle \Sigma; \pi, p : E[e_1 \mid\!\triangleright\ e_2] \rangle \longrightarrow \langle \Sigma, q; \pi, q : e_1, p : E[e_2] \rangle$	By rule fork
Subsubcase nu	
$e = E[\nu(x_1, x_2).\ e_1],\ e' = E[[\text{Read}(d)/x_1][\text{Write}(d)/x_2]e_1],\ d \notin \Sigma$	Suppose
$\langle \Sigma; \pi, p : [\nu(x_1, x_2).\ e_1] \rangle \longrightarrow \langle \Sigma, d; \pi, p : E[[\text{Read}(d)/x_1][\text{Write}(d)/x_2]e_1] \rangle$	By rule nu
Subsubcase rw	
$e = E[\text{rd}(c_1, x.e_1)],\ e' = E[[(!v, c_1)_1/x]e_1],$ and $c_2 \rightsquigarrow c_1,$ or	
$\quad e = E[\text{wr}(v, c_2)],\ e' = E[()],$ and $c_2 \rightsquigarrow c_1$	
Subsubsubcase $e = E[\text{rd}(c_1, x.e_1)],\ e' = E[[(!v, c_1)_1/x]e_1],$ and $c_2 \rightsquigarrow c_1$	
$\exists\ (q : E[\text{wr}(v, c_2)]) \in \pi$	By $c_2 \rightsquigarrow c_1$
$\langle \Sigma; \pi, p : E[\text{rd}(c_1, x.e_1)] \longrightarrow \langle \Sigma; \pi, p : E[[(!v, c_1)_1/x]e_1] \rangle \rangle$	By rule rw
Subsubsubcase $e = E[\text{wr}(v, c_2)],\ e' = E[()],$ and $c_2 \rightsquigarrow c_1$	
$\exists\ (q : E[\text{rd}(c_1, x.e_1)]) \in \pi$	By $c_2 \rightsquigarrow c_1$
$\langle \Sigma; \pi, p : E[\text{wr}(v, c_2)] \longrightarrow \langle \Sigma; \pi, p : E[()] \rangle \rangle$	By rule rw
Subsubcase cw	
$e = E[\text{ch}(c_1, x_1.e_1, c_2, x_2.e_2)],\ e' = E[[(!v, c_1, c_2)_1/x_i]e_i],\ c \rightsquigarrow c_i,\ i \in \{1, 2\},$ or	
$\quad e = E[\text{wr}(v, c)],\ e' = E[()],\ c \rightsquigarrow c_i,\ i \in \{1, 2\}$	
Subsubsubcase $e = E[\text{ch}(c_1, x_1.e_1, c_2, x_2.e_2)],\ e' = E[[(!v, c_1, c_2)_1/x_i]e_i],$	
$\quad c \rightsquigarrow c_i,\ i \in \{1, 2\}$	
$\exists\ (q : E[\text{wr}(v, c)]) \in \pi$	By $c \rightsquigarrow c_i$

$\langle \Sigma; \pi, p : E[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)] \rangle \longrightarrow \langle \Sigma; \pi, p : E[[(!v, c_1, c_2)_1 / x_i] e_i] \rangle$ By rule cw

Subsubsubcase $e = E[\mathsf{wr}(v, c)]$, $e' = E[()]$, $c \rightsquigarrow c_i$, $i \in \{1, 2\}$

$\exists\, (q : E[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)]) \in \pi$ By $c \rightsquigarrow c_i$

$\langle \Sigma; \pi, p : E[\mathsf{wr}(v, c)] \rangle \longrightarrow \langle \Sigma; \pi, p : E[()] \rangle$ By rule cw

Subcase $\exists \langle \Sigma'; \pi' \rangle$ s.t. $\langle \Sigma; \pi \rangle \longrightarrow \langle \Sigma'; \pi' \rangle$

$\langle \Sigma; \pi, p : e \rangle \longrightarrow \langle \Sigma'; \pi', p : e \rangle$ By rules local and congr

Subcase $\langle \Sigma; p : e \rangle$ **term** and $\langle \Sigma; \pi \rangle$ **term**

$\Sigma_1 = \mathrm{RdChans}(\pi, p : e)$ and $\Sigma_2 = \mathrm{WrChans}(\pi, p : e)$ Suppose

$\{(c_1, c_2) \mid c_1 \in \Sigma_1, c_2 \in \Sigma_2, c_2 \rightsquigarrow c_1\} = \varnothing$ or

$\{(c_1, c_2) \mid c_1 \in \Sigma_1, c_2 \in \Sigma_2, c_2 \rightsquigarrow c_1\} \neq \varnothing$

Subsubcase $\{(c_1, c_2) \mid c_1 \in \Sigma_1, c_2 \in \Sigma_2, c_2 \rightsquigarrow c_1\} = \varnothing$

$\langle \Sigma; \pi, p : e \rangle$ **term** By rule Cterm

Subsubcase $\{(c_1, c_2) \mid c_1 \in \Sigma_1, c_2 \in \Sigma_2, c_2 \rightsquigarrow c_1\} \neq \varnothing$

$\exists\, c_2 \rightsquigarrow c_1$ s.t. $c_1 \in \Sigma_1, c_2 \in \Sigma_2$ Above

$p : v$ or $p : E[\mathsf{rd}(c_1, x.e)]$ or $p : E[\mathsf{ch}(c_1, x_1.e_1, c_3, x_2.e_2)]$ or

$p : E[\mathsf{ch}(c_3, x_1.e_1, c_1, x_2.e_2)]$ or $p : E[\mathsf{wr}(v, c_2)]$ By definition of **lterm**

Subsubsubcase $p : v$ Impossible

Subsubsubcase $p : E[\mathsf{rd}(c_1, x.e)]$

$\exists\, q : E[\mathsf{wr}(v, c_2)] \in \pi$ By $c_2 \rightsquigarrow c_1$

$\langle \Sigma; \pi, p : E[\mathsf{rd}(c_1, x.e)] \rangle \longrightarrow \langle \Sigma; \pi, p : E[[(!v, c_1)_1 / x] e] \rangle$ By rule rw

Subsubsubcase $p : E[\mathsf{ch}(c_1, x_1.e_1, c_3, x_2.e_2)]$

$\exists\, q : E[\mathsf{wr}(v, c_2)] \in \pi$ By $c_2 \rightsquigarrow c_1$

$\langle \Sigma; \pi, p : E[\mathsf{ch}(c_1, x_1.e_1, c_3, x_2.e_2)] \rangle \longrightarrow \langle \Sigma; \pi, p : E[[(!v, c_1, c_3)_1 / x_1] e_1] \rangle$ By rule cw

Subsubsubcase $p : E[\mathsf{ch}(c_3, x_1.e_1, c_1, x_2.e_2)]$

$\exists\, q : E[\mathsf{wr}(v, c_2)] \in \pi$ By $c_2 \rightsquigarrow c_1$

$\langle \Sigma; \pi, p : E[\mathsf{ch}(c_3, x_1.e_1, c_1, x_2.e_2)] \rangle \longrightarrow \langle \Sigma; \pi, p : E[[(!v, c_1, c_3)_1 / x_2] e_2] \rangle$ By rule cw

Subsubsubcase $p : E[\mathsf{wr}(v, c_2)]$

$\exists\, q : E[\mathsf{rd}(c_1, x.e)] \in \pi$ or $\exists\, q : E[\mathsf{ch}(c_1, x_1.e_1, c_3, x_2.e_2)] \in \pi$ or

$\exists\, q : E[\mathsf{ch}(c_3, x_1.e_1, c_1, x_2.e_2)] \in \pi$ By $c_2 \rightsquigarrow c_1$

$\langle \Sigma; \pi, p : E[\mathsf{wr}(v, c_2)] \rangle \longrightarrow \langle \Sigma; \pi, p : E[()] \rangle$ By rule rw

□

B.2 Preservation

Preservation for the functional fragment of ILC (local preservation) is standard.

Lemma B.3 (Local Preservation). *If $\Psi \vdash e : U$ and $e \longrightarrow e'$, then there exists $\Psi' \supseteq \Psi$ such that $\Psi \vdash e' : U$.*

Proof. By structural induction on the derivation of $e \longrightarrow e'$. □

To state preservation on configurations, we first state several auxiliary results, which follow the formulation of Gay and Vasconcelos [25]. Lemma B.4 shows that typing of configurations is preserved under configuration equivalence.

Lemma B.4 (Preservation Modulo Equivalence). *If $\Psi \vdash C : \Phi$ and $C \equiv C'$, then $\Psi \vdash C' : \Phi$.*

Proof. By structural induction on $\Psi \vdash C : \Phi$. □

Lemma B.5 shows that a subterm of a well-typed evaluation context is typeable with a subset of the type contexts.

Lemma B.5 (Typeability of Subterms). *If $\mathcal{D}$ is a derivation of $\Psi; \Delta; \Gamma \vdash E[e] : U$ (written $\mathcal{D} :: \Psi; \Delta; \Gamma \vdash E[e] : U$), then*

1. *there exists $\Psi_1, \Psi_2; \Delta_1, \Delta_2; \Gamma_1, \Gamma_2$ and V such that $\Psi = \Psi_1, \Psi_2$, $\Delta = \Delta_1, \Delta_2$, $\Gamma = \Gamma_1, \Gamma_2$,*
2. *$\mathcal{D}$ has a subderivation $\mathcal{D}'$ (written $\mathcal{D}' \sqsubseteq \mathcal{D}$) concluding $\Psi_1; \Delta_1; \Gamma_1 \vdash e : V$,*
3. *the position of $\mathcal{D}'$ in $\mathcal{D}$ corresponds to the position of the hole in E (written $E[\mathcal{D}' \sqsubseteq \mathcal{D}]$).*

Proof. By structural induction on the structure of E. □

Lemma B.6 shows that the subterm of a well-typed evaluation context can be replaced.

Lemma B.6 (Replacement (Evaluation Contexts)). *If*

1. $\mathcal{D} :: \Psi_1, \Psi_2; \Delta_1, \Delta_2; \Gamma_1, \Gamma_2 \vdash E[e] : U$,
2. $\mathcal{D}' \sqsubseteq \mathcal{D}$ *such that* $\mathcal{D}' :: \Psi_2; \Delta_2; \Gamma_2 \vdash e : V$,
3. $E[\mathcal{D}' \sqsubseteq \mathcal{D}]$,
4. $\Psi_3; \Delta_3; \Gamma_3 \vdash e' : V$,
5. $\Psi_1, \Psi_3; \Delta_1, \Delta_3; \Gamma_1, \Gamma_3$ *is defined,*

then $\Psi_1, \Psi_3; \Delta_1, \Delta_3; \Gamma_1, \Gamma_3 \vdash E[e'] : U$.

Proof. By structural induction on the structure of E. □

Finally, Lemmas B.7, B.8, B.9, B.10 show that typing of terms is preserved by substitution.

Lemma B.7 (Substitution (Unrestricted)). *If*

1. $\Psi_1; \Delta_1; \Gamma_1, x : A \vdash e : U$,
2. $\Psi_2; \Delta_2; \Gamma_2 \vdash e' : A$,
3. $\Psi_1, \Psi_2; \Delta_1, \Delta_2; \Gamma_1, \Gamma_2$ *is defined,*

then $\Psi_1, \Psi_2; \Delta_1, \Delta_2; \Gamma_1, \Gamma_2 \vdash [e'/x]e : U$.

Proof. By structural induction on the derivation of $\Psi_1; \Delta_1; \Gamma_1, x : A \vdash e : U$. □

Lemma B.8 (Substitution (Affine)). *If*

1. $\Psi_1; \Delta_1, x : X; \Gamma_1 \vdash e : U$,
2. $\Psi_2; \Delta_2; \Gamma_2 \vdash e' : X$,
3. $\Psi_1, \Psi_2; \Delta_1, \Delta_2; \Gamma_1, \Gamma_2$ *is defined,*

then $\Psi_1, \Psi_2; \Delta_1, \Delta_2; \Gamma_1, \Gamma_2 \vdash [e'/x]e : U$.

Proof. By structural induction on the derivation of $\Psi_1; \Delta_1, x : X; \Gamma_1 \vdash e : U$. □

Lemma B.9 (Substitution (Read Endpoint)). *If*

1. $\Psi; \Delta, x : \mathsf{Rd}\, S; \Gamma \vdash e : U$,
2. $\Psi, d : S; \Delta; \Gamma$ *is defined,*

then $\Psi, d : S; \Delta; \Gamma \vdash [\mathsf{Read}(d)/x]e : U$.

Proof. By structural induction on the derivation of $\Psi; \Delta, x : \mathsf{Rd}\, S; \Gamma \vdash e : U$. □

Lemma B.10 (Substitution (Write Endpoint)). *If*

1. $\Psi; \Delta; \Gamma, x : \mathsf{Wr}\, S \vdash e : U$,
2. $\Psi, d : S, \Psi; \Delta; \Gamma$ *is defined,*

then $\Psi, d : S; \Delta; \Gamma \vdash [\mathsf{Write}(d)/x]e : U$.

Proof. By structural induction on the derivation of $\Psi; \Delta; \Gamma, x : \mathsf{Wr}\, S \vdash e : U$. □

Theorem B.11 (Preservation). *If* $\Psi \vdash C : \Phi$ *and* $C \longrightarrow C'$, *then there exists* $\Psi' \supseteq \Psi$ *and* $\Phi' \supseteq \Phi$ *such that* $\Psi' \vdash C' : \Phi'$.

Proof. By structural induction on the derivation of $C \longrightarrow C'$.

Case

$$\frac{e_1 \longrightarrow e_2}{\langle \Sigma; \pi, p : E[e_1] \rangle \longrightarrow \langle \Sigma; \pi, p : E[e_2] \rangle}\ \text{local}$$

$\Psi \vdash \langle \Sigma; \pi, p : E[e_1] \rangle : \Phi$ s.t. $\Phi = \Phi_\pi, p : U$, $\Psi = \Psi_1, \Psi_2$, and $\mathcal{D} :: \Psi_1, \Psi_2 \vdash E[e_1] : U$	Assumption
$\exists\, \mathcal{D}' \sqsubseteq \mathcal{D}$ s.t. $\mathcal{D}' :: \Psi_2 \vdash e_1 : V$ and $E[\mathcal{D}' \sqsubseteq \mathcal{D}]$	By Lemma B.5
$\Psi_2 \vdash e_2 : V$	By i.h. and Lemma B.3
$\Psi_1, \Psi_2 \vdash E[e_2] : U$	By Lemma B.6

$\Psi \vdash E[e_2] : U$	By above equalities
$\Psi \vdash \langle \Sigma; \pi \rangle : \Phi_\pi$	Above
$\Psi \vdash \langle \Sigma; \pi, p : E[e_2] \rangle : (\Phi_\pi, p : U)$	By rule cons
$\Psi \vdash \langle \Sigma; \pi, p : E[e_2] \rangle : \Phi$	By above equalities
$\Psi' = \Psi$ and $\Phi' = \Phi$	Suppose
$\Psi' \vdash \langle \Sigma; \pi, p : E[e_2] \rangle : \Phi'$	By above equalities

Case

$$\frac{q \notin \Sigma}{\langle \Sigma; \pi, p : E[e_1 \mid\triangleright e_2] \rangle \longrightarrow \langle \Sigma, q; \pi, q : e_1, p : E[e_2] \rangle}\ \text{fork}$$

$\Psi \vdash \langle \Sigma; \pi, p : E[e_1 \mid\triangleright e_2] \rangle : \Phi$ s.t. $\Phi = \Phi_\pi, p : U$, $\Psi = \Psi_1, \Psi_2$, and $\mathcal{D} :: \Psi_1, \Psi_2 \vdash E[e_1 \mid\triangleright e_2] : U$	Assumption
$\exists\, \mathcal{D}' \sqsubseteq \mathcal{D}$ s.t. $\mathcal{D}' :: \Psi_2 \vdash e_1 \mid\triangleright e_2 : V_2$ and $E[\mathcal{D}' \sqsubseteq \mathcal{D}]$	By Lemma B.5
$\Psi_2 \vdash e_1 : V_1$	By inversion on fork
$\Psi_2 \vdash e_2 : V_2$	By inversion on fork
$\Psi_1, \Psi_2 \vdash E[e_2] : U$	By Lemma B.6
$\Psi \vdash E[e_2] : U$	By above equalities
$\Psi \vdash \langle \Sigma; \pi \rangle : \Phi_\pi$	Above
$\Psi \vdash \langle \Sigma, q; \pi \rangle : \Phi_\pi$	By $q \notin \Sigma$
$\Psi \vdash \langle \Sigma, q; \pi, q : e_1 \rangle : (\Phi_\pi, q : V_1)$	By rule cons
$\Psi \vdash \langle \Sigma, q; \pi, q : e_1, p : E[e_2] \rangle : (\Phi_\pi, q : V_1, p : U)$	By rule cons
$\Psi \vdash \langle \Sigma, q; \pi, q : e_1, p : E[e_2] \rangle : \Phi, q : V_1$	By above equalities
$\Psi' = \Psi$ and $\Phi' = \Phi, q : V_1$	Suppose
$\Psi' \vdash \langle \Sigma, q; \pi, q : e_1, p : E[e_2] \rangle : \Phi'$	By above equalities

Case

$$\frac{C_1 \equiv C_1' \qquad C_1' \longrightarrow C_2' \qquad C_2' \equiv C_2}{C_1 \longrightarrow C_2}\ \text{congr}$$

$\Psi \vdash C_1 : \Phi$	Assumption
$C_1 \equiv C_1'$	Given
$\Psi \vdash C_1' : \Phi$	By Lemma B.4
$\Psi' \supseteq \Psi$ and $\Phi' \supseteq \Phi$	Suppose
$\Psi' \vdash C_2' : \Phi'$	By i.h.
$\Psi' \vdash C_2 : \Phi'$	By Lemma B.4

Case

$$\frac{d \notin \Sigma}{\langle \Sigma; \pi, p : E[\nu(x_1, x_2).\, e] \rangle \longrightarrow \langle \Sigma, d; \pi, p : E[[\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e] \rangle}\ \text{nu}$$

$\Psi \vdash \langle \Sigma; \pi, p : E[\nu(x_1, x_2).\, e] \rangle : \Phi$ s.t. $\Phi = \Phi_\pi, p : U$, $\Psi = \Psi_1, \Psi_2$ and $\mathcal{D} :: \Psi_1, \Psi_2 \vdash E[\nu(x_1, x_2).\, e] : U$	Assumption
$\exists\, \mathcal{D}' \sqsubseteq \mathcal{D}$ s.t. $\mathcal{D}' :: \Psi_2 \vdash \nu(x_1, x_2).\, e : V$ and $E[\mathcal{D}' \sqsubseteq \mathcal{D}]$	By Lemma B.5
$\Psi_2; \Gamma; \Delta \vdash e : U$ where $\Gamma; \Delta = x_1 : \mathsf{Rd}\, S; x_2 : \mathsf{Wr}\, S$	By inversion on nu
$d : S \vdash \mathsf{Read}(d) : \mathsf{Rd}\, S$	By rule rdend

$d : S \vdash \mathsf{Write}(d) : \mathsf{Rd}\, S$	By rule wrend
$\Psi_3 \vdash [\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e : U$ where $\Psi_3 = \Psi_2, d : S$	By Lemmas B.9 and B.10
$\Psi_1, \Psi_3 \vdash E[[\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e] : U_p$	By Lemma B.6
$\Psi, \Psi_4 \vdash E[[\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e] : U_p$ where $\Psi_4 = c_1 : \mathsf{Rd}\, S, c_2 : \mathsf{Wr}\, S$	By above equalities
$\Psi, \Psi_4 \vdash \langle \Sigma; \pi \rangle : \Phi_\pi$	Above
$\Psi, \Psi_4 \vdash \langle \Sigma; \pi, p : E[[\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e] \rangle : (\Phi_\pi, p : U)$	By rule cons
$\Psi, \Psi_4 \vdash \langle \Sigma; \pi, p : E[[\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e] \rangle : \Phi$	Above
$\Psi' = \Psi, \Psi_4$ and $\Phi' = \Phi$	Suppose
$\Psi' \vdash \langle \Sigma; \pi, p : E[[\mathsf{Read}(d)/x_1][\mathsf{Write}(d)/x_2]e] \rangle : \Phi'$	By above equalities

Case

$$\frac{c_2 \rightsquigarrow c_1}{\langle \Sigma; \pi, p : E_1[\mathsf{rd}(c_1, x.e)], q : E_2[\mathsf{wr}(v, c_2)] \rangle \longrightarrow \langle \Sigma; \pi, p : E_1[[(!v, c_1)_1/x]e], q : E_2[()] \rangle}\ \text{rw}$$

$\Psi \vdash \langle \Sigma; \pi, p : E_1[\mathsf{rd}(c_1, x.e)], q : E_2[\mathsf{wr}(v, c_2)] \rangle : \Phi$ s.t. $\Phi = \Phi_\pi, p : U, q : V$, $\Psi = \Psi_1, \Psi_2, \mathcal{D}_p :: \Psi_1, \Psi_2 \vdash E_1[\mathsf{rd}(c_1, x.e)] : U$, $\Psi = \Psi_3, \Psi_4,$ and $\mathcal{D}_q :: \Psi_3, \Psi_4 \vdash E_2[\mathsf{wr}(v, c_2)] : V$	Assumption
$\exists\, \mathcal{D}'_p \sqsubseteq \mathcal{D}_p$ s.t. $\mathcal{D}'_p :: \Psi_2 \vdash \mathsf{rd}(c_1, x.e) : U'$ and $E_1[\mathcal{D}'_p \sqsubseteq \mathcal{D}_p]$	By Lemma B.5
$\exists\, \mathcal{D}'_q \sqsubseteq \mathcal{D}_q$ s.t. $\mathcal{D}'_q :: \Psi_4 \vdash \mathsf{wr}(v, c_2) : \mathbb{1}$ and $E_2[\mathcal{D}'_q \sqsubseteq \mathcal{D}_q]$	By Lemma B.5
$c_2 \rightsquigarrow c_1$ s.t. $\Psi(c_2) = \mathsf{Wr}\, S$ and $\Psi(c_1) = \mathsf{Rd}\, S$	Given
$\Psi_2; \Delta; \cdot \vdash e : U'$ where $\Delta = ⓦ, x : !\, S \otimes \mathsf{Rd}\, S$	By inversion on rd
$\vdash v : S$	By inversion on wr
$\vdash !\, v : !\, S$	By rule bang
$\vdash (!v, c_1)_1 : !\, S \otimes \mathsf{Rd}\, S$	By rule apair
$\Psi_2; ⓦ; \cdot \vdash [(!v, c_1)_1/x]e : U'$	By Lemma B.8
$\Psi_1, \Psi_2 \vdash E_1[[(!v, c_1)_1/x]e] : U$	By Lemma B.6
$\Psi \vdash E_1[[(!v, c_1)_1/x]e] : U$	By above equalities
$\Psi \vdash \langle \Sigma; \pi \rangle : \Phi_\pi$	Above
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1)_1/x]e] \rangle : (\Phi_\pi, p : U)$	By rule cons
$\Psi_4 \vdash () : \mathbb{1}$	By rule unit
$\Psi_3, \Psi_4 \vdash E_2[()] : V$	By Lemma B.6
$\Psi \vdash E_2[()] : V$	By above equalities
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1)_1/x]e], q : E_2[()] \rangle : (\Phi_\pi, p : U, q : V)$	By rule cons
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1)_1/x]e], q : E_2[()] \rangle : \Phi$	By above equalities
$\Psi' = \Psi$ and $\Phi' = \Phi$	Suppose
$\Psi' \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1)_1/x]e], q : E_2[()] \rangle : \Phi'$	By above equalities

Case

$$\frac{c \rightsquigarrow c_i \qquad i \in \{1, 2\}}{\langle \Sigma; \pi, p : E_1[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)], q : E_2[\mathsf{wr}(v, c)] \rangle \longrightarrow \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1/x_i]e_i], q : E_2[()] \rangle}\ \text{cw}$$

$\Psi \vdash \langle \Sigma; \pi, p : E_1[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)], q : E_2[\mathsf{wr}(v, c)] \rangle : \Phi$ s.t. $\Phi = \Phi_\pi, p : U, q : V$, $\Psi = \Psi_1, \Psi_2, \mathcal{D}_p :: \Psi_1, \Psi_2 \vdash E_1[\mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2)] : U$, $\Psi = \Psi_3, \Psi_4,$ and $\mathcal{D}_q :: \Psi_3, \Psi_4 \vdash E_2[\mathsf{wr}(v, c)] : V$	Assumption
$\exists\, \mathcal{D}'_p \sqsubseteq \mathcal{D}_p$ s.t. $\mathcal{D}'_p :: \Psi_2 \vdash \mathsf{ch}(c_1, x_1.e_1, c_2, x_2.e_2) : U'$ and $E_1[\mathcal{D}'_p \sqsubseteq \mathcal{D}_p]$	By Lemma B.5
$\exists\, \mathcal{D}'_q \sqsubseteq \mathcal{D}_q$ s.t. $\mathcal{D}'_q :: \Psi_4 \vdash \mathsf{wr}(v, c_2) : \mathbb{1}$ and $E_2[\mathcal{D}'_q \sqsubseteq \mathcal{D}_q]$	By Lemma B.5

$c \rightsquigarrow c_1$ s.t. $\Psi(c) = \mathsf{Wr}\, S,\ \Psi(c_1) = \mathsf{Rd}\, S,\ \Psi(c_2) = \mathsf{Rd}\, T$ or	
$\quad c \rightsquigarrow c_2$ s.t. $\Psi(c) = \mathsf{Wr}\, T,\ \Psi(c_1) = \mathsf{Rd}\, S,\ \Psi(c_2) = \mathsf{Rd}\, T$	Given
Subcase $c \rightsquigarrow c_1$	
$\Psi_2; \Delta; \cdot \vdash e : U'$ where $\Delta = ⓦ, x_1 : !\, S \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T$	By inversion on choice
$\vdash v : S$	By inversion on wr
$\vdash !\, v : !\, S$	By rule bang
$\vdash (!v, c_1, c_2)_1 : !\, S \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T$	By rule apair
$\Psi_2; ⓦ; \cdot \vdash (!v, c_1, c_2)_1 / x_1]e_1 : U'$	By Lemma B.8
$\Psi_1, \Psi_2 \vdash E_1[[(!v, c_1, c_2)_1 / x_1]e_1] : U$	By Lemma B.6
$\Psi \vdash E_1[[(!v, c_1, c_2)_1 / x_1]e_1] : U$	By above equalities
$\Psi \vdash \langle \Sigma; \pi \rangle : \Phi_\pi$	Above
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_1]e_1] \rangle : (\Phi_\pi, p : U)$	By rule cons
$\Psi_4 \vdash () : \mathbb{1}$	By rule unit
$\Psi_3, \Psi_4 \vdash E_2[()] : V$	By Lemma B.6
$\Psi \vdash E_2[()] : V$	By above equalities
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_1]e_1], q : E_2[()] \rangle : (\Phi_\pi, p : U, q : V)$	By rule cons
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_1]e_1], q : E_2[()] \rangle : \Phi$	By above equalities
$\Psi' = \Psi$ and $\Phi' = \Phi$	Suppose
$\Psi' \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_1]e_1], q : E_2[()] \rangle : \Phi'$	By above equalities
Subcase $c \rightsquigarrow c_2$	
$\Psi_2; \Delta; \cdot \vdash e : U'$ where $\Delta = ⓦ, x_2 : !\, T \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T$	By inversion on choice
$\vdash v : T$	By inversion on wr
$\vdash !\, v : !\, T$	By rule bang
$\vdash (!v, c_1, c_2)_1 : !\, T \otimes \mathsf{Rd}\, S \otimes \mathsf{Rd}\, T$	By rule apair
$\Psi_2, ⓦ; \cdot \vdash (!v, c_1, c_2)_1 / x_2]e_2 : U'$	By Lemma B.8
$\Psi_1, \Psi_2 \vdash E_1[[(!v, c_1, c_2)_1 / x_2]e_2] : U$	By Lemma B.6
$\Psi \vdash E_1[[(!v, c_1, c_2)_1 / x_2]e_2] : U$	By above equalities
$\Psi \vdash \langle \Sigma; \pi \rangle : \Phi_\pi$	Above
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_2]e_2] \rangle : (\Phi_\pi, p : U)$	By rule cons
$\Psi_4 \vdash () : \mathbb{1}$	By rule unit
$\Psi_3, \Psi_4 \vdash E_2[()] : V$	By Lemma B.6
$\Psi \vdash E_2[()] : V$	By above equalities
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_2]e_2], q : E_2[()] \rangle : (\Phi_\pi, p : U, q : V)$	By rule cons
$\Psi \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_2]e_2], q : E_2[()] \rangle : \Phi$	By above equalities
$\Psi' = \Psi$ and $\Phi' = \Phi$	Suppose
$\Psi' \vdash \langle \Sigma; \pi, p : E_1[[(!v, c_1, c_2)_1 / x_2]e_2], q : E_2[()] \rangle : \Phi'$	By above equalities

□

C Confluence

The following lemmas state structural invariants over write effects and read endpoints of a well-typed configuration: at most one process owns the write token ⓦ, and every read endpoint is a non-duplicable (affine) resource.

Lemma C.1 (Unique writer process). *If C is a well-typed configuration with process pool π, then there exists at most one process in π that owns the write token ⓦ (i.e., has ⓦ in its affine context).*

Proof. By structural induction over the typing derivation for C. □

Lemma C.2 (Unique reader process). *If C is a well-typed configuration with process pool π, and c is a read endpoint in this configuration, then there exists at most one process in π where c appears.*

Proof. By structural induction over the typing derivation for C. □

Theorem C.3 (Single-step confluence). *For all well-typed configurations C, if $C \longrightarrow C_1$ and $C \longrightarrow C_2$ then there exists renaming a function f such that either:*

1. *$C_1 = f(C_2)$, or*
2. *there exists C_3 such that $C_1 \longrightarrow C_3$ and $f(C_2) \longrightarrow C_3$.*

Proof. By induction on the pair of steps $\langle C \longrightarrow C_1\,,\ C \longrightarrow C_2 \rangle$.
We consider the following cases:

Case congruence
If either step uses congr, we apply the inductive hypothesis.

Case independent processes
If both steps advance distinct processes, using any of the rules local, fork and nu, we produce C_3 by combining those two (independent) steps.

Case one process
If both steps advance the same process, we show that this is deterministic (up to naming) by constructing the naming function f such that $C_2 = f(C_1)$. Most cases are straightforward since they perform no nondeterministic choices. The only source of nondeterminism is the name choices, in rules nu and fork. In each case, we map the name choice from the second step to that of the first step.

Case interaction
If either step uses rw or cw, we rely on Lemmas C.1 and C.2 to show that both steps use either rw or cw, and that the reader-writer process pair is unique.

□

By composing multiple uses of this theorem we prove multi-step confluence. However, to carry forth this composition, we need a more general notion of single-step confluence, which is parameteric in a renaming function for the initial configurations.

Theorem C.4 (Single-step confluence, generalized). *For all well-typed configurations C and renaming functions f, if $C \longrightarrow C_1$ and $f(C) \longrightarrow C_2$ then there exists renaming function g such that either:*

1. *$C_1 = g(C_2)$, or*
2. *there exists C_3 such that $C_1 \longrightarrow C_3$ and $g(C_2) \longrightarrow C_3$.*

Proof. Analogous to the proof of Theorem C.3 (single-step confluence). □

We prove a full confluence theorem that is generalized similarly, by accepting a renaming function f to produce a new function g:

Theorem C.5 (Full confluence). *For all well-typed configurations C, and renaming functions f, if $C \longrightarrow^* C_1$ and $f(C) \longrightarrow^* C_2$ and C_1 **term** and C_2 **term** then there exists a renaming function g such that $C_1 = g(C_2)$.*

Proof. By induction on the reduction sequence pair $\langle C \longrightarrow^* C_1\,,\ f(C) \longrightarrow^* C_2 \rangle$. Because of single-step confluence, we know that if either reduction sequence is empty, then the other must be empty, and that if either takes a step, the other must take a step.

Case empty
When empty, we have the resulting renaming function g via single-step confluence.

Case step
We consider the case where each reduction consists of at least one step: $C \longrightarrow C_1'$ and $C_1' \longrightarrow^* C_1$ and $f(C) \longrightarrow C_2'$ and $C_2' \longrightarrow^* C_2$. By single-step confluence, we have that there exists g_0 such that $g_0(C_2') = C_1'$. By the inductive hypothesis, we have that there exists g such that $C_1 = g(C_2)$.

□

D ILC Implementation of execUC

The full implementation of the UC execution experiment is given in Figure 10.

```
data Crupt = CruptP | CruptQ | CruptNone

corruptOrNot :: ∀ ... . A_p → Nat → [Bit] → Bool → ···
let corruptOrNot p k bits iscrupt toZ toF toA toQ frZ frF frA frQ=
      if iscrupt then
          let _ = rd frZ in error "Z can't wr to corrupt"
        |▷ fwd toA frF
        |▷ fwd toA frQ
        |▷ fwd toF frA
      else
        p k bits toZ toF toQ frZ frF frQ

execUC :: A_z→_w A_p × A_q → A_f → A_a → Crupt → Nat → [Bit] → Bit
let execUC z (p,q) f a crupt k r =
  ν (rZ2P, wZ2P), (rP2Z, wP2Z)
  , (rZ2Q, wZ2Q), (rQ2Z, wQ2Z)
  , (rP2F, wP2F), (rF2P, wF2P)
  , (rQ2F, wQ2F), (rF2Q, wF2Q)
  , (rF2A, wF2A), (rA2F, wA2F)
  , (rA2Z, wA2Z), (rZ2A, wZ2A)
  , (rP2A, wP2A), (rA2P, wA2P)
  , (rQ2A, wQ2A), (rA2Q, wA2Q)
  , (rP2Q, wP2Q), (rQ2P, wQ2P)
  . let (rf,ra,rp,rq,rz) = splitBits r in
      f k rf crupt wF2P wF2Q wF2A rP2F rQ2F rA2F
    |▷ a k ra crupt wA2Z wA2F wA2P wA2Q rZ2A rF2A rP2A rQ2A
    |▷ corruptOrNot p k rp (crupt == CruptP) wP2Z wP2F wP2A wP2Q rZ2P rF2P rA2P rQ2P
    |▷ corruptOrNot q k rq (crupt == CruptQ) wQ2Z wQ2F wQ2A wQ2P rZ2Q rF2Q rA2Q rP2Q
    |▷ z k rz wZ2P wZ2Q wZ2A rP2Z rQ2Z rA2Z
```

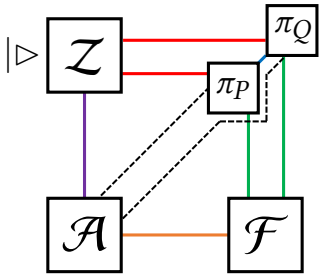

Figure 10. Full implementation of execUC. The channels follow a uniform naming scheme. The read end of a channel is prefixed with r- and the write end of a channel is prefixed with w-. The channel rZ2P denotes the read end of communications from the environment z to the party p. First, the random bitstring is split amongst each of the five parties. Then, the functionality, the adversary, and both protocol parties are spawned in a child process (given the appropriate channels and parameters), and the process continues as the environment process. Notice that parties are run in wrapper functions, which alter their behavior depending on whether or not they are corrupted. If a party is corrupted, then the adversary masquerades as the party.

E Extending ILC with Trapdoor Permutations

UC Commitments are realized from cryptographic primitives, such as trapdoor permutations, which require extensions to ILC. The new syntactic forms are kgen, tdp, inv, and hc with the static and dynamic semantics shown in Figure 11. The semantics are written in terms of the cryptographic objects themselves.

The key generation function keygen takes as input a random bitstring and outputs a random public key v_{pk} and a trapdoor v_{td}. The trapdoor permutation function tdp takes as inputs a key v_{pk} and a bitstring v_{in} and outputs a bitstring v_{out}. The inv function takes as inputs a key-trapdoor pair (v_{pk}, v_{td}) and a bitstring v_{in} and outputs a bitstring v_{out}. The hardcore predicate function hc takes as input a key v_{pk} and outputs a single bit.

We can use these to implement a special pseudorandom number generator $G_{pk} : \{0,1\}^k \to \{0,1\}^{4k}$ that has a trapdoor property, i.e., it is easy to compute, but difficult to invert except with special information called the "trapdoor."

$$G_{pk}(r) = \left(\mathbf{f}_{pk}^{(3n)}(r), \mathbf{B}(\mathbf{f}_{pk}^{(3n-1)}(r)), \ldots, \mathbf{B}(\mathbf{f}_{pk}(r)), \mathbf{B}(r)\right)$$

Here, $\mathbf{f}_{pk}$ is a trapdoor permutation over $\{0,1\}^k$, with $\mathbf{f}_{pk}^{(i)}(r)$ denoting the i^{th}-fold application of $\mathbf{f}_{pk}$, and $\mathbf{B}$ is a hardcore predicate for $\mathbf{f}_{pk}$. In ILC, this can be implemented as:

```
iterate :: ∀ a . Int → (a → a) → a → a
prg :: [Bit] → [Bit] → Nat → [Bit]
let prg pk r k =
  letrec aux j =
    if j ≤ 0 then [hc r]
    else hc (iterate j (tdp pk) r) : aux pk r (j − 1) in
  iterate (3 ∗ k) (tdp pk) r ++ aux pk r (3 ∗ k − 1)
```

F Universally Composable Commitment Protocol

In this section we give the full elaboration of our UC commitment instantiation. The specification functionality is given in the body in Figure 1, along with the protocol implementation in Section 6.5. Our development follows closely from the psuedocode in the UC literature [19], which we show here in Algorithm 1. The protocol relies on the CRS functionality which we define here in Figure 15. To briefly summarize what is going: the setup CRS samples a random string σ and two trapdoor pseudorandom generators (prgs pk_0, pk_1). To commit to the bit b, the commiter produces a string y that is the result of applying one or the other of the prgs, and if $b = 1$ additionally applying xor with σ. The intuitive explanation why this is hiding is that without the trapdoor, it is difficult to tell whether a random $4k$-bit string is in the range of either prg. To open the commitment, the committer simply reveals the preimage and the receiver checks which of the two cases applies. The intuitive explanation why this is binding is that it is difficult to find a pair $y, y \oplus \sigma$ that are respectively in the range of both prgs.

The UC proof consists of two simulators, one for the ideal world and one for the real world. The ideal world simulator, given in Figure 17 is ported directly from the UC literature [19], while the non-standard real world simulator, given in Figure 18, is required because our protocol emulation definition requires simulation in both directions. The key to the ideal world simulator is to allow the simulator to generate its own "fake" CRS, for which it stores the trapdoors. The string σ is not truly random, but instead is the result of combining two evaluations of the prgs. The ideal world simulator consists of two cases, depending on which of the parties is corrupt.

In the case that the committer P is corrupt, the simulator needs to be able to *extract* the committed value. The simulator is activated when $\mathcal{Z}$ sends a message (Commit′ y); in the real world, this is relayed by the dummy adversary to Q, who outputs Committed back to the environment. Hence to achieve the same effect in the ideal word, the simulator must send (Commit b) to $\mathcal{F}_{\text{Com}}$. To extract b from y, the simulator makes use of the prg trapdoor check which one has y in its range. It is necessary to argue by cryptographic reduction that this simulation is sound. To show this, we would define an alternative execution where the prg is substituted for a truly random function (i.e., a random oracle). If an environment $\mathcal{Z}$ could distinguish between these two worlds, then we could adapt the execution to distinguish the prg from random, violating the prg assumption.

In the case that the receiver Q is corrupt, the simulator needs to *equivocate*. The simulator is activated when $\mathcal{Z}$ inputs (Commit b) to P, after which $\mathcal{F}_{\text{Com}}$ sends Committed to the simulator. In the real world, the environment receives a commitment message (Commit′ y) from corrupted Q for some seemingly-random y. To achieve the same effect, the simulator must choose y. However, the simulator is next activated when the $\mathcal{Z}$ inputs (Open b) to P, after which the simulator learns b from $\mathcal{F}_{\text{Com}}$. However, in the real world the environment receives a valid opening (Opened′ b r) that is consistent with y and with the value chosen by the environment. Thus the simulator must initially choose y so that it can later be opened to either value b may take. The simulator achieves this by choosing σ and y ahead of time while generating the fake CRS. The reduction step is the same, and involves replacing prg with a true random function.

Recall that the motivation for the real world simulator is to rule out degenerate protocols that diverge in some way. For every well behaved environment such that the ideal world is PPT, we need to demonstrate an adversary in the real world that is also PPT. Fortunately, the real world simulator, shown in Figure 18 is much simpler than ideal world simulator. Essentially the simulator runs a copy of the honest protocol

$$\text{Expressions} \quad e ::= \mathsf{kgen}(e) \mid \mathsf{tdp}(e_1, e_2) \mid \mathsf{inv}(e_1, e_2) \mid \mathsf{hc}(e)$$

$\boxed{\Delta; \Gamma \vdash e : U}$ Under affine context Δ and unrestricted context Γ, expression e has type U.

$$\frac{\Delta; \Gamma \vdash e : [\mathsf{Bit}]}{\Delta; \Gamma \vdash \mathsf{kgen}(e) : [\mathsf{Bit}] \times [\mathsf{Bit}]}\ \text{kgen} \qquad \frac{\Delta_1; \Gamma \vdash e_1 : [\mathsf{Bit}] \qquad \Delta_2; \Gamma \vdash e_2 : [\mathsf{Bit}]}{\Delta_1, \Delta_2; \Gamma \vdash \mathsf{tdp}(e_1, e_2) : [\mathsf{Bit}]}\ \text{tdp}$$

$$\frac{\Delta_1; \Gamma \vdash e_1 : [\mathsf{Bit}] \times [\mathsf{Bit}] \qquad \Delta_2; \Gamma \vdash e_2 : [\mathsf{Bit}]}{\Delta_1, \Delta_2; \Gamma \vdash \mathsf{inv}(e_1, e_2) : [\mathsf{Bit}]}\ \text{inv} \qquad \frac{\Delta; \Gamma \vdash e : [\mathsf{Bit}] \to \mathsf{Bit}}{\Delta; \Gamma \vdash \mathsf{hc}(e) : \mathsf{Bit}}\ \text{hc}$$

$\boxed{e_1 \longrightarrow e_2}$ Expression e_1 reduces to e_2.

$$\frac{\mathbf{Gen}(v_r) = (v_{pk}, v_{td})_\infty \qquad v_{pk}, v_{td} \in \{0,1\}^k}{\mathsf{kgen}(v_r) \longrightarrow (v_{pk}, v_{td})_\infty}\ \text{kgen} \qquad \frac{\mathbf{f}(v_{pk}, v_{in}) = v_{out} \qquad \mathbf{f} : \{0,1\}^k \to \{0,1\}^k \to \{0,1\}^k}{\mathsf{tdp}(v_{pk}, v_{in}) \longrightarrow v_{out}}\ \text{tdp}$$

$$\frac{\mathbf{Inv}((v_{pk}, v_{td})_\infty, v_{in}) = v_{out} \qquad \mathbf{Inv} : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^k \to \{0,1\}^k}{\mathsf{inv}((v_{pk}, v_{td})_\infty, v_{in}) \longrightarrow v_{out}}\ \text{inv} \qquad \frac{\mathbf{B}(v_{pk}) = v \qquad \mathbf{B} : \{0,1\}^k \to \{0,1\}}{\mathsf{hc}(v_{pk}) \longrightarrow v}\ \text{hc}$$

Figure 11. Extending ILC with trapdoor permutations. The semantics are parameterized by a security parameter k.

for each of the corrupted parties. The simulation that results in this case is identical.

Protocol 1: Universally Composable Commitment

Public strings:
  σ: Random string in $\{0,1\}^{4n}$
  pk_0, pk_1: Keys for generator $G_k : \{0,1\}^n \to \{0,1\}^{4n}$
Commit(b):
  $r \leftarrow \{0,1\}^n$
  $y := G_{pk_b}(r)$
  if $b = 1$ then $y := y \oplus \sigma$
  Send (Commit, y) to receiver.
  Upon receiving (Commit, y) from A, B outputs (Receipt).
Decommit(x):
  Send (b, r) to receiver.
  Receiver checks $y = G_{pk_b}(r)$ for $b = 0$, or $y = G_{pk_b}(r) \oplus \sigma$ for $b = 1$. If verification succeeds, then B outputs (Open, b).

```
dummy :: ∀ a ... . Nat → [Bit] → Crupt → ··· a
let dummyA k bits crupt toZ toF toP toQ toQasP frZ frF frP frQ toPasQ=
  let fwd2Z () c = loop (λ m . wr (X2Z m) → toZ) c in
      loop (λ x . match x with
                  | A2F m ⇒ wr m → toF
                  | A2P m ⇒ if crupt == CruptP
                            then wr m → toQasP
                            else wr m → toP) frZ
    |▷ fwd2Z () frF
    |▷ fwd2Z () frP
    |▷ fwd2Z () frQ
```

Figure 12. Dummy adversary. The dummy adversary forwards messages from the environment to either the functionality (if the message has constructor A2F) or the party p (if the message has constructor A2P). Similarly, the dummy adversary forwards messages from the functionality or the procotol parties to the environment.

```
dummyP :: ∀ a b ... . Nat → [Bit] → Wr a → ··· b
let dummyP k r toZ toF toQ frZ frF frQ = fwd toF frZ |▷ fwd toZ frF
```

Figure 13. Dummy party. The dummy party simply relays information between the environment and the functionality.

```
fCrs :: ∀ a ... . Nat → [Bit] → Crupt → ··· a
let fCrs k bits crupt toP toQ toA frP frQ frA =
  let (σ, bits) = sample (4*k) bits in
  let (r0, bits) = sample k bits in
  let (r1, bits) = sample k bits in
  let pk0 = kgen k r0 in
  let pk1 = kgen k r1 in
  let pub = PublicStrings σ pk0 pk1 in
  let replyCrs to fr = loop (λ _ . wr pub → to) fr in
      replyCrs toP frP
    |▷ replyCrs toQ frQ
    |▷ replyCrs toA frA
```

Figure 14. Ideal functionality for common reference string.

```
bCrs :: ∀ a ... . Nat → [Bit] → Crupt → ··· a
let bCrs k bits crupt toP toQ toA frP frQ frA =
  let (r0, bits) = sample k bits in
  let (r1, bits) = sample k bits in
  let pk0 = kgen k r0 in
  let pk1 = kgen k r1 in
  let σ = xors (prg pk0 r0) (prg pk1 r1)
  let pub = PublicStrings σ pk0 pk1 in
  let replyCrs to fr = loop (λ _ . wr pub → to) fr in
      replyCrs toP frP
    |▷ replyCrs toQ frQ
    |▷ replyCrs toA frA
```

Figure 15. Bad ideal functionality for common reference string.

```
fCom :: Nat → [Bit] → Crupt → ··· —∘R 𝟙
let fCom k bits crupt toP toQ toA frP frQ frA =
  let (!(Commit b), frP) = rd frP in
    wr Receipt → toQ ;
    let (!Open, frP) = rd frP in
      wr (Opened b) → toQ
```

Figure 16. Ideal functionality for one-time bit commitment.

```
simI :: Nat → [Bit] → Crupt → ··· —∘ 𝟙
let simI k bits crupt toZ toF toP toQ frZ frF frP frQ =
  let (pk0,td0) = kgen k in
  let (pk1,td1) = kgen k in
  let (r0, bits) = sample k bits in
  let (r1, bits) = sample k bits in
  let σ = xors (prg pk0 r0) (prg pk1 r1) in
    match crupt with
    | CruptP ⇒
      let (!GetCRS, frZ) = rd frZ in
        wr (X2Z (PublicStrings σ pk0 pk1)) →toZ ;
        let (!(A2P (Commit' y)), frZ) = rd frZ in
          if check td0 pk0 y then
            wr (Commit 0) → toP
          else
            if check td1 pk1 (xors y σ) then
              wr (Commit 1) → toP
            else error "Fail" ;
          let (!(A2P (Open' b r)), frZ) = rd frZ in
            if b == 0 && y == prg pk0 r ||
              b == 1 && y == xors (prg pk1 r) σ
            then wr Open → toP
            else error "Fail"
    | CruptQ ⇒
      let (!GetCRS, frZ) = rd frZ in
        wr (X2Z (PublicStrings σ pk0 pk1)) → toZ ;
        let (!Receipt, frQ) = rd frQ in
        let y = prg pk0 r0 in
          wr (X2Z (Commit' y)) → toZ ;
          let (!(Opened b'), frQ) = rd frQ in
            if (b' == 0) then
              wr (X2Z (Opened' r0)) → toZ
            else
              wr (X2Z (Opened' r1)) → toZ
    | CruptNone ⇒ error "Fail"
```

Figure 17. Ideal world simulator for UC commitment.

```
simR :: Nat ⟶ [Bit] ⟶ Crupt ⟶ ··· ⊸ 𝟙
let simR k bits crupt toZ toF toP toQ frZ frF frP frQ =
  match crupt with
  | CruptP ⇒
    let (!(Commit b), frZ) = rd frZ in
      wr GetCRS ⟶ toF ;
      let (!(PublicStrings σ pk0 pk1), frF) = rd frF in
      let r = take k bits in
      let y = if b == 0 then prg pk0 r else xors (prg pk0 r) σ in
        wr (Commit' y) ⟶ toQ ;
        let (!(Open), frZ)= rd frZ in
          wr (Open' b r) ⟶ toQ
  | CruptQ ⇒
    let (!(Commit' y), frQ) = rd frQ in
      wr Receipt ⟶ toZ ;
      let (!(Open' b r), frQ) = rd frQ in
        wr (Opened b) ⟶ toZ
  | CruptNone ⇒ error "Fail"
```

Figure 18. Real world simulator for UC commitment.