

# IPDL: A Probabilistic Dataflow Logic for Cryptography

Xiong Fan\*    Joshua Gancher†    Greg Morrisett†    Elaine Shi†    Kristina Sojakova†

December 6, 2019

## Abstract

While there have been many successes in verifying cryptographic security proofs of noninteractive primitives such as encryption and signatures, less attention has been paid to interactive cryptographic protocols. Interactive protocols introduce the additional verification challenge of *concurrency*, which is notoriously hard to reason about in a cryptographically sound manner.

When proving the (approximate) observational equivalence of protocols, as is required by simulation based security in the style of Universal Composability (UC), a *bisimulation* is typically performed in order to reason about the nontrivial control flows induced by concurrency. Unfortunately, bisimulations are typically very tedious to carry out manually and do not capture the high-level intuitions which guide informal proofs of UC security on paper. Because of this, there is currently a large gap of formality between proofs of cryptographic protocols on paper and in mechanized theorem provers.

We work towards closing this gap through a new methodology for iteratively constructing bisimulations in a manner close to on-paper intuition. We present this methodology through Interactive Probabilistic Dependency Logic (IPDL), a simple calculus and proof system for specifying and reasoning about (a certain subclass of) distributed probabilistic computations. The IPDL framework exposes an *equational* logic on protocols; proofs in our logic consist of a number of rewriting rules, each of which induce a single low-level bisimulation between protocols.

We show how to encode simulation-based security in the style of UC in our logic, and evaluate our logic on a number of case studies; most notably, a semi-honest secure Oblivious Transfer protocol, and a simple multiparty computation protocol robust to Byzantine faults. Due to the novel design of our logic, we are able to deliver mechanized proofs of protocols which we believe are comprehensible to cryptographers without verification expertise. We provide a mechanization in Coq of IPDL and all case studies presented in this work.

## 1 Introduction

With new decentralized computing paradigms such as blockchains and cloud outsourcing, the community’s pace at designing and rolling out new, efficient cryptographic protocols has accelerated. Many of the protocols being deployed involve rich building blocks such as commitment schemes, zero-knowledge proofs, oblivious transfer, and multi-party computation. As a result, it has become increasingly important to scale up the effort of proving cryptographic protocols secure. An important and necessary step towards this goal is to enable *easy-to-use* computer-aided proof systems for general and possibly complex cryptographic protocols.

---

\*University of Maryland. Email: [xfan@cs.umd.edu](mailto:xfan@cs.umd.edu). Part of this work was done while at Cornell.

†Cornell University. Email: [jrg358@cornell.edu](mailto:jrg358@cornell.edu), [elaine@cs.cornell.edu](mailto:elaine@cs.cornell.edu), [greg.morrisett@cornell.edu](mailto:greg.morrisett@cornell.edu), [sojakova.kristina@gmail.com](mailto:sojakova.kristina@gmail.com)

Existing systems for mechanically verifying cryptography are either symbolic in nature [BSCS18, MSCB13, DY83, Cre08] and therefore do not exactly match the computational reduction-style of reasoning that underlies the mathematical foundation of cryptography, or do not provide much support for encoding general protocols (e.g., EasyCrypt [BGHZ11], CryptVerif [Bla06], and FCF [PM15]). Moreover, except for a few recent/concurrent endeavors [CSV19, LSBM19], most prior systems do not provide good support for reasoning about general multi-party cryptographic protocols and their composition.

In the cryptography literature, Universal Composition (UC) [Can01] is the *de facto* framework for modeling distributed cryptographic protocols and conduct compositional security reasoning. To prove a protocol secure in the UC framework, one must prove it *approximately equivalent* to a particular idealization. This is typically done through a number of low-level (approximate) equivalences to intermediate protocols, each of which either simplifies the protocol in some way or applies a cryptographic hardness assumption.

In order to formally prove one of these low-level protocol equivalences correct, the formal methods community typically employs a *bisimulation* argument (e.g., as seen in EasyUC [CSV19], FCF [PM15], and CryptHOL [LSBM19]), which requires one to construct a particular relational invariant between (distributions of) states of the two protocols which satisfies certain behavioral properties. Unfortunately, it is quite unsatisfactory to use bisimulations directly to reason about cryptographic protocols, since bisimulations induce overly technical proofs which do not match cryptographers’ style of reasoning.

## 1.1 Our Contributions

Our goal is to design a new technique for proving approximate equivalences of cryptographic protocols that is *easy-to-use* and matches the style of on-paper cryptographic proofs. We do so by constructing a new methodology based on *channel dependency* which does not require the end user to manually construct bisimulations. Instead, the end user interacts with an *equational logic* in order to conduct cryptographic proofs. In this work, we introduce Interactive Probabilistic Dependency Logic (IPDL), one such equational logic for cryptographic protocols.

For simplicity, we currently only consider a certain important subclass of protocols (those which do not make use of nontrivial control flow); however, we believe our proof technique can be readily extended to more complicated scenarios. More detail about this limitation (and others) can be found in Section 9.

**A probabilistic logic capturing reactions and dependencies** IPDL encodes a distributed computation using *channels* and *reactions*: when all inputs to a reaction have been collected, the reaction is triggered and a value is written to its output channel. This style of encoding features *dependency among channels* (both dependency in *timing* and *value*) as a first-class citizen, since each reaction effectively captures the dependency of the output events on the input events. Since many cryptographic proofs are fundamentally about removing dependencies (e.g., removing a message from a ciphertext), we believe our dependency-centric perspective enables proofs which are often preferable to manually constructing bisimulations. Our perspective can be seen as following in the spirit of equational security [MT13].

In IPDL, a protocol  $\pi$  is expressed as the collection of code for each protocol participant. Channels that are not internally consumed by  $\pi$  form the interface of  $\pi$  to the outside world, including interfaces to an external environment who provides inputs to the protocol and receives outputs and interfaces to the adversary. In this way, IPDL allows us to capture UC-style protocol proofs: to show that  $\pi_R$  implements  $\pi_I$ , we can define a simulator  $\text{Sim}$  in IPDL which converts attacks on

$\pi_R$  to attacks in  $\pi_I$ . We can then ask in IPDL whether the execution of  $\pi_R$  is (approximately) equivalent to the execution of  $\pi_I||\text{Sim}$ , where the notion of equivalence implicitly and universally quantifies over the environment that connects to the external interfaces of the protocol.

**A compositional semantics for IPDL** Proofs in IPDL can be obtained by applying a set of high-level rewrite rules defined over channels and dependencies. The main judgement in IPDL is that  $\Sigma \rightsquigarrow_\varepsilon \Sigma'$ , meaning that (informally) no environment interacting with  $\Sigma$  can distinguish it from  $\Sigma'$  with probability better than  $\varepsilon$ . We prove the equational logic of IPDL sound for approximate protocol equivalence through corresponding each rewrite rule to a low-level bisimulation.

The semantics for protocols in our setting is quite subtle, since concurrency is usually formalized using nondeterminism. In the presence of cryptography, however, reasoning about nondeterminism is undesirable, since it intuitively corresponds to *demonic choice*, which leads to an unrealistic attacker model. We take an approach similar to CryptHOL [LSBM19] and UC [Can01], in which all internal computations which usually are resolved through nondeterminism are instead resolved by the protocol agents themselves. Unlike previous approaches, however, the user of our system is not exposed to the technical formalities this solution induces in resolving this internal choice.

**Implementation in Coq** We have implemented the IPDL logic in Coq as well as mechanized proofs for all of the case studies considered in this paper. Our Coq implementation is publicly available at <https://github.com/ipdl/ipdl>. To enable succinct mechanized proofs, we provide a library of tactics which automatically apply our rewriting rules, hiding low-level details of the proof from the user.

**Case studies** To illustrate the usefulness of IPDL, we present case studies for several cryptographic protocols, including 1) realizing a secure channel from an authenticated channel under trusted setup and an encryption scheme with suitable security; 2) the ElGamal encryption system; 3) a UC-secure rock-paper-scissors protocol which employs UC-secure commitments; and 4) an oblivious transfer (OT) scheme realized from trapdoor permutations. All case studies have been mechanized in Coq and come with our open source repository. We demonstrate in Section 7 that IPDL delivers proofs which scale well with the complexity of the protocol in question.

In summary, we provide a mechanized proof system for cryptographic protocols expressed naturally as an equational logic on protocols. IPDL is distinctive in that channel dependency is captured as a first-class notion. Such a dependency-centric design allows every step of a cryptographic proof to be expressed using a simple and concise rewrite rule. We believe that our dependency-based perspective has the potential to deliver concise and easy-to-understand proofs. As an example, for the most sophisticated case study we consider, Oblivious Transfer, the full definitions and proofs contain 926 lines of Coq code; we argue this is a quite reasonable number, given the complexity of the protocol.

The full version, containing all IPDL examples and proofs, can be found at <https://github.com/ipdl/ipdl>.

## 2 IPDL by Example: Authenticated to Secure Channel

As an introductory example, we will consider a simple protocol that constructs a one-time use secure communication network between a sender S and a receiver R from an authenticated network and a trusted setup for a symmetric encryption key. (This example is similar to the main example in [Mau11].) For simplicity, we do not model message size, but instead only consider messages of

some arbitrary but fixed length. It is straightforward to extend this example to support a dynamic message length.

The construction is modeled as an interactive protocol  $\pi_R$  between participants S, R, the actual network  $\mathcal{F}_{\text{auth}}$ , a key functionality  $\mathcal{F}_{\text{key}}$ , and an implicit external context. The context is thought of as serving two roles: an *environment*, supplying and receiving external inputs and outputs from the parties, and an *attacker*, who interacts with the attack surface of the protocol. Channels meant for the attacker are written in **red**, while those for the external environment are written in **blue**. The sender receives an input message  $m$  from the external environment, a key  $k$  from the trusted setup, and sends the value  $c \stackrel{\$}{\leftarrow} \text{Enc}(k, m)$  to the network. The network leaks this ciphertext  $c$  to the attacker, and waits for a confirmation **ans<sub>R</sub>** from the attacker before delivering the ciphertext to the receiver, who may decrypt the ciphertext using  $k$ , and output the decrypted message.

In Figure 1, we see the protocol  $\pi_R$  encoded in IPDL. (Strictly speaking, we see the *uncorrupted execution* of  $\pi_R$ . This terminology is described in Section 4. In this section, we use the term “protocol” loosely, but define it formally later.) Protocols in IPDL are collections of *inputs* and *reactions*, which interact through *channels*. Channels are used for external input, output, as well as internal hidden computations. The first line, **in<sub>S</sub> : msg inp**, names a hidden channel **in<sub>S</sub>** of type **msg**, and declares it an external input to the protocol. All channels may only be assigned to once. Lines 2-4 implement the trusted setup  $\mathcal{F}_{\text{key}}$  through three reactions: the first samples from the distribution  $\text{Rnd}(\mathcal{K})$  and assigns the result to a local channel **keygen**, while the next two simply copy this result to **key<sub>S</sub>** and **key<sub>R</sub>**. Reactions may fire only when all input to the reaction is available and the output channel of the reaction has not been assigned to yet. When this is the case, we sample from the corresponding distribution and assign the return value to the output channel.

Line 5 corresponds to the code of the sender: when the input message  $m$  and the key  $k$  are available, the local channel **send<sub>R</sub><sup>S</sup>** is probabilistically assigned to from the distribution  $\text{Enc}(k, m)$ . This ciphertext is then copied to the channel **leak** on the next line. This channel **leak** is tagged with a *visibility label* **vis**, meaning that the external adversary is able to read from this channel. Channels without a visibility label are implicitly tagged with **hid**, meaning that the channel is hidden from the context.

Next, we wait for an answer from the adversary on channel **ans<sub>R</sub>**. (If no type is present, it is implicitly of unit type.) This channel models a network adversary that can control if and when the message gets delivered. Once the adversary has given **ans<sub>R</sub>**, we deliver the ciphertext on **send<sub>R</sub><sup>S</sup>** to the sender through the channel **deliv**. The sender may now decrypt the ciphertext, using the input from **deliv** and the key **key<sub>R</sub>**.

Our main goal of IPDL is to keep the formal framework surrounding protocols as simple as possible, while still maintaining enough expressive power to do interesting cryptographic proofs. This is done by restricting our scope in two ways: first, as discussed above, channels may only be written once. This keeps the value of all channels unambiguous, and guarantees that our protocols will terminate. Second, all behaviors on channels are constrained to only be a function of the input values of the channel. Thus, we do not consider protocols in which a channel’s value depends on the order in which input arrives.

Reactions are generally of the form  $\Gamma \vdash c \ell : \tau \leftarrow D$ , read as “once all input channels in  $\Gamma$  have been set, if  $c$  has not been set, sample from  $D$  and write its value of type  $\tau$  to  $c$ .” As discussed above,  $\ell \in \{\text{hid}, \text{vis}\}$  controls whether the channel is visible to the outside world. We model non-probabilistic assignments through the syntactic sugar  $\Gamma \vdash c \ell : \tau := e$ , which is an abbreviation for  $\Gamma \vdash c \ell : \tau \leftarrow \delta(e)$ , where  $\delta(e)$  is the point mass distribution which assigns probability 1 to the value at  $e$ .

1.  $\text{in}_S : \text{msg inp}$
2.  $\cdot \vdash \text{keygen} : \text{key} \leftarrow \text{Rnd}(\mathcal{K})$
3.  $k : \text{keygen} : \text{key} \vdash \text{key}_S : \text{key} := k$
4.  $k : \text{keygen} : \text{key} \vdash \text{key}_R : \text{key} := k$
5.  $m : \text{in}_S : \text{msg}, k : \text{key}_S : \text{key}$   
 $\vdash \text{send}_R^S : \text{ct} \leftarrow \text{Enc}(m, k)$
6.  $c : \text{send}_R^S : \text{ct} \vdash \text{leak vis} := c$
7.  $\text{ans}_R \text{ inp}$
8.  $c : \text{send}_R^S : \text{ct}, \text{ans}_R \vdash \text{deliv} : \text{ct} := c$
9.  $c : \text{deliv} : \text{ct}, k : \text{key}_R : \text{key}$   
 $\vdash \text{out}_R \text{ vis} : \text{msg} := \text{Dec}(c, k)$

Figure 1: The (uncorrupted) execution of the secure network protocol  $\pi_R$  expressed in IPDL. Each line is either a locally controlled reaction of the form  $\Gamma \vdash c : \tau \leftarrow D$ , or an input of the form  $c : \tau \text{ inp}$ . Locally controlled reactions are hidden by default; visible reactions are marked vis.

**Reasoning in IPDL** The main feature in IPDL is that protocol events are naturally structured according to their *dependency* on other events. This is in contrast to more conventional program logics, which impose a linear order on the statements in programs. Structuring our protocols in terms of their dependency allows us to reason easily about equivalence of protocols.

Since we know exactly what distribution the values of a channel may take given its input, we are able to manipulate channels not only by rewriting their output distributions to equivalent ones, but by rewriting their *dependency* to equivalent ones. Consider the channel  $\text{out}_R$  in Figure 1: it decrypts a ciphertext  $c$  coming from the channel  $\text{deliv}$ , where  $\text{deliv}$  is simply copied from the channel  $\text{send}_R^S$ . Additionally, we know from the protocol that  $\text{key}_S$  and  $\text{key}_R$  are equal, when both are defined. We may reflect this information in the protocol by performing a sequence of rewrites to obtain that  $\text{out}_R$  is equivalent to the reaction

$$c : \text{send}_R^S : \text{ct}, k : \text{key}_S : \text{key}, \text{ans}_R \\ \vdash \text{out}_R \text{ vis} : \text{msg} := \text{Dec}(c, k),$$

thus essentially performing a partial evaluation of the protocol in order to propagate the value of  $\text{send}_R^S$  to  $\text{out}_R$ . We may then perform additional rewrites on the protocol to obtain that this reaction is equivalent to

$$m : \text{in}_S : \text{msg}, k : \text{key}_S : \text{key}, \text{ans}_R \\ \vdash \text{out}_R \text{ vis} : \text{msg} := m,$$

by using the correctness property of the encryption scheme. Now, we see that  $\text{out}_R$  no longer uses the value of  $k$ : in certain cases, we may then soundly *remove* dependencies from reactions, producing the following simplified reaction for  $\text{out}_R$ :

$$m : \text{in}_S : \text{msg}, \text{ans}_R \vdash \text{out}_R \text{ vis} : \text{msg} := m.$$

**Protocol Simulation** We express protocol security in the style of UC [Can01] by using IPDL to express and prove equivalence of protocols. In Figure 2 we have an idealized version of a secure channel. The sender and receiver do not communicate using an encryption scheme, but instead in cleartext. However, the network functionality only reports the message `query` to the adversary, and does not leak any information about the ciphertext. Similarly to  $\pi_R$ , the network here waits for the adversary to send `ansI` before delivering the message.

$$\begin{aligned}
& \text{in}_S : \text{msg inp} \\
& m : \text{in}_S : \text{msg} \vdash \text{send}_I^S : \text{msg} := m \\
& m : \text{send}_I^S : \text{msg} \vdash \text{query vis} \\
& \text{ans}_I \text{ inp} \\
& m : \text{send}_I^S : \text{msg}, \text{ans}_I \vdash \text{deliv} : \text{msg} := m \\
& m : \text{deliv} : \text{msg} \vdash \text{out}_R \text{ vis} : \text{msg} := m
\end{aligned}$$

Figure 2: The idealized secure network  $\pi_I$  expressed in IPDL.

We wish to argue that  $\pi_R$  somehow performs the same job as  $\pi_I$ . Note that the input and output channels for the two parties are the same across  $\pi_R$  and  $\pi_I$ ; the only input and visible channels that differ are those which correspond to the respective adversaries. In  $\pi_R$ , the adversary learns the ciphertext from `leak`, and chooses to deliver that ciphertext using `ansR`; in  $\pi_I$ , the adversary only learns the message has been sent using `query`, and chooses to deliver the message using `ansI`.

In order for  $\pi_R$  to be as effective as  $\pi_I$ , we need a way of *converting* all adversarial attacks on  $\pi_R$  to attacks on  $\pi_I$ . If we can invent a converter `Sim` (also called the *simulator*) which converts attacks on  $\pi_R$  to attacks on  $\pi_I$ , then we can ask in IPDL whether  $\pi_R$  is equivalent to `Sim|| $\pi_I$` , where `||` is the parallel composition operator in IPDL. Details on our composition operator are given in Section 3. If such a `Sim` exists, we say that  $\pi_R$  *realizes*  $\pi_I$ . Crucially, this simulator only has access to the channels designated for the adversary. Following UC, we may also *corrupt* certain parties in both protocols, and require the corrupted protocols also satisfy this realization property. More details on simulators, corruption, and protocol realization is in Section 4.

In the case when neither party is corrupted, the simulator is standard: for `Sim` to convert an attack on  $\pi_R$  to an attack on  $\pi_I$ , it will receive a `query` message from  $\pi_I$  and need to deliver a corresponding `leak` message to  $\pi_R$ . It does not have access to the key nor the message, so it will generate a fresh key  $k'$  from  $\text{Rnd}(\mathcal{K})$  and output a ciphertext from  $\text{Enc}(0, k')$ , where 0 is any distinguished message from the message space. More details on this simulator are given in Section 5.

**Approximate compositional reasoning** In order to show that  $\pi_R$  is equivalent to `Sim|| $\pi_I$` , we must apply a cryptographic assumption about the encryption scheme; namely, that over the randomness of  $k$ , ciphertexts coming from  $\text{Enc}(m, k)$  look identical to ciphertexts coming from  $\text{Enc}(0, k)$  for all  $m$ . In Figure 3, we express this cryptographic assumption as an approximate equivalence of two reaction sets `OSSR` and `OSSI` (`OSS` standing for *one-shot security*). In order to apply this assumption to  $\pi_R$  and `Sim|| $\pi_I$` , we first must rewrite  $\pi_R$  to be equivalent to  $\pi'_R||\text{OSS}_R$ . Thus, we “factor out” the encryption performed in  $\pi_R$  to `OSSR`. This rewrite is made possible by our simplifications on channel `outR` from above, wherein we use the correctness property of the encryption scheme to essentially give the plaintext value directly to the receiver.



Once this factoring is done, we make use of a general rule in IPDL which states that the equivalence notion  $\rightsquigarrow$  is a congruence: since  $\text{OSS}_R$  is (approximately) equivalent to  $\text{OSS}_I$  by assumption, we know that  $\pi_R \rightsquigarrow \pi'_R || \text{OSS}_R$  is approximately equivalent to  $\pi'_R || \text{OSS}_I$ , where the encryption is replaced with a dummy ciphertext. We may then use the logic to deduce that  $\pi'_R || \text{OSS}_I$  is equivalent to  $\text{Sim} || \pi_I$ , which concludes our proof. More detail about this proof (as well as a security proof when the sender is malicious) is given in Section 5.

In Section 3, we describe IPDL and its equational logic in detail; in Section 4, we present a general framework for expressing cryptographic security in IPDL, including both semi-honest and malicious corruption models; in Sections 5 and 6, we present multiple case studies of cryptographic proofs in IPDL. Finally, in Section 7, we discuss our mechanization of IPDL in Coq.

$$\begin{array}{ll}
\text{in}_{\text{OSS}} : \text{msg inp} & \text{in}_{\text{OSS}} : \text{msg inp} \\
\cdot \vdash \text{keygen} : \text{key} \leftarrow \text{Rnd}(\mathcal{K}) & \cdot \vdash \text{keygen} : \text{key} \leftarrow \text{Rnd}(\mathcal{K}) \\
m : \text{in} : \text{msg}, k : \text{keygen} : \text{key} & m : \text{in} : \text{msg}, k : \text{keygen} : \text{key} \\
\vdash \text{out}_{\text{OSS}} \text{ vis} : \text{ct} \leftarrow \text{Enc}(k, m) & \vdash \text{out}_{\text{OSS}} \text{ vis} : \text{ct} \leftarrow \text{Enc}(k, 0)
\end{array}$$

Figure 3: Definition of one-shot security of an encryption scheme in IPDL. The top reaction set,  $\text{OSS}_R$ , is approximately equivalent to  $\text{OSS}_I$ . (Our definition of approximate equivalence is given in Section 3.)

### 3 Logic

Protocols in our logic are modeled as *reaction sets*, whose syntax is shown in Figure 4. Reaction sets operate on a set of *channels*, which may model I/O communication with the environment as well as internal computation through hidden channels. Input channels are defined through the syntax  $c \text{ inp} : \tau$ , which declares the channel  $c$  as an external input of type  $\tau$ . Output and hidden channels are modeled through *reactions* of the form  $x_1 : c_1 : \tau_1, \dots, x_n : c_n : \tau_n \vdash c \ell : \tau \leftarrow D$ , which says that if for all  $i \in \{1, \dots, n\}$ ,  $c_i$  is bound to a value  $x_i$  of type  $\tau_i$ , and if  $c$  is not yet bound to a value, then sample from the distribution  $D$ , ranging over values of type  $\tau$ , and bind the result to  $c$ . Thus, once a value is set to a channel, it remains set on the channel for the entire execution. The left side  $x_1 : c_1 : \tau_1, \dots, x_n : c_n : \tau_n$  is called the *context* of the channel  $c$ . The label  $\ell$  may be either *out* or *hid*, denoting if the channel is a hidden internal computation, or an output to an external environment. Reactions written without a label are implicitly marked hidden.

Our logic is parameterized over a universe of types  $\tau$  including products and probability distributions. All probabilistic reasoning in IPDL is factored out into an equational theory  $\mathcal{E}$  supporting

$$\begin{array}{ll}
\tau & := \dots \mid \tau \times \tau \mid \mathcal{D}(\tau) & \text{types} \\
\ell & := \text{out} \mid \text{hid} & \text{output/hidden channels} \\
\Gamma & := \emptyset \mid x : c; \Gamma & \text{reaction context} \\
\Sigma & := \emptyset \mid \Gamma \vdash c \ell \leftarrow D; \Sigma & \text{locally controlled reactions} \\
& \quad \mid c : \tau \text{ inp}; \Sigma & \text{inputs}
\end{array}$$

Figure 4: Main Syntax of IPDL.

the judgement  $\Gamma \vDash_{\mathcal{E}} D = D'$ , meaning that for all valuations of the variables in  $\Gamma$ , the distributions  $D$  and  $D'$  are equivalent. A channel  $c$  is *defined* in reaction set  $\Sigma$  if it appears in  $\Sigma$  as an input, output, or hidden channel. Given a reaction set  $\Sigma$ , we define the sets  $\text{out}(\Sigma)$  to be the output channels set in  $\Sigma$ , and similarly for  $\text{in}(\Sigma)$  and  $\text{hid}(\Sigma)$ . We assume the following well-formedness conditions on reaction sets  $\Sigma$  and contexts  $\Gamma$ :

1. All variable names  $x_i$  in  $\Gamma$  are distinct.
2. No duplicate channels exist in  $\Sigma$ .
3. All appearances of a channel in  $\Sigma$  have the same type.

In particular, the above conditions imply that the sets  $\text{out}(\Sigma)$ ,  $\text{in}(\Sigma)$ , and  $\text{hid}(\Sigma)$  are pairwise disjoint.

We now define a parallel composition operator for reaction sets. Two reaction sets  $\Sigma$  and  $\Sigma'$  are *compatible* if their output channels are disjoint, and no hidden channel in one reaction set appears in the other. For compatible  $\Sigma$  and  $\Sigma'$ , we define the union  $\Sigma \cup \Sigma'$  to be the union of all the hidden and output channels in  $\Sigma \cup \Sigma'$ , plus all the input channels of  $\Sigma$  and/or  $\Sigma'$  that do not simultaneously appear as outputs of the other reaction set. If  $c$  is an output channel of a reaction set  $\Sigma$ , we denote by  $\Sigma \text{ hide } c$  the reaction set that defines  $c$  to be a hidden, rather than an output channel. Finally, for compatible  $\Sigma$  and  $\Sigma'$  we denote by  $\Sigma \parallel \Sigma'$  the union of  $\Sigma$  and  $\Sigma'$  followed by a successive hiding of all the output channels of  $\Sigma$  and/or  $\Sigma'$  that simultaneously appear as inputs to the other reaction set.

### 3.1 Semantics

We model IPDL protocols in a manner similar to CryptHOL [LSBM19] and UC [Can01]. Similarly to an IPDL protocol, a protocol agent has a set of typed *input* and *output* channels. Input channels are partitioned into two kinds: *token input channels* and *listener input channels*; the latter allow a certain form of broadcast in our semantics. A *listener input*  $I_L$  is a pair of a listener input channel  $I$  and a value  $v$  of the type  $\tau$  associated to the channel  $I$ ; similarly, we define token inputs  $I_T$  and outputs  $O$  to be pairs of channels and values on that channel.

Unlike an IPDL protocol, a general protocol agent does not have any explicit hidden channels. Instead, the behavior of a protocol agent is completely determined by a set  $S$  of states (with a distinguished start state  $s$ ) and two transition functions  $\mu$  and  $\nu$ :

- The transition function  $\mu$  for listener inputs associates to each state  $e$  and each listener input  $I_L$  a distribution  $\mathcal{D}(S)$  on states. Intuitively, this models the fact that the protocol agent learns the value  $v$  sent on the channel  $I$  but does not take control of the computation.
- The transition function  $\nu$  for token inputs associates to each state  $e$  and each token input  $I_T$  a distribution  $\mathcal{D}(S \times (1 + O))$  on states and, optionally, outputs. This models the fact that the protocol agent learns the value  $v$  sent on the channel  $I$  and takes control of the computation, optionally executing an output.

Similarly to IPDL programs, we say that two protocol agents are *compatible* if their respective token input channels and their respective output channels are disjoint, and all appearances of the same channel common to both have the same type. The same channel  $C$  of type  $\tau$  can appear, *e.g.*, as a listener input to both protocol agents, as a listener input to one and a token input to the other, or as an input (of either kind) to one and an output of the other. A *protocol*  $\sigma$  is a finite family of pairwise compatible protocol agents. Every channel appearing in  $\sigma$  can be classified as one of three kinds:



- a *free input channel* is a channel that does not appear as an output of any protocol agent of  $\sigma$
- a *free output channel* is a channel that does not appear as a token input of any protocol agent of  $\sigma$
- an *internal channel* is a channel that appears as both an output and a token input channel of some protocol agent(s) of  $\sigma$

Two protocols are *comparable* if their respective sets of free input channels, free output channels, and internal channels coincide. Two protocols  $\sigma$  and  $\sigma'$  are *compatible* if the protocol agents selected from each protocol are pairwise compatible, no internal channel of  $\sigma$  appears in  $\sigma'$  and vice versa, and no free output of  $\sigma$  appears as a free output of  $\sigma'$ . In this case the union  $\sigma \parallel \sigma'$  of  $\sigma$  and  $\sigma'$  forms a new protocol.

In contrast with other formalisms such as Task PIOA [CCK<sup>+</sup>18], this form of modeling protocols does not involve any nondeterministic choice; instead, the control flow of the protocol is entirely determined by the protocol agents themselves who decide which message to send next. To deterministically execute a protocol  $\sigma$  consisting of  $n$  protocol agents  $P_1, \dots, P_n$ , we assume a *context*  $\mathcal{C}$ , specified by the following data:

- A protocol agent  $A$  (“adversary”) that has no listener inputs, whose token input channels are a subset of the free output channels of  $\sigma$ , and whose output channels are a subset of the free input channels of  $\sigma$ . This in particular means that  $A$  does not have access to any communication happening on the channels internal to  $\sigma$ .
- An *initial distribution*  $\mathcal{D}(S \times (1 + O))$  on the states and possible outputs of  $A$ .
- A *yield transition function*  $\xi$  associating to each state  $e$  and each (token) input  $I_T$  of  $A$  a distribution  $\mathcal{D}(S \times (1 + O))$  on the states and possible outputs of  $A$ .

The adversary  $A$  thus adaptively generates free inputs to  $\sigma$  while observing the communication on the free outputs of  $\sigma$ . The initial distribution specifies how to start the execution of  $\sigma$ , whereas the yield transition function  $\xi_i$  specifies how to continue the execution of  $\sigma$  in the case when some protocol agent (but the adversary does not learn which) received a token input but has not followed by an output (hence *yielded* to the adversary). The effects of the initial distribution and the yield transition functions do not factor into the trace of the execution since they do not exist as inputs to  $A$ .

If  $\sigma$  has a free input channel that does not appear as an output channel of  $A$ , then no communication on this channel is possible. On the other hand, if there is a free output channel of  $\sigma$  that does not appear as an input channel to  $\sigma$ , then any message sent on this channel will necessarily terminate the protocol. Likewise, if an agent  $P_i$  yields to  $A$  and  $A$  does not follow by an output, the protocol terminates. This is also the case if the initial distribution does not induce an output (this case is the least interesting but we allow it for consistency).

A context  $\mathcal{C}$  for a protocol  $\sigma$  deterministically induces a distribution on executions of length  $\leq k$ , and thus on execution traces of length  $\leq k$  (which we denote by  $\text{trDist}_k(\sigma)$ ), by executing  $\sigma$  for  $k$  steps for each  $k \in \mathbb{N}$ . Following [CCK<sup>+</sup>06a, CCK<sup>+</sup>06b, CCK<sup>+</sup>18], we compare protocols by comparing their trace distributions: let  $\varepsilon(\tau, \mathcal{C}, n)$  be a mapping from protocols, contexts, and natural numbers to real numbers in  $[0, 1]$ . Then, given two comparable protocols  $\sigma$  and  $\sigma'$ , we say that  $\sigma \rightsquigarrow_\varepsilon \sigma'$  if for all compatible protocols  $\tau$  and contexts  $\mathcal{C}$  for  $\sigma \parallel \tau$  (or  $\sigma' \parallel \tau$ ), and all  $k \in \mathbb{N}$ ,  $d(\text{trDist}_k(\sigma), \text{trDist}_k(\sigma')) \leq \varepsilon(\tau, \mathcal{C}, k)$  where  $d(D_1, D_2) = \sup_x |D_1(x) - D_2(x)|$  is the TV-distance on distributions.

We call this  $\varepsilon(\tau, \mathcal{C}, k)$  an *error*, and will write 0 for the error that maps all protocols, contexts, and trace lengths to zero. It is easy to see that this notion of protocol equivalence is an equivalence relation and a congruence.

### 3.2 Reaction sets as protocol agents

We present a translation from reaction sets  $\Sigma$  to protocol agents  $\llbracket \Sigma \rrbracket$ , as described above. As an intermediate step, we first describe a natural semantics for reaction sets as a certain kind of probabilistic I/O automaton [CCK<sup>+</sup>06a, CCK<sup>+</sup>06b, CCK<sup>+</sup>18]. The state space is a function mapping each channel  $c : \tau$  of  $\Sigma$  to a value of type  $1 + \tau$ , *i.e.*, the channel  $c$  has either not been set yet or is set to a value of its assigned type. The start state is when each channel is undefined. Given a state  $s$ , a channel  $c$  defined by the reaction  $x_1 : c_1 : \tau_1, \dots, x_n : c_n : \tau_n \vdash c \ell : \tau \leftarrow D$  induces a distribution on states as follows. If  $c$  is already set in  $s$ , it leaves  $s$  unchanged; if any of the channels  $c_1, \dots, c_n$  are not yet set, it leaves  $s$  unchanged; otherwise we draw  $v$  from the distribution  $D(v_1, \dots, v_n)$ , where  $v_i$  is the value of channel  $c_i$ , and set the value of  $c$  to  $v$ . A *schedule* is a sequence of hidden or output channels and any schedule deterministically induces a distribution on states.

To turn this automaton into a protocol agent, we proceed as follows. The listener input channels are the input channels of the reaction sets. The effect of an input  $v$  arriving on an input channel  $i$  of type  $\tau$  has the effect of setting the value of  $i$  in the state  $s$  to  $v$  if it has not yet been set, and leaving  $s$  unchanged otherwise. For each output channel  $o$ , we have one token input  $get\_o$ , and one output  $o$ . To describe the effect of  $get\_o$  on a state  $s$ , we proceed as follows. Call a schedule *proper* for the pair  $(s, o)$  if it has the following properties:

- Each channel in the schedule is a hidden channel that  $o$  transitively depends on.
- Whenever  $o$  depends on a hidden channel  $h$ , then either  $h$  has already been set in  $s$  or  $h$  is in the schedule.
- Whenever  $h_2$  transitively depends on a hidden channel  $h_1$  and  $h_2$  is in the schedule, then either  $h_1$  is already set in  $o$  or  $h_1$  precedes  $h_2$  in the schedule.

Proper schedules have some important properties:

- If all the visible channels that  $o$  transitively depends on have been set in  $s$ , a proper schedule guarantees that  $o$  will be set after the schedule has been executed.
- Without loss of generality we can assume that every channel  $h$  appears at most once in a proper schedule.
- If  $h_1$  immediately precedes  $h_2$  in a proper schedule, and  $h_2$  does not depend on  $h_1$ , then we can swap the order of  $h_1$  and  $h_2$  in the schedule without altering the resulting distribution on states.
- Any two proper schedules induce the same distribution on states.

To execute  $get\_o$ , we first check whether  $o$  is already set in the state  $s$ ; if so, we return the value of  $o$ . Otherwise we check whether every visible channel that  $o$  transitively depends on has been set in  $s$ ; if not, we leave the state unchanged and yield as there is no way for us to compute  $o$ . Otherwise we select an arbitrary proper schedule for  $(s, o)$ ; if such a schedule does not exist (perhaps due to a circular dependency on channels such as  $h \vdash h$ ) then we again leave the state unchanged and yield. Otherwise we execute this proper schedule followed by  $o$  to induce a distribution on states, and for each state  $s'$  in the support we output the value of the channel  $o$  in  $s'$ .

$$\begin{array}{c}
\frac{\Sigma, \Sigma' \text{ permutations}}{\Sigma \rightsquigarrow \Sigma'} \text{ [PERM-}\Sigma\text{]} \qquad \frac{\Gamma, \Gamma' \text{ permutations}}{\Sigma\{\Gamma \vdash c \ell \leftarrow D\} \rightsquigarrow \Sigma\{\Gamma' \vdash c \ell \leftarrow D\}} \text{ [PERM-}\Gamma\text{]} \\
\\
\frac{\Sigma \rightsquigarrow_{\varepsilon} \Sigma' \text{ and } \Sigma' \rightsquigarrow_{\varepsilon'} \Sigma''}{\Sigma \rightsquigarrow_{\varepsilon+\varepsilon'} \Sigma''} \text{ [TRANS]} \qquad \frac{\Gamma \vDash_{\varepsilon} D = D'}{\Sigma\{\Gamma \vdash c \ell \leftarrow D\} \rightsquigarrow \Sigma\{\Gamma \vdash c \ell \leftarrow D'\}} \text{ [EXT]} \\
\\
\frac{h \notin \Sigma}{\Sigma\{\Gamma \vdash h \text{ hid} \leftarrow D; \Gamma', \Gamma, x : h \vdash c \ell \leftarrow D'\} \rightsquigarrow \Sigma\{\Gamma', \Gamma \vdash c \ell \leftarrow (x \leftarrow D; D')\}} \text{ [FOLD]} \\
\\
\frac{x \notin FV(D')}{\Sigma\{\Gamma, x : c \vdash c' \ell \leftarrow D; \Gamma', y : c' \vdash c'' \ell' \leftarrow D'\} \rightsquigarrow \Sigma\{\Gamma, x : c \vdash c' \ell \leftarrow D; \Gamma', x : c, y : c' \vdash c'' \ell' \leftarrow D'\}} \text{ [DEP-TRANS]} \\
\\
\frac{}{\Sigma\{\Gamma \vdash c \ell := e; \Gamma', \Gamma, x : c \vdash c' \ell' \leftarrow D\} \rightsquigarrow \Sigma\{\Gamma \vdash c \ell := e; \Gamma', \Gamma, x : c \vdash c' \ell' \leftarrow D[x := e]\}} \text{ [SUBST]} \\
\\
\frac{x \notin FV(D')}{\Sigma\{\Gamma \vdash h \text{ hid} \leftarrow D; \Gamma', \Gamma, x : h \vdash c \ell \leftarrow D'\} \rightsquigarrow \Sigma\{\Gamma \vdash h \text{ hid} \leftarrow D; \Gamma', \Gamma \vdash c \ell \leftarrow D'\}} \text{ [HIDDEN RESOURCE]} \\
\\
\frac{h \notin \Sigma}{\Sigma\{\Gamma \vdash h \text{ hid} \leftarrow D\} \rightsquigarrow \Sigma} \text{ [ADD/REM]} \\
\\
\frac{h \notin \Sigma}{\Sigma\{\Gamma \vdash c \ell \leftarrow D\} \rightsquigarrow \Sigma\{\Gamma \vdash h \text{ hid} \leftarrow D; x : h \vdash c \ell := x\}} \text{ [RENAME]} \\
\\
\frac{\Sigma \rightsquigarrow_{\varepsilon} \Sigma' \quad P \text{ compatible with } \Sigma \text{ and } \Sigma'}{\Sigma \cup P \rightsquigarrow_{\varepsilon} (\|P\|, \dots) \Sigma' \cup P} \text{ [CONGR]} \qquad \frac{\Sigma \rightsquigarrow_{\varepsilon} \Sigma}{\text{hide}(\Sigma, c) \rightsquigarrow_{\varepsilon} \text{hide}(\Sigma', c)} \text{ [HIDING]}
\end{array}$$

Figure 5: The IPDL Proof System.

### 3.3 Proof System

We present a proof system for reasoning about reaction sets. The main judgement of the logic is  $\Sigma \rightsquigarrow_{\varepsilon} \Sigma'$ , meaning that the protocol represented by  $\Sigma$  is indistinguishable from  $\Sigma'$ , up to error  $\varepsilon$ . We read  $\Sigma \rightsquigarrow_{\varepsilon} \Sigma'$  as  $\Sigma$  *rewrites to*  $\Sigma'$  (and vice versa) with error  $\varepsilon$ . Recall from Section 3.1 that we measure cryptographic error with a *function*, which takes as input an automata and natural number and outputs a real number in  $[0, 1]$ . The automata we give as an argument to the error function corresponds to the simulator of the approximate rewrite. We say the judgement  $\Sigma \rightsquigarrow_{\varepsilon} \Sigma'$  is *valid* if  $\llbracket \Sigma \rrbracket \rightsquigarrow_{\varepsilon} \llbracket \Sigma' \rrbracket$ , as defined in Section 3.1. When the error is zero, we leave it out of the judgement; i.e.,  $\Sigma \rightsquigarrow \Sigma'$  is an abbreviation for  $\Sigma \rightsquigarrow_0 \Sigma'$ .

The proof system of IPDL is defined by the rules in Figure 5. The rules [PERM- $\Sigma$ ] and [PERM- $\Gamma$ ] state that reaction sets and contexts of individual reactions are equivalent up to reordering, and in particular imply that  $\rightsquigarrow$  is reflexive, and that  $\|$  is symmetric and associative. Rule [EXT] states that we may replace one distribution by another if the underlying equational theory proves them to be the same. The [TRANS] rule states that approximate equivalence is transitive, except that we accumulate error.

The [FOLD] rule states that sampling a channel  $c$  from a monadic bind ( $x \stackrel{\$}{\leftarrow} D; D'$ ) where  $x$  is free in  $D'$  is equivalent to generating a fresh hidden channel  $h$  for  $x$ , sampling  $D$  in this fresh channel, and passing the result to  $c$ . The context for  $h$  may be any subcontext of  $c$  that allows  $D$  to be defined. This rule connects the syntax of the underlying monadic language of distributions to IPDL. In particular, we may use the [FOLD] rule to decompose a channel which carries a product

distribution into two different channels.

The [DEP-TRANS] rule states that dependence between channels is transitive. The rule [SUBST] deals with non-probabilistic assignments: it says that if channel  $c$  is assigned to value  $e$  under context  $\Gamma$ , then any channel with a context that includes  $c$  and its context may propagate this value  $e$  in for  $c$ .

The rule [HIDDEN RESOURCE] states that if a hidden channel  $h$  depends on the context  $\Gamma$ , and  $c$  is a channel with context  $\Gamma'$  such that all of  $\Gamma$  appears in  $\Gamma'$ , then we may add or remove  $h$  from  $\Gamma'$  at will. Rule [ADD/REM] states that unused hidden channels may be removed from the reaction set. Rule [RENAME] allows us rename internal computations using hidden channels. Rule [CONGR] states that  $\rightsquigarrow_\varepsilon$  is a congruence for  $\cup$ ; we appropriately modify the error  $\varepsilon$  to take into account the common context  $P$ . Note that this implicitly allows us to reason about execution time, since the error  $\varepsilon$  may take into account the context  $P$ . Finally, rule [HIDING] allows us to hide output channels of the protocol.

We have proven our logic to be sound:

**Theorem 1.** *If the judgement  $\Sigma \rightsquigarrow_\varepsilon \Sigma'$  is derivable in IPDL, then it is valid; i.e., that  $\llbracket \Sigma \rrbracket \rightsquigarrow_\varepsilon \llbracket \Sigma' \rrbracket$ .*

The proof of Theorem 1 can be found in the appendix, in Section A. We stress that the soundness proof quantifies over adversaries which have the power of general probabilistic (polytime) algorithms, and are not limited to the IPDL language.

## 4 Encoding Protocol Security

The primary function of our logic is to show when two reaction sets,  $\Sigma$  and  $\Sigma'$ , behave (nearly) identically on their external I/O channels. This problem is orthogonal to many of the tasks performed by protocol security frameworks, such as Universal Composability [Can01]. We now detail how protocol security may be encoded in our logic in a way compatible with UC.

An *interface* for a reaction set  $\Sigma$  is a set of I/O channels between  $\Sigma$  and the outside context. Each security protocol  $\pi$  is defined to be a set of *parties*  $\{P_i \mid i \in 1 \dots n\}$  and a set of *trusted functionalities*  $\{F_i \mid i \in 1 \dots k\}$ , both of which are specified as reaction sets. Protocols may be used to express real, runnable code which performs cryptography, or they may be used to express idealized executions which replace cryptography with trusted third parties. Each protocol  $\pi$  has two kinds of I/O interfaces: the first is the *external interface*  $\text{Ext}(\pi)$ , which is used for specifying high-level I/O behaviors of the protocol; the second is the *attacker interface*  $\text{Att}(\pi)$ , used for specifying the *attacks* by which the outside world can influence the protocol. The external interface is supplied by the parties, while the attacker interface can be supplied by both parties and trusted functionalities.

We now specify how these interfaces are defined for protocols. First, we describe in more detail the role of parties and trusted functionalities:

**Parties** Each party  $P_i$  participates in a protocol as an independent agent with possibly distinct roles in the protocol. The I/O channels of each party  $P_i$  are partitioned into the following disjoint sets:

- $\text{Fun}(P_i)$ , the interface between  $P_i$  and the set of trusted functionalities in the protocol;
- $\text{Att}(P_i)$ , the interface between  $P_i$  and the *attacker*, and
- $\text{Ext}(P_i)$ , the set of *external* I/O channels of  $P_i$ .

The *functionality interface*  $\text{Fun}(P_i)$  specifies the channels along which a party may invoke a trusted functionality. These may be used for a variety of tasks, such as sending a message over a network, sampling a key in a trusted manner, invoking an abstract commitment, etc. The *attacker interface*  $\text{Att}(P_i)$  specifies the ways in which an attacker may influence  $P_i$ . (Strictly speaking, the attacker is not separate from the context which influences  $\text{Ext}(P_i)$ . The attacker interface is used to constrain how the simulator is constructed. This is discussed in more detail below.) The way in which the attacker interface is constructed depends on the *corruption model*, which we describe below in Section 4.2. The *external interface* is the intended external I/O behavior of the party in the protocol. (In UC-style protocol composition, the caller and callee of a subprotocol communicate along this external interface.) Note that we *do not* allow parties to communicate directly with one another; they may only indirectly communicate through functionalities which act as a network.

**Functionalities** Each functionality  $F_i$  specifies a trusted third party which is uncorruptable. Functionalities are used for specifying everything that is not a party, including the network by which parties communicate. The I/O channels of each functionality are partitioned into disjoint sets:

- $\text{Att}(F_i)$ , the interface between  $F_i$  and the attacker;
- $\text{Fun}(F_i)$ , the interface between  $F_i$  and other functionalities.
- $\text{Par}(F_i)$ , the interface between  $F_i$  and the parties in the protocol.

Channels in  $\text{Fun}(F_i)$  may only be matched with their dual channel in  $\text{Fun}(F_j)$  for some  $j$ . (Thus, functionalities may not directly communicate with the attacker interface of another functionality.) Channels in  $\text{Par}(F_i)$  may only be matched with channels in  $\text{Fun}(P)$  for some  $P$ . The channels in  $\text{Fun}(P)$ ,  $\text{Fun}(F)$  and  $\text{Par}(F)$  for all  $P$  and  $F$  are called *internal*. A protocol  $\pi$  consisting of parties  $P_i$  and functionality  $F_i$  is *fully specified* if all internal channels are matched with their dual. If this is the case, composing all parties and functionalities together using the hiding composition operator  $\parallel$  will eliminate all internal channels. This operation, called the *execution of  $\pi$* , is denoted by  $\Sigma(\pi)$ . When we do so, we are left with the attacker interface  $\text{Att}(\pi)$  and the external interface  $\text{Ext}(\pi)$  consisting of the union of all attacker and external interfaces in  $\pi$  respectively.

## 4.1 Comparing protocols

Given this setup, we now describe how we may compare two protocols through their external and attacker interfaces. There are two ways to do so, corresponding to *malicious* and *semi-honest* security. Common to both is the notion of a *simulation*: given two protocols  $\pi$  and  $\pi'$  with identical external interfaces, a simulator  $\text{Sim}$  is a reaction set which *converts* the attacker interface of  $\pi'$  to that of  $\pi$  by closing off all channels in  $\text{Att}(\pi')$ , and exposing exactly those channels in  $\text{Att}(\pi)$ . (Thus,  $\text{Sim}$  *supplies* inputs to  $\text{Att}(\pi')$ , while it *consumes* inputs from  $\text{Att}(\pi)$ ; dually,  $\text{Sim}$  consumes outputs from  $\text{Att}(\pi')$ , and supplies outputs to  $\text{Att}(\pi)$ .)

Two protocols  $\pi$  and  $\pi'$  are *comparable* if they have the same number of parties and the  $i$ th party of each protocol has the same external interface. If this is the case, we say that  $\pi$  *reduces from  $\pi'$*  if there exists a simulator  $\text{Sim}$  such that in IPDL,  $\Sigma(\pi) \rightsquigarrow \Sigma(\pi') \parallel \text{Sim}$ , where  $\Sigma(\pi)$ , the execution of  $\pi$ , is equal to the composition of all parties and functionalities in  $\pi$ . (Note that both attacker and external interfaces of  $\pi$  and  $\pi' \parallel \text{Sim}$  are the same.)

## 4.2 Corrupting protocols

When proving security of a protocol  $\pi$  relative to an idealized protocol  $\pi'$ , we do so relative to a corruption model, which specifies the way in which parties can be corrupted by the adversary. We consider two corruption models here: *semi-honest* and *malicious* security. In both cases, given a corruption model  $\mathcal{C} \in \{\text{sh}, \text{mal}\}$  and an index set  $J \subseteq \{1, \dots, n\}$  of parties to corrupt, we construct *derived protocols*  $\pi_{\mathcal{C}}^J$  and  $\pi'_{\mathcal{C}}^J$ , and say that  $\pi$   $\mathcal{C}$ -realizes  $\pi'$  if for all  $J$ ,  $\pi_{\mathcal{C}}^J$  reduces from  $\pi'_{\mathcal{C}}^J$ . (Note that there is a separate proof of reduction for each index set  $J$ , so there exists a simulator  $\text{Sim}^J$  for each  $J$ . Differing choices of  $J$  will give in general very different information to the simulator, and may require differing proof techniques to prove security.)

Below, we detail both corruption models, and show how to derive protocols in these two models.

### 4.2.1 Semi-honest security

In semi-honest security, the adversary has no control over the behavior of the corrupted party, but learns its internal state including any messages sent and received. This is modeled by appropriately instrumenting the code of each party  $P_i, i \in J$  to leak all probabilistic channels as well as any input channels. (A channel is probabilistic when it is not of the form  $\Gamma \vdash c \ell := e$ ; i.e., when it is associated to a nontrivial probability distribution.) Thus, the attacker interface of  $P_i$  grows to include these leaked channels.

### 4.2.2 Malicious Security

In malicious security, the adversary has full control over the party. This is modeled by performing the following modifications to  $\pi$ , for each  $i \in I$ :

- Remove the code for  $P_i$  from  $\pi$ .
- Move all channels in  $\text{Fun}(P_i)$  to the adversary interface on their respective functionality.
- Disregard all channels in  $\text{Ext}(P_i)$  and  $\text{Att}(P_i)$ .

The most surprising modification is perhaps the third one, stating that we do not consider the channels in  $\text{Ext}(P_i)$  and  $\text{Att}(P_i)$  if  $P_i$  is corrupted. The intuition behind this is that we do not need to reason about the context's outputs to the corrupted party along either the external or attacker interface, since the context and corrupted party are one and the same; for a similar reason, we need not reason about the inputs from a corrupted party to the context.

## 4.3 Comparison with UC

Our formalism is restricted compared to full UC; in UC, corruption may happen dynamically, while in our setting we currently reason only about static corruption. The other major difference from UC is how we model the attacker. In UC, the adversary and environment play two differing but similar roles: the adversary provides an attack on a system, while the environment attempts to distinguish two protocols through their external channels.

We choose to model a simpler but equally expressive model, where the “adversary” and “environment” are both played by the context. The simulator, which in full UC is simply another adversary, is instead a “converter” of attacks between one protocol and another. This converter is constrained by the choice of channels which appear in the attacker interfaces of the two protocols. This viewpoint is supported by two theorems in the theory of UC: the first is black-box simulatability [Can01], which states that it suffices to consider UC-simulators which behave the same for all adversaries, and have



only black-box access to the adversary; the second is the dummy adversary theorem [Can01], which states that it suffices to consider a “dummy” UC-adversary which simply forwards all messages to and from the UC-environment. By combining these two theorems, we collapse the “adversary” and “environment” together, and only consider simulators which attach to the attacker interfaces of the two protocols (which formerly would have been meant for the UC-adversary.)

## 5 IPDL Proofs for Running Example

We continue our discussion of the running example from Section 2 in this section. We first study the case when neither party is corrupted. Remember that Figure 1 and 2 describe the real and ideal protocol executions in this case. In the no-corruption case, when the simulator Sim receives message **query** from the protocol  $\pi_I$ , it internally generates a symmetric key  $k$  and sends  $\text{Enc}(k, 0)$  along channel **leak** to  $\pi_R$ . The message  $\text{ans}_R$  is then forwarded to  $\pi_I$  as message  $\text{ans}_I$ . Thus, we get the simulator in Figure 6.

$$\begin{array}{c}
 \text{query inp} \\
 q : \text{query} \vdash \text{keygen} : \text{key} \\
 k : \text{keygen} : \text{key} \vdash \text{leak vis} : \text{ct} := \text{Enc}(0, k) \\
 \text{ans}_R \text{ inp} \\
 \text{ans}_R \vdash \text{ans}_I := \text{ans}_R
 \end{array}$$

Figure 6: Definition of the simulator for the secure channel example when neither party is corrupted.

The security argument is based on semantic security of the symmetric key encryption, as formalized in Figure 3, which says the encryption of a specified message is indistinguishable from encryption of message 0, which means the adversary cannot distinguish whether the ciphertext comes from simulator Sim or the real protocol  $\pi_R$ . The proof follows our description given in Section 2.

In order to get a feel for conducting a proof in IPDL, we will detail the most important step of this proof. Our goal is to obtain a factoring  $\pi_R \rightsquigarrow \pi'_R ||| \text{OSS}_R$ . We first simplify out the  $\text{key}_S$  and  $\text{key}_R$  channels of  $\pi_R$ . Note that both channels simply copy the value of  $\text{keygen}$  onwards, and are only used by the  $\text{send}_R^S$  and  $\text{out}_R$  channels, respectively. We thus use the [DEP-TRANS] rules to move  $\text{keygen}$  into the contexts of both  $\text{send}_R^S$  and  $\text{out}_R$ . Now, we may use the [SUBST] rule to substitute the value of  $\text{key}_S$  into  $\text{send}_R^S$ , and similarly for  $\text{key}_R$  and  $\text{out}_R$ . At this point, we have the following channels (suppressing unused variable names and types):

$$\begin{array}{c}
 k : \text{keygen} \vdash \text{key}_S := k \\
 k : \text{keygen} \vdash \text{key}_R := k \\
 m : \text{in}_S, k : \text{keygen}, \text{key}_S \vdash \text{send}_R^S \leftarrow \text{Enc}(m, k) \\
 c : \text{deliv}, k : \text{keygen}, \text{key}_R \vdash \text{out}_R \text{ vis} := \text{Dec}(c, k).
 \end{array}$$

Because the context of  $\text{key}_S$  is fully contained within  $\text{send}_R^S$  and  $\text{send}_R^S$  does not use the value of  $\text{key}_S$ , we now may remove  $\text{key}_S$  from the context of  $\text{send}_R^S$ . At this point,  $\text{key}_S$  is a hidden channel not used anywhere, so we use the [ADD/REM] rule to remove it from the reaction set. We do the same for  $\text{key}_R$ , removing it from the context of  $\text{in}_S$  and then from the protocol.

At this point, we need to reason about the value of  $\text{Dec}(c, k)$  in  $\text{out}_R$ . We will first rename the channel  $\text{send}_R^S$  to  $\text{out}_{\text{OSS}}$  using [RENAME], in anticipation for later in the proof. We first use the [DEP-TRANS] rule to add  $\text{keygen}$  to the dependencies of  $\text{leak}$ , which we may do since  $\text{leak}$  depends on  $\text{out}_{\text{OSS}}$  which depends on  $\text{keygen}$ . The relevant channels are now below:

$$\begin{aligned} & \cdot \vdash \text{keygen} \leftarrow \text{Rnd}(\mathcal{K}) \\ & m : \text{in}_S, k : \text{keygen} \vdash \text{out}_{\text{OSS}} \leftarrow \text{Enc}(m, k) \\ & c : \text{out}_{\text{OSS}}, k : \text{keygen}, \text{ans}_R \vdash \text{out}_R \text{ vis} := \text{Dec}(c, k) \\ & k : \text{keygen}, c : \text{out}_{\text{OSS}} \vdash \text{leak} \text{ vis} := c \end{aligned}$$

We need to somehow connect the value of  $\text{Enc}(m, k)$  to the value of  $\text{Dec}(c, k)$ . This is typically done using the [FOLD] rule. However, we cannot fold the value of  $\text{out}_{\text{OSS}}$  directly into  $\text{out}_R$ , since the channel  $\text{leak}$  also needs to use this ciphertext. This is complicated by the fact that  $\text{out}_R$  needs to wait on  $\text{ans}_R$  from the adversary, but  $\text{leak}$  needs to happen before. Thus, the first thing we do is separate out the computation of  $\text{Dec}(c, k)$  from the firing of  $\text{out}_R$ . This is done by rewriting  $\text{Dec}(c, k) \equiv (m \leftarrow \text{Dec}(c, k); \text{ret } m)$  (by the basic monadic identities), and using [FOLD] to unfold the computation of  $\text{Dec}(c, k)$ . This gives us a new channel,  $\text{tmp}$ , such that we get the two below reactions:

$$\begin{aligned} & c : \text{out}_{\text{OSS}}, k : \text{keygen} \vdash \text{tmp} := \text{Dec}(c, k) \\ & m : \text{tmp}, \text{ans}_R \vdash \text{out}_R \text{ vis} := m \end{aligned}$$

(We also used the [DEP-TRANS] rule in reverse to remove the dependencies of  $\text{tmp}$  from  $\text{out}_R$ ). At this point, we can see that the dependencies of  $\text{tmp}$  and  $\text{leak}$  are the same, so we will use the [FOLD] rule to combine them together. This will construct a new hidden channel,  $P$ , that outputs the pair  $(c, \text{Dec}(c, k))$  given the ciphertext  $c$  and key  $k$ . By properly removing channels from  $\text{leak}$  and  $\text{tmp}$ , we now get that the only channel which uses  $\text{out}_{\text{OSS}}$  is this new channel  $P$ . Thus, we may fold the  $\text{Enc}(m, k)$  distribution into  $\text{out}_{\text{OSS}}$ . After removing some redundant dependencies using [DEP-TRANS] in reverse, the situation is shown below:

$$\begin{aligned} & \cdot \vdash \text{keygen} \leftarrow \text{Rnd}(\mathcal{K}) \\ & m : \text{in}_S, k : \text{keygen} \\ & \vdash P \leftarrow (c \leftarrow \text{Enc}(m, k); \text{ret } (c, \text{Dec}(c, k))) \\ & p : P \vdash \text{leak} \text{ vis} := p.1 \\ & p : P, \text{ans}_R \vdash \text{out}_R \text{ vis} := p.2 \end{aligned}$$

At this point, we can see that by the correctness of decryption,  $(c \leftarrow \text{Enc}(m, k); \text{ret } (c, \text{Dec}(c, k))) \equiv (c \leftarrow \text{Enc}(m, k); \text{ret } (c, m))$ ; thus we use the [EXT] rule to perform this rewrite. By carrying out the above operations in reverse (i.e., unfolding the monadic bind in  $P$  to obtain  $\text{out}_{\text{OSS}}$ , and substituting in the value of  $P$  into  $\text{leak}$  and  $\text{out}_R$ ), we finally get a simplified protocol  $\pi_R^{\text{simp}}$  containing the channel

$$m : \text{in}_S, \text{ans}_R \vdash \text{out}_R := m$$

as shown in Section 2. The corresponding mechanized proof that  $\pi_R \rightsquigarrow \pi_R^{\text{simp}}$  is shown in the Appendix, in Figure 11.

What the above protocol simplification allows us to do is express the real protocol in a way such that *only the ciphertext* is needed to carry out the protocol, and not the decryption key. In this simplified protocol  $\pi_R^{\text{simp}}$ , the channel  $\text{out}_{\text{OSS}}$  is used to generate the ciphertext. In order to

carry out the security argument, we may now *factor*  $\pi_R^{\text{simp}}$  into the reaction set  $\pi'_R||\text{OSS0}$ . This  $\pi'_R$  is equal to  $\pi_R^{\text{simp}}$ , except that we remove the `keygen` channel and turn the channel `outOSS` into an input, which is output by `OSS0`. This factored protocol  $\pi'_R$  can be thought of as the simulator for the encryption scheme.

At this point, we now use the [CONGR] rule to rewrite to  $\pi'_R||\text{OSS1}$ , where the ciphertext is now equal to  $\text{Enc}(0, k)$ . It is now a straightforward series of substitutions to see that  $\pi'_R||\text{OSS1} \rightsquigarrow \text{Sim}||\pi_I$ . Altogether, by applying [TRANS] we receive that  $\pi_R \rightsquigarrow_{\varepsilon_{\text{OSS}}(\llbracket \pi'_R \rrbracket, \cdot, \cdot)} \pi_I$ , where  $\varepsilon_{\text{OSS}}$  is the error incurred from rewriting from `OSS0` to `OSS1`. We are able to concretely track which simulator we used for the rewrite involving `OSS0` and `OSS1`. By inspecting this simulator  $\pi'_R$ , we see that no inefficient computations are performed; thus, we remain computationally sound.

We also analyze the case when the sender is corrupted. We follow the framework in Section 4 to derive corrupted protocols  $\pi_{R_{\text{mal}}}^{\{S\}}$  and  $\pi_{I_{\text{mal}}}^{\{S\}}$ . In the real protocol, when we corrupt the sender we expose the channels `keyS` and `sendRS` to the sender as an input and output channel respectively. In the ideal protocol, the adversary obtains the channel `sendIS`.

Thus when the sender is corrupted, the simulator needs to provide the real-world adversary with a key by running the key generation algorithm itself. When the real-world adversary outputs a ciphertext on `sendRS`, the simulator decrypts it using this key and sends the decrypted message along `sendIS`. (We may assume wlog that the adversary is constrained to send well-formed ciphertexts.) The simulator then handles `query`, `leak`, `ansR`, and `ansI` as before, except now the value of `leak` is not generated randomly but comes from `sendRS` from the adversary. Since the adversary is the one providing the ciphertext and the simulator knows the key, we do not need to apply the security assumption about the encryption scheme here. The proof is a straightforward sequence of rewriting steps.

## 6 Case Studies

In this section, we describe more case studies.

### 6.1 The ElGamal Cryptosystem

We replicate in IPDL the standard security of the ElGamal cryptosystem [EIG85], which constructs a CPA-secure public key encryption scheme from the Decisional Diffie Hellman (DDH) assumption, which states that two certain distributions, `DDH0` and `DDH1` are computationally indistinguishable. This assumption is interpreted as an axiom  $\Sigma_{\text{DDH0}} \rightsquigarrow_{\varepsilon_{\text{DDH}}} \Sigma_{\text{DDH1}}$  in IPDL. We define the security of the ElGamal scheme similar to Figure 3, comparing a functionality which takes a message as input and encrypts it to one which disregards the message and instead encrypts a dummy message. Since we are proving semantic security, we additionally insert reactions which leak the public key of the encryption scheme to the adversary. Conducting this proof in IPDL requires us to “factor” the real ElGamal execution  $\Sigma_{\text{real}}^{\text{EG}}$  into two reaction sets  $\Sigma_{\text{factor}}||\Sigma_{\text{DDH0}}$ . We then apply the [CONGR] rule to rewrite  $\Sigma_{\text{DDH0}}$  into  $\Sigma_{\text{DDH1}}$ . The security of the cryptosystem follows, since doing this rewrite essentially rerandomizes the ciphertext so that it no longer depends on the message.

### 6.2 MPC in the $\mathcal{F}_{\text{comm}}$ -hybrid Model

We prove secure a two party secure multiparty computation protocol for Rock-Paper-Scissors (RPS) using an idealized commitment scheme. In the ideal protocol, each party sends its input to the ideal functionality who computes the result of the computation and sends it back to the two parties. The real protocol for RPS is implemented using two trusted functionalities  $\mathcal{F}_{\text{comm}}$  implementing

commitment schemes: (1) Both parties send (`commit`, input) to  $\mathcal{F}_{\text{comm}}$ ; (2) When  $\mathcal{F}_{\text{comm}}$  collects (`commit`, input) from both parties, it sends `committed` to both parties and `ready` to adversary; (3) Upon receiving `commit` from  $\mathcal{F}_{\text{comm}}$ , both party send `open` to  $\mathcal{F}_{\text{comm}}$ ; (4) Upon receiving `open` from one party and `ok` from adversary,  $\mathcal{F}_{\text{comm}}$  send one party’s input to the other party. The ideal functionality can be appropriately modified to allow adv to schedule message delivery

We prove security in the malicious corruption model. There are no adversarial channels exposed when neither party is corrupted; thus, in this case, to prove security we let the simulator be empty, and prove directly that the real protocol rewrites under  $\rightsquigarrow$  to the ideal one. Since the only exposed channels are the parties’ external channels, this amounts to proving functional correctness of the real protocol; this proof is a straightforward application of the rewriting rules.

When one of the parties (wlog, party B) is corrupted, the simulator sends a fake commitment message to the adversary, who then responds with their input. The simulator forwards this input to the trusted functionality in the ideal world, who responds back with the result of the computation. Now that the simulator has party B’s input as well as the result of the computation, it can deduce party A’s input and that input to the adversary. Once this simulator is specified in IPDL, it is a straightforward exercise to see that the real world simplifies to the same reaction set as the ideal world with this adversary. We describe more details on the RPS protocol in Appendix B.

### 6.3 Semi-honest Oblivious Transfer Protocol

Our final case study is the classic oblivious transfer (OT) protocol of [GMW87], which constructs an OT protocol from any hard-core predicate associated to a trapdoor permutation family.<sup>1</sup> The protocol consists of two parties, the transmitter and receiver, communicating over an authenticated network. The network is modeled similarly to Section 2, where we instrument the network with explicit `leak` channels. The transmitter takes as input two bits,  $m_0$  and  $m_1$ , while the receiver takes an input an index bit  $i$ . The goal is for the receiver to learn  $m_i$ , while the receiver learns nothing.

We briefly describe the 1-out-of-2 OT protocol with message space  $\{0, 1\}$ , using trapdoor permutation family  $\mathcal{F}$  associated with hardcore predicate  $b(\cdot)$ . On input of two message  $m_0, m_1 \in \{0, 1\}$ , the transmitter samples a random permutation  $(f, f^{-1})$  from family  $\mathcal{F}$ , and sends the permutation description  $f$  to receiver. On input index  $i$  and permutation  $f$  from transmitter, the receiver computes  $y_i = f(r_i)$  and  $y_{1-i} = r_{1-i}$ , where  $(r_i, r_{1-i})$  are sampled randomly from the domain of permutation  $f$ . The receiver then sends  $(y_i, y_{1-i})$  to the transmitter. Next, transmitter first uses trapdoor  $f^{-1}$  to invert  $(y_i, y_{1-i})$  and xor the message  $(m_i, m_{1-i})$  with hardcore predicate  $b(\cdot)$ , i.e.  $c_i = m_i \oplus b(f^{-1}(y_i))$ . Upon message  $(c_i, c_{1-i})$  from transmitter, the receiver can recover message  $m_i$  by computing  $m_i = c_i \oplus b(r_i)$ .

The protocol is secure in the semi-honest corruption model, as described in Section 4. The protocol is modeled similarly to our previous example, where the two ideal parties send their inputs  $(m_0, m_1)$  and  $i$  to a trusted functionality who securely communicates  $m_i$  to the ideal receiver. We focus on the case where the receiver is corrupted. Following Section 4, this is modeled by instrumenting the code of both the real and ideal world receivers in order to leak to the adversary all secret state and input from other parties. The interesting aspect of this case study is its complexity; while this OT protocol has been studied before in formal cryptography frameworks [CCK<sup>+</sup>05], we have found that our abstraction of dependency leads to a very manageable proof. While expressible on paper, this proof is further enhanced by the proof automation capabilities of our mechanization, which we present in Section 7.

<sup>1</sup>An on-paper description of our formal proof can be seen at <https://github.com/ipdl/ipdl>.

## 7 Mechanization

In this section, we describe our mechanization of IPDL as a logic shallowly embedded in Coq, publicly available at <https://github.com/ipdl/ipdl>. All IPDL proofs in this paper have been mechanized. The soundness proof of IPDL using probabilistic automata is not yet mechanized.

Our embedding of IPDL in Coq is *shallow*, in the sense that individual reactions of IPDL are encoded using Coq functions. By doing so, we are able to turn our rewriting rules of Figure 5 into rewriting tactics in Coq. In order to assist proof automation, we also make heavy use of LTAC, Coq’s native proof automation mechanism. We use a shallow embedding in order to encode IPDL reactions into Coq. By doing so, we are able to write reactions that look very close to their on-paper counterpart: for example, the reaction  $c : \text{send}_R^S : \text{ct}, k : \text{key}_S : \text{key}, \text{ans}_R \vdash \text{out}_R \text{ vis} : \text{msg} := \text{Dec}(c, k)$  is encoded in Coq using the concrete syntax

```
[:: ("sendRS", tyCtx); ("keyS", tyKey);
  ("ansR", tyUnit)] ~>
  ("outR", tyMsg) dvis
  (fun ct k _ => dec (c, k)).
```

Each rule in our logic corresponds (roughly) to a single Coq tactic. The main proof goal in Coq is written  $S1 \rightsquigarrow S2$ , corresponding to the  $\rightsquigarrow$  judgment.<sup>2</sup> Most of our tactics work on either side of this proof goal. For example, the tactic `r_move_at pos c1 c2` takes a boolean flag `pos` indicating which reaction set to operate on (either `leftc` or `rightc`, borrowing this design from FCF [PM15]) and two channel names `c1` and `c2`, and swaps the position of `c1` with `c2`. Doing so automatically applies the [PERM- $\Sigma$ ] rule. We have a similar tactic for permuting the elements of a reaction’s context.

While doing these proofs, we found the most often used tactic was `autosubst_at pos c1 c2`, which given a non-probabilistic channel `c1` automatically fills the context of `c2` with the the context of `c1`, substitutes the value of `c1` into `c2`, and removes `c1` from the context of `c2`. This rule is often paired with the tactic `remove_at pos c`, which removes unused hidden actions from the reaction set.

The tactic `r_ext_at pos c D` (corresponding to the rule [EXT]) replaces the value of `c` with the distribution `D`. When this tactic is called, a subgoal is generated that states for all possible interpretations of the context of `c`, the distribution currently in `c` is equal to the distribution `D`. We use a custom library for representing probability distributions. Our library has the feature that representations of distributions are naturally *extensional*, meaning that equality of the distribution coincides with Coq equality. This feature is not essential for the mechanization, but is very important for reducing the complexity of the proof engineering.

Our proof development is overviewed in Figure 7. Only the file `Logic.v` contains trusted definitions; all others are either derived tactics, auxiliary lemmas, or our case studies. The longest proof is of the Oblivious Transfer protocol, at 816 lines total; our other proofs are small in size, at about 300-400 lines each. An example excerpt from our mechanization of the running example from Section 2 is given in the Appendix, in Figure 11. Our mechanized proof sizes are either favorable or on par with relevant work. For protocols, we compare with the recent prior work EasyUC [CSV19] (which uses EasyCrypt as the proof assistant) and the independent work CryptHOL [LSBM19]. EasyUC encoded a secure channel example with more than 10,000 lines of code, while CryptHOL encoded one-time pad with 347 lines of code. For non-interactive primitives, we have 244 lines of

<sup>2</sup>We do not currently reason about approximate error in the mechanization; i.e., our mechanization can be seen as manipulating statements of the form  $\exists \varepsilon, s \rightsquigarrow_\varepsilon s'$ .

code for ElGamal encryption whereas FCF [PM15] is at 249 lines (although FCF does not provide an easy-to-use logic for expressing protocols). Our most complicated example, OT, was not implemented by these other systems — however, we can compare with the *on-paper* proof in the Task-PIOA paper [CCK<sup>+</sup>18] which proved OT in 24 pages (*c.f.* ours is 816 lines of actual code). We conclude that our framework delivers proofs of succinct size<sup>3</sup>.

Section	LoC	Section	LoC
Core Logic	338	RPS	372
Tactics and Lemmas	802	ElGamal	244
Secure Channel	307	OT	816

Figure 7: Lines of code for our IPDL Coq development.

## 8 Related Work

**Symbolic techniques for protocol verification.** A long line of research in the verification of cryptographic protocols has been done in the so-called *Dolev-Yao model* [DY83], where all cryptography is assumed to be perfectly secure. While admitting simple and effective proof methods and supported by multiple mechanized provers ([BSCS18, MSCB13, Cre08]), this technique is limited by requiring complicated *computational soundness* arguments, which prove that a symbolic attacker is equivalent in power to a corresponding computational one which interacts with instantiations of the idealized cryptographic primitives.

A promising direction ([BCL12], [BCL14], [BCEO19]) in this space is initiated by Bana and Comon, where the attacker is not limited by interacting with idealized cryptography, but instead constrained by a number of axioms which state what the attacker is not able to do. By considering the most powerful attacker which does not violate the axioms, it is easy to see such attackers are equivalent to their computational counterparts, if the axioms are sound. The soundness of these axioms are subtle and require careful expert analysis.

In comparison, our work aims to match the *computational reduction* model of reasoning from the ground-up, and therefore match the mathematical foundations of cryptography. While systems based on symbolic systems like Bana and Comon would require an additional soundness proof outside the system, systems based on computational cryptography are intrinsically sound. This difference reflects a tradeoff between the two approaches: while symbolic systems generally support automatic reasoning, computational systems support more general forms of reasoning. Because of this tradeoff, it is likely most profitable to combine aspects of both systems into one, such as partially automating computational verification tools using symbolic techniques [BGJ<sup>+</sup>19]. We additionally point out that a design goal of IPDL is to enable *human-readable* formal proofs (either on paper or in a theorem prover), while symbolic systems generally favor computer-generated proofs.

**Computational mechanized security proofs.** To avoid some of the limitations of the symbolic approach, a separate line of work instead chooses to reason directly in the computational model; *i.e.*, with explicit reasoning principles about computation time, computational error, and probability. CryptoVerif [Bla06] reasons in the computational model about security protocols in a manner

<sup>3</sup>We did not directly compare succinctness with Bana and Comon [BCL12, BCL14, BCEO19], Micciancio and Tessaro [MT13], and other related systems which either do not have an implementation or we did not find an open-source implementation.



similarly to symbolic systems such as ProVerif [BSCS18], thus mainly aims for automated proofs for a well-chosen set of cryptographic primitives.

There are a number of successful verification tools for computational cryptographic proofs, such as EasyCrypt [BGHZ11] and FCF [PM15], which focus on noninteractive primitives and thus do not reason about concurrency explicitly. These tools are capable of nontrivial probabilistic reasoning which is currently not captured directly in IPDL, such as explicit reasoning about the security parameter, and up-to-bad reasoning for bounding the kinds of attacks the adversary may perform. It is interesting future work to combine these advanced forms of cryptographic reasoning with IPDL.

**Frameworks for cryptographic protocols.** In the cryptography literature, Universal Composability [Can01] and Constructive Cryptography [Mau11] are the two dominant definitional frameworks for simulation-based security. Several automata-based frameworks also exist, such as [BPW07] and [CCK<sup>+</sup>18], which, while similar in spirit, aim for a more formal treatment. Additionally, some works use process calculi to model computational cryptographic protocols, such as [MRST01]. Our basic semantic framework is closest to that of UC, since we semantically model a protocol as a configuration of multiple agents, each of which may send at most one message after becoming activated. This requirement that only one message is sent when each protocol agent is activated is naturally captured through a linear typing judgement in ILC [LHM19], a recent effort to formalize the semantics of UC in a process calculus. While ILC does not yet offer formal proofs, it would be a natural next step to formalize the semantics of IPDL using ILC.

An interesting alternative framework is given in Micciancio and Tessaro [MT13] (hereafter M&T), where they use *complete partial orders* to represent cryptographic protocols as the least fixed point of a recursive set of equations. Each communication channel induces a stream of values, which are recursively interrelated through the protocol’s execution. We follow a number of intuitions made by M&T: both IPDL and M&T choose to view communication channels as pure values which may be related through equational rewrites. In both systems, these rewrites are sound under the assumption that the protocols in question do not make use of timing information, but only the values taken on the channels. However, we prove soundness in a more general semantic framework which allows for general adversaries which do make use of timing information. We point out that the monotonicity requirement in M&T (that further inputs can only create more outputs) is naturally captured through our semantic framework, since protocol agents in our framework only operate on single messages at a time.

**Computational proofs of simulation-based security.** Two recent concurrent works, EasyUC [CSV19] and CryptHOL [LSBM19], deliver mechanized computational proofs of simulation-based security of cryptographic protocols. EasyUC does so by directly encoding the coroutine-based semantics of UC in EasyCrypt. As the authors acknowledged in their paper, encoding proofs of cryptographic protocols in EasyUC is highly non-trivial, which is largely due to EasyCrypt not natively supporting reasoning about concurrency. The EasyUC work points out a need for a DSL for specifying and reasoning about cryptographic protocols in EasyCrypt; future work on IPDL could potentially allow it to fill this gap.

CryptHOL is a formalization of the Constructive Cryptography framework [Mau11] in Isabelle. While superficially similar to our work, CryptHOL (and EasyUC) serve a complementary purpose to IPDL: while the former works serve to formalize the semantics of cryptographic protocols in a mechanized theorem prover, IPDL explores the space of appropriate abstractions for conducting cryptographic proofs. We believe that the reaction-based equational logic of IPDL could potentially benefit proofs in EasyUC and CryptHOL, both of which employ handcrafted bisimulations to

conduct equivalence proofs.

## 9 Limitations and Future Work

Given our focus on simplicity and ease-of-use, IPDL currently imposes some restrictions on the cryptographic proofs it can model. We believe, however, that the scope of IPDL nonetheless covers a wide and interesting class of commonly encountered cryptographic proofs.

First, IPDL currently does not explicitly encode and reason about the security parameter. This does not harm any computational reasoning, since proofs in IPDL consist of a finite number of rewrites. Thus, the error incurred in an IPDL proof will be upper bounded by  $n\varepsilon$ , where  $n$  is the size of the proof, and  $\varepsilon$  is the largest error incurred during the proof. By thinking of the security parameter as implicit in the IPDL protocol, we see that the error incurred by any proof is negligible, assuming that  $\varepsilon$  is. However, if extensions to IPDL enable a way to encode loops, then this reasoning breaks down and an explicit security parameter is required.

Second, IPDL currently does not capture protocols where events are triggered based on nontrivial properties about their dependencies. Thus, we do not yet consider protocols which make nontrivial use of control flow, such as certain consensus protocols. We note that this has previously been argued by Micciancio and Tessaro [MT13], who have a similar restriction in their work, to nevertheless represent an expressive class of cryptographic protocols. This restriction also means that for UC-style simulation proofs, we currently support only straightline, blackbox notions of simulation.

Third, it would not be too hard to extend IPDL to support a parametrized number of parties and a parameterized number of executions (as long as the number of parties or sessions are fixed a-priori). We plan to extend IPDL in this direction in the near future. We point out that since IPDL adopt UC-flavor security definitions, just like UC [Can01], the security notion allows us to reason about a single protocol instance while guaranteeing concurrent or sequential composition with polynomially unbounded instances/sessions.

We stress that we do not require that the adversary be an IPDL program — the adversary is assumed to be a probabilistic polynomial-time algorithm. It would be an interesting direction to extend IPDL and allow the syntax to describe a *family* of channels, with a uniform description of their values and dependencies. This would allow us to describe protocols and security games which operate over many queries, up to some (publicly known) parameter  $q$ . Security analysis in this setting would require explicit security parameter analysis, since handling an arbitrary number of queries is similar to handling loops.

We currently capture only static security, in which parties are corrupted a-priori, rather than becoming so in the middle of the protocol. We leave adaptive corruption to later work.

An exciting future direction is to integrate IPDL with an underlying battle-hardened cryptographic proof system (such as EasyCrypt) which may enable more expressiveness, thus achieving ease-of-use and generality simultaneously. Other exciting future directions include to provide a greater degree of proof automation, compiling IPDL programs to executable code (e.g., in C) and proving the correctness of the compilation. As discussed in Section 8, there are a number of directions in which our work can potentially intersect with other concurrent work in formalized cryptography proofs. Integrating our method for constructing bisimulations into EasyUC [CSV19] or CryptHOL [LSBM19] may allow for more accessible proofs for a certain class of protocols. Additionally, it would be interesting to prove our bisimulation methodology sound in ILC [LHM19], a formalization of UC semantics which is likely compatible with our own semantics.

## References

- [BCEO19] Gergei Bana, Rohit Chadha, Ajay Kumar Eeralla, and Mitsuhiro Okada. Verification methods for the computationally complete symbolic attacker based on indistinguishability. *ACM Transactions on Computational Logic (TOCL)*, 21(1):2, 2019.
- [BCL12] Gergei Bana and Hubert Comon-Lundh. Towards unconditional soundness: Computationally complete symbolic attacker. In *International Conference on Principles of Security and Trust*, pages 189–208. Springer, 2012.
- [BCL14] Gergei Bana and Hubert Comon-Lundh. A computationally complete symbolic attacker for equivalence properties. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 609–620. ACM, 2014.
- [BGHZ11] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. Computer-aided security proofs for the working cryptographer. In Phillip Rogaway, editor, *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 71–90, Santa Barbara, CA, USA, August 14–18, 2011. Springer, Heidelberg, Germany.
- [BGJ<sup>+</sup>19] Gilles Barthe, Benjamin Grégoire, Charlie Jacomme, Steve Kremer, and Pierre-Yves Strub. Symbolic methods in computational cryptography proofs. 2019.
- [Bla06] Bruno Blanchet. A computationally sound mechanized prover for security protocols. In *2006 IEEE Symposium on Security and Privacy*, pages 140–154, Berkeley, CA, USA, May 21–24, 2006. IEEE Computer Society Press.
- [BPW07] Michael Backes, Birgit Pfitzmann, and Michael Waidner. The reactive simulatability (rsim) framework for asynchronous systems. *Information and Computation*, 205(12):1685–1720, 2007.
- [BSCS18] Bruno Blanchet, Ben Smyth, Vincent Cheval, and Marc Sylvestre. Proverif 2.00: Automatic cryptographic protocol verifier, user manual and tutorial, 2018.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science*, pages 136–145, Las Vegas, NV, USA, October 14–17, 2001. IEEE Computer Society Press.
- [CCK<sup>+</sup>05] Ran Canetti, Ling Cheung, Dilsun Kaynar, Moses Liskov, Nancy Lynch, Olivier Pereira, and Roberto Segala. Using probabilistic i/o automata to analyze an oblivious transfer protocol. *Cryptology ePrint Archive*, Report 2005/452, 2005. <http://eprint.iacr.org/2005/452>.
- [CCK<sup>+</sup>06a] Ran Canetti, Ling Cheung, Dilsun Kaynar, Moses Liskov, Nancy Lynch, Olivier Pereira, and Roberto Segala. Task-structured probabilistic i/o automata. In *2006 8th International Workshop on Discrete Event Systems*, pages 207–214. IEEE, 2006.
- [CCK<sup>+</sup>06b] Ran Canetti, Ling Cheung, Dilsun Kaynar, Moses Liskov, Nancy Lynch, Olivier Pereira, and Roberto Segala. Time-bounded task-pioas: A framework for analyzing security protocols. In *International Symposium on Distributed Computing*, pages 238–253. Springer, 2006.

- [CCK<sup>+</sup>18] Ran Canetti, Ling Cheung, Dilsun Kaynar, Moses Liskov, Nancy Lynch, Olivier Pereira, and Roberto Segala. Task-structured probabilistic i/o automata. *Journal of Computer and System Sciences*, 94:63–97, 2018.
- [Cre08] Cas JF Cremers. The scyther tool: Verification, falsification, and analysis of security protocols. In *International Conference on Computer Aided Verification*, pages 414–418. Springer, 2008.
- [CSV19] Ran Canetti, Alley Stoughton, and Mayank Varia. Easyuc: Using easycrypt to mechanize proofs of universally composable security. In *32nd IEEE Computer Security Foundations Symposium*, 2019. <https://eprint.iacr.org/2019/582>.
- [DY83] Danny Dolev and Andrew Yao. On the security of public key protocols. *IEEE Transactions on information theory*, 29(2):198–208, 1983.
- [ElG85] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE transactions on information theory*, 31(4):469–472, 1985.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 218–229. ACM, 1987.
- [LHM19] Kevin Liao, Matthew A. Hammer, and Andrew Miller. Ilc: A calculus for composable, computational cryptography. In *PLDI*, 2019. <https://eprint.iacr.org/2019/402>.
- [LSBM19] Andreas Lochbihler, S. Reza Sefidgar, David Basin, and Ueli Maurer. Formalizing constructive cryptography using crypthol. In *32nd IEEE Computer Security Foundations Symposium*, 2019. <http://www.andreas-lochbihler.de/pub/lochbihler2019csf.pdf>.
- [Mau11] Ueli Maurer. Constructive cryptography—a new paradigm for security definitions and proofs. In *Joint Workshop on Theory of Security and Applications*, pages 33–56. Springer, 2011.
- [MRST01] John Mitchell, Ajith Ramanathan, Andre Scedrov, and Vanessa Teague. A probabilistic polynomial-time calculus for analysis of cryptographic protocols:(preliminary report). *Electronic Notes in Theoretical Computer Science*, 45:280–310, 2001.
- [MSCB13] Simon Meier, Benedikt Schmidt, Cas Cremers, and David Basin. The tamarin prover for the symbolic analysis of security protocols. In Natasha Sharygina and Helmut Veith, editors, *Computer Aided Verification*, pages 696–701, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [MT13] Daniele Micciancio and Stefano Tessaro. An equational approach to secure multi-party computation. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 355–372. ACM, 2013.
- [PM15] Adam Petcher and Greg Morrisett. The foundational cryptography framework. In *International Conference on Principles of Security and Trust*, pages 53–72. Springer, 2015.

## A Soundness proof for IPDL

We proceed by case analysis on the proof rules. In each case, we need to provide a bisimulation that relates distributions of states in the first protocol to distribution of states in the second, such that related distributions of states exhibit the same behavior. (This distribution bisimulation is similar to that proposed in [LSBM19].)

We proceed by case analysis on the proof rules:

- Cases [PERM- $\Sigma$ ], [PERM- $\Gamma$ ], [EXT], [TRANS], [SUBST] are clear, since the two protocol agents considered are equal.
- Case [ADD/REM]: The corresponding bisimulation relates a Dirac distribution on state  $s_1$  of the reaction set on the left to a Dirac distribution on state  $s_2$  of the reaction set on the right iff the value of  $h$  in  $s_1$  is undefined and the values of  $s_1$  and  $s_2$  for all remaining channels agree.
- Case [RENAME]: The corresponding bisimulation relates a Dirac distribution on state  $s_1$  of the reaction set on the left to a Dirac distribution on state  $s_2$  of the reaction set on the right iff the value of  $h$  in  $s_2$  is equal to the value of  $c$  in both  $s_1$  and  $s_2$ , and the values of  $s_1$  and  $s_2$  for all remaining channels agree.
- Case [HIDDEN RESOURCE]: The corresponding bisimulation relates a distribution  $\eta$  on states of the reaction set on the left to a Dirac distribution on state  $s$  of the reaction set on the right iff either
  - $\eta$  is a Dirac distribution on the same state  $s$ , or
  - $\eta$  is the result of applying  $h$  to the state  $s$
- Case [FOLD]: The corresponding bisimulation relates a distribution  $\eta_1$  on states of the reaction set on the left to a distribution  $\eta_2$  on states of the reaction set on the right iff either
  - $\eta_1$  is a Dirac distribution on state  $s_1$ ,  $\eta_2$  is a Dirac distribution on state  $s_2$ , the values of  $h$  and  $c$  in  $s_1$  are undefined, the value of  $c$  in  $s_2$  is undefined, and  $s_1$  and  $s_2$  agree on all remaining channels, or
  - there are Dirac distributions on states  $s_1$  and  $s_2$ , respectively, that satisfy the conditions of the previous point, such that  $\eta_1$  is the result of applying  $h$  and then  $c$  to  $s_1$  and  $\eta_2$  is the result of applying  $c$  to  $s_2$ .
- Case [CONGR]: This follows from the observation that the protocol consisting of the single IPDL program  $P \cup Q$  is equivalent to the protocol consisting of the two IPDL programs  $P$  and  $Q$ . The bisimulation for this observation relates a Dirac distribution on state  $s_1$  to the Dirac distribution on the combined state  $(s_2, s_3)$  iff  $s_1$  agrees with  $s_2$  and also with  $s_3$  on any common channels.
- Case [HIDING]: This follows from the observation that given a protocol  $\sigma$  compatible with the protocol agent  $P \text{ hide } c$ , we can extend  $\sigma$  by another protocol agent  $Q$  whose (token) inputs are requests to execute outputs  $o$  of  $P$  (*i.e.*, inputs of the form  $get\_o$ ). If  $o$  does not depend on  $c$ ,  $Q$  forwards this request to  $P$ . Otherwise it first asks  $P$  to execute  $h$  and if it receives a value back from  $P$ , it forwards to  $P$  the request to execute  $o$ . If instead of a value for  $h$  the protocol  $Q$  receives another execution request, *e.g.*, for a channel  $l$ , this means the previous attempt to execute  $h$  was unsuccessful and the effort to execute  $o$  is abandoned.

$\text{comm}_P : \text{play inp}$ $c_P : \text{comm}_P : \text{play} \vdash \text{committed}_P \text{ vis}$ $\text{open}_P \text{ inp}$ $c_P : \text{comm}_P : \text{play}, \text{open}_P \vdash \text{val}_P \text{ vis} : \text{play} := \text{comm}_P$	$\text{in}_P : \text{play inp}$ $i_P : \text{in}_P : \text{play} \vdash \text{comm}_P \text{ vis} : \text{play} := \text{in}_P$ $\text{committed}_{\bar{P}} \text{ inp}$ $c_P : \text{comm}_P : \text{play}, c'_{\bar{P}} : \text{committed}_{\bar{P}} \vdash \text{open}_P \text{ vis}$ $\text{val}_{\bar{P}} : \text{play inp}$ $i_P : \text{in}_P : \text{play}, v_{\bar{P}} : \text{val}_{\bar{P}} : \text{play}$ $\vdash \text{out}_P \text{ vis} : \text{ans} := f_{\text{rps}}(\text{in}_P, \text{val}_{\bar{P}})$
--	---

Figure 8: The real protocol for RPS. Left: the commitment functionalities  $\mathcal{F}_{\text{comm}}$ , for parties  $P \in \{A, B\}$ . Right: the code for party  $P$ . We use the notation  $\bar{P}$  to denote  $B$  if  $P = A$ , and  $A$  if  $P = B$ .

## B More on RPS Protocol

The ideal protocol is implemented in IPDL as having each party  $P \in \{A, B\}$  listen on input channel  $\text{in}_P$ , forward this input to  $\text{send}_P$ , receive a reply on  $\text{recv}_P$ , and forward this value on  $\text{out}_P$ ; correspondingly, the ideal functionality listens on both  $\text{send}_A$  and  $\text{send}_B$ , and sends the value  $f_{\text{rps}}(x, y)$  on both  $\text{recv}_A$  and  $\text{recv}_B$ , where  $x$  and  $y$  are the two parties' inputs. We say the real protocol is secure when it reduces from the ideal protocol, in the terminology of Section 4. In particular, this means input privacy (no party can learn more about another party's input than what the function outputs), and input independence (no party can correlate its input with the input of another party).

The real protocol for RPS is implemented using two trusted functionalities implementing commitment schemes. In UC terminology, this means we are operating in the  $\mathcal{F}_{\text{comm}}$ -hybrid model. The commitment functionality operates by receiving a value  $x \in \{\text{rock}, \text{paper}, \text{scissors}\}$  on input  $\text{comm}_P$ , and sending an indication to the other party  $\bar{P}$  that a value has been committed. Once  $P$  sends  $\text{open}_P$  to the functionality, it then reveals the commitment to the other party via the  $\text{val}_P$  channel.

Informally, the protocol is secure since the commitment functionalities enforce this input independence property described above. The input privacy requirement for this protocol is degenerate, since each party can deduce the other's private input, given its own private input and the output value of the protocol.

There are no adversarial channels exposed when neither party is corrupted; thus, in this case, to prove security we let the simulator be empty, and prove directly that the real protocol rewrites under  $\rightsquigarrow$  to the ideal one. Since the only exposed channels are the parties' external channels, this amounts to proving functional correctness of the real protocol; this proof is a straightforward application of the rewriting rules.

When one of the parties (wlog, party B) is corrupted, the simulator gets access to the real world channels  $\text{comm}_B$  (coming from the adversary);  $\text{committed}_A$  (coming from the simulator);  $\text{open}_B$  (coming from the adversary); and  $\text{val}_A$  (coming from the simulator). In the ideal world, it gets access to the channels  $\text{send}_B$  (coming from the simulator) and  $\text{recv}_B$  (coming from the functionality.) This situation is shown in Figure 9.






Figure 9: Illustration of Corrupted Bob

The simulator operates as follows: it may immediately fire  $\text{committed}_A$ , indicating that the other party has committed its input. On input  $\text{comm}_B$  from the adversary, the simulator forwards this input to  $\text{send}_B$  in the ideal protocol. The simulator receives the value  $\text{open}_B$  from the adversary, but does nothing with it. After receiving  $\text{recv}_B$  from the ideal protocol,  $\text{Sim}$  can compute and send  $\text{val}_A := f_{\text{rps}}^{-1}(\text{recv}_B, \text{send}_B)$ , where the function  $f_{\text{rps}}^{-1}$  computes A's input based on the result  $\text{recv}_B$  and B's input  $\text{send}_B$ .

$\text{comm}_B : \text{play inp}$   
 $c_B : \text{comm}_B : \text{play} \vdash \text{send}_B : \text{play} := \text{comm}_B$   
 $\cdot \vdash \text{committed}_A$   
 $\text{recv}_B : \text{ans inp}, \text{open}_B \text{ inp}$   
 $s_B : \text{send}_B : \text{play}, r_B : \text{recv}_B : \text{ans} \vdash \text{val}_A : \text{play} := f_{\text{rps}}^{-1}(s_B, r_B)$

Figure 10: Formalization of Simulator Sim

Once we describe the simulator, this proof is similar to the honest case and only requires one to compute that the real protocol and the ideal protocol with the simulator simplify to the same reaction set.

```

Theorem noCorr_real_simp : realNoCorr < ~ > realSimp.
  rewrite /realNoCorr /realSimp /rlist_comp_hide; vm_compute RChans; simpl.
  autosubst_at leftc "keyR" "outR".
  remove_at leftc "keyR".
  autosubst_at leftc "keyS" "sendSR".
  remove_at leftc "keyS".

  autosubst_at leftc "deliv" "outR".
  remove_at leftc "deliv".

  rename_at leftc "sendSR" "outCPA".
  arg_move_at leftc "outCPA" "inS" 0.
  trans_at leftc "outCPA" "leak" "keygen" tyKey.

  arg_move_at leftc "outR" "keygen" 0.
  r_ext_at leftc "outR" (rbind [:: ("keygen", tyKey); ("outCPA", tyCtx)]
    [:: ("ans", tyUnit)] ("tmp", tyMsg) ("outR", tyMsg)
    (fun x y => ret (dec y x)) (fun (x : msg) _ _ => ret
      intros; unlock rbind; rewrite //=.
      msimp; done.
    unfold_at leftc "outR".

  trans_rev_at leftc "tmp" "outR" "keygen".
  trans_rev_at leftc "tmp" "outR" "outCPA".
  pair_at leftc "leak" "tmp" "X".
  trans_rev_at leftc "X" "leak" "outCPA".
  trans_rev_at leftc "X" "tmp" "outCPA".
  trans_at leftc "outCPA" "X" "inS" tyMsg.
  r_move_at leftc "outCPA" 0.
  r_move_at leftc "X" 1.
  arg_move_at leftc "X" "outCPA" 0.
  fold_at leftc.
  r_ext_at leftc "X" (fun m k => n2 <- enc m k; ret (n2, m)).
  intros.
  apply mbind_eqP.
  intros; msimp.
  erewrite dec_correct.
  apply: erefl.
  done.
  unfold_bind2_at leftc "X" "outCPA" tyCtx.
  unfold_at leftc "X".
  autosubst_at leftc "X" "leak".
  autosubst_at leftc "X" "tmp".
  remove_at leftc "X".
  autosubst_at leftc "tmp" "outR".
  remove_at leftc "tmp".
  trans_rev_at leftc "outCPA" "leak" "keygen".
  trans_rev_at leftc "outCPA" "leak" "inS".
  hid_str_at leftc "outCPA" "outR".
  hid_str_at leftc "keygen" "outR".
  align.
  reflexivity.
Qed.

```

Figure 11: Example proof from our mechanization, showing that  $\pi_R \rightsquigarrow \pi_R^{\text{simp}}$ , from Section 2. Our library is generic over a type for names of channels (here, strings) and a universe of types (here, including a type for messages, ciphertexts, and keys). Each tactic invocation soundly edits the current goal state, thus enabling the user to reason about protocols through interactive transformations.