

Concrete Security Bounds for Simulation-Based Proofs of Multi-Party Computation Protocols

KRISTINA SOJAKOVA, Vrije Universiteit Amsterdam, Netherlands

MIHAI CODESCU, Research Group of the NLNet Project IPDL, Romania

JOSHUA GANCHER, Northeastern University Boston, United States of America

The *concrete security* paradigm aims to give precise bounds on the probability that an adversary can subvert a cryptographic mechanism. This is in contrast to asymptotic security, where the probability of subversion may be *eventually* small, but large enough in practice to be insecure. Fully satisfactory concrete security bounds for Multi-Party Computation (MPC) protocols are difficult to attain, as they require reasoning about the *running time* of cryptographic adversaries and reductions.

In this paper we close this gap by introducing a new foundational approach that allows us to automatically compute concrete security bounds for MPC protocols. We take inspiration from the meta-theory of IPDL, a prior approach for formally verified distributed cryptography, to support reasoning about the runtime of protocols and adversarial advantage. For practical proof developments, we implement our approach in Maude, an extensible logic for equational rewriting.

We carry out four case studies of concrete security for simulation-based proofs. Most notably, we deliver the first formal verification of the GMW MPC protocol over N parties. To our knowledge, this is the first time that formally verified concrete security bounds are computed for a proof of an MPC protocol in the style of Universal Composability. Our tool provides a layer of abstraction that allows the user to write proofs at a high level, which drastically simplifies the proof size. For comparison, a case study that in prior works required 2019 LoC only takes 567 LoC, thus reducing proof size by 72%.

CCS Concepts: • **Security and privacy** → **Logic and verification**; • **Theory of computation** → *Equational logic and rewriting*.

Additional Key Words and Phrases: multi-party computation, equational reasoning, formal verification

ACM Reference Format:

Kristina Sojakova, Mihai Codescu, and Joshua Gancher. 2025. Concrete Security Bounds for Simulation-Based Proofs of Multi-Party Computation Protocols. 1, 1 (July 2025), 36 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

Advanced distributed cryptographic protocols such as Multi-Party Computation (MPC) have the potential to enable new, privacy-preserving modes of computing. However, their inherent complexity introduces new risks, including the risk that the protocol itself (or protocol optimizations employed by implementations) is insecure. The commonly-used security definition for MPC is simulation-based security in the style of Universal Composability, or UC [Canetti 2000]. UC provides strong security guarantees that are robust under an embedding of the protocol into a larger distributed

Authors' addresses: Kristina Sojakova, k.sojakova@vu.nl, Vrije Universiteit Amsterdam, Netherlands; Mihai Codescu, mscodescu@gmail.com, Research Group of the NLNet Project IPDL, Romania; Joshua Gancher, j.gancher@northeastern.edu, Northeastern University Boston, United States of America.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Association for Computing Machinery.

XXXX-XXXX/2025/7-ART \$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

system. However, UC-style proofs for MPC protocols are generally very difficult, as they require complex bisimulation-based security arguments in conjunction with low-level runtime analysis.

While a number of prior approaches for verifying MPC protocols have been proposed [Almeida et al. 2017; Defrawy and Pereira 2019; Gancher et al. 2023b; Haagh et al. 2018; Stoughton and Varia 2017], most of these approaches deliver standalone security definitions not compatible with UC [Almeida et al. 2017; Defrawy and Pereira 2019; Haagh et al. 2018; Stoughton and Varia 2017], which harms composability. In contrast, the IPDL system [Gancher et al. 2023b] can deliver UC-style security results for MPC protocols via a convenient *equational* style of reasoning. Prior work has shown that IPDL can reason about realistic protocols, including a two-party variant of the GMW [Goldreich et al. 1987] MPC protocol, along with various Oblivious Transfer [Beaver 1995] protocols.

However, a number of security-critical caveats remain. None of the prior approaches for verifying MPC adequately reason about the *runtime* of adversaries and simulators constructed during the proof. Reasoning about runtime is essential for security, since nearly all cryptographic proofs contain *reductions* of the form “the probability that A breaks protocol P is bounded by the probability that the reduced adversary $R(A)$ breaks indistinguishability assumption Q ”. Here $R(A)$ is an adversary whose interaction with Q mimics the interaction of the original adversary A with P . If $R(A)$ ’s runtime is not adequately bounded, then the probability that $R(A)$ breaks the assumption Q may be 1, rendering the security result essentially meaningless. This problem is even more pronounced for large protocols such as MPC, since the constructed simulators also become increasingly complex.

In this work, we address this issue for MPC protocols. Our method supports *concrete security* [Bellare et al. 1997], which gives precise bounds on the probability $\epsilon(t)$ that an attacker running in time t violates the security guarantees of the system. Our novel strategy to obtain these bounds is to bound the size of the IPDL program context, and analyze the runtime of a Turing Machine that interprets IPDL programs. This strategy allows us to obtain practical bounds for MPC: indeed, *to our knowledge, there is no other formally verified proof of an MPC protocol that reasons about concrete security and runtime of adversaries/simulators*.

To enable fast and scalable formal proofs, we implemented our proof system in the equational rewrite tool Maude [Clavel et al. 2007] as an extension of SpeX [Tutu 2022] and equipped it with a Domain-Specific Language (DSL) for concise proofs that automatically compile down to lower-level Maude code. Using our implementation, we carry out four different case studies. To highlight the scalability of our approach, we present a new, fully mechanized proof of *simulation-based security* for the *Multi-Party GMW* MPC protocol defined over an arbitrary Boolean circuit and for arbitrarily many parties. Crucially, our tool *automates* the computation of concrete security bounds, which would otherwise be infeasible to carry out for a protocol of this size.

1.1 Contributions

In this work, we develop:

- a proof system that automatically computes concrete security bounds for composable simulation-based proofs of MPC protocols;
- a formally verified proof of the *GMW* MPC protocol with N parties, which to our knowledge marks the first time that formally verified concrete security bounds have been computed for a UC-style proof of an MPC protocol;
- an accompanying Maude implementation that automatically computes the aforementioned bounds and a DSL for writing proofs that hides most low-level details from the user. This dramatically simplifies the proof size: e.g., excluding definitions, the proof of the *Coin Flip* case study in [Gancher et al. 2023b] takes 1905 LoC, whereas we deliver it in 256 LoC. Our

proofs also significantly outperform their Coq equivalents from [Gancher et al. 2023b] in terms of runtime: e.g., 5 seconds vs. a few minutes for the *Coin Flip* case study.

Limitations of our system. Since we build upon the process calculus in [Gancher et al. 2023b], we are only able to support protocols expressible in IPDL. Specifically, IPDL only considers protocols with static communication topologies and static security. Furthermore, it does not consider protocols that exhibit threshold behavior such as consensus protocols. As IPDL is targeted towards UC security, it also does not consider proofs that use rewinding.

Limitations of our proof effort. Our proof of the GMW protocol assumes that party N is semi-honest and party $N + 1$ is honest. While this is not fully general, all other cases are either trivial (all parties honest/all parties semi-honest) or can be essentially reduced to the aforementioned case. We note that the GMW protocol is not secure against a malicious adversary.

Structure of paper. In Section 3 we briefly review simulation-based security, present IPDL, and illustrate our DSL on a simple example. In Section 4, we give the syntax and semantics of our cost-aware proof system for simulation-based security, and present our soundness results. In Section 5, we outline our proof of the N -Party GMW protocol and briefly describe the other case studies. We conclude by indicating some directions for future work.

2 RELATED WORK

This paper is part of a long line of formal verification efforts for MPC and similar protocols. Some works target domain-specific, standalone security notions for protocols [Almeida et al. 2017; Defrawy and Pereira 2019; Haagh et al. 2018; Stoughton and Varia 2017], while others [Barbosa et al. 2021; Canetti et al. 2019; Gancher et al. 2023b; Lochbihler et al. 2019] aim to construct general frameworks for *simulation-based security* in the style of UC. While runtime analysis is in a sense required for sound cryptographic reasoning, almost all of the above works (with the exception of [Barbosa et al. 2021] and [Gancher et al. 2023b]) declare reasoning about runtime out of scope, and instead defer to the reader to ensure that all relevant cryptographic reductions have a reasonable running time.

The line of work proposed in [Gancher et al. 2023b] shows that an equational proof strategy is useful for proving MPC and related protocols secure. However, in lieu of reasoning about runtime, [Gancher et al. 2023b] relies on symbolic bounds, which provide some measure of complexity for a syntactic simulation context. However, [Gancher et al. 2023b] does not give a low-level computational semantics to protocols, nor does it carry out any cryptographic reductions. It is therefore unclear what such a symbolic bound means for the runtime of adversaries and simulators. In this paper, we build upon the process calculus introduced in [Gancher et al. 2023b] with a new cost-aware proof system and semantics that reason explicitly about the runtime of cryptographic adversaries and simulators. In particular, we prove *concrete* security bounds of the form $\epsilon(t)$, where t is the running time of the adversary.

The system by Barbosa et al. [Barbosa et al. 2021] extends EasyCrypt [Barthe et al. 2011] by a cost-aware Hoare logic and uses it to analyze a secure channel protocol. However, this runtime analysis only applies to sequential programs, while the process calculus of [Gancher et al. 2023b] natively handles concurrency. Thus, while [Barbosa et al. 2021] considers UC, their computational model assumes that the protocol follows a stack discipline, which is not a good fit for MPC.

Squirrel [Baelde et al. 2024] and CryptoVerif [Blanchet 2008] are two popular tools that deliver concrete security bounds for large classes of cryptographic protocols. However, the bulk of the proof for an MPC protocol consists of manipulations such as inlining the computation from one channel to another, or removing parts of the protocol that have become unused after simplification.

The process calculus of [Gancher et al. 2023b] was explicitly designed around such congruences, which do not have counterparts in Squirrel or CryptoVerif. This design choice enabled us to carry out the largest formally verified UC-style proof of an MPC protocol that we are aware of.

Overture [Skalka and Near 2024] is a recently-proposed system for proving security properties of MPC protocols via checking secrecy and integrity hyperproperties. While it offers automatic proofs, it does not prove UC-style properties, nor analyze the runtime of simulators. Owl [Gancher et al. 2023a] uses an information-flow type system to deliver modular proofs of protocols that use cryptographic mechanisms (rather than reason *about* them, which is required for MPC). Resource-aware session types (RAST) [Das et al. 2018] allow one to embed runtime guarantees into session-typed protocols. While RAST and our work both target concurrent process calculi, the protocols we consider do not naturally carry session types, and instead have a low-level computational interpretation in terms of Turing Machines. In particular, it is unclear how to embed MPC into RAST, or use session types in general to perform cryptographic security proofs. GAuV [Xie et al. 2024] uses graph transformations to automate proofs of semi-honest security for concrete instances of the BGW protocol (*i.e.*, for a fixed number of parties and a fixed circuit).

3 OVERVIEW: SIMULATION-BASED SECURITY AND IPDL

We now briefly describe simulation-based security, give an overview of the IPDL process calculus from [Gancher et al. 2023b], and illustrate our approach on a simple running example.

3.1 Simulation-Based Security

Simulation-based security relates the behavior of a protocol to that of an idealization, where cryptographic mechanisms are replaced by trusted resources secure by construction. In this sense, an idealization is a specification for the behavior that the real-world protocol aims to approximate with cryptographic methods. For example, a real-world protocol utilizes encryption to securely send a message from Alice to Bob over a public network. The idealization instead relies on a trusted third party that securely obtains the message from Alice and forwards it to Bob.

The simulation-based paradigm, as employed *e.g.* in UC and Constructive Cryptography [Maurer 2012] is very powerful and unifies various other security notions such as secrecy and integrity. In UC-style security, a protocol is an interactive system consisting of *parties*, *e.g.* Alice and Bob, and *functionalities*, *e.g.* a public network that forwards messages from Alice to Bob, or a key generation mechanism that randomly generates a secret key and securely delivers it to Alice and Bob.

Formal proofs in this setting amount to showing *observational equivalences* $P \approx Q$ between a “real” protocol P , and an “ideal” protocol Q . These proofs typically take the form of a sequence of exact (=) and approximate (\approx) equality steps:

$$P = P_1 \approx Q_1 = P_2 \approx Q_2 = \dots = P_n \approx Q_n = Q$$

As observed in [Gancher et al. 2023b], a typical approximate equality step $Q_i \approx P_i$ consists of replacing the left-hand side of an *indistinguishability assumption* $G_i \approx H_i$ by its right-hand side in a common context R_i ; *i.e.*, P_i arises as $R_i[G_i]$ and Q_i as $R_i[H_i]$. An example of such an assumption is IND-CPA, which states that encryptions of adversarially chosen messages are indistinguishable from encryptions of zeros (in the absence of a decryption oracle).

3.2 Background: IPDL

IPDL [Gancher et al. 2023b] is a process calculus for distributed cryptographic protocols (*e.g.*, MPC) that enables one to prove simulation-based security results similar to those analyzed in UC. IPDL *protocols* are composed of mutually interacting *reactions*, which are sequential monadic programs that probabilistically compute an *expression*. In the context of a protocol, a reaction operates on a

unique *channel* and may read from other channels, thereby utilizing computations coming from other reactions.

Aside from a formal process calculus, the original paper [Gancher et al. 2023b] defines two *equational* proof systems: an *exact* equational logic for proving perfect equivalences between protocols, and an *approximate* logic for proving computational indistinguishability results. The exact equational logic is proven sound in terms of *bisimulations*, while the approximate logic is proven sound in terms of *computational reductions*.

Following [Gancher et al. 2023b], we assume a user-defined *signature* that specifies the base types and the (probabilistic) functions we have at our disposal:

DEFINITION 1 (SIGNATURE). A signature Σ consists of:

- type constants t ,
- function symbols $f : \sigma \rightarrow \tau$, and
- distribution symbols $d : \sigma \rightarrow \tau$.

We summarize the syntax of IPDL in Figure 1. Data types and expressions are standard. Here 1 denotes the unit type and \checkmark the canonical inhabitant of 1 . The expression $\text{app}_{\sigma \rightarrow \tau} f e$ denotes the application of the function symbol $f : \sigma \rightarrow \tau$ declared in the signature Σ to an expression e . Similarly, the reaction $\text{samp}_{\sigma \rightarrow \tau} d e$ denotes the application of the distribution symbol $d : \sigma \rightarrow \tau$ declared in the signature Σ to an expression e .

The reaction $\text{read}(c : \tau)$ denotes the read of a value of type τ from the channel c . We also have branching (if e then R_1 else R_2) and the standard monadic operations of return ($\text{ret } e$) and bind ($x : \sigma \leftarrow R; S$). At the protocol level, we have the trivial protocol 0 , the single-channel protocol $o := R$ that assigns a reaction R to the channel o , the parallel composition $P \parallel Q$ of two protocols, and the spawning new $o : \tau$ in P of a new internal channel o of type τ for use in P .

In our version of the IPDL syntax, references $\text{var}(x : \tau)$ to variables and $\text{read}(c : \tau)$ to channels include a typing annotation. We will need these later on when encoding an IPDL construct as a sequence of symbols on a Turing Machine tape; knowing the type τ will allow us to allocate the correct number of bits for the variable x or the channel c .

Variables	x, y, z	
Channels	i, o, c	
Channel Sets	I, O	$::= \{c_1, \dots, c_n\}$
Data Types	τ, σ	$::= t \mid 1 \mid \text{Bool} \mid \tau_1 \times \tau_2$
Expressions	e	$::= \text{var}(x : \tau) \mid \checkmark \mid \text{true} \mid \text{false} \mid \text{app}_{\sigma \rightarrow \tau} f e \mid (e_1, e_2) \mid \text{fst}_{\sigma \times \tau} e \mid \text{snd}_{\sigma \times \tau} e$
Reactions	R, S	$::= \text{ret } e \mid \text{samp}_{\sigma \rightarrow \tau} d e \mid \text{read}(c : \tau) \mid \text{if } e \text{ then } R_1 \text{ else } R_2 \mid x : \sigma \leftarrow R; S$
Protocols	P, Q	$::= 0 \mid o := R \mid P \parallel Q \mid \text{new } o : \tau \text{ in } P$
Type Contexts	Γ	$::= \cdot \mid \Gamma, x : \tau$
Channel Contexts	Δ	$::= \cdot \mid \Delta, c : \tau$

Fig. 1. Syntax of IPDL.

Typing of protocols in IPDL has the form $\Delta \vdash P : I \rightarrow O$, where Δ is a channel context assigning types to channel names, and I, O are disjoint sets of *input* and *output* channels, respectively. Each output channel in O must be assigned a reaction inside P . The exact equational logic of IPDL is parameterized by a finite set of axioms of the form $\Delta \vdash P_1 = P_2 : I \rightarrow O$, where $\Delta \vdash P_1 : I \rightarrow O$

$$\boxed{\Delta \vdash P = Q : I \rightarrow O}$$

$$\frac{\Delta \vdash P : I \cup O_2 \rightarrow O_1 \quad \Delta, o : \tau \vdash Q : I \cup O_1 \rightarrow O_2 \cup \{o\}}{\Delta \vdash P \parallel (\text{new } o : \tau \text{ in } Q) = \text{new } o : \tau \text{ in } (P \parallel Q) : I \rightarrow O_1 \cup O_2} \text{COMP-NEW}$$

$$\frac{\Delta \vdash P : I \rightarrow O \quad \Delta \vdash Q : I \cup O \rightarrow \emptyset}{\Delta \vdash P \parallel Q = P} \text{ABSORB}$$

$$\frac{\Delta; \cdot \vdash R : I \rightarrow \sigma \quad \Delta; x : \sigma \vdash S : I \rightarrow \tau}{\Delta \vdash (\text{new } c : \sigma \text{ in } o := x \leftarrow \text{read } c; S \parallel c := R) = (o := x \leftarrow R; S)} \text{FOLD-BIND}$$

$$\frac{\Delta; \cdot \vdash (x \leftarrow R_1; y \leftarrow R_1; \text{ret } (x, y)) = (x \leftarrow R_1; \text{ret } (x, x))}{\Delta \vdash (o_1 := R_1 \parallel o_2 := x_1 \leftarrow \text{read } o_1; R_2) = (o_1 := R_1 \parallel o_2 := x_1 \leftarrow R_1; R_2)} \text{SUBST}$$

$$\frac{\Delta; \cdot \vdash R_1 : I \rightarrow \tau_1 \quad \Delta; \cdot \vdash R_2 : I \rightarrow \tau_2 \quad \Delta; \cdot \vdash (x_1 \leftarrow R_1; R_2) = R_2 : I \rightarrow \tau_2}{\Delta \vdash (o_1 := R_1 \parallel o_2 := x_1 \leftarrow \text{read } o_1; R_2) = (o_1 := R_1 \parallel o_2 := R_2)} \text{DROP}$$

Fig. 2. Selected rules for exact equality of IPDL protocols.

and $\Delta \vdash P_2 : I \rightarrow O$. We use these axioms to express example-specific functional assumptions, e.g., the correctness of an encryption/decryption scheme. Figure 2 shows a few illustrative rules of the exact fragment of IPDL.

Rule COMP-NEW allows us to pull a sub-protocol P outside the scope of a channel declaration if the bound channel name does not appear in P . Rule ABSORB allows us to discard a sub-protocol that has become unused after simplification. Rule SUBST says that if the computation R_1 assigned to a channel o_1 is deterministic, then we may replace every occurrence of $\text{read } o_1$ by R_1 . Rule FOLD-BIND states that if we only read from channel c once in the context of a protocol, we can soundly replace $\text{read } c$ by the computation assigned to c , even if this computation is probabilistic. Finally, rule DROP allows us to drop a vacuous dependency on channel o_1 from channel o_2 if the computation R_1 assigned to o_1 does not introduce additional dependencies to o_2 .

IPDL protocols come with a natural operational semantics. We slightly generalize the semantics given in [Gancher et al. 2023b] to support dynamic-length bitstrings. To this end, we use a special placeholder symbol \bullet and by abuse of terminology we refer to strings $v \in \{0, 1, \bullet\}^*$ as bitstrings.

DEFINITION 2 (INTERPRETATION). An interpretation $\llbracket - \rrbracket$ for a signature Σ associates to:

- each type symbol t a subset $\subseteq \{0, 1, \bullet\}^{|t|}$ of bitstrings of length $|t| \geq 0$;
- each function symbol $f : \sigma \rightarrow \tau$ a function $\llbracket f \rrbracket$ from $\llbracket \sigma \rrbracket$ to $\llbracket \tau \rrbracket$;
- each distribution symbol $d : \sigma \rightarrow \tau$ a function $\llbracket d \rrbracket$ from $\llbracket \sigma \rrbracket$ to distributions on $\llbracket \tau \rrbracket$.

In the above, we generalize the interpretation $\llbracket - \rrbracket$ to all types in the obvious way. To handle partial computations, we follow [Gancher et al. 2023b] and augment the syntax of IPDL protocols to contain intermediate bitstring values

$$\text{Protocols } P, Q ::= o := v \mid \dots$$

We give semantics to IPDL protocols via two main small-step rules, see Figure 3, where we write $1[P]$ for the distribution with unit mass at the protocol P , and freely use a distribution in place of a

reaction or a protocol to indicate the obvious lifting of the corresponding construct to distributions on protocols. As reactions are sequential monadic programs, they admit a straightforward small-step semantics $R \rightarrow \eta$, which we omit. Big-step operational semantics for protocols $P \Downarrow \eta$ performs output and internal steps in an arbitrary order until no more steps are possible, resulting in a unique distribution η on protocols.

$$\begin{array}{c}
 \boxed{P \xrightarrow{o := v} Q} \\
 \\
 \frac{P \xrightarrow{o := v} P'}{P \parallel Q \xrightarrow{o := v} P' \parallel Q[\text{read } o := \text{val } v]} \qquad \frac{Q \xrightarrow{o := v} Q'}{P \parallel Q \xrightarrow{o := v} P[\text{read } o := \text{val } v] \parallel Q'} \\
 \\
 \frac{}{(o := \text{val } v) \xrightarrow{o := v} (o := v)} \qquad \frac{P \xrightarrow{o := v} P' \quad o \neq c}{(\text{new } c : \tau \text{ in } P) \xrightarrow{o := v} (\text{new } c : \tau \text{ in } P')} \\
 \\
 \boxed{P \rightarrow \eta} \\
 \\
 \frac{R \rightarrow \eta}{(o := R) \rightarrow (o := \eta)} \qquad \frac{P \rightarrow \eta}{P \parallel Q \rightarrow \eta \parallel Q} \qquad \frac{Q \rightarrow \eta}{P \parallel Q \rightarrow P \parallel \eta} \\
 \\
 \frac{P \rightarrow \eta}{(\text{new } c : \tau \text{ in } P) \rightarrow (\text{new } c : \tau \text{ in } \eta)} \qquad \frac{P \xrightarrow{c := v} P'}{(\text{new } c : \tau \text{ in } P) \rightarrow 1[\text{new } c : \tau \text{ in } P']}
 \end{array}$$

Fig. 3. Small-step operational semantics for IPDL protocols.

3.3 Example: Authenticated-To-Secure Channel

We now revisit the running example of [Gancher et al. 2023b]. Alice wants to securely communicate n messages to Bob using an authenticated channel which leaks all messages to the adversary. To do so, we will assume they share a pre-shared key, which enables them to encrypt and decrypt all messages. We show how to encode this protocol and its proof in our DSL. First, we declare the number of sessions as a parameter n to our case study, with the intended interpretation that n is a function of the security parameter λ :

parameter $n : \text{nat}$.

3.3.1 The Assumptions. We use a version of the IND-CPA assumption, which states that encoding n context-chosen messages with the same secret key is indistinguishable from encrypting zeros. For simplicity, we assume a type of messages with constant length, so that the constant zeros need not depend on the length. We express the assumption as an *axiom* about the equality between two protocols:

approx-assumption CPA :

```

(fam In[i < n] :: msg)
(fam Enc[i < n] :: ctxt)

```

```

inputs: fam In[i < n] |=
new Key : key in
  (Key ::= samp gen_key ||
   (family Enc[i < n] ::=
     m : msg <- read In[i] ;
     k : key <- read Key ;
     samp enc((m, k))))
~
new Key : key in
  (Key ::= samp gen_key ||
   (family Enc[i < n] ::=
     m : msg <- read In[i] ;
     k : key <- read Key ;
     samp enc((zeros, k)))) .

```

The CPA axiom is parameterized by two *families* of channels: $\text{In}[i < n]$, for n input channels carrying messages, and $\text{Enc}[i < n]$, for n output channels carrying ciphertexts. The left side of the CPA assumption samples a key on channel Key and, for each i , encrypts $\text{In}[i]$ under the key. We do this by declaring a **family** of protocols — one for each i less than the parameter n — which reads from $\text{In}[i]$, reads from Key , and samples from the distribution of probabilistic encryptions under that key and message. We declare the Key channel as *internal* using **new** so that the outside context cannot read it. The right side of the CPA assumption is similar, but encrypts zeros (defined to be a constant) rather than the message. Importantly, the family $\text{Enc}[i < n]$ on the right hand side still reads from $\text{In}[i]$, since the two protocols must have the same (logical) *timing* behaviors between the channels.

In addition to the CPA assumption, we also have the assumption about the encryption scheme's *correctness*: encrypting and decrypting must return the same message. We encode this in our DSL as an assumption similar to CPA, but since we assume that the encryption scheme is *perfectly* correct, with zero probability of error, we use the declaration **protocol-assumption** rather than **approx-assumption**:

```

protocol-assumption enc-dec-correctness :
  (chn In :: msg) (chn Key :: key)
  (chn Enc :: ctxt) (chn Dec :: msg)
inputs: chn In, chn Key |=
  (Enc ::= m : msg <- read In;
   k : key <- read Key;
   samp enc((m, k))) ||
  (Dec ::= c : ctxt <- read Enc;
   k : key <- read Key;
   return dec((c, k)))
=
  (Enc ::= m : msg <- read In;
   k : key <- read Key;
   samp enc((m, k))) ||
  (Dec ::= i : msg <- read In;

```


return i).

3.3.2 The Protocol. The Authenticated-To-Secure Channel protocol *Real* now takes the following form:

```

protocol Real =
  new Key : key in
  newfamily Recv[i < n] : ctxt in
  newfamily Send[i < n] : ctxt in
  (Keygen || Alice || Channel || Bob)
  where Alice =
    (family Send[i < n] ::=
      m : msg <- read In[i] ;
      k : key <- read Key ;
      samp enc((m, k)))
  and Bob =
    (family Out[i < n] ::=
      c : ctxt <- read Recv[i] ;
      k : key <- read Key ;
      return dec((c, k)))
  and Channel =
    (family Leak[i < n] ::= read Send[i]) ||
    (family Recv[i < n] ::=
      c : ctxt <- read Send[i] ;
      ok : unit <- read Ok[i] ;
      return c)
  and Keygen =
    (Key ::= samp gen_key) .

```

The body of the protocol is a parallel composition of the key generating functionality *Keygen*, the two parties Alice and Bob, and the authenticated channel functionality *Channel*. Alice encrypts each input with the shared key stored on the internal channel *Key*, samples a ciphertext from the resulting distribution, and sends the result to the authenticated channel functionality on the channel *Send[i]*. Bob reads the ciphertext forwarded to him from the authenticated channel functionality on the channel *Recv[i]*, decrypts it with the shared key, and outputs the plaintext on the channel *Out[i]*. The channels *Send[i]* and *Recv[i]* are connected by the *Channel* functionality, which allows the adversary to read/schedule messages via *Leak[i]* and *Ok[i]*.

Proving Protocols Secure. Our protocol is named *Real* because we will compare it to an *ideal* version, where Alice and Bob communicate directly through a secure channel, without the need for encryption. As in UC, we do this by proving that the *Real* protocol is indistinguishable from the *Ideal* protocol composed with a *simulator* *Sim* that can emulate the *Leak[i]* messages without knowledge of the secret messages:

$$\text{Real} = \text{Ideal} \parallel \text{Sim}.$$

The key advantage of IPDL [Gancher et al. 2023b] is that it enables proofs through *equational reasoning* principles. Using basic identities of protocols (e.g., inlining definitions of channels into

other channels) and assumptions (e.g., CPA), one proves a protocol secure by progressively rewriting the real protocol into its idealization (plus the simulator).

We illustrate some key steps of the proof. The ideal functionality for our secure channel example has two output channels per session i : the adversarial output channel `LeakMsgRcvdIdAdv[i]` reads the message from `In[i]` and lets the adversary know that a message has been received – by passing a term of the unit type – but divulges nothing about the value of the message. The channel `Out[i]` first waits to receive a confirmation on the adversarial input channel `OkMsgAdvId[i]` that gives the green light to the functionality to process the message. It subsequently reads the message from `In[i]` on behalf of Alice, and outputs it on behalf of Bob:

```

Ideal =
  (family LeakMsgRcvdIdAdv[i < n] ::=
    m : msg <- read In[i];
    return ()) ||
  (family Out[i < n] ::=
    okMsg : unit <- read OkMsgAdvId[i];
    m : msg <- read In[i];
    return m)

```

The simulator turns the adversarial interface of the real protocol into the adversarial interface of the ideal functionality, thereby converting any adversary for the real protocol into an adversary for the functionality. In our example, the channels `LeakMsgRcvdIdAdv[i]` and `OkCtxtAdvNet[i]` are the inputs to the simulator, while the channels `LeakCtxtNetAdv[i]` and `OkMsgAdvId[i]` are the outputs.

Hence, upon receiving the information from the ideal functionality that a message has been received, the simulator must conjure up a ciphertext to leak to the adversary. This is accomplished by randomly generating a secret key and encrypting the chosen message zeros in each session. Upon receiving the approval from the adversary for the generated ciphertext, the simulator gives the approval to the ideal functionality to output the message:

```

Sim =
  new Key : key in
    (Key ::= samp gen_key(())) ||
    (family LeakCtxtNetAdv[i < n] ::=
      x : unit <- read LeakMsgRcvdIdAdv[i];
      k : key <- read Key;
      samp enc((zeros()), k))) ||
    (family OkMsgAdvId[i < n] ::=
      okCtxt : unit <- read OkCtxtAdvNet[i];
      return okCtxt)

```

As the channel `OkMsgAdvId[i]` carries a deterministic computation, we want to inline the computation into `Out[i]` to yield the following:

```

family Out[i] i < n ::=
  okCtxt : unit <- read OkCtxtAdvNet[i];
  m : msg <- read In[i];
  return m

```

This form of substitution is justified by the rule **SUBST** from Figure 2, which says that if the computation R_1 assigned to a channel o_1 is deterministic, then we may replace every occurrence of $\text{read } o_1$ by R_1 . However, for the rule **SUBST** to apply, we must first massage the protocol into a form where $\text{OkMsgAdvId}[i]$ appears immediately next to $\text{Out}[i]$. In the implementation of [Gancher et al. 2023b], this required tedious manual transformations that, *e.g.*, permute channels inside a parallel composition or move channels in and out of scope of a new channel declaration (lines 87–104 of `MultiChan.v`). In our tool, all the necessary massaging is performed automatically, and we can simply write:

```
subst fam OkMsgAdvId into fam Out then
subst fam LeakMsgRcvdIdAdv into fam LeakCtxtNetAdv then
absorb fam LeakMsgRcvdIdAdv then
absorb fam OkMsgAdvId
```

A crucial step in simplifying the real protocol is to conceptually separate the encryption and decryption actions from the message-passing by introducing new internal channels $\text{Enc}[i]$ and $\text{Dec}[i]$ along with their definitions:

```
add internal family Enc  $i < n$  typed: ctxt
  assigned:  $m : \text{msg} \leftarrow \text{read In}[i];$ 
              $k : \text{key} \leftarrow \text{read Key};$ 
             samp  $\text{enc}((m, k))$  then
add internal family Dec  $i < n$  typed: msg
  assigned:  $c : \text{ctxt} \leftarrow \text{read Enc}[i];$ 
              $k : \text{key} \leftarrow \text{read Key};$ 
             return  $\text{dec}((c, k))$ 
```

We can now modify the channels $\text{Send}[i]$ and $\text{Out}[i]$ to read from $\text{Enc}[i]$ and $\text{Dec}[i]$ directly:

```
sym from change fam Send with
   $e : \text{ctxt} \leftarrow \text{read Enc}[i];$ 
  return  $e$ 
in currentProtocol(
  subst fam Enc into fam Send ) then
sym from change fam Out with
   $\text{okCtxt} : \text{unit} \leftarrow \text{read OkCtxtAdvNet}[i];$ 
   $d : \text{msg} \leftarrow \text{read Dec}[i];$ 
  return  $d$ 
in currentProtocol(
  subst fam Dec into fam Out )
```

We can now invoke the correctness assumption in each individual session to cancel the effect of encryption followed by decryption. In [Gancher et al. 2023b], the generalization to n sessions was stated as a separate lemma with a nontrivial manual proof (lines 67 – 103 in `CPA.v`). In our code, the corresponding proof looks like this:

```
by induction on  $i$  with variable  $x$  (
  use assumption enc-dec-correctness
  on chn Dec[ $x$ ], chn Enc[ $x$ ] )
```

The above code snippet proves by induction on $i < n$ that if the channels $\text{Dec}(i)$ with $i < x$ rewrite from the original formulation that performs the decryption to the new formulation that simply reads off the message $\text{In}[i]$, so does the channel $\text{Dec}[x]$.

On the other hand, our CPA assumption is applied just once across all sessions upon encountering the following protocol snippet:

```
new Key : key in
  ((Key ::= samp gen_key()) ||
   (family Enc[i] i < n ::=
     m : msg <- read In[i];
     k : key <- read Key;
     samp enc((m, k))))
```

We invoke the approximate CPA assumption as shown below:

```
use approx assumption cpa
```

This yields the following protocol snippet:

```
new Key : key in
  ((Key ::= samp gen_key()) ||
   (family Enc[i] i < n ::=
     m : msg <- read In[i];
     k : key <- read Key;
     samp enc((zeros()), k))))
```

A final step in the proof *folds* the internal channels $\text{Enc}[i]$ and $\text{Dec}[i]$ that we introduced earlier into the rest of the protocol:

```
fold fam Enc into fam LeakCtxtNetAdv then
fold fam Dec into fam Out
```

This is justified by the rule FOLD-BIND in Figure 2, which states that if we only read from channel c once, we can soundly replace $\text{read } c$ by the computation assigned to c , even if this computation is probabilistic.

Concrete Security Bounds. Such a proof could be carried out in the IPDL logic alone. However, that proof would only guarantee asymptotic security with respect to the custom symbolic bounds defined in [Gancher et al. 2023b]. In this work, we aim for *concrete* security with respect to a low-level Turing Machine semantics: a precise probabilistic bound for the difference in the probability that an adversary can distinguish Real from Ideal $\parallel \text{Sim}$.

After encoding the proof in the DSL, our tool computes the following bounds:

```
indistinguishability assumption cpa :
count: 1
context: n * |msg| * 6 + n * |ctxt| * 3 + n * 96 + 12
```

Here **count** denotes the number of times the CPA assumption was applied, and **context** bounds the maximal size of the program context in which it was applied. The expressions $|msg|$ and $|ctxt|$ denote the concrete length of bitstrings needed to represent the two types.

Given these two quantities, our main theorem allows us to derive the following concrete security bound on the distinguishing advantage for our protocol:

$$\left| \Pr[\text{Adv} \rightleftharpoons \text{Real} = 1] - \Pr[\text{Adv} \rightleftharpoons \text{Ideal} \mid \text{Sim} = 1] \right| \leq \varepsilon_{\text{cpa}}.$$

Here, $\text{Adv} \rightleftharpoons P$ denotes the interaction of the adversary Adv with a protocol P . The value ε_{cpa} is the maximal distinguishing advantage for the CPA assumption against any adversary with computational “cost” at most $\mathcal{P}(\text{context})$. Here \mathcal{P} is a fixed polynomial, and “cost” bounds the runtime of the adversary and the number of states in its Turing Machines (among other quantities). In other words, our theorem exactly bounds the security error in our protocol’s proof by the error present in the IND-CPA game against the reduction.

Our novel strategy for computing concrete security bounds is to implement an *interpreter* for IPDL programs as a Turing Machine, and compute its cost as the aforementioned polynomial \mathcal{P} . The polynomial takes the size of the interpreted program as its argument, and bounds the number of Turing Machine states and transitions that the resulting interpretation will need.

3.4 The Main Theorem

Before giving technical details, we now discuss our main result informally. Roughly speaking, if P is approximately equal to Q , then the advantage that an adversary has in distinguishing P and Q is a reasonable combination of the distinguishing advantages against each indistinguishability assumptions by an adversary whose computational resources are only slightly larger than those of the original adversary.

THEOREM (SOUNDNESS OF APPROXIMATE EQUALITY OF PROTOCOLS, INFORMAL). *There exists a polynomial $\mathcal{P}(x, y, z)$ with the following property. Given:*

- *finitely many built-in functions that can be computed with cost at most $C_{\text{sem}} \in \mathbb{N}$ and can be approximated by probabilistic Turing Machines with error at most $\eta_{\text{sem}} \in \mathbb{Q}_{\geq 0}$;*
- *indistinguishability assumptions $\vdash P^1 \approx Q^1, \dots, \vdash P^n \approx Q^n$;*
- *a proof of indistinguishability $\vdash P \approx Q$ with output bounds **count** _{i} and **context** _{i} for the i -th indistinguishability assumption;*
- *an adversary Adv for P/Q that computes with cost at most $C_{\text{adv}} \in \mathbb{N}$;*
- *axiom bounds $\varepsilon^1, \dots, \varepsilon^n \in \mathbb{Q}_{\geq 0}$ with the property that for any adversary Adv^i for P^i/Q^i such that Adv^i computes with cost at most $\mathcal{P}(C_{\text{sem}}, C_{\text{adv}}, \text{context}_i)$, we have*

$$\left| \Pr[\text{Adv}^i \stackrel{\llbracket - \rrbracket}{\rightleftharpoons} P^i = 1] - \Pr[\text{Adv}^i \stackrel{\llbracket - \rrbracket}{\rightleftharpoons} Q^i = 1] \right| \leq \varepsilon^i,$$

we have

$$\left| \Pr[\text{Adv} \stackrel{\llbracket - \rrbracket}{\rightleftharpoons} P = 1] - \Pr[\text{Adv} \stackrel{\llbracket - \rrbracket}{\rightleftharpoons} Q = 1] \right| \leq \sum_{i=1}^n \text{count}_i * \varepsilon^i.$$

In other words, we can bound the probability in distinguishing P from Q by the sum of the maximal probabilities of violating an indistinguishability axiom. The actual theorem we prove (Thm. 1) is slightly more general, as we allow the distribution symbols to be general distributions that are only *approximated* by Turing Machines, which requires an *error term* to appear in the theorem.

The probability ε^i must work for *every* adversary with cost at most $\mathcal{P}(C_{\text{sem}}, C_{\text{adv}}, \text{context}_i)$; thus, the larger the adversary cost, the looser the bound ε^i must be. Indeed, if $\mathcal{P}(C_{\text{sem}}, C_{\text{adv}}, \text{context}_i) = \text{context}_i^{500}$, the bound is still polynomial but not particularly meaningful. Crucially, our proof is *constructive*: we are able to *compute* the polynomial as

$$\mathcal{P}(x, y, z) = y^2 + 8yz + 15z^2 + (|\Sigma_f| + |\Sigma_d| + 1)x + 34y + 47z + O(1).$$

This polynomial, where $|\Sigma_f|$ and $|\Sigma_d|$ are the number of function and distribution symbols in our signature, serves as a precise bound on the reduction overhead incurred when using a cryptographic assumption. We are able to achieve such a concrete bound precisely by using a low-level computational semantics for adversaries as Turing Machines. The exact polynomial, which we include in the appendix, is computed once and for all, and counts the precise number of steps that our TMs take, along with binding the TM's number of states, tapes, and symbols (to encode a protocol on a tape, we use additional symbols besides 0, 1). This is in contrast to almost all prior cryptographic work targeting MPC, which either reasons about runtime informally (without being able to compute the polynomial bound on reduction overhead) or subverts reasoning about runtime via error terms that contain reductions themselves [Brzuska et al. 2018]. A notable exception is Barbosa et al. [Barbosa et al. 2021], which adds a Hoare logic for running time to EasyCrypt, at the cost of being limited to imperative programs (and thus imperative encodings of protocols), and manual proof effort for each runtime bound.

3.4.1 Overview of Results. In this work, we build a new cost-aware proof system on top of the existing exact equational logic of IPDL. The rest of our framework diverges significantly. In particular:

- We carry out explicit cryptographic reductions instead of using symbolic adversaries. We represent an adversary (Section 4.3) as a tuple of (essentially arbitrary) probabilistic algorithms for scheduling interactions, updating the adversary's internal state, querying the protocol for output channel values, and assigning new input channel values. When absorbing a program context Q into the adversary, we explicitly extend the adversary's state by the encoding of Q as a sequence of symbols on the Turing Machine tape.
- We deliver *concrete security bounds* (Section 4.1 and 4.2) instead of symbolic ones. In particular, the bound induced by invoking an approximate congruence rule, which allows us to conclude $P \parallel Q \approx P' \parallel Q$ from $P \approx P'$, is the length $\|Q\|$ of the aforementioned Turing Machine encoding (plus some overhead).
- We give a natural definition of computational indistinguishability (Sec. 4.4) that does not involve reasoning about syntactic contexts (as in prior work [Gancher et al. 2023b]). Informally, we define two families of protocols to be indistinguishable if for any polynomial $p(\lambda)$ and negligible function $\eta(\lambda)$, there exists a negligible function $\varepsilon(\lambda)$ such that for any sufficiently large λ and any adversary Adv with cost bounded by $p(\lambda)$, the distinguishing advantage of Adv is bounded by $\varepsilon(\lambda)$.
- We carry out an explicit analysis of errors induced by a probabilistic Turing Machine that ends up in a non-accepting state with a negligible probability. This probability makes an appearance in the concrete bounds we derive (Thm. 1).

4 COST-AWARE SYNTAX AND SEMANTICS FOR IPDL

In this section, we extend the IPDL logic [Gancher et al. 2023b] to handle *cost-aware* proofs; that is, proofs which guarantee precise concrete security bounds. Our main theorem for concrete security bounds assumes a sound ambient theory for the strict fragment of IPDL.

While the exact fragment of IPDL is exactly what is desired for cryptographic proofs, the approximate fragment is missing a crucial point of reasoning. In particular, its soundness proof is in terms of *symbolic* bounds, which abstract away the underlying cost semantics of IPDL protocols; in short, it does not reason about runtime. Because of this, the prior approximate logic of IPDL does not prove the same class of security results generally accepted by the cryptographic community.

For the rest of this section, we assume a fixed signature Σ with type constants $t_1, \dots, t_{|\Sigma_t|}$.

4.1 Approximate Congruence

Our cost-aware equational theory consists of three layers. Firstly, we have *approximate congruence*, see Figure 4, which applies a program context to a single *approximate axiom*, resulting in the judgment $\Delta \vdash P \cong Q : I \rightarrow O \text{ ctxt } \psi$ for axiom # k . We assume a set of n approximate axioms of the form $\Delta^k \vdash P^k \approx Q^k : I^k \rightarrow O^k$ for $1 \leq k \leq n$, where $\Delta^k \vdash P^k : I^k \rightarrow O^k$ and $\Delta^k \vdash Q^k : I^k \rightarrow O^k$. These axioms capture cryptographic assumptions on computational indistinguishability.

Here, the ctxt parameter tracks the increase in the adversary's resources incurred by the proof. A typical proof step in the exact fragment transforms the protocol into a form where an approximate axiom applies. We subsequently carry out an approximate congruence step, where we use the approximate axiom to replace a small protocol fragment nested inside an arbitrary program context by its computationally indistinguishable counterpart.

The program context is formally a part of the adversary, and as such it must be resource-bounded for the indistinguishability assumption to apply. Some nesting patterns do not effect any change on the adversary's resources: for example, a simple renaming of channels (rule *EMBED*); the formal addition of an unused channel i to the protocol's inputs I (rule *INPUT-UNUSED*), in which case any value assigned by the adversary to channel i will leave the protocol unchanged; or the introduction of an internal channel $o : \tau$ (rule *CONG-NEW*), in which case the adversary will never query o because internal channels are only visible in the scope of their declaration.

On the other hand, composing two approximately equal protocols $P \approx P'$ with another protocol Q requires the adversary to simulate the interaction of the program context Q with P versus P' . In other words, the adversary *absorbs* Q and the protocol becomes part of the new adversary's code. In particular, the number of symbols needed for encoding the adversary's code on a Turing Machine tape increases, and the parameter ψ approximates this increase. As rule *CONG-COMP* shows, composition with protocol Q incurs $\|Q\| + 3$ additional symbols: $\|Q\|$ symbols for encoding Q ; a parallel composition symbol to combine the original code with the code for Q ; and two parenthesis symbols “(”, “)” for enclosing the composition. We emphasize that the exact numbers here are not crucial; what matters is that we eventually deliver a (reasonable) polynomial in λ .

We show how to compute the Turing Machine bound of an IPDL construct in Section A. This bound, and consequently the ctxt parameter ψ , is not a natural number but a function $\psi(t_1, \dots, t_{|\Sigma_I|}) : \mathbb{N}^{|\Sigma_I|} \rightarrow \mathbb{N}$ that is *monotonically increasing in each argument*. When encoding a protocol Q as a sequence of symbols on a Turing Machine tape, we invariably encounter variables x of type τ . At this point, we do not know how many bits we will need to encode values of type τ , because the type constants $t \in \Sigma$ are yet uninterpreted. Instead, we leave the size of each type constant as a variable to the function ψ , which will later be instantiated by the appropriate natural number according to $\|-\|$.

4.2 Approximate and Asymptotic Equality

In the *approximate equality* of protocols, see Figure 5, we chain together a sequence of strict equalities and approximate congruence transformations to obtain the judgment $\Delta \vdash P \approx Q : I \rightarrow O \text{ count } \xi \text{ ctxt } \psi$. The parameter ξ counts the number of axiom invocations for each of the n axioms. One application of the k -th approximate axiom incurs a count that maps k to 1 and all other axioms to 0. The use of transitivity requires us to add up the respective values of ξ per each axiom (rule *TRANS*). Even though each individual axiom invocation introduces a negligible error, summing up exponentially many negligible errors might not be negligible, which is why we need to keep track of the number of times each axiom is applied. The parameter ψ tracks the maximum size of a program context in which each axiom is applied.

$$\boxed{\Delta \vdash P \cong Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}$$

$$\frac{}{\Delta^k \vdash P^k \cong Q^k : I^k \rightarrow O^k \text{ ctxt } 0 \text{ for axiom \# } k} \text{AXIOM}$$

$$\frac{c \notin I \cup O \quad \Delta \vdash P \cong Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash P \cong Q : I \cup \{c\} \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{INPUT-UNUSED}$$

$$\frac{\phi : \Delta_1 \rightarrow \Delta_2 \quad \Delta_2 \vdash P \cong Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta_1 \vdash \phi^*(P) \cong \phi^*(Q) : \phi^*(I) \rightarrow \phi^*(O) \text{ ctxt } \psi \text{ for axiom \# } k} \text{EMBED}$$

$$\frac{\Delta \vdash P \cong P' : I \cup O_2 \rightarrow O_1 \text{ ctxt } \psi \text{ for axiom \# } k \quad \Delta \vdash Q : I \cup O_1 \rightarrow O_2}{\Delta \vdash P \parallel Q \cong P' \parallel Q : I \rightarrow O_1 \cup O_2 \text{ ctxt } (\psi + \|Q\| + 3) \text{ for axiom \# } k} \text{CONG-COMP}$$

$$\frac{\Delta, o : \tau \vdash P \cong P' : I \rightarrow O \cup \{o\} \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash (\text{new } o : \tau \text{ in } P) \cong (\text{new } o : \tau \text{ in } P') : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{CONG-NEW}$$

Fig. 4. Approximate congruence of IPDL protocols.

$$\boxed{\Delta \vdash P \approx Q : I \rightarrow O \text{ count } \xi \text{ ctxt } \psi}$$

$$\frac{\Delta \vdash P = Q : I \rightarrow O}{\Delta \vdash P \approx Q : I \rightarrow O \text{ count } (i \mapsto 0) \text{ ctxt } (i \mapsto 0)} \text{STRICT}$$

$$\frac{\Delta \vdash P \cong Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash P \approx Q : I \rightarrow O \text{ count } (k \mapsto 1, i \neq k \mapsto 0) \text{ ctxt } (k \mapsto \psi, i \neq k \mapsto 0)} \text{APPROX-CONG}$$

$$\frac{\Delta \vdash P_1 \approx P_2 : I \rightarrow O \text{ count } \xi \text{ ctxt } \psi}{\Delta \vdash P_2 \approx P_1 : I \rightarrow O \text{ count } \xi \text{ ctxt } \psi} \text{SYM}$$

$$\frac{\Delta \vdash P_1 \approx P_2 : I \rightarrow O \text{ count } \xi_1 \text{ ctxt } \psi_1 \quad \Delta \vdash P_2 \approx P_3 : I \rightarrow O \text{ count } \xi_2 \text{ ctxt } \psi_2}{\Delta \vdash P_1 \approx P_3 : I \rightarrow O \text{ count } (i \mapsto \xi_1(i) + \xi_2(i)) \text{ ctxt } (i \mapsto \max(\psi_1(i), \psi_2(i)))} \text{TRANS}$$

Fig. 5. Approximate equality for IPDL protocols.

Finally, analogously to [Gancher et al. 2023b], we define the *asymptotic equality* of two protocol families, see Figure 6, as functions of the security parameter $\lambda \in \mathbb{N}$. Informally speaking, if two protocol families are asymptotically equal, then any *resource-bounded* adversary cannot distinguish them with greater than negligible error. Formally, we assume a finite set of *approximate axiom families* of the form $\{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$, where $\{P_\lambda\}, \{Q_\lambda\}$ are two protocol families with pointwise-identical typing judgments. If the axiom families comprising our asymptotic theory are clear from the context, we will omit them from the asymptotic equality judgment.

$$\boxed{\{\Delta_\lambda^1 \vdash P_\lambda^1 \approx Q_\lambda^1 : I_\lambda^1 \rightarrow O_\lambda^1\}_\lambda, \dots, \{\Delta_\lambda^n \vdash P_\lambda^n \approx Q_\lambda^n : I_\lambda^n \rightarrow O_\lambda^n\}_\lambda \vdash \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_\lambda}$$

$$\frac{\begin{array}{c} \forall \lambda, \Delta_\lambda^i \vdash P_\lambda^i \approx Q_\lambda^i : I_\lambda^i \rightarrow O_\lambda^i, i = 1, \dots, n \vdash \Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda \text{ count } \xi_\lambda \text{ ctxt } \psi_\lambda \\ \forall i, \xi_{(\cdot)}^i = O(\text{poly}(\lambda)) \quad \forall i, \psi_{(\cdot)}^i = O(\text{poly}(\lambda, t_1, \dots, t_{|\Sigma_t|})) \end{array}}{\{\Delta_\lambda^1 \vdash P_\lambda^1 \approx Q_\lambda^1 : I_\lambda^1 \rightarrow O_\lambda^1\}_\lambda, \dots, \{\Delta_\lambda^n \vdash P_\lambda^n \approx Q_\lambda^n : I_\lambda^n \rightarrow O_\lambda^n\}_\lambda \vdash \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_\lambda}$$

Fig. 6. Asymptotic equality for IPDL protocol families.

For any fixed λ , we obtain an approximate theory by selecting from each axiom family the particular axiom corresponding to λ . Similarly, from each of the two protocol families we select the protocol corresponding to λ , which gives us two concrete protocols to equate approximately. We recall that an approximate equality judgment is tagged by a pair of parameters ξ and ψ , and for each axiom $1 \leq i \leq n$ we have $\xi^i \in \mathbb{N}$ and $\psi^i : \mathbb{N}^{|\Sigma_t|} \rightarrow \mathbb{N}$, where $|\Sigma_t|$ is the number of type constants declared in our ambient signature Σ . Fixing the axiom i and letting the security parameter λ vary thus gives us two functions $\mathbb{N} \rightarrow \mathbb{N}$ and $\mathbb{N}^{|\Sigma_t|+1} \rightarrow \mathbb{N}$, and we require that these be bounded by polynomials in the appropriate number of variables. Unlike [Gancher et al. 2023b], we do not impose a bound on the channel context Δ , which can be particularly burdensome to check in a formal tool. Instead, our definition of an adversary ensures that it only interacts with the protocol via a bounded number of channels.

4.3 Adversaries for IPDL Protocols

To seamlessly account for the possible renaming of channel names, an adversary for protocols of type $\Delta \vdash I \rightarrow O$ is allowed to operate in a larger context Δ' that subsumes the original context Δ via an *embedding* $\phi : \Delta' \rightarrow \Delta$. Here ϕ is an injective, type-preserving mapping that specifies how to rename channels in Δ to fit in the larger context Δ' . In this larger context, we specify the adversarial input channels I' that the adversary will query for a value, and the adversarial output channels O' that the adversary will assign values to. The adversarial inputs I' will be a subset of the protocol outputs O , appropriately translated along ϕ . Dually, the protocol inputs I , appropriately translated along ϕ , will be a subset of the adversarial outputs O' . In the interaction between the adversary and the protocol, every query for a value of a channel $o \in I'$ will extract the value of the channel $o \in \phi^*(O)$ as computed by the the protocol, and pass it on to the adversary. Conversely, an input on channel $i \in \phi^*(I)$ to the protocol occurs after the adversary computes the value of the channel $i \in O'$.

Since our adversaries will be resource-bounded, we need to bind the number of interactions or *rounds* between the adversary and the protocol. In each round, the adversary examines its internal state to determine the type of interaction to perform next, and steps to a new state. This transition function is a partial probabilistic function of type $\text{St} \rightarrow (\{\perp\} \cup I' \cup O') \times \text{St}$. That is, for any internal state s the adversary probabilistically decides among: 1) no interaction, coupled with stepping to a new state s' ; 2) querying a channel $o \in I'$, coupled with stepping to a new state s' ; 3) an assignment to a channel $i \in O'$, coupled with stepping to a new state s' ; or 4) halting, in which case the game between the adversary and the protocol ends without a decision Boolean. We use this last option to capture probabilistic computations that only succeed up to a negligible error.

If the adversary queries channel $o \in I'$ and receives a value v as a response to the query, it updates its internal state according to an input assignment function of type $\llbracket \tau \rrbracket \times \text{St} \rightarrow \text{St}$, where τ is the type of the channel o in Δ' . That is, for any value $v \in \llbracket \tau \rrbracket$ and any state s , the adversary

steps to a new state s' that records the value v as a result of the query. If the adversary chooses a value assignment to a channel $i \in O'$, the value v – if any – is determined by an output valuation function of type $\text{St} \rightarrow \llbracket \tau \rrbracket \cup \{\perp\}$, where τ is the type of the channel i in Δ' . After completing the designated number of rounds, the adversary converts its internal state to a final decision Boolean according to a function of type $\text{St} \rightarrow \{0, 1\}$.

To bind the complexity of the aforementioned operations, we implement them as Turing Machines¹. For convenience, we allow TMs with multiple tapes. As is standard, in the initial configuration all tapes except the first are fully blank. The internal state of the adversary is typically encoded as a bitstring, containing e.g., register values together with the sequence of instructions to be executed, if we wish to view the adversary as an essentially arbitrary probabilistic program. For our purposes, it is convenient to allow additional symbols besides 0, 1 on our TM tapes: when justifying the COMP-CONG rule of our theory, we suitably compose the adversary with the common context Q . The protocol Q thus becomes integrated into the new adversary's code. Instead of encoding protocols as bitstrings, we will suitably enrich our baseline set of symbols so that we can faithfully capture IPDL code.

DEFINITION 3 (ADVERSARY). Fix an interpretation $\llbracket - \rrbracket$ for Σ . An adversary for protocols $\Delta \vdash I \rightarrow O$ is a tuple $(\Delta', I', O', \phi, \#_{\text{rnd}}, \#_{\text{tape}}, \text{Symb}, \text{St}, s_\star, T, \{I_o\}_{o \in I'}, \{O_i\}_{i \in O'}, D)$, where

- Δ' is a channel context;
- $I' \subseteq \Delta'$ is a set of channels that the adversary can query for a value;
- $O' \subseteq \Delta'$ is a set of channels to which the adversary can assign a value;
- $\phi : \Delta' \rightarrow \Delta$ is an embedding of Δ into Δ' ;
- $\#_{\text{rnd}} \geq 1$ is the number of rounds the adversary will perform;
- $\#_{\text{tape}} \geq 1$ is the number of TM tapes at our disposal;
- Symb is a finite set of additional symbols that will be used to encode the adversary's internal state;
- $\text{St} \subseteq (\{0, 1\} \sqcup \text{Symb})^l$ is a set of strings of a fixed length $l \geq 1$ consisting of symbols drawn from the disjoint union of the sets $\{0, 1\}$ and Symb ;
- $s_\star \in \text{St}$ is the initial state;
- T is a probabilistic TM that computes a partial function $\text{St} \rightarrow (\{\perp\} \cup I' \cup O') \times \text{St}$, with $\#_{\text{tape}}$ -many tapes using symbols from the set $\{0, 1\} \sqcup \{\perp\} \sqcup (I' \cup O') \sqcup \text{Symb}$,
- I_o where $o : \tau \in \Delta$ is a TM that computes a function $\text{St} \times \llbracket \tau \rrbracket \rightarrow \text{St}$, with $\#_{\text{tape}}$ -many tapes using symbols from the set $\{0, 1\} \sqcup \text{Symb}$,
- O_i where $i : \tau \in \Delta$ is a TM that computes a function $\text{St} \rightarrow \llbracket \tau \rrbracket \cup \{\perp\}$, with $\#_{\text{tape}}$ -many tapes using symbols from the set $\{0, 1\} \sqcup \{\perp\} \sqcup \text{Symb}$,
- D is a TM that computes a function $\text{St} \rightarrow \{0, 1\}$, with $\#_{\text{tape}}$ -many tapes using symbols from the set $\{0, 1\} \sqcup \text{Symb}$.

We furthermore require that $I' \subseteq \phi^\star(O)$, $\phi^\star(I) \subseteq O'$, and $\phi^\star(O) \cap O' = \emptyset$.

The probabilistic TM T that computes the transition function can terminate in a non-accepting state with probability > 0 . We will be interested in families of adversaries where this error as a function of the security parameter is negligible.

DEFINITION 4 (ADVERSARIAL ERROR). We say that an adversary Adv has error up to $\varepsilon \in \mathbb{Q}_{\geq 0}$, written $\text{err}(\text{Adv}) \leq \varepsilon$, if for any state $s \in \text{St}$ the transition function $T(s)$ is undefined with probability

¹When using Turing Machines to compute (probabilistic) functions, we only consider (probabilistic) Turing Machines that have a finite runtime $N \in \mathbb{N}$; i.e., for every input in the domain, after N (probabilistic) transitions the TM ends up in a configuration where no further transitions are possible. That is, the TM has either reached an accepting state or it has halted after reading a symbol for which no transition is possible in the current state.

$\leq \epsilon$. In other words, when T is run with the initial tape contents s , it halts in a non-accepting state with probability $\leq \epsilon$.

To ensure that the adversary does not have access to computationally expensive functions such as the discrete logarithm, we need to impose a bound on its computational resources. We will be interested in families of adversaries where the bound as the function of the security parameter is polynomial.

DEFINITION 5 (RESOURCE-BOUNDED ADVERSARIES). We say that an adversary Adv has cost at most $K \in \mathbb{N}$, written $|\text{Adv}| \leq K$, if:

- $\#_{\text{rnd}}, \#_{\text{tape}} \leq K$;
- $|I'| \leq K$ and for each $o \in I'$ with $o : \tau \in \Delta'$, we have $|\tau| \leq K$,
- $|O'| \leq K$ and for each $i \in O'$ with $i : \tau \in \Delta'$, we have $|\tau| \leq K$,
- $|\text{Symb}| \leq K$;
- the length k of a state $s \in \text{St}$ is $\leq K$;
- the number of states of each² TM T, I_o, O_i, D is $\leq K$;
- the runtime of each TM T, I_o, O_i, D is $\leq K$.

DEFINITION 6 (INTERACTION). Fix an interpretation $\llbracket - \rrbracket$ for Σ . Let Adv be an adversary for protocols $\Delta \vdash I \rightarrow O$ and let $\Delta \vdash P : I \rightarrow O$. We define $\text{Adv} \xrightarrow{\llbracket - \rrbracket} P$ to be the probability sub-distribution on Booleans induced by the algorithm in Figure 7.

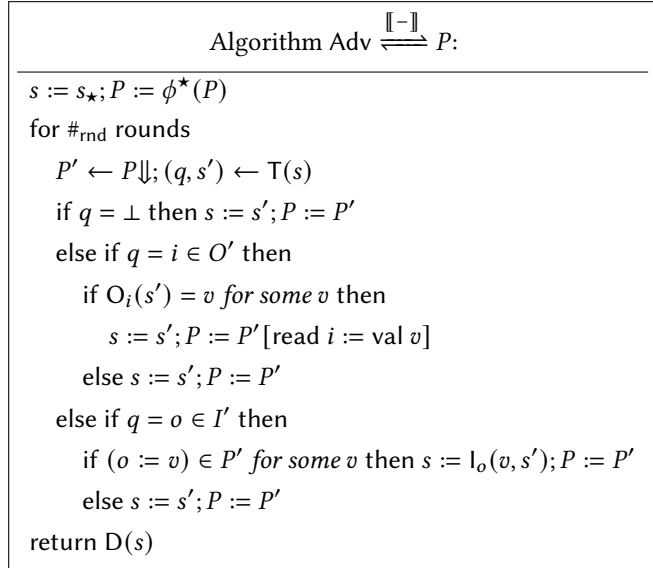


Fig. 7. Interaction of an an adversary Adv with an IPDL protocol P .

In Figure 7, the adversary interacts with the protocol through the specified number of rounds. The algorithm maintains a state variable s along with a protocol variable P , which we respectively

²Instead of having a separate Turing Machine I_o for each channel $o \in I'$ we could have required a single Turing Machine that performs the computation across all channels in I' , and analogously for O_i . However, this is unnecessary as the number of channels in both I' and O' is $O(\text{poly}(\lambda))$, and the current formulation is more convenient for our purposes.

initialize to the initial state s_* and the original protocol P , appropriately embedded in Δ' . In each round, the protocol P probabilistically evolves to a new protocol P' . Independently, the adversary probabilistically computes the type of interaction $q \in \{\perp\} \cup I' \cup O'$ together with a new state s' according to $T(s)$. If $q = \perp$, in which case no interaction takes place, the state and the protocol are updated to s' and P' . If $q = i$ for some $i \in O'$, we compute $O_i(s')$ to see if in the adversary's current state s' the channel i carries a value v . If this is the case, we update the state to s' while computing a new protocol $P'[\text{read } i := \text{val } v]$. Otherwise we update the state and the protocol to s' and P' . Finally, if $q = o$ for some $o \in I'$, we examine the protocol P' to see if the output channel o carries a value v . If this is the case, we compute a new adversary state $I_o(v, s')$, while updating the protocol to P' . Otherwise we update the state and the protocol to s' and P' . After completing the prescribed number of rounds, we obtain a decision Boolean³ $D(s)$ based on the adversary's current state.

4.4 Computational Indistinguishability

A family of interpretations is PPT (probabilistic polynomial-time) if it assigns polynomial lengths to type symbols t , and PPT-computable functions to function symbols f and distribution symbols d . A small caveat is that a random distribution on a subset $S \subseteq \{0, 1\}^n$ of bitstrings is in general computable by a probabilistic Turing Machine only up to a small error ϵ , which is the probability that the TM does not end up in an accepting state. In effect, the TM computes a distribution μ on $S \cup \{\perp\}$ with $\mu(\perp) = \epsilon$. To relate μ to our original distribution on S , we introduce the following:

DEFINITION 7 (APPROXIMATING DISTRIBUTIONS). *Let $S \subseteq \{0, 1\}^n$ be a subset of bitstrings of a fixed length. We say that a distribution μ_1 on $S \cup \{\perp\}$ approximates a distribution μ_2 on S with error $0 \leq \epsilon \leq 1$ if there are distributions η_1, η_2 on S such that $\mu_1 = (1 - \epsilon)\eta_1 + \epsilon 1[\perp]$ and $\mu_2 = (1 - \epsilon)\eta_1 + \epsilon\eta_2$.*

A function $\epsilon : \mathbb{N} \rightarrow \mathbb{Q}_{\geq 0}$ is negligible if it is eventually smaller than the inverse of any polynomial: for any $n \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that for all $\lambda \geq N$ we have $\epsilon(\lambda) \leq \frac{1}{\lambda^n}$.

DEFINITION 8 (PPT FAMILY OF INTERPRETATIONS). *We say that a family $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$ of interpretations for Σ is PPT if there is a polynomial $p(\lambda)$, a negligible function $\eta(\lambda)$, and a natural number $N \in \mathbb{N}$ such that the following holds:*

- For all type symbols t , $|t|_\lambda \leq p(\lambda)$ if $\lambda \geq N$.
- For all function symbols $f : \sigma \rightarrow \tau$, the function $\llbracket f \rrbracket_\lambda$ from bitstrings $\llbracket \sigma \rrbracket_\lambda$ to bitstrings $\llbracket \tau \rrbracket_\lambda$ is computable by a deterministic Turing Machine TM_λ with symbols 0, 1. Both the number of states and the runtime of TM_λ are $\leq p(\lambda)$ if $\lambda \geq N$.
- For all distribution symbols $d : \sigma \rightarrow \tau$, the function $\llbracket d \rrbracket_\lambda$ from bitstrings $\llbracket \sigma \rrbracket_\lambda$ to distributions on bitstrings $\llbracket \tau \rrbracket_\lambda$ is computable up to an error $\eta(\lambda)$ by a probabilistic Turing Machine TM_λ with symbols 0, 1. Specifically, for every $v \in \llbracket \sigma \rrbracket_\lambda$, the distribution $\text{TM}_\lambda(v)$ on $\llbracket \tau \rrbracket_\lambda \cup \{\perp\}$ approximates $\llbracket d \rrbracket_\lambda(v)$ with error $\leq \eta(\lambda)$. Both the number of states and the runtime of TM_λ are $\leq p(\lambda)$ if $\lambda \geq N$.

We are now ready to give the definition of computational indistinguishability. We will be only interested in indistinguishability with respect to a PPT family of interpretations.

DEFINITION 9 (COMPUTATIONAL INDISTINGUISHABILITY). *Consider a family $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$ of interpretations for Σ . Let $\{\Delta_\lambda \vdash P_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\Delta_\lambda \vdash Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$ be two protocol families with identical typing judgments. We say that $\{P_\lambda\}$ and $\{Q_\lambda\}$ are indistinguishable under $\{\llbracket - \rrbracket_\lambda\}$,*

³Strictly speaking, the interaction $\text{Adv} \stackrel{\llbracket - \rrbracket}{\longleftrightarrow} P$ is only a sub-distribution on Booleans, since $T(s)$ may halt without a result. As the probability of this happening will be negligible, we gloss over this technical point here.

written

$$\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}} \models \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$$

if for any polynomial $p(\lambda)$ and negligible function $\eta(\lambda)$, there exists a negligible function $\varepsilon(\lambda)$ and an $N \in \mathbb{N}$ such that for any $\lambda \geq N$ and any adversary Adv for protocols $\Delta_\lambda \vdash I_\lambda \rightarrow O_\lambda$ with respect to the interpretation $\llbracket - \rrbracket_\lambda$ such that $|\text{Adv}| \leq p(\lambda)$ and $\text{err}(\text{Adv}) \leq \eta(\lambda)$, we have⁴

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket_\lambda} P_\lambda = 1] - \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket_\lambda} Q_\lambda = 1] \right| \leq \varepsilon(\lambda).$$

If the family of interpretations is clear from the context, we may omit it from the computational indistinguishability judgment.

4.5 Concrete and Asymptotic Security

We can now present our main result in full:

THEOREM 1 (SOUNDNESS OF APPROXIMATE EQUALITY OF PROTOCOLS). *There exists a polynomial $\mathcal{P}(x, y, z)$ such that for any*

- *interpretation $\llbracket - \rrbracket$ for Σ for which there are $C_{\text{sem}} \in \mathbb{N}$ and $\eta_{\text{sem}} \in \mathbb{Q}_{\geq 0}$ such that*
 - *for all type symbols t , $|t| \leq C_{\text{sem}}$;*
 - *for all function symbols f , $\llbracket f \rrbracket$ is computable by a TM with symbols $0, 1$ such that the number of states and the runtime are $\leq C_{\text{sem}}$; and*
 - *for all distribution symbols d , $\llbracket d \rrbracket_\lambda$ is computable up to an error η_{sem} by a probabilistic TM with symbols $0, 1$ such that the number of states and the runtime are $\leq C_{\text{sem}}$;*
- *approximate axioms $\Delta^1 \vdash P^1 \approx Q^1 : I^1 \rightarrow O^1, \dots, \Delta^n \vdash P^n \approx Q^n : I^n \rightarrow O^n$;*
- *derivation $\Delta \vdash P \approx Q : I \rightarrow O$ count ξ ctxt ψ ;*
- *adversary Adv for protocols of type $\Delta \vdash I \rightarrow O$ such that $|\text{Adv}| \leq C_{\text{adv}}$ and $\text{err}(\text{Adv}) \leq \eta_{\text{adv}}$ for some $C_{\text{adv}} \in \mathbb{N}$ and $\eta_{\text{adv}} \in \mathbb{Q}_{\geq 0}$;*
- *context bounds $C_{\text{ctxt}}^1, \dots, C_{\text{ctxt}}^n \in \mathbb{N}$ such that $\psi^i(|t_1|, \dots, |t_{|\Sigma_t|}|) \leq C_{\text{ctxt}}^i$; and*
- *axiom bounds $\varepsilon^1, \dots, \varepsilon^n \in \mathbb{Q}_{\geq 0}$ with the property that for any adversary Adv^i for protocols $\Delta^i \vdash I^i \rightarrow O^i$ such that*

$$|\text{Adv}^i| \leq \mathcal{P}(C_{\text{sem}}, C_{\text{adv}}, C_{\text{ctxt}}^i)$$

and $\text{err}(\text{Adv}^i) \leq \max(\eta_{\text{sem}}, \eta_{\text{adv}})$, we have

$$\left| \Pr[\text{Adv}^i \xrightarrow{\llbracket - \rrbracket} P^i = 1] - \Pr[\text{Adv}^i \xrightarrow{\llbracket - \rrbracket} Q^i = 1] \right| \leq \varepsilon^i,$$

we have

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} P = 1] - \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} Q = 1] \right| \leq \sum_{i=1}^n \xi^i * (\varepsilon^i + 2 * C_{\text{ctxt}}^i * \eta_{\text{sem}}).$$

We briefly explain where the error term $2 * C_{\text{ctxt}}^i * \eta_{\text{sem}}$ comes from. The original error η_{sem} is multiplied by C_{ctxt}^i because C_{ctxt}^i provides an upper bound on how many times we perform probabilistic samplings when executing the absorbed protocol, and hence serves as a proxy for how much the total error accumulates. Finally, we multiply the overall error term by two because computing the distinguishing advantage of an adversary A for protocols $P_1 \parallel Q$ versus $P_2 \parallel Q$, we accumulate errors twice: once when comparing $A \rightleftharpoons (P_1 \parallel Q)$ against $B \rightleftharpoons P_1$ (where B is the reduced adversary obtained from A by absorbing the program context Q), and then again when

⁴Instead of comparing probabilities for 1, we could have likewise used 0: the probability $\Pr[b = 0]$ that the decision Boolean b is 0 is only negligibly different from $1 - \Pr[b = 1]$, since the probability that the game ends without a decision Boolean is negligible. This follows from the fact that the error is negligible and the number of rounds is $O(\text{poly}(\lambda))$.

comparing $B \rightleftharpoons P_2$ against $A \rightleftharpoons (P_2 \parallel Q)$. Our proof of Theorem 1 relies on two intermediate results. The first one says that strict equality of protocols implies perfect indistinguishability against any adversary (not just a resource-bounded one); for proof see Section D.

LEMMA 1 (PERFECT INDISTINGUISHABILITY). *For any interpretation $\llbracket - \rrbracket$ for Σ , derivation $\Delta \vdash P = Q : I \rightarrow O$, and adversary Adv for protocols $\Delta \vdash I \rightarrow O$, we have*

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} P = 1] - \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} Q = 1] \right| = 0.$$

The next result is the *absorption lemma*, the proof of which we sketch in Section F, which allows us to absorb a protocol into an adversary at the cost of correspondingly increasing its cost and error. This is precisely where the polynomial $\mathcal{P}(x, y, z)$ in Theorem 1 comes from:

LEMMA 2 (ABSORPTION). *There exists a polynomial $\mathcal{P}(x, y, z) \geq y$ such that for any*

- *interpretation $\llbracket - \rrbracket$ for Σ for which there are $C_{\text{sem}} \in \mathbb{N}$, $\eta_{\text{sem}} \in \mathbb{Q}_{\geq 0}$ such that*
 - *for all type symbols t , $|t| \leq C_{\text{sem}}$,*
 - *for all function symbols f , $\llbracket f \rrbracket$ is computable by a TM with symbols 0, 1 such that the number of states and the runtime are $\leq C_{\text{sem}}$, and*
 - *for all distribution symbols d , $\llbracket d \rrbracket$ is computable up to error η_{sem} by a probabilistic TM with symbols 0, 1 such that the number of states and the runtime are $\leq C_{\text{sem}}$,*
- *adversary Adv for protocols of type $\Delta \vdash I \rightarrow O_1 \cup O_2$ such that $|\text{Adv}| \leq C_{\text{adv}}$ and $\text{err}(\text{Adv}) \leq \eta_{\text{adv}}$ for some $C_{\text{adv}} \in \mathbb{N}$ and $\eta_{\text{adv}} \in \mathbb{Q}_{\geq 0}$,*
- *protocol $\Delta \vdash Q : I \cup O_1 \rightarrow O_2$,*

we have an adversary $\text{Adv}_{\mathcal{R}}$ for protocols $\Delta \vdash I \cup O_2 \rightarrow O_1$ with

$$|\text{Adv}_{\mathcal{R}}| \leq \mathcal{P}(C_{\text{sem}}, C_{\text{adv}}, \|Q\|(|t_1|, \dots, |t_{|\Sigma_t|}|))$$

and $\text{err}(\text{Adv}_{\mathcal{R}}) \leq \max(\eta_{\text{sem}}, \eta_{\text{adv}})$ such that for any protocol $\Delta \vdash P : I \cup O_2 \rightarrow O_1$ we have

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} P \parallel Q = 1] - \Pr[\text{Adv}_{\mathcal{R}} \xrightarrow{\llbracket - \rrbracket} P = 1] \right| \leq \|Q\|(|t_1|, \dots, |t_{|\Sigma_t|}|) * \eta_{\text{sem}}.$$

The high-level idea behind the proof of Theorem 1 is to restructure the derivations of approximate equality so that all invocations of the rule **EMBED** are carried out first, followed by applications of the rule **INPUT-UNUSED**, which are in turn followed by invocations of the rule **CONG-COMP**, and finally by applications of the rule **CONG-NEW**. The new layered form of our approximate judgments is shown in Figures 9 and 10.

Crucially, we collapse a sequence of applications of the **CONG-COMP** rule with common contexts Q_1, \dots, Q_n into a single application with the combined common context $Q_1 \parallel \dots \parallel Q_n$. This is necessary because every time we absorb a protocol into the adversary, we increase the adversary's resources: e.g., the number of rounds that $\text{Adv}_{\mathcal{R}}$ takes is more than double the original number of rounds. Thus, if we carried out this process λ -many times, we would see an exponential increase in the adversary's resources.

Our second result (Theorem 2; for proof see Section E) concerns asymptotic security and serves as a sanity check for the concrete bounds we derived. Our asymptotic theory is said to be sound if each of its axioms is sound:

DEFINITION 10. *We say that an approximate axiom family $\{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$ is sound with respect to a family of interpretations $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$ if $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}} \models \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$.*

THEOREM 2 (SOUNDNESS OF ASYMPTOTIC EQUALITY OF PROTOCOLS). *For any*

- protocol families $\{\Delta_\lambda \vdash P_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\Delta_\lambda \vdash Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$ with identical typing judgments;
- PPT family of interpretations $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$ for Σ ; and
- an asymptotic theory that is sound with respect to $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$;

we have that

$$\vdash \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$$

implies

$$\models \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}.$$

5 CASE STUDIES

We now briefly describe our four case studies, focusing on the Multi-Party GWM protocol. In Fig. 8, we summarize our case studies by lines of code (second number in columns 2 and 3) and runtime to verify (column 4). Whenever applicable, we compare against the corresponding case study from [Gancher et al. 2023b] implemented in Coq (first number in columns 2 and 3). Our largest case study currently takes 12 minutes to verify. This number can likely be reduced with performance optimizations and/or by splitting the proof across multiple files that can be verified in parallel⁵.

Case Study	Proofs (LoC)		Defs (LoC)		Runtime
	Coq	DSL	Coq	DSL	
Auth-To-Secure Channel [Maurer 2012]	128	61	97	143	1.38s
DHKE-OTP [Barbosa et al. 2021]	532	164	183	371	3.3s
Multi-Party Coin Flip [Blum 1983a]	1905	256	114	311	5.1s
Multi-Party GMW [Goldreich et al. 1987]	-	9102	-	5738	12.08m

Fig. 8. Case Studies: size of proofs and definitions, and runtime to verify.

Authenticated-To-Secure Channel. Our smallest case study constructs a secure channel from an authenticated one. Alice wants to communicate q messages to Bob using an authenticated channel. The authenticated channel is not secure: it leaks each message to the adversary, and waits to receive an *ok* message back from the adversary before delivering the in-flight message. Thus, the adversary cannot modify any of the messages but can read and delay them for any amount of time. To transmit information securely, Alice sends encryptions of her messages, which Bob decrypts using a shared key not known to the adversary. The indistinguishability assumption states that the encryption scheme is CPA-secure.

One-Time Pad (OTP) From Diffie-Hellman Key Exchange. In symmetric-key encryption, the sender (Alice) and the receiver (Bob) need to agree on a shared secret key. One such key-agreement protocol is the Diffie-Hellman Key Exchange (DHKE), which assumes a cyclic group G of a prime order with generator g . The indistinguishability assumption is the *decisional Diffie-Hellman (DDH)* assumption: as long as the exponents k, l are generated uniformly, even if the adversary knows the values g^k and g^l , they will be unable to distinguish $(g^k)^l$ from a uniformly generated element of G .

We subsequently use the DHKE protocol to turn an authenticated channel into a one-time pad (OTP) that delivers a single secret message from Alice to Bob. Our proof is modular: in the first step, we establish that the DHKE protocol can be replaced with its idealization. In the second step, we prove that the resulting OTP protocol itself reduces to an idealization.

⁵The case studies and the sources can be found at <https://github.com/concrete-bounds-for-mpc-proofs/concrete-bounds-for-mpc-simulation-proofs>.

Multi-Party Coin Flip. We implement a protocol (originally due to [Blum 1983b]) where $N + 2$ parties labeled $0, \dots, N - 1$ reach a Boolean consensus. At the start of the protocol, each party commits to a randomly-generated Boolean. After all parties have committed, each party opens its commit. The consensus is the Boolean sum of all commits. We prove the protocol secure against a malicious attacker in the case when party N is corrupt, party $N + 1$ is honest, and any other party is arbitrarily honest or corrupt. As this protocol is perfectly secure, there are no indistinguishability assumptions.

Multi-Party GMW Protocol. In the multi-party GMW protocol [Goldreich et al. 1987], $N + 2$ parties securely compute the value of a given Boolean circuit built out of *xor*-, *and*-, and *not* gates. The inputs to the circuit are divided among the parties, and no party has access to the inputs of any other. Each party maintains its share of the value v computed by each gate, and summing up the shares of all parties yields v . We prove the protocol secure in the case when party N is semi-honest, party $N + 1$ is honest, and any other party is arbitrarily honest or semi-honest.

The protocol consists of the $N + 2$ parties, plus an instance of the 1-Out-Of-4 Oblivious Transfer (OT) protocol for each gate and each pair of parties $n < m$, where n is the sender and m is the receiver. The code for each party is separated into three parts: in the initial phase, each party computes and distributes everyone's shares for each of its inputs. In the inductive phase, each party computes their share of each gate by induction on the ambient circuit. At last, in the final phase, parties send their shares of each output wire to one another and add them up to compute the result.

The shares of each gate are computed as follows. In the case of an *input* gate, the parties use the corresponding input share from the initial phase. In the case of a *not* gate, parties $0, \dots, N$ simply copy their share of the incoming wire, whereas party $N + 1$ negates its share. If the gate is an *xor* gate, the resulting share is the sum of the shares of the incoming two wires. The case of an *and* gate is the most complex. The sum of everybody's shares must equal $(x_0 \oplus \dots \oplus x_{N+1}) * (y_0 \oplus \dots \oplus y_{N+1})$, where x_n, y_n are the respective shares of party n on the incoming two wires. We have

$$(x_0 \oplus \dots \oplus x_{N+1}) * (y_0 \oplus \dots \oplus y_{N+1}) = \bigoplus_i \bigoplus_j x_i * y_j$$

Parties n and m engage in a 1-Out-Of-4 OT exchange to compute $(x_n * y_m) \oplus (x_m * y_n)$. There are four possible combinations of values that x_m, y_m can take, and party n computes the value of $(x_n * y_m) \oplus (x_m * y_n)$ for each. This offers party m four messages to choose from, and he selects the one corresponding to the actual values of x_m, y_m . A small caveat: in the exchange as described above, party m would still be able to infer the value of party n 's shares in certain cases: e.g., if $x_m = 0$ and $y_m = 1$, party m gets the share x_n as the result of the exchange. To prevent this, party n encodes her messages by masking them with a random Boolean b that only she knows. To offset for the presence of this Boolean, she includes it in her own share $b \oplus (x_n * y_n)$.

We assume four opaque protocols representing four implementations of the 1-Out-Of-4 Oblivious Transfer, one for each of the possible combinations of an honest/semi-honest sender and receiver. We also assume four cryptographic hardness axioms, stating that each OT implementation is approximately equal to an ideal 1-Out-Of-4 OT functionality. In the first step of our proof, we replace each OT implementation by its the corresponding functionality. The rest of the proof is carried out in the exact fragment of our proof system.

Since the last party is by assumption honest, the simulator does not have access to its inputs. Therefore, any computation that depends on the value of the inputs belonging to the last party must be eliminated. In particular, all shares of the last party must be eliminated. Instead, the simulator only computes shares for parties $0, \dots, N$ in the inductive part, and replaces every mention of the

last party's share in its leakage by the quantity $x \oplus (\bigoplus_{i \leq N} x_i)$. Here x is the value carried by the gate as computed by the ideal functionality, and leaked to the simulator by a semi-honest party.

In this case study, the main challenge for formal verification is how to effectively carry out proofs by induction nested three levels deep: one for the arbitrary Boolean circuit, and two for each of the arbitrarily many parties communicating with another party. Additionally, we have to account for an essentially arbitrary combination of honest/semi-honest parties, as all but the last two parties can be honest or semi-honest. Every one of these factors reflects into the actual code of the protocol, and has to be taken into account when computing concrete security bounds. The final concrete bound computed by our tool is as follows:

indistinguishability assumption HH2HH :

count: $|\{(n, m, k) \text{ s.t. } n < N + 2, m < N + 2, k < K,$
 $\text{when isHonest}(n) \text{ and isHonest}(m)\}|$

context:

$N * K * \max(N * 434 + 990, N * 434 + 1011, N * 607 + 1329, N * 434 + 983) +$
 $N * N * K * \max(|1\text{OutOf4OTReal-Honest-Honest}|, 543) + N * N * K * 218 +$
 $N * K * \max(|1\text{OutOf4OTReal-Honest-Honest}|, 543) * 4 + N * K * 932 +$
 $K * \max(N * 434 + 990, N * 434 + 1011, N * 607 + 1329, N * 434 + 983) * 2 +$
 $N * N * \maxValue(I, N + 2) * 300 + N * N * 68 + K * 992 + N * 352 +$
 $K * \max(|1\text{OutOf4OTReal-Honest-Honest}|, 543) * 3 +$
 $(K - 1) * \max(|1\text{OutOf4OTReal-Honest-Honest}|, 543) +$
 $N * \maxValue(I, N + 2) * 1169 + \maxValue(I, N + 2) * 1138 + 509$

indistinguishability assumption SHH2SHH : ...

indistinguishability assumption HSH2SHS : ...

indistinguishability assumption SHSH2SHSH : ...

Here K is the number of wires in the circuit, HH2HH is the indistinguishability assumption stating that the 1-Out-Of-4 Oblivious Transfer protocol when both the sender and the receiver are honest can be soundly replaced by the corresponding functionality, and $\maxValue(I, N + 2)$ selects the maximum of $I(0), \dots, I(N + 1)$, where $I(n)$ denotes the number of inputs to party n . We omit the bounds for the other indistinguishability assumptions, as these are entirely analogous. As expected, the bound **context** is $O(N * N * K)$, since we have one instance of an OT protocol for each of the K wires and each pair of parties. As we can see from the the bound **count**, the total number of times we invoke an indistinguishability assumption is $(N + 2) * (N + 2) * K$, since we replace each of the OT protocols with its idealization. The indistinguishability assumption applicable to a particular OT instance is conditional on whether the sender and receiver are honest or semi-honest, respectively.

DSL Implementation. Our DSL serves as a layer of abstraction over the implementation internals, hiding low-level proof details from the user. For each rule, we have language construct in the DSL that handles its application. For example, an application of the rule **SUBST** in Fig. 2 is written in the DSL as **subst** o1 **into** o2. Internally, this translates into a call of a Maude strategy, which applies the substitution rule over a protocol without explicitly writing applications of congruence or exchange rules. Moreover, in many cases these strategies are generated on the fly, instead of requiring the user to manually specify them.

6 FUTURE WORK

We present a promising formalism for developing verified cryptographic proofs for MPC protocols. We compute concrete security bounds by bounding the size of the program context in which indistinguishability assumptions are applied, and analyzing the cost of the interpreter that executes the program context. Currently, our main source of over-approximation is the runtime of this interpreter, which is currently quadratic due to using linear scans for inputs (see Section F); however, reducing the overhead to near-linear is possible using standard techniques.

Our tool can model arbitrary (static) corruption scenarios, including protocols that tolerate k/N corruptions (either semi-honest or malicious); formalizing a protocol that demonstrates this capability is one avenue for future work. In addition to verifying even more sophisticated MPC protocols (e.g., garbled circuits), we want to develop new algorithms for automatically synthesizing cryptographic simulators, rather than requiring them to be manually encoded. Finally, we aim to extend our system to handle more expressive classes of protocols, such as those exhibiting threshold behavior (e.g., consensus protocols).

REFERENCES

- J. Almeida, M. Barbosa, G. Barthe, François Dupressoir, B. Grégoire, Vincent Laporte, and Vitor Pereira. 2017. A Fast and Verified Software Stack for Secure Function Evaluation. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (2017).
- David Baelde, Caroline Fontaine, Adrien Koutsos, Guillaume Scerri, and Theo Vignon. 2024. A Probabilistic Logic for Concrete Security. In *2024 IEEE 37th Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, Los Alamitos, CA, USA, 324–339. <https://doi.org/10.1109/CSF61375.2024.00046>
- Manuel Barbosa, Gilles Barthe, Benjamin Grégoire, Adrien Koutsos, and Pierre-Yves Strub. 2021. Mechanized Proofs of Adversarial Complexity and Application to Universal Composability. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (Virtual Event, Republic of Korea) (CCS '21)*. Association for Computing Machinery, New York, NY, USA, 2541–2563. <https://doi.org/10.1145/3460120.3484548>
- G. Barthe, B. Grégoire, S. Heraud, and Santiago Zanella Béguelin. 2011. Computer-Aided Security Proofs for the Working Cryptographer. In *CRYPTO*.
- Donald Beaver. 1995. Precomputing Oblivious Transfer. In *Annual International Cryptology Conference*. Springer, 97–109.
- Mihir Bellare, Anand Desai, Eron Jokipii, and Phillip Rogaway. 1997. A Concrete Security Treatment of Symmetric Encryption. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*. IEEE, 394–403.
- Bruno Blanchet. 2008. A Computationally Sound Mechanized Prover for Security Protocols. *IEEE Trans. Dependable Secur. Comput.* 5, 4 (2008), 193–207. <https://doi.org/10.1109/TDSC.2007.1005>
- Manuel Blum. 1983a. Coin Flipping by Telephone a Protocol for Solving Impossible Problems. *SIGACT News* 15, 1 (1983), 23–27. <https://doi.org/10.1145/1008908.1008911>
- Manuel Blum. 1983b. Coin flipping by telephone a protocol for solving impossible problems. *ACM SIGACT News* 15, 1 (1983), 23–27.
- Chris Brzuska, Antoine Delignat-Lavaud, Cédric Fournet, Konrad Kohbrok, and Markulf Kohlweiss. 2018. State separation for code-based game-playing proofs. In *Advances in Cryptology—ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part III* 24. Springer, 222–249.
- Ran Canetti. 2000. Universally Composable Security: A New Paradigm for Cryptographic Protocols. Cryptology ePrint Archive, Report 2000/067. <https://ia.cr/2000/067>.
- Ran Canetti, Alley Stoughton, and Mayank Varia. 2019. EasyUC: Using EasyCrypt to Mechanize Proofs of Universally Composable Security. In *32nd IEEE Computer Security Foundations Symposium*. <https://eprint.iacr.org/2019/582>.
- Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn L. Talcott (Eds.). 2007. *All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic*. Lecture Notes in Computer Science, Vol. 4350. Springer. <https://doi.org/10.1007/978-3-540-71999-1>
- Ankush Das, Jan Hoffmann, and Frank Pfenning. 2018. Work Analysis with Resource-Aware Session Types. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. 305–314.
- Karim M. El Defrawy and Vitor Pereira. 2019. A High-Assurance Evaluator for Machine-Checked Secure Multiparty Computation. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (2019).
- Joshua Gancher, Sydney Gibson, Pratap Singh, Samvid Dharanikota, and Bryan Parno. 2023a. Owl: Compositional Verification of Security Protocols via an Information-Flow Type System. In *44th IEEE Symposium on Security and Privacy, SP 2023*,

- San Francisco, CA, USA, May 21-25, 2023. IEEE, 1130–1147. <https://doi.org/10.1109/SP46215.2023.10179477>
- Joshua Gancher, Kristina Sojakova, Xiong Fan, Elaine Shi, and Greg Morrisett. 2023b. A Core Calculus for Equational Proofs of Cryptographic Protocols. *Proc. ACM Program. Lang.* 7, POPL (2023), 866–892. <https://doi.org/10.1145/3571223>
- Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play Any Mental Game. In *Proceedings of the 19th annual ACM symposium on Theory of Computing*. ACM, 218–229.
- Helene Haagh, Aleksandr Karbyshev, Sabine Oechsner, Bas Spitters, and Pierre-Yves Strub. 2018. Computer-Aided Proofs for Multiparty Computation with Active Security. *2018 IEEE 31st Computer Security Foundations Symposium (CSF)* (2018), 119–131.
- Andreas Lochbihler, S. Reza Sefidgar, David Basin, and Ueli Maurer. 2019. Formalizing Constructive Cryptography using CryptHOL. In *32nd IEEE Computer Security Foundations Symposium*. <http://www.andreas-lochbihler.de/pub/lochbihler2019csf.pdf>.
- Ueli Maurer. 2012. Constructive Cryptography – A New Paradigm for Security Definitions and Proofs. In *Theory of Security and Applications*, Sebastian Mödersheim and Catuscia Palamidessi (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 33–56.
- Christian Skalka and Joseph Near. 2024. Language-Based Security for Low-Level MPC. In *Proceedings of the 26th International Symposium on Principles and Practice of Declarative Programming*. 1–14.
- Alley Stoughton and Mayank Varia. 2017. Mechanizing the Proof of Adaptive, Information-Theoretic Security of Cryptographic Protocols in the Random Oracle Model. *2017 IEEE 30th Computer Security Foundations Symposium (CSF)* (2017), 83–99.
- Ionut Tăutu. 2022. SpeX: A Rewriting-Based Formal Specification Environment. In *Recent Trends in Algebraic Development Techniques - 26th IFIP WG 1.3 International Workshop, WADT 2022, Aveiro, Portugal, June 28-30, 2022, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 13710)*, Alexandre Madeira and Manuel A. Martins (Eds.). Springer, 163–178. https://doi.org/10.1007/978-3-031-43345-0_8
- Xingyu Xie, Yifei Li, Wei Zhang, Tuowei Wang, Shizhen Xu, Jun Zhu, and Yifan Song. 2024. GAuV: A Graph-Based Automated Verification Framework for Perfect Semi-Honest Security of Multiparty Computation Protocols. *IACR Cryptol. ePrint Arch.* (2024).

A TURING MACHINE BOUNDS

The Turing Machine bound of a type τ is straightforward:

$$\begin{aligned}
 \|t_i\| &:= t_i \\
 \|1\| &:= 0 \\
 \|\text{Bool}\| &:= 1 \\
 \|\tau_1 \times \tau_2\| &:= \|\tau_1\| + \|\tau_2\|
 \end{aligned}$$

The encoding of variables x of type τ uses the symbols “(”, “var”, “:”, “)” in addition to the de Bruijn index of the variable x , encoded as a single symbol, and the encoding of the type annotation τ . For expressions \checkmark , true, false, we use the corresponding symbols “ \checkmark ”, “true”, “false” and the two parenthesis symbols “(”, “)”. For an application $f\ e$ of a function of type $\sigma \rightarrow \tau$, we use the symbols “(”, “app”, “ \rightarrow ”, “)” in addition to the function symbol f , encoded as a single symbol, and the encodings of the two type annotations σ, τ and the expression e . To encode a pair (e_1, e_2) , we will only need the encodings of the two expressions e_1 and e_2 . Finally, to encode first and second projections, we will use the symbols “(”, “fst”, “snd”, “ \times ”, “of”, “)” in addition to the encodings of

the two type annotations σ, τ and the expression e .

$$\begin{aligned}
\|\text{var}(x : \tau)\| &:= \|\tau\| + 5 \\
\|\checkmark\| &:= 3 \\
\|\text{true}\| &:= 3 \\
\|\text{false}\| &:= 3 \\
\|\text{app}_{\sigma \rightarrow \tau} f e\| &:= \|\sigma\| + \|\tau\| + \|e\| + 5 \\
\|(e_1, e_2)\| &:= \|e_1\| + \|e_2\| \\
\|\text{fst}_{\sigma \times \tau} e\| &:= \|\sigma\| + \|\tau\| + \|e\| + 5 \\
\|\text{snd}_{\sigma \times \tau} e\| &:= \|\sigma\| + \|\tau\| + \|e\| + 5
\end{aligned}$$

For a `ret` e , we use the symbols “(”, “ret”, “)” in addition to the encoding of the expression e . For a `sampling` `samp` d e from a distribution of type $\sigma \rightarrow \tau$, we use the symbols “(”, “samp”, “ \rightarrow ”, “)” in addition to the distribution symbol d , encoded as a single symbol, and the encodings of the two type annotations σ, τ and the expression e . For a `read` c from a channel of type τ , we use the symbols “(”, “read”, “:”, “)” in addition to the de Bruijn index of the channel c , encoded as a single symbol, and the encoding of the type annotation τ . Furthermore, we will need one extra symbol: one of “input-to-query”, “input-queried”, “input-not-to-query”. When encoding a protocol $Q : I \cup O_1 \rightarrow O_2$ coming from the `COMP-CONG` rule, we use “input-to-query” or “input-queried” if we are reading from a channel $o_1 \in O_1$, according to whether we have already queried the channel o_1 , and “input-not-to-query” otherwise.

For a conditional `if` e `then` R_1 `else` R_2 , we use the symbols “(”, “if”, “then”, “else”, “)” in addition to the encodings of the expression e and the two reactions R_1, R_2 . Finally, to encode a `bind`, we use the symbols “{”, “_”, “:”, “ \leftarrow ”, “;”, “}” in addition to the encodings of the type annotation σ and the two reactions R and S . The symbol “_” is used in lieu of the bound variable name x and stands for de Bruijn index 0.

$$\begin{aligned}
\|\text{ret } e\| &:= \|e\| + 3 \\
\|\text{samp}_{\sigma \rightarrow \tau} d e\| &:= \|\sigma\| + \|\tau\| + \|e\| + 5 \\
\|\text{read}(c : \tau)\| &:= \|\tau\| + 6 \\
\|\text{if } e \text{ then } R_1 \text{ else } R_2\| &:= \|e\| + \|R_1\| + \|R_2\| + 5 \\
\|x : \sigma \leftarrow R; S\| &:= \|\sigma\| + \|R\| + \|S\| + 6
\end{aligned}$$

To encode the zero protocol 0, we use the single symbol “0”. For an assignment $o := R$, we use the symbols “[”, “:=”, “react”, “]” in addition to the de Bruijn index of the channel c , encoded as a single symbol, and the encoding of the reaction R . For a parallel composition $P \parallel Q$, we use the symbols “(”, “||”, “)” in addition to the encodings of the two protocols P and Q . Finally, for the declaration of a new channel `new` $o : \tau$ `in` P , we use the symbols “new”, “_”, “:”, “in”, “wen” in addition to the encodings of the typing annotation τ and the protocol P . The symbol “_” is used in lieu of the bound channel name c and stands for de Bruijn index 0. The symbol “wen” indicates the end of the binding scope.

$$\begin{aligned}
\|0\| &:= 1 \\
\|o := R\| &:= \|R\| + 5 \\
\|P \parallel Q\| &:= \|P\| + \|Q\| + 3 \\
\|\text{new } c : \tau \text{ in } P\| &:= \|\tau\| + \|P\| + 5
\end{aligned}$$

B ENCODING PROTOCOLS ON A TURING MACHINE TAPE

To encode protocols on a Turing Machine tape, we make use of the following sets of symbols:

- Punc with symbols “<”, “>”, “(”, “)”, “{”, “}”, “[”, “]”, “_”, “:”, “.”, “;”, “→”, “⇒”, “←”, “×”, “:=”, “||”, “•”,
- KeyWords with symbols “var”, “√”, “true”, “false”, “app”, “fst”, “snd”, “of”, “ret”, “smp”, “read”, “if”, “then”, “else”, “0”, “new”, “in”, “wen”, “input-to-query”, “input-queried”, “input-not-to-query”.

We also need a finite set of de Bruijn indices in lieu of channel and variable names. To derive an upper bound on how many indices we will need, we statically count the maximum depth of variable and channel declarations in the protocol, giving us a *variable-index bound* and a *channel-index bound*, respectively.

To avoid an infinite loop, an adversary executing the absorbed protocol will need to keep track of which channels have already been queried for a value. We store this information inside the protocol in the form of an annotation: for each channel read $\text{read}(c : \tau)$, we denote whether the channel c has already been queried for a value, if applicable. By erasing the annotations from a query-annotated reaction or protocol, we obtain the underlying IPDL construct.

Given an ambient interpretation $\llbracket - \rrbracket$ for the signature Σ , we now show how to encode IPDL constructs as a sequence of symbols on a Turing Machine tape. For types, the encoding $\text{Enc}[\tau]$ consists of the symbol “.” repeated $|\tau|$ -many times. For expressions, we have the encoding below, where $+$ denotes string concatenation. We recall that each variable name is represented as a de Bruijn index, and is in particular a natural number.

$$\begin{aligned}
 \text{Enc}[v] &:= v \\
 \text{Enc}[\text{var}(x : \tau)] &:= “(” + “\text{var}” + x + “:” + \text{Enc}[\tau] + “)” \\
 \text{Enc}[\sqrt{}] &:= “(” + “\sqrt{” + “)” \\
 \text{Enc}[\text{true}] &:= “(” + “\text{true}” + “)” \\
 \text{Enc}[\text{false}] &:= “(” + “\text{false}” + “)” \\
 \text{Enc}[\text{app}_{\sigma \rightarrow \tau} f e] &:= “(” + “\text{app}” + \text{Enc}[\sigma] + “\rightarrow” + \text{Enc}[\tau] + f + \text{Enc}[e] + “)” \\
 \text{Enc}[(e_1, e_2)] &:= \text{Enc}[e_1] + \text{Enc}[e_2] \\
 \text{Enc}[\text{fst}_{\sigma \times \tau} e] &:= “(” + “\text{fst}” + \text{Enc}[\sigma] + “\times” + \text{Enc}[\tau] + “\text{of}” + \text{Enc}[e] + “)” \\
 \text{Enc}[\text{snd}_{\sigma \times \tau} e] &:= “(” + “\text{snd}” + \text{Enc}[\sigma] + “\times” + \text{Enc}[\tau] + “\text{of}” + \text{Enc}[e] + “)”
 \end{aligned}$$

The encoding $\text{Enc}[a]$ of an annotation is the corresponding symbol. For reactions, we have the following encoding, where we recall that each channel name is represented as a de Bruijn index, and is in particular a natural number.

$$\begin{aligned}
 \text{Enc}[\text{val } v] &:= “<” + v + “>” \\
 \text{Enc}[\text{ret } e] &:= “(” + “\text{ret}” + \text{Enc}[e] + “)” \\
 \text{Enc}[\text{smp}_{\sigma \rightarrow \tau} d e] &:= “(” + “\text{smp}” + \text{Enc}[\sigma] + “\rightarrow” + \text{Enc}[\tau] + d + \text{Enc}[e] + “)” \\
 \text{Enc}[\text{read}[a](c : \tau)] &:= “(” + “\text{read}” + \text{Enc}[a] + c + “:” + \text{Enc}[\tau] + “)” \\
 \text{Enc}[\text{if } e \text{ then } R_1 \text{ else } R_2] &:= “(” + “\text{if}” + \text{Enc}[e] + “\text{then}” + \text{Enc}[R_1] + “\text{else}” + \text{Enc}[R_2] + “)” \\
 \text{Enc}[x : \sigma \leftarrow R; S] &:= “\{” + “_” + “:” + \text{Enc}[\sigma] + “\leftarrow” + \text{Enc}[R] + “;” + \text{Enc}[S] + “\}”
 \end{aligned}$$

Finally, for protocols we have the encoding below.

$$\begin{aligned}
\text{Enc}[0] &:= "0" \\
\text{Enc}[o := v] &:= "[" + o + " := " + v + "]" \\
\text{Enc}[o := R] &:= "(" + o + " := " + \text{"react"} + \text{Enc}[R] + ")" \\
\text{Enc}[P \parallel Q] &:= "(" + \text{Enc}[P] + " \parallel " + \text{Enc}[Q] + ")" \\
\text{Enc}[\text{new } c : \tau \text{ in } P] &:= \text{"new"} + \text{"_"} + \text{" : " + Enc}[\tau] + \text{" in " + Enc}[P] + \text{"wen"}
\end{aligned}$$

To avoid having to shift the tape contents when executing IPDL protocols on a Turing Machine tape, we will make use of the white-space symbol " ", which we consider as distinct from the symbol *blank*. The former will be used as a placeholder so that our protocol encoding remains at a constant length throughout the execution. For this reason, we extend our notion of encoding to allow extra white-spaces around the encoding of an expression e or a query-annotated reaction R occurring inside a query-annotated protocol P .

C LAYERED APPROXIMATE JUDGEMENTS

The layered form of our approximate judgments is shown in Figures 9 and 10.

LEMMA 3. *We have $\Delta \vdash P \approx Q : I \rightarrow O$ count k ctxt l iff $\Delta \vdash P \approx_5 Q : I \rightarrow O$ count k ctxt l for any protocols $\Delta \vdash P : I \rightarrow O$ and $\Delta \vdash Q : I \rightarrow O$.*

D PERFECT INDISTINGUISHABILITY: PROOF OF LEMMA 1

LEMMA (PERFECT INDISTINGUISHABILITY). *For any interpretation $\llbracket - \rrbracket$ for Σ , derivation $\Delta \vdash P = Q : I \rightarrow O$, and adversary Adv for protocols $\Delta \vdash I \rightarrow O$, we have*

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} P = 1] - \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} Q = 1] \right| = 0.$$

PROOF. Fix an adversary Adv as in Definition 3. By assumption, we have a proof $\Delta \vdash P = Q : I \rightarrow O$, which means we also have a proof that $\Delta' \vdash \phi^*(P) = \phi^*(Q) : \phi^*(I) \rightarrow \phi^*(O)$. The soundness theorem for strict equality of protocols applied to this proof gives us a bisimulation \sim such that $1[\phi^*(P)] \sim 1[\phi^*(Q)]$. Now let \sim_{adv} be a binary relation on sub-distributions on pairs where the first element is an adversary state and the second is a protocol of type $\Delta \vdash I \rightarrow O$, defined as follows:

- $(s, \eta) \sim_{\text{adv}} (s, \varepsilon)$ if $s \in \text{St}$ and $\eta \sim \varepsilon$, where we use a distribution in place of a protocol to indicate the obvious lifting to sub-distributions on pairs of the the aforementioned form, and
- $1[\perp] \sim_{\text{adv}} 1[\perp]$, where \perp indicates that the security game between the adversary and the protocol halted without a decision Boolean.

Let $\mathcal{L}_{\sim_{\text{adv}}}$ be the closure of \sim_{adv} under joint convex combinations. Explicitly, $\mathcal{L}_{\sim_{\text{adv}}}$ is defined by

$$\left(\sum_i c_i \eta_i \right) \mathcal{L}_{\sim_{\text{adv}}} \left(\sum_i c_i \varepsilon_i \right)$$

for coefficients $c_i > 0$ with $\sum_i c_i = 1$ and distributions $\eta_i \sim_{\text{adv}} \varepsilon_i$. We now establish a loop invariant for the algorithm in Figure 7. Before starting the first round, the initial distributions are suitably related: by assumption, we have $1[\phi^*(P)] \sim 1[\phi^*(Q)]$, which means that

$$1[(s_\star, \phi^*(P))] \mathcal{L}_{\sim_{\text{adv}}} 1[(s_\star, \phi^*(Q))]$$

as the two distributions are already related under \sim_{adv} . Now assume that we have two sub-distributions related by $\mathcal{L}_{\sim_{\text{adv}}}$. We prove that performing a single round yields sub-distributions that are again related by $\mathcal{L}_{\sim_{\text{adv}}}$. It suffices to show this for the case $(s, \eta) \sim_{\text{adv}} (s, \varepsilon)$, where $s \in \text{St}$

$$\begin{array}{c}
\boxed{\Delta \vdash P \cong_0 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \\
\\
\frac{\Delta^k \vdash P^k \approx Q^k : I^k \rightarrow O^k}{\Delta^k \vdash P^k \cong_0 Q^k : I^k \rightarrow O^k \text{ ctxt } 0 \text{ for axiom \# } k} \text{AXIOM} \\
\\
\boxed{\Delta \vdash P \cong_1 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \\
\\
\frac{\Delta \vdash P \cong_0 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash P \cong_1 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{SUB} \\
\\
\frac{\phi : \Delta_1 \rightarrow \Delta_2 \quad \Delta_2 \vdash P \cong_0 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta_1 \vdash \phi^*(P) \cong_1 \phi^*(Q) : \phi^*(I) \rightarrow \phi^*(O) \text{ ctxt } \psi \text{ for axiom \# } k} \text{EMBED} \\
\\
\boxed{\Delta \vdash P \cong_2 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \\
\\
\frac{\Delta \vdash P \cong_1 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash P \cong_2 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{SUB} \\
\\
\frac{i \notin I \cup O \quad \Delta \vdash P \cong_2 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash P \cong_2 Q : I \cup \{i\} \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{INPUT-UNUSED} \\
\\
\boxed{\Delta \vdash P \cong_3 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \\
\\
\frac{\Delta \vdash P \cong_2 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash P \cong_3 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{SUB} \\
\\
\frac{\Delta \vdash P \cong_2 P' : I \cup O_2 \rightarrow O_1 \text{ ctxt } \psi \text{ for axiom \# } k \quad \Delta \vdash Q : I \cup O_1 \rightarrow O_2}{\Delta \vdash P \parallel Q \cong_3 P' \parallel Q : I \rightarrow O_1 \cup O_2 \text{ ctxt } \psi + \|Q\| + 3 \text{ for axiom \# } k} \text{CONG-COMP-LEFT} \\
\\
\boxed{\Delta \vdash P \cong_4 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \\
\\
\frac{\Delta \vdash P \cong_3 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash P \cong_4 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{SUB} \\
\\
\frac{\Delta, o : \tau \vdash P \cong_4 P' : I \rightarrow O \cup \{o\} \text{ ctxt } \psi \text{ for axiom \# } k}{\Delta \vdash (\text{new } o : \tau \text{ in } P) \cong_4 (\text{new } o : \tau \text{ in } P') : I \rightarrow O \text{ ctxt } \psi \text{ for axiom \# } k} \text{CONG-NEW}
\end{array}$$

Fig. 9. Layered approximate judgements for protocols.

$$\boxed{\Delta \vdash P \cong_5 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom } \# k}$$

$$\frac{\Delta \vdash P = P' : I \rightarrow O \quad \Delta \vdash P' \cong_4 Q' : I \rightarrow O \text{ ctxt } \psi \text{ for axiom } \# k \quad \Delta \vdash Q' = Q : I \rightarrow O}{\Delta \vdash P \cong_5 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom } \# k} \text{ SUB}$$

$$\boxed{\Delta \vdash P \approx_5 Q : I \rightarrow O \text{ count } \xi \text{ ctxt } \psi}$$

$$\frac{\Delta \vdash P = Q : I \rightarrow O}{\Delta \vdash P \approx_5 Q : I \rightarrow O \text{ count } (i \mapsto 0) \text{ ctxt } (i \mapsto 0)} \text{ STRICT}$$

$$\frac{\Delta \vdash P \cong_5 Q : I \rightarrow O \text{ ctxt } \psi \text{ for axiom } \# k}{\Delta \vdash P \approx_5 Q : I \rightarrow O \text{ count } (k \mapsto 1, i \neq k \mapsto 0) \text{ ctxt } (k \mapsto \psi, i \neq k \mapsto 0)} \text{ APPROX-CONG}$$

$$\frac{\Delta \vdash P_1 \approx_5 P_2 : I \rightarrow O \text{ count } \xi \text{ ctxt } \psi}{\Delta \vdash P_2 \approx_5 P_1 : I \rightarrow O \text{ count } \xi \text{ ctxt } \psi} \text{ SYM}$$

$$\frac{\Delta \vdash P_1 \approx_5 P_2 : I \rightarrow O \text{ count } \xi_1 \text{ ctxt } \psi_1 \quad \Delta \vdash P_2 \approx_5 P_3 : I \rightarrow O \text{ count } \xi_2 \text{ ctxt } \psi_2}{\Delta \vdash P_1 \approx_5 P_3 : I \rightarrow O \text{ count } (i \mapsto \xi_1(i) + \xi_2(i)) \text{ ctxt } (i \mapsto \max(\psi_1(i), \psi_2(i)))} \text{ TRANS}$$

Fig. 10. Layered approximate judgements for protocols, continued.

and $\eta \sim \varepsilon$. We first compute the distributions $\eta' := \eta \Downarrow$ and $\varepsilon' := \varepsilon \Downarrow$. By definition of \sim we have $\eta' \sim \varepsilon'$. Independently, we probabilistically compute the type of interaction to perform together with a new adversary state s' . If no interaction has been chosen, the resulting distributions are (s', η') and (s', ε') . We have

$$(s', \eta') \mathcal{L}_{\sim_{\text{adv}}} (s', \varepsilon')$$

as desired, as the two distributions are already related under \sim_{adv} . If the interaction is an input on channel i , we compute $O_i(s')$ to see if in the adversary's current state s' the channel i carries a value. If this is not the case, the resulting distributions are (s', η') and (s', ε') . Here we again have $(s', \eta') \mathcal{L}_{\sim_{\text{adv}}} (s', \varepsilon')$, as desired. On the other hand, if the channel i carries a value v , the resulting distributions are $(s', \eta'[\text{read } i := \text{val } v])$ and $(s', \varepsilon'[\text{read } i := \text{val } v])$. Now because $\eta' \sim \varepsilon'$, by definition of \sim we have $\eta'[\text{read } i := \text{val } v] \sim \varepsilon'[\text{read } i := \text{val } v]$. Thus we have

$$(s', \eta'[\text{read } i := \text{val } v]) \mathcal{L}_{\sim_{\text{adv}}} (s', \varepsilon'[\text{read } i := \text{val } v])$$

as desired, as the two distributions are already related under \sim_{adv} . Finally, if the interaction is a query for an output channel o , we recall that the valuation property of the bisimulation \sim allows us to jointly partition the distributions $\eta' \sim \varepsilon'$ into a joint convex combination

$$\eta' = \sum_i c_i \eta'_i \sim \sum_i c_i \varepsilon'_i = \varepsilon'$$

with $c_i > 0$ and $\sum_i c_i = 1$ such that

- the respective components $\eta'_i \sim \varepsilon'_i$ are again related, and
- $\eta'_i|_{\text{val}(o)} = v_\perp = \varepsilon'_i|_{\text{val}(o)}$ for the same $v_\perp \in \{\perp\} \cup \llbracket \tau \rrbracket$ where $o : \tau$ in Δ' .

Therefore, it suffices to consider the respective components $\eta'_i \sim \varepsilon'_i$ with the same v_\perp . If v_\perp is \perp , then the resulting distributions are (s', η'_i) and (s', ε'_i) . Here we again have $(s', \eta'_i) \mathcal{L}_{\sim_{\text{adv}}} (s', \varepsilon'_i)$, as desired. On the other hand, if v_\perp is a value v , then the resulting distributions are $(l_o(v, s'), \eta'_i)$ and $(l_o(v, s'), \varepsilon'_i)$. Thus we have

$$(l_o(v, s'), \eta'_i) \mathcal{L}_{\sim_{\text{adv}}} (l_o(v, s'), \varepsilon'_i)$$

as desired, as the two distributions are already related under \sim_{adv} . This proves that after completing the required number of rounds, we end up with two sub-distributions related by $\mathcal{L}_{\sim_{\text{adv}}}$. It is now easy to see that they induce the same sub-distribution on decision Booleans. It suffices to prove this for the case $(s, \eta) \sim_{\text{adv}} (s, \varepsilon)$, where $s \in \text{St}$ and $\eta \sim \varepsilon$. But the state s is the same for both distributions, so the resulting distribution on decision Booleans is $1[D(s)]$. This finishes the proof. \square

E SOUNDNESS OF ASYMPTOTIC EQUALITY: PROOF OF THEOREM 2

THEOREM (SOUNDNESS OF ASYMPTOTIC EQUALITY OF PROTOCOLS). *For any*

- *protocol families $\{\Delta_\lambda \vdash P_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\Delta_\lambda \vdash Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$ with identical typing judgments;*
- *PPT family of interpretations $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$ for Σ ; and*
- *an asymptotic theory that is sound with respect to $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$;*

we have that

$$\vdash \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$$

implies

$$\models \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}.$$

PROOF. Given n approximate axioms, the top-level asymptotic equality judgement

$$\vdash \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}}$$

gives us functions

$$\xi : \mathbb{N} \rightarrow \{0, \dots, n-1\} \rightarrow \mathbb{N}$$

$$\psi : \mathbb{N} \rightarrow \{0, \dots, n-1\} \rightarrow \mathbb{N}^{|\Sigma_t|+1} \rightarrow \mathbb{N}$$

such that for each $\lambda \in \mathbb{N}$, we have

$$\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda \text{ count } \xi_\lambda \text{ ctxt } \psi_\lambda \quad (\star)$$

and for each $1 \leq i \leq n$, we have

$$\xi_{(\cdot)}^i = O(\text{poly}(\lambda))$$

$$\psi_{(\cdot)}^i = O(\text{poly}(\lambda, t_1, \dots, t_{|\Sigma_t|}))$$

In particular, there are polynomials $p_{\text{count}}^i(\lambda)$ with $N_{\text{count}}^i \in \mathbb{N}$ such that $\xi_\lambda^i \leq p_{\text{count}}^i(\lambda)$ if $\lambda \geq N_{\text{count}}^i$, and polynomials $p_{\text{ctxt}}^i(\lambda, t_1, \dots, t_{|\Sigma_t|})$ with $N_{\text{ctxt}}^i \in \mathbb{N}$ such that

$$\psi_\lambda^i(t_1, \dots, t_{|\Sigma_t|}) \leq p_{\text{ctxt}}^i(\lambda, t_1, \dots, t_{|\Sigma_t|})$$

if $\lambda \geq N_{\text{ctxt}}^i$ and $t_1, \dots, t_{|\Sigma_t|} \geq N_{\text{ctxt}}^i$.

Since $\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}}$ is PPT, we have a polynomial $C_{\text{sem}}(\lambda)$, a negligible function $\eta_{\text{sem}}(\lambda)$, and a $N_{\text{sem}} \in \mathbb{N}$ such that:

- *for all type symbols t , $|t|_\lambda \leq C_{\text{sem}}(\lambda)$ if $\lambda \geq N_{\text{sem}}$,*
 - *for all function symbols f , $\llbracket f \rrbracket_\lambda$ is computable by a deterministic Turing Machine TM_λ with symbols $0, 1$, and both the number of states and the runtime of TM_λ are $\leq C_{\text{sem}}(\lambda)$ if $\lambda \geq N_{\text{sem}}$,*
- and*

- for all distribution symbols d , $\llbracket d \rrbracket_\lambda$ is computable up to an error $\eta_{\text{sem}}(\lambda)$ by a probabilistic Turing Machine TM_λ with symbols 0, 1, and both the number of states and the runtime of TM_λ are $\leq C_{\text{sem}}(\lambda)$ if $\lambda \geq N_{\text{sem}}$.

To prove

$$\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}} \models \{\Delta_\lambda \vdash P_\lambda \approx Q_\lambda : I_\lambda \rightarrow O_\lambda\}_{\lambda \in \mathbb{N}},$$

assume a polynomial $C_{\text{adv}}(\lambda)$ and a negligible function $\eta_{\text{adv}}(\lambda)$. Since the ambient asymptotic theory is sound under the ambient family of interpretations, we have the computational indistinguishability assumption

$$\{\llbracket - \rrbracket_\lambda\}_{\lambda \in \mathbb{N}} \models \{\Delta_\lambda^i \vdash P_\lambda^i \approx Q_\lambda^i : I_\lambda^i \rightarrow O_\lambda^i\}_{\lambda \in \mathbb{N}}.$$

Define

$$C_{\text{ctxt}}^i(\lambda) := p_{\text{ctxt}}^i(\lambda, C_{\text{sem}}(\lambda) + N_{\text{ctxt}}^i, \dots, C_{\text{sem}}(\lambda) + N_{\text{ctxt}}^i),$$

and apply the computational indistinguishability assumption above to the polynomial

$$\mathcal{P}(C_{\text{sem}}(\lambda), C_{\text{adv}}(\lambda), C_{\text{ctxt}}^i(\lambda))$$

and the negligible function $\max(\eta_{\text{sem}}(\lambda), \eta_{\text{adv}}(\lambda))$.

This yields a negligible function $\varepsilon^i(\lambda)$ with an $N^i \in \mathbb{N}$ such that for any $\lambda \geq N^i$ and any adversary Adv^i for protocols $\Delta_\lambda^i \vdash I_\lambda^i \rightarrow O_\lambda^i$ under the interpretation $\llbracket - \rrbracket_\lambda$ with the property that

$$|\text{Adv}^i| \leq \mathcal{P}(C_{\text{sem}}(\lambda), C_{\text{adv}}(\lambda), C_{\text{ctxt}}^i(\lambda))$$

and $\text{err}(\text{Adv}^i) \leq \max(\eta_{\text{sem}}(\lambda), \eta_{\text{adv}}(\lambda))$, we have

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket_\lambda} P_\lambda^i = 1] - \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket_\lambda} Q_\lambda^i = 1] \right| \leq \varepsilon^i(\lambda).$$

We can now define our desired negligible function as

$$\varepsilon(\lambda) := \sum_{i=1}^n \xi_\lambda^i * (\varepsilon^i(\lambda) + 2 * C_{\text{ctxt}}^i(\lambda) * \eta_{\text{sem}}(\lambda))$$

The negligibility of $\varepsilon(\lambda)$ follows easily: if $\lambda \geq N_{\text{count}}^i$, then

$$\xi_\lambda^i * (\varepsilon^i(\lambda) + 2 * C_{\text{ctxt}}^i(\lambda) * \eta_{\text{sem}}(\lambda)) \leq p_{\text{count}}^i(\lambda) * (\varepsilon^i(\lambda) + 2 * C_{\text{ctxt}}^i(\lambda) * \eta_{\text{sem}}(\lambda)),$$

thus it suffices to show that this latter function is negligible. But this is immediate from the negligibility of $\varepsilon^i(\lambda)$, the negligibility of $\eta_{\text{sem}}(\lambda)$, and the fact that $p_{\text{count}}^i(\lambda)$ and $C_{\text{ctxt}}^i(\lambda)$ are polynomials. Define

$$N := \max(N_{\text{sem}}, N_{\text{ctxt}}^1, \dots, N_{\text{ctxt}}^n, N^1, \dots, N^n).$$

Now assume $\lambda \geq N$ and take any adversary Adv for protocols $\Delta_\lambda \vdash I_\lambda \rightarrow O_\lambda$ under the interpretation $\llbracket - \rrbracket_\lambda$, such that $|\text{Adv}| \leq C_{\text{adv}}(\lambda)$ and $\text{err}(\text{Adv}) \leq \eta_{\text{adv}}(\lambda)$. We aim to show that

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket_\lambda} P_\lambda = 1] - \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket_\lambda} Q_\lambda = 1] \right| \leq \sum_{i=1}^n \xi_\lambda^i * (\varepsilon^i(\lambda) + 2 * C_{\text{ctxt}}^i(\lambda) * \eta_{\text{sem}}(\lambda))$$

But this is precisely the conclusion of Theorem 1 applied to the derivation (\star) . It thus suffices to prove the hypotheses of Theorem 1. Among these, the only non-trivial assumption is

$$\psi_\lambda^i(|t_1|, \dots, |t_{|\Sigma_i|}|) \leq C_{\text{ctxt}}^i(\lambda).$$

We show this via the following sequence of inequalities:

$$\begin{aligned}
& \psi_\lambda^i(|t_1|_\lambda, \dots, |t_{|\Sigma_t|}|_\lambda) \leq \\
& \psi_\lambda^i(C_{\text{sem}}(\lambda), \dots, C_{\text{sem}}(\lambda)) \leq \\
& \psi_\lambda^i(C_{\text{sem}}(\lambda) + N_{\text{ctxt}}^i, \dots, C_{\text{sem}}(\lambda) + N_{\text{ctxt}}^i) \leq \\
& p_{\text{ctxt}}^i(\lambda, C_{\text{sem}}(\lambda) + N_{\text{ctxt}}^i, \dots, C_{\text{sem}}(\lambda) + N_{\text{ctxt}}^i) = C_{\text{ctxt}}^i(\lambda).
\end{aligned}$$

The first inequality follows since the function $\psi_\lambda^i : \mathbb{N}^{|\Sigma_t|} \rightarrow \mathbb{N}$ is monotonically increasing in each argument and $|t_1|_\lambda, \dots, |t_{|\Sigma_t|}|_\lambda \leq C_{\text{sem}}(\lambda)$ by assumption since $\lambda \geq N_{\text{sem}}$. The second inequality is again monotonicity of ψ_λ^i , and the third follows from the definition of p_{ctxt}^i since $\lambda \geq N_{\text{ctxt}}^i$ and $C_{\text{sem}}(\lambda) + N_{\text{ctxt}}^i \geq N_{\text{ctxt}}^i$. \square

F ABSORPTION: PROOF SKETCH FOR LEMMA 2

LEMMA (ABSORPTION). *There exists a polynomial $\mathcal{P}(x, y, z) \geq y$ such that for any*

- *interpretation $\llbracket - \rrbracket$ for Σ for which there are $C_{\text{sem}} \in \mathbb{N}$, $\eta_{\text{sem}} \in \mathbb{Q}_{\geq 0}$ such that*
 - *for all type symbols t , $|t| \leq C_{\text{sem}}$,*
 - *for all function symbols f , $\llbracket f \rrbracket$ is computable by a TM with symbols 0, 1 such that the number of states and the runtime are $\leq C_{\text{sem}}$, and*
 - *for all distribution symbols d , $\llbracket d \rrbracket$ is computable up to error η_{sem} by a probabilistic TM with symbols 0, 1 such that the number of states and the runtime are $\leq C_{\text{sem}}$,*
- *adversary Adv for protocols of type $\Delta \vdash I \rightarrow O_1 \cup O_2$ such that $|\text{Adv}| \leq C_{\text{adv}}$ and $\text{err}(\text{Adv}) \leq \eta_{\text{adv}}$ for some $C_{\text{adv}} \in \mathbb{N}$ and $\eta_{\text{adv}} \in \mathbb{Q}_{\geq 0}$,*
- *protocol $\Delta \vdash Q : I \cup O_1 \rightarrow O_2$,*

we have an adversary $\text{Adv}_{\mathcal{R}}$ for protocols $\Delta \vdash I \cup O_2 \rightarrow O_1$ with

$$|\text{Adv}_{\mathcal{R}}| \leq \mathcal{P}(C_{\text{sem}}, C_{\text{adv}}, \|Q\|(|t_1|, \dots, |t_{|\Sigma_t|}|))$$

and $\text{err}(\text{Adv}_{\mathcal{R}}) \leq \max(\eta_{\text{sem}}, \eta_{\text{adv}})$ such that for any protocol $\Delta \vdash P : I \cup O_2 \rightarrow O_1$ we have

$$\left| \Pr[\text{Adv} \xrightarrow{\llbracket - \rrbracket} P \parallel Q = 1] - \Pr[\text{Adv}_{\mathcal{R}} \xrightarrow{\llbracket - \rrbracket} P = 1] \right| \leq \|Q\|(|t_1|, \dots, |t_{|\Sigma_t|}|) * \eta_{\text{sem}}.$$

SKETCH. Let context $:= \|Q\|(|t_1|, \dots, |t_{|\Sigma_t|}|)$, and let O_1^{\min} be the minimal set of query inputs to Q from among O_1 . In other words, O_1^{\min} contains precisely those channels of O_1 that Q reads from. The reason for replacing O_1 with O_1^{\min} is that we do not have a bound on the size of the former, but the size of the latter is bounded by the number of occurrences of the query annotation “input-to-query”, and this is in turn bounded by context. Let

$$\text{Adv} := (\Delta', I', O', \phi, \#_{\text{rnd}}, \#_{\text{tape}}, \text{Symb}, \text{St}, s_\star, T, \{l_o\}_{o \in I'}, \{O_i\}_{i \in O'}, D)$$

be the adversary for protocols of type $\Delta \vdash I \rightarrow O_1 \cup O_2$. Define the reduced adversary $\text{Adv}_{\mathcal{R}}$ for protocols of type $\Delta \vdash I \cup O_2 \rightarrow O_1$ as follows:

$$\text{Adv}_{\mathcal{R}} := (\Delta', I'_{\mathcal{R}}, O'_{\mathcal{R}}, \phi, \#_{\text{rnd}}^{\mathcal{R}}, \#_{\text{tape}}^{\mathcal{R}}, \text{Symb}^{\mathcal{R}}, \text{St}^{\mathcal{R}}, s_\star^{\mathcal{R}}, T^{\mathcal{R}}, \{l_o^{\mathcal{R}}\}_{o \in I'_{\mathcal{R}}}, \{O_i^{\mathcal{R}}\}_{i \in O'_{\mathcal{R}}}, D^{\mathcal{R}})$$

Here:

- the set of *inputs* is $I'_{\mathcal{R}} := (I' \cup \phi^\star(O_1)) \cup \phi^\star(O_1^{\min})$;
- the set of *outputs* is $O'_{\mathcal{R}} := O' \cup \phi^\star(O_2)$;
- the number of *rounds* is $\#_{\text{rnd}}^{\mathcal{R}} := \#_{\text{rnd}} * \text{context} + \text{context}^2 + 2 * \#_{\text{rnd}} + \text{context} + 1$;
- the number of *tapes* is $\#_{\text{tape}}^{\mathcal{R}} := \#_{\text{tape}} + 1$;

- the set of *symbols* is

$$\text{Symb}^{\mathcal{R}} := \{r \mid 0 \leq r \leq \#_{\text{rnd}}\} \sqcup \text{ProtEncSymb} \sqcup \{\text{"}\Rightarrow\text{"}, \text{"}\#\text{"}\} \sqcup \text{Symb},$$

where ProtEncSymb is the disjoint union of the following sets: $\{\text{"}\text{"}\}$, Punc , KeyWords , the set Σ_f of function symbols declared in Σ , the set Σ_d of distribution symbols declared in Σ , the set

$$\{n \mid 0 \leq n < \text{variable-index bound}(Q)\},$$

and the set

$$\begin{aligned} &\{n \mid 0 \leq n < \text{channel-index bound}(Q)\} \cup \\ &\{m + n \mid m \in \phi^*(I \cup O_1^{\min} \cup O_2) \text{ and } 0 \leq n \leq \text{channel-index bound}(Q)\}; \end{aligned}$$

- the set of *states* is

$$\text{St}^{\mathcal{R}} := \{\text{"}(\text{"} + r + b + \text{prot} + \text{"}\Rightarrow\text{"} + \text{adv} + \text{"})\},$$

where

- the round counter $0 \leq r \leq \#_{\text{rnd}}$ denotes the number of rounds remaining,
- the Boolean $b \in \{0, 1\}$ indicates whether we are processing the adversary code (0) or the absorbed protocol code (1),
- $\text{adv} \in \text{St}$ is the original adversary code, and
- prot is the absorbed protocol code,
- the *initial state* $s_{\star}^{\mathcal{R}}$ sets:
 - $r := \#_{\text{rnd}}$,
 - $b := 1$,
 - $\text{adv} := s_{\star}$, and
 - prot is the encoding of the protocol Q with every channel read from O_1^{\min} annotated with “input-to-query”.

The TM $T^{\mathcal{R}}$ executes the encoded protocol code by first searching for a channel read annotated with “input-to-query”; if it finds one, it updates the annotation to “input-queried” and performs the query. Otherwise all inputs from O_1^{\min} have been queried. We thus enter a second phase, where we search for a reaction that computes; if we find one, we perform the computation. Otherwise the protocol computation terminates and we set the bit b to 0. Our proof yields the following polynomial:

$$\begin{aligned} P(x, y, z) := & y^2 + 8yz + 15z^2 + (|\Sigma_f| + |\Sigma_d| + 1)x + 34y + 47z + \\ & (|\text{Punc}| + |\text{KeyWords}| + |\Sigma_f| + |\Sigma_d| + 161) \end{aligned}$$

□

ACKNOWLEDGMENTS

This work is partially supported by the NGI0 Core Fund (<https://nlnet.nl/core>), a fund established by NLnet (<https://nlnet.nl>) with financial support from the European Commission’s Next Generation Internet programme (<https://ngi.eu>), under the aegis of DG Communications Networks, Content and Technology under grant agreement Nr. 101092990.