

Universal Composability is Robust Compilation

MARCO PATRIGNANI*, University of Trento, Italy

ROBERT KÜNNEMANN, CISPA Helmholtz Center for Information Security, Germany

RIAD S. WAHBY†, Carnegie Mellon University, USA

ETHAN CECCHETTI, University of Wisconsin–Madison, USA

This paper discusses the relationship between two frameworks: universal composability (UC) and robust compilation (RC). In cryptography, UC is a framework for the specification and analysis of cryptographic protocols with a strong compositionality guarantee: UC protocols remain secure even when composed with other protocols. In programming language security, RC is a novel framework for determining secure compilation by proving whether compiled programs are as secure as their source-level counterparts no matter what target-level code they interact with. Presently, these disciplines are studied in isolation, though we argue that there is a deep connection between them and exploring this connection will benefit both research fields.

This paper formally proves the connection between UC and RC and then it explores the benefits of this connection (focussing on perfect, rather than computational UC). For this, this paper first identifies which conditions must programming languages fulfil in order to possibly attain UC-like composition. Then, it proves UC of both an existing and a new commitment protocol as a corollary of the related compilers attaining RC. Finally, it mechanises these proofs in Deepsec, obtaining symbolic guarantees that the protocol is indeed UC.

Our connection lays the groundwork towards a better and deeper understanding of both UC and RC, and the benefits we showcase from this connection provide evidence of scalable mechanised proofs for UC.

In order to differentiate various sub-parts of the UC and RC frameworks, we use syntax highlighting to a degree that colourblind and black&white readers can benefit from [78].

*UC talks about ideal functionalities (typeset in a **verbatim**, emerald font) and protocol implementations (typeset in a **sans-serif**, orange font). RC deals with source languages (typeset in an **italics**, blue font) and target ones (typeset in a **bold**, red font). Elements that are common to each framework are typeset in a black, sans-serif font for UC and in a black, italicised font for RC to avoid repetition.*

For a better experience, please view this paper in colour.

CCS Concepts: • **Security and privacy** → **Cryptography**; • **Theory of computation** → **Program semantics**;

Additional Key Words and Phrases: secure compilation, programming language semantics, universal composability, cryptography, composition, proof techniques, backtranslation

ACM Reference Format:

Marco Patrignani, Robert Künnemann, Riad S. Wahby, and Ethan Cecchetti. 2022. Universal Composability is Robust Compilation. *ACM Trans. Program. Lang. Syst.* 1, 1, Article 1 (January 2022), 64 pages. <https://doi.org/10.1145/3436809>

*Part of this work was conducted while at MPI-SWS, at Cisca, and at Stanford University.

†Part of this work was conducted while at Stanford University.

Authors' addresses: Marco Patrignani, disi, University of Trento, Trento, Italy, mp@cs.stanford.edu; Robert Künnemann, CISPA Helmholtz Center for Information Security, Saarbrücken, Germany, robert.kuennemann@cispa.de; Riad S. Wahby, Carnegie Mellon University, Pittsburgh, PA, USA, riad@cmu.edu; Ethan Cecchetti, University of Wisconsin–Madison, Wisconsin, USA, cecchetti@wisc.edu.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

0164-0925/2022/1-ART1

<https://doi.org/10.1145/3436809>

1 INTRODUCTION

In cryptography, universal composability (UC) is a framework for the specification and analysis of cryptographic protocols with a key guarantee about compositionality [24, 39, 40, 60, 69, 75]. If a cryptographic protocol is proven UC, that protocol behaves like some high-level, secure-by-construction *ideal functionality* no matter what the protocol interacts with. As such, if that protocol is used within a larger protocol, in order to reason about the latter, we can replace the former protocol with its ideal functionality and just reason about the rest. In other words, UC protocols are secure *even when composed* with other protocols.

In programming language security, robust compilation (RC) [12, 14] is a framework for studying secure compilation as the *robust* preservation of classes of hyperproperties [45]. That is, for each class \mathcal{H} of hyperproperties (e.g., safety, hypersafety, subset-closed hyperproperties etc.), RC provides a criterion stating that if attained by a compiler, then the compiler preserves class \mathcal{H} from the source programs it inputs to the target ones it produces. Moreover, this preservation is done robustly, i.e., *no matter* what code is linked against the source and target programs.

These two frameworks (Section 2) belong to different research fields and seem to deal with different notions, however we demonstrate that they are deeply connected. Both frameworks are concerned with abstract notions which are generally *deemed* secure, and more concrete notions whose security must be *proven against arbitrary attackers*. In UC, the abstract notions are called ideal functionalities (F) while concrete ones are called protocols (π). In RC, the abstract notions are called source programs (P), for the source language of the compiler $\llbracket \cdot \rrbracket$, while concrete ones are target programs (P), for the target language of the compiler. Moreover, proving that a protocol UC-realises an idea functionality ($\pi \vdash_{\text{UC}} F$), or that a compiler attains a RC criterion for class \mathcal{H} ($\vdash \llbracket \cdot \rrbracket : \mathcal{H}$), is done by showing that any attack at the concrete level (protocol or target program) can be simulated at the abstract level (functionality or source program).

This paper is the first to witness (and formally prove in Isabelle) the connection between these two frameworks (Section 3). Specifically, we connect UC to the RC criterion that corresponds to *Robust Hyperproperties Preservation* (dubbed *RHP*). We detail all assumptions typically made by the two frameworks and clarify which additional (simple) assumptions need to hold for this connection to hold. Effectively, we demonstrate that proving $\pi \vdash_{\text{UC}} F$ is equivalent to: code F as a program P , code π as a program P , and prove that $\vdash \llbracket \cdot \rrbracket : RHP$ for the compiler that translates P into P .

After presenting this connection, this paper discusses the benefits that both frameworks gain from it. Admittedly, the benefits we identify, and reap, benefit the UC world more than the RC one, in fact the paper focusses on the possibility of obtaining rigorous, scalable, mechanised proofs of UC from RC proofs. For this, we note that UC results are stated in terms of the semantics of Interactive Turing Machines (ITMs), while RC ones are stated in terms of arbitrary source and target languages. Thus, we first identify which conditions arbitrary programming languages must fulfil in order to be used to attain UC proofs from RC ones (Section 4). These conditions formalise the behaviour of the linking operators of such languages, and they impose some constraint on the languages' semantics. In addition, these conditions let us derive common UC corollaries that simplify the mechanisation of RC proofs.

We then present one language that we prove to fulfil these conditions: the *reactive, interactive λ -calculus* (RILC, Section 5). RILC extends the interactive λ -calculus of Liao et al. [73] with a module system and a trace semantics, which are needed in order to carry out RC proofs. The module system is needed to clearly demarcate the boundary between protocol and attacker (resp. functionality and simulator). The trace semantics describes the behaviour of a protocol linked with an attacker from the perspective of an external observer (i.e., the environment perspective, in UC terms), as

78 commonly done in reactive systems formalisations [18]. Together, these additions yield a language
 79 where we can express all the concepts required by UC, but still perform proofs in a *RC* style,
 80 leveraging the formal language semantics.

81 Next, we demonstrate how the semantics-based proof techniques used for *RC* proofs can be
 82 leveraged to obtain rigorous, scalable UC proofs (Section 6). We discuss structural properties and
 83 automation of both UC and *RC* proofs in the literature. This includes formalisations of UC that
 84 are mechanized and allow for rigorous proofs, but could so far not exploit the connection to *RC*.
 85 For this we take a protocol for one-time commitments π_{comm} [41, 73] and prove that it UC-realises
 86 some ideal functionality F_{comm} . We prove this both with static and with dynamic corruption models;
 87 to the best of our knowledge the former is a known result [41, 73] but the latter is novel. By
 88 relying on our connection, we code the ideal functionality and the protocols as programs (P_{comm}
 89 and P_{comm} respectively) in RILC. We prove that the compiler $[\cdot]$, which translates P_{comm} to P_{comm}
 90 attains *RHP*, so from our connection, π_{comm} UC-realises F_{comm} . We carry out this proof for both the
 91 static and the dynamic corruption models, and both proofs rely on simplifications of a trace-based
 92 backtranslation, a semantics-based proof technique often used in *RC* proofs [11, 12, 14, 51, 79, 82, 83,
 93 85]. While similar proof approaches for UC have been used in the literature (e.g., in computer-aided
 94 cryptography), our connection is the first to formally bridge the gap between the UC setting
 95 using ITMs and that in which other results are stated. The proofs that we show imply *RC* proof
 96 techniques, but they nevertheless follow a pattern that is already known in the literature.

97 Finally, we mechanise both these proofs in DEEPSEC [44], a fully automated protocol verification
 98 tool for privacy properties. We translate P_{comm} and P_{comm} in the languages of DEEPSEC (which is
 99 very close to RILC) and use the tool to prove that the programs are trace equivalent, so in turn, the
 100 compiler from the former to the latter is *RHP*.

101 The connection between UC and programming languages semantics had already been con-
 102 jectured, but the programming-languages counterpart had been (wrongly) identified [58] to be
 103 fully-abstract compilation [1, 80] (interestingly, the first formal criterion for secure compilation).
 104 We discuss this conjecture and other elements of our connection in Section 8, before presenting
 105 related work (Section 9) and concluding (Section 10).

106 Concretely, this paper makes the following contributions:

- 107 • it proves in Isabelle that (under some simple conditions), UC coincides with *RHP*, the hyper-
 108 property preservation instance of *RC*;
- 109 • in doing so, it disproves an existing conjecture [58] claiming that UC is a different secure
 110 compilation criterion called fully-abstract compilation;
- 111 • it provides the conditions for arbitrary languages to be used for our connection;
- 112 • it formalises RILC, the reactive, interactive λ -calculus and proves it fulfils the aforementioned
 113 conditions (together with other useful properties);
- 114 • it proves a one-time commitment protocol to UC-realise some ideal functionality under both
 115 static and dynamic corruption models by proving the compiler from functionality to protocol
 116 is *RHP*;
- 117 • it mechanises these proofs in DEEPSEC, obtaining symbolic UC guarantees.

118 For the sake of simplicity, we delegate much of the auxiliary formalism to the technical report [86],
 119 where the interested reader can find full language formalisation, auxiliary lemmas and proofs.
 120 Mechanised proofs for the main results (in Isabelle) and for the commitment protocol (in DEEPSEC)
 121 can be found at our project page:

122 <https://uc-is-sc.github.io/>

123 2 BACKGROUND: UC AND RC

124 This section presents the UC (Section 2.1) and RC (Section 2.2) frameworks and the details that we
125 rely on to define our connection between them.

126 *A Note on Abstraction.* There exist several UC frameworks and a lot of programming languages
127 that can fit the RC framework. In the following, the details of both frameworks (e.g., the choice of
128 interactive Turing machines or the language semantics) are deliberately abstract and handled in an
129 axiomatic manner. Such an abstract treatment lets us derive the connection in the most general
130 sense, but it may leave the knowledgeable reader concerned about some technical details. These
131 technical details are handled and discussed later on, when we show how concrete UC frameworks
132 and what real languages fulfil these axioms and provide a concrete instance of our connection.

133 2.1 Universal Composability

134 In UC [37, 39] the involved parties form a collection of *Interactive Turing Machines* (ITM), that
135 is, Turing Machines that operate concurrently and that have shared tapes between each other.
136 Crucially: in a collection of ITMs, only a single machine operates at a time, so computation happens
137 under sequential consistency.

138 *ITMs.* The framework of UC talks about entities in two distinct worlds: the *real* and the *ideal*
139 ones. The entities of each world comprise exactly the same roles: one encodes the protocol of
140 interest, one models the adversary and one that models an environment providing inputs and
141 observing behaviour. In the real world, these entities are: the protocol π , the adversary A and the
142 environment Z . In the ideal world, these entities are: the ideal functionality F , the simulator S and
143 the (same) environment Z . In either world, the three entities form a *network* that executes according
144 to the semantics of ITMs.

145 *ITMs Semantics.* ITMs have a deterministic semantics whose source of randomness comes from
146 a specific tape called the random tape that any machine in the network can read from. The final
147 outcome of the interaction of a network of ITMs (as perceived by the environment) is (wlog) a
148 single bit 0/1 [37, §1.1]. Thus, the semantics of a network of ITMs (e.g., π, A, Z) can be expressed
149 as a random variable $\text{EXEC}(\pi, A, Z)$ whose image is a single bit.

150 Oftentimes, e.g., in UC proof arguments, it is useful to describe the sequence of messages
151 exchanged by the entities in the network. We call this sequence of messages a *trace*, to uniform the
152 lexicon with the one of RC (described in Section 2.2 below). Thus, the trace of a network of ITMs
153 (e.g., π, A, Z) can be expressed as a random variable $\text{EXEC T}(\pi, A, Z)$ whose image is a sequence of
154 messages.

Definition 1 (ITMs Semantics in UC).

$\text{EXEC}(\cdot, \cdot, \cdot) : 0/1$

$\text{EXEC T}(\cdot, \cdot, \cdot) : \text{trace of messages}$

155 The precise definition of EXEC T and EXEC differs from framework to framework, so in Section 3.2.1
156 we capture their essence with a set of axioms that abstract from the framework-specific (gory)
157 details. We verified that these axioms hold for the UC-like frameworks [36, 37, 73], as we report in
158 Section 8.1.

159 The semantics of ITMs plays a role in the security argument of UC, which is about the real world
160 emulating the ideal one. This means indistinguishability of executions i.e., of the random variable
161 EXEC when run on a real and on an ideal network.

162 *UC-Emulation and Indistinguishability.* UC distinguishes different forms of emulation, the two
 163 most important ones being *perfect* and *computational* emulation [61, 62]. In this work we focus on
 164 perfect emulation, and we leave discussing computational emulation for future work. In perfect
 165 emulation, the real and ideal worlds are considered indistinguishable (\approx) if their semantics are
 166 equally distributed. In general, for two random variables X, Y in the same observation space E (bits
 167 in our case), this means that $X \approx Y \iff \forall e \in E. \Pr[X = e] = \Pr[Y = e]$, where $\Pr[c]$ indicates
 168 the probability of condition c .

169 *UC Security.* We now have all the elements to define the UC security argument (Definition 2). In
 170 UC, a protocol π is secure (denoted with $\pi \vDash_{\text{UC}} F$) if it refines an ideal functionality F that is secure
 171 by construction (aka, the protocol π UC-emulates the functionality F). UC ensures that π is as
 172 secure as F by requiring that any attack on π can be rewritten into an attack on F . This is expressed
 173 via simulation: π emulates F if for any adversary A , there is a simulated adversary S such that, for
 174 any environment Z , the networks (Z, π, A) and (Z, F, S) are indistinguishable.

Definition 2 (UC Perfect Emulation).

$$\pi \vDash_{\text{UC}} F \stackrel{\text{def}}{=} \forall A. \exists S. \forall Z. \text{EXEC}(Z, A, \pi) \approx \text{EXEC}(Z, S, F)$$

175 2.2 Robust Compilation

176 Robust compilation criteria [12, 14] are defined for a generic notion of a source language S and a
 177 target language T so long as they provide a few elements that we now describe.

178 *Formal Languages, Traces and Semantics.* In *RC*, languages must have a notion of partial programs
 179 P (or, components) and of attackers to the programs (or, program contexts) A . The two can be
 180 linked together $A \rightsquigarrow P^1$ and the result is a another program W that we call *complete*. Intuitively,
 181 a complete program is one that has no external dependencies. Notice however that a complete
 182 program can still be linked with other programs that use the code of the former program.² For
 183 example, a program may be the implementation of an encryption library and it can be linked with
 184 an attacker that interacts with said library, yielding a complete program. This complete program
 185 can then be linked within a larger program, say modelling a web browser, that uses the encryption
 186 library. While this setup with complete programs is not canonical in the *RC* setting (which is built
 187 on programming language semantics theory), we justify it later, when reasoning about composition
 188 of programs (and attackers, and complete programs) in Section 4.1.1.

189 Languages must provide a representation of their behaviour through traces $\bar{\tau}$ which are an
 190 infinite sequence of security-relevant events. The trace model considered by Abate et al. [14] is
 191 that of infinite traces performed against an interface *common* to both source and target languages
 192 (e.g., syscalls).³ The trace model also considers prefixes ($\bar{\mu}$), i.e., finite traces.

To talk meaningfully about cryptographic operations, we expect the language semantics to take
 a randomness parameter in input. This, in turn, changes the trace model and traces become a pair
 of a sequence of actions ($\bar{\tau}$, as before) together with the probability for that sequence to happen (ρ).
 Formally, we indicate a trace and a prefix as follows:

$$\text{Interaction Trace } \bar{\tau} ::= \text{infinite sequence of messages} \qquad \text{Trace } \bar{t} ::= (\bar{\tau}, \rho)$$

¹Linking is often denoted as $A[P]$, we use a different notation here to drive a cleaner visual connection with UC.

²We refrain from using the word ‘whole’ since in programming language semantics it has a slightly different connotation than the one of this paper. Essentially, canonical whole programs cannot be extended anymore with other code, while our complete programs can (as we show in Section 4).

³Having a common trace model across language is a common assumption in compiler works [14, 72], though existing work has shown how to lift this assumption [12, 83]. We keep this assumption in our development for the sake of simplicity.

Interaction Prefix $\bar{\mu} ::=$ finite sequence of messages

Prefix $\bar{m} ::= (\bar{\mu}, \rho)$

193 Technically, the trace model may be something different than what we present here, so long as it is
194 something that agrees with the UC trace model.

To connect traces and prefixes, we rely on the prefixing operation, denoted with \leq . Let $\bar{\mu} \leq \bar{\tau}$ be the canonical prefixing operation between prefix $\bar{\mu}$ and trace $\bar{\tau}$, i.e., the finite sequence of messages of $\bar{\mu}$ is contained in the infinite sequence of messages of $\bar{\tau}$. Since traces (and prefixes) are augmented with probabilities, we lift the prefixing operator \leq to our trace model by requiring the trace part to be a prefix and the probability part to be the same.

$$(\bar{\mu}, \rho) \leq (\bar{\tau}, \rho') \stackrel{\text{def}}{=} \bar{\mu} \leq \bar{\tau} \text{ and } \rho = \rho'$$

195 We leave the semantics of our languages abstract, and indicate the reduction relation of the
196 semantics with \rightsquigarrow . However, there must be a way to indicate the behaviour of programs in terms
197 of the set of traces generated by a program according to the semantics of the language:

$$\text{Behav}(W) \stackrel{\text{def}}{=} \{\bar{t} \mid W \rightsquigarrow \bar{t}\}$$

198 and indicate that two programs W_1 and W_2 have the same behaviours as follows:

$$W_1 \simeq W_2 \stackrel{\text{def}}{=} \text{Behav}(W_1) = \text{Behav}(W_2)$$

Since in this paper we assume a common trace model, shared between all languages, note that \simeq can also be used to compare whether programs in different languages have the same behaviour or not. The \simeq relation is reflexive, symmetric and transitive, i.e., it is an equivalence. As an example, below are the traces that describe the behaviour of a program that returns a single random bit once queried:

$$\{(\text{Query?} \cdot \text{Reply } 0!, 1/2), (\text{Query?} \cdot \text{Reply } 1!, 1/2)\}$$

199 The semantics of the languages that we consider (which includes the semantics of most pro-
200 gramming languages, including ITMs) is supposed to have a simple property that, for lack of a
201 better word, we call ‘being ok’. Intuitively, for any program W , a semantics is *ok* if, given that it
202 can produce all prefixes of a trace, then it produces the trace too, with the same probability.

Definition 3 (Ok Semantics).

$$\vdash \rightsquigarrow : \text{ok} \stackrel{\text{def}}{=} \forall W, \bar{t}. \text{ if } (\text{ if } \forall \bar{m}. \bar{m} \leq \bar{t} \text{ then } W \rightsquigarrow \bar{m}) \text{ then } W \rightsquigarrow \bar{t}$$

203 *Hyperproperties.* Hyperproperties [45] are a formal representation of predicates on programs, i.e.,
204 they are predicates on sets of traces. They capture many security-relevant properties including con-
205 ventional safety and liveness (i.e., predicates on traces), as well as properties like non-interference
206 (i.e., predicates on sets of traces). The class of hyperproperties we focus on (for now) is *arbitrary*
207 *hyperproperties* (H).

208 *Compilers.* The last element we need to introduce for RC are compilers. Given a source language
209 S and a target language T , the compiler that translates S components P into T ones is denoted with
210 $\llbracket S \rrbracket_T^S$. When the languages of the compiler are not relevant, we simply write $\llbracket \cdot \rrbracket$.

211 We take compilers to be partial functions, and indicate a compiler to be undefined for certain
212 inputs i as $\llbracket i \rrbracket = \perp$. The criteria we define (and use) in this paper only hold for those inputs for
213 which a compiler is defined.

214 *RHP*. Below we present *RHP* (Definition 4), a criterion that is fulfilled by compilers that pre-
 215 serve subset-closed hyperproperties robustly. Intuitively, a compiler attains *RHP* if the behaviour
 216 ($\mathbf{Behav}(\cdot)$) of the compiled program ($\llbracket P \rrbracket$) linked (\bowtie) with an arbitrary target attacker (A) is the
 217 same as the behaviour ($\mathit{Behav}(\cdot)$) of the source program (P) linked (\bowtie) with a source attacker (A).
 218 If we unfold the definition of behaviour, we obtain that a compiler attains *RHP* if and only if any
 219 prefix (\bar{m}) produced (\rightsquigarrow) by the compiled program linked with an arbitrary target attacker can be
 220 replicated (\rightsquigarrow) by the source program linked with a source attacker.

Definition 4 (Robust HP Preservation).

$$\llbracket \cdot \rrbracket \vdash RHP \stackrel{\text{def}}{=} \mathbf{A}. \exists \mathbf{A}. \mathbf{A} \bowtie \llbracket P \rrbracket \simeq \mathbf{A} \bowtie P$$

$$\text{or equivalently: } \forall P, \mathbf{A}. \exists \mathbf{A}. \mathbf{Behav}(\mathbf{A} \bowtie \llbracket P \rrbracket) = \mathit{Behav}(\mathbf{A} \bowtie P)$$

$$\text{or equivalently: } \forall P, \mathbf{A}. \exists \mathbf{A}. \forall \bar{m}. \mathbf{A} \bowtie \llbracket P \rrbracket \rightsquigarrow \bar{m} \text{ iff } \mathbf{A} \bowtie P \rightsquigarrow \bar{m}$$

221 By looking at the statement of *RHP* it is not clear that amounts to the preservation of all H . Such
 222 a result has been provided by Abate et al. [14], as they prove that *RHP* is equivalent to the statement
 223 of Definition 5 that more clearly amounts to the preservation of all H . We refer the interested
 224 reader to the work of Abate et al. [14] for a proof of the equivalence of the two definitions.

Definition 5 (Robust HP-Preserving Compiler).

$$\llbracket \cdot \rrbracket \vdash RHP \stackrel{\text{def}}{=} \forall H. \forall P. \text{ if } (\forall \mathbf{A}. \mathit{Behav}(\mathbf{A} \bowtie P) \in H) \text{ then } (\forall \mathbf{A}. \mathbf{Behav}(\mathbf{A} \bowtie \llbracket P \rrbracket) \in H)$$

225 Having defined the two frameworks of interest, we now describe the connection between them.

226 3 CONNECTING THE FRAMEWORKS

227 This section first discusses our connection visually and informally (Section 3.1) before formally
 228 proving it (Section 3.2).

229 3.1 Informal Connection




Fig. 1. Connecting UC and *RHP*, visually.

230 Figure 1 presents a visual depiction of our connection. There, (almost) each element in a frame-
 231 work has a corresponding counterpart in the other. Protocols π correspond to compiled programs
 232 $\llbracket P \rrbracket$, ideal functionalities F correspond to source programs P and the environment Z corresponds to
 233 the prefix \bar{m} . Concrete adversaries A correspond to target attackers A and existentially-quantified
 234 simulators S correspond to existentially-quantified source attackers A . However, the connection
 235 between certain elements is non-trivial, and we discuss these cases below.

236 In the following, we use the word *system* to mean one of the four pairs at the bottom of each
 237 triangle in the figure (i.e., what in *RC* we called a complete program). As such, the system is always

238 composed of an adversary (or an attacker, or a simulator) and another entity: the protocol, the
239 functionality, the source program or its compiled counterpart.

240 *The Role of Environments.* The connection between the environment Z and the prefix \bar{m} is most
241 apparent when thinking about their role, which is observing the system and providing input to it.

242 A difference between environments and prefixes is their security role—at least at the intuitive
243 level. In UC, the environment is a *benign* entity, so it is not supposed to provide messages that
244 deviate from the protocol and its observation is assumed to be fair. In *RC*, a prefix is just an
245 objective indication that some reduction has happened in the system, so a prefix cannot be benign
246 or malicious.

247 Fortunately, from the formal perspective, the environment in UC is an ITM that has no way to
248 be restricted to just benign behaviour. Instead, since the environment is universally quantified,
249 it can behave both in a benign and in a malicious way. For this reason, formally, the (malicious)
250 adversary role can be subsumed by the environment, and in UC it is therefore possible to replace
251 the adversary with a dummy forwarder (this is the dummy attacker theorem, as we discuss in
252 Section 4.2).

253 On the other hand, there are languages in which prefixes are not an objective indication of a
254 reduction happening, but they are generated by a computing entity that the system is reacting
255 to: these are called *reactive languages*. Our connection works with both reactive and non-reactive
256 languages, and we discuss how using reactive ones simplifies certain results (such as porting the
257 dummy attacker theorem from UC to *RC*) in Section 8.3. However, we focus on reactive languages
258 since non-reactive ones can have infinite traces, which leads to weaker results than UC [76]

259 *Traces VS Single Bit.* A seeming discrepancy in the connection is the single-bit differentiation
260 output of UC contrasted with a full prefix in *RC*. We argue this is merely a cosmetic difference: in
261 UC all parties exchange messages – whose concatenation forms indeed a prefix – and the bit is a
262 function of those messages (as captured by Axiom 2). It is therefore just for convenience (and a
263 technicality due to dealing with probability distributions) that the environment outputs a single bit
264 in order to tell whether it is interacting with the real or ideal world.

265 *What are Compilers in UC?* In the UC framework there is no element that corresponds to
266 compilers. Instead, the translation of high-level **functionalities** into **protocols** is human made.
267 As such, the typical UC translation risks being imprecise and error prone, and it cannot be automated
268 for large functionalities, or between functionalities with recurring parts.

269 On the other hand, in *RC* the compilers we consider are not what one expects: they do not
270 traverse the program in input and translate its subparts. Instead, they perform a syntactic check
271 on the whole source program: if it is a specific one (i.e., the ideal functionality), then the related
272 protocol is produced by the compiler. In turn, this means that the compilers we consider are not
273 total functions on their input, but partial ones, defined only on those source programs that are
274 ideal functionalities. Notation-wise, we identify such partial functions as $\llbracket P \rrbracket \rightarrow P$.

275 Note that in this work we still do not provide a compiler for the systematic translation of
276 functionalities into protocols. We envision such a compiler exists, for a protocol-specification
277 language, and leave a discussion of the kind of compiler we consider for now for Section 3.2.4.

278 3.2 Formally Proving the Connection

279 The proof strategy we adopt to formally prove the connection is depicted in Figure 2. Our main result
280 is the dashed co-implication on top, which represents the equivalence of UC and *RC* (Theorem 3
281 in Section 3.2.4). We break down that equivalence into two other equivalences, the sloped, full
282 co-implications on the bottom (Theorem 1 and Theorem 2).

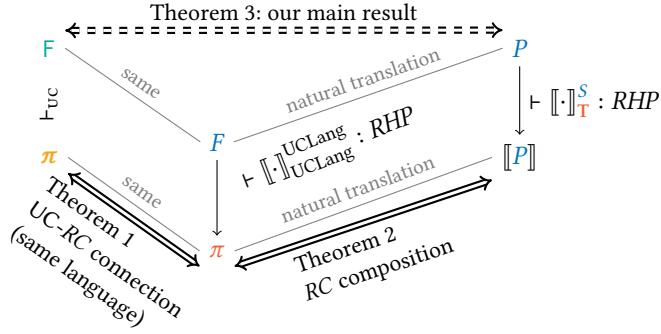


Fig. 2. Visual depiction of the proof of Theorem 3.

283 First we must establish the formal semantics of UC. Despite there being a number of flavours
 284 of UC, we define an abstract language UCLang that captures the essence of all these flavours. We
 285 define the semantics of UCLang via a set of axioms (Section 3.2.1); later (Section 8.1), we show that
 286 many UC flavours respect these axioms.

287 With UCLang, we can prove the connection between UC and RC provided that we are using
 288 UCLang in the RC statements (Theorem 1 in Section 3.2.2). That is, the compiler we consider must
 289 have the same source and target language: UCLang, and we indicate this compiler as $[[\cdot]]_{UCLang}^{UCLang}$.
 290 Since the semantics of ITMs (in UC) and of UCLang (in RC) is the same, in the figure we indicate
 291 that F (resp. π) and F (resp. π) are the ‘same’.

292 Once this connection is proven, we lift the restriction on the languages of the compiler and
 293 consider any compiler between an arbitrary source and an arbitrary target language (Section 3.2.3).
 294 We assume a “natural translation”, i.e., a way to translate programs between S and UCLang and
 295 between T and UCLang that preserves the same behaviour. We argue that such natural translations
 296 already exist, since real-world protocols are not written in ITMs but in programming languages. So
 297 finally, by composing these natural translations with $[[\cdot]]_{UCLang}^{UCLang}$ we obtain $[[\cdot]]_T^S$. By relying on existing
 298 results about the composition of translations [66] we obtain that $[[\cdot]]_T^S$ is also *RHP* (Theorem 2).

299 **3.2.1 Axioms for UC Semantics.** The first axiom relates the UC semantics of ITMs to the one of
 300 UCLang, a programming language semantics in the style of RC, and it defines what we mean when
 301 an adversary and a program produce a trace (Axiom 1).⁴ In this axiom (as well as in subsequent ones)
 302 we assume a function $z(\cdot) : \bar{m} \rightarrow Z$ that assigns each possible trace a *canonical environment*. A
 303 canonical environment for \bar{m} produces exactly this prefix and then halts the execution with the final
 304 bit 1. Such a canonical environment is a known concept in the programming languages research
 305 literature, it is the environment obtained from the backtranslation of a trace [11, 14, 65, 79, 81, 83].
 306 We will discuss this in more detail in Section 6.1.2.

Axiom 1 (UC and UCLang Semantics).

$$A \rightsquigarrow P \rightsquigarrow \bar{m} \text{ iff } \Pr[\text{EXEC T}(z(\bar{m}), A, P) = \bar{\mu}] = \rho > 0 \quad \text{where } \bar{m} = (\bar{\mu}, \rho)$$

307 (Formalized in UC.relation_uc.)

⁴This axiom is typeset in black since it is valid for connecting source and ideal functionality semantics as well as trace and real protocol semantics. Technically both pairs talk about the same semantics since source and target languages here are UCLang and real and ideal world talk about ITMs.

308 Next, prefixes must hold enough information to extract whether the environment has decided
 309 to output a final bit b or not (\perp). In UC [37], the environment terminates the trace with the final
 310 bit, hence any prefix must have it. Abstracting away from the encoding, we assume an extraction
 311 function $\beta : \bar{m} \rightarrow \{0, 1, \perp\}$.

Axiom 2 (Finite Traces Contain the Final Bit).

$$\Pr[\text{EXEC}(Z, A, \pi) = b] = \sum_{\beta(\bar{m})=b} \Pr[\text{EXEC T}(Z, A, \pi) = \bar{\mu}]$$

312 (Formalized in UC.prefix_iff_trace_is_prefix.)

313 To be sure that the canonical environment does not introduce undesired behaviour, we require
 314 that it is correct (Axiom 3). We define correctness to mean $z(\bar{m})$ can produce \bar{m} if any environment
 315 can, and that it distinguishes \bar{m} from all other executions by only outputting the final bit $\beta(\bar{m})$ only
 316 when the environment sees exactly the prefix it is looking for. The side conditions required for this
 317 axiom ensure that the canonical environment exists for a non-probabilistic environment Z that can
 318 produce the prefix at hand.

Axiom 3 (Canonical Environment Correctness). For all non-probabilistic Z that produce some
 prefix \bar{m} , i.e., $\exists A_e, \pi_e. \Pr[\text{EXEC T}(Z, A_e, \pi_e) = \bar{m}] > 0$, we have

$$\begin{aligned} \Pr[\text{EXEC T}(Z, A, \pi) = \bar{m}] &= \Pr[\text{EXEC T}(z(\bar{m}), A, \pi) = \bar{m}] \\ &= \Pr[\text{EXEC}(z(\bar{m}), A, \pi) = \beta(\bar{m})] \end{aligned}$$

319 (Formalized in UC.construct_canonical_env and UC.construct_canonical_env2.)

320 The construction of this canonical environment is straightforward in most UC frameworks: it
 321 matches each incoming input against the expected input in the prefix and hard-codes its output.
 322 Thus, intuitively, the canonical environment reduces the environment to its input/output behaviour.

323 In the axiom we assume Z is non-probabilistic, meaning that the probability of that prefix depends
 324 only on the randomness used in A and π . This is not a real limitation, since non-probabilistic
 325 environments are complete (Axiom 4).

Axiom 4 (Non-Probabilistic Environments are Complete).

if $\text{EXEC}(Z, A, \pi) \not\approx \text{EXEC}(Z, S, F)$ then \exists non-probabilistic $Z_n. \text{EXEC}(Z_n, A, \pi) \not\approx \text{EXEC}(Z_n, S, F)$

326 (Formalized in UC.nonprobabilistic_envs_are_complete.)

327 For UC, Canetti [37, p. 47] states this is the case. Intuitively, if there is a distinguishing environ-
 328 ment, then there is at least one set of random choices that we can condition the environment's
 329 randomness on such that the environment still succeeds in distinguishing both systems. Since the
 330 environment is quantified after the adversary and protocol (Definition 2), we can hard-code these
 331 choices into the environment, making it deterministic.

332 **3.2.2 Connecting UC and RHP Between UCLang.** With the semantics of UC defined via the axioms,
 333 we can prove a simplified version of our connection. If a protocol UC-realises a functionality, then
 334 the UCLang-compiler from that functionality to that protocol is *RHP* (Theorem 1).

Theorem 1 (UC and *RHP* Coincide for UCLang).

$$\llbracket P \rrbracket_{\text{UCLang}}^{\text{UCLang}} \vdash_{\text{UC}} P \text{ iff } \vdash \llbracket \cdot \rrbracket_{\text{UCLang}}^{\text{UCLang}} : \text{RHP}$$

PROOF. (Proven in Isabelle/HOL: Theorem RHPtoUC and Theorem UCtoRHP.)

In both directions, we proceed by contradiction, we only discuss the $RHP \Rightarrow UC$ one since the structure of the two directions is the same. At the intuitive level, the proof relies on the axioms of Section 3.2.1 in order to switch between the EXECUT representation of ITMs semantics (of the UC framework) to the \rightsquigarrow representation of the same semantics (of the RC framework).

$RHP \Rightarrow UC$ By contradiction, we assume the negation of UC (Definition 2), so program P is not emulated by its compiled counterpart $\llbracket P \rrbracket$.

This gives us an adversary A such that for any simulator S , there exists an environment Z that can distinguish real and ideal interactions (negation of Definition 2, point 1).

Thus, without loss of generality, we have a prefix \bar{m} whose generating execution is more probable in the real world than in the ideal one (Axiom 2).

This execution can also be generated by feeding \bar{m} to the canonical environment (Axiom 3), so we eliminate the environment in favour of the canonical one.

We now use Axiom 1 twice, on the real and on the ideal world, in order to translate point 1 into RC. Thus, we obtain the RC counterparts of the UC elements: $A = A$ and $A = S$ such that this holds:

$$A \bowtie \llbracket P \rrbracket \rightsquigarrow \bar{m}$$

but at the same this does not hold:

$$A \bowtie P \rightsquigarrow \bar{m}$$

This is in direct contrast with the RHP assumption (Definition 4): contradiction. \square

3.2.3 Connecting UC and RHP Between Any Language. UC and RHP coincide for programs written in UCLang, but what about other languages? To answer this, we need to take a protocol in UCLang and essentially translate it into any other language without changing its meaning: we call this a *natural translation*.

We argue this natural translations are already done in practice, at an abstract level, and in the following, we formalise the meaning of natural translations. In fact, while UC proofs are carried out between ITMs, real protocol implementations are done in real programming languages. Thus, cryptographers already know how to naturally translate between UCLang and any other programming language.

Two programs P and P are the natural translation of each other if they have the same behaviour robustly.

Definition 6 (Program Natural Translation).

$$P_1 \lesssim P_2 \stackrel{\text{def}}{=} \forall A_1. \exists A_2. Behav(A_1 \bowtie P_1) = Behav(A_2 \bowtie P_2)$$

$$P \sim P \stackrel{\text{def}}{=} P \lesssim P \text{ and } P \lesssim P$$

The notion of natural translation can then be lifted to languages by ensuring that there are natural translation for all of their programs.⁵

Definition 7 (Language Natural Translation).

$$L_1 \lesssim L_2 \stackrel{\text{def}}{=} \forall P_1 \in L_1. \exists P_2 \in L_2. P_1 \sim P_2$$

$$L_1 \sim L_2 \stackrel{\text{def}}{=} L_1 \lesssim L_2 \text{ and } L_2 \lesssim L_1$$

⁵So, the natural translation of two languages can witnessed by the existence of a galois insertion between components of each language whose abstraction and concretisation functions are essentially two RHP compilers between the languages. Moreover, this insertion must induce equivalence between the behaviours of the insertion-related programs. Since it is not the focus of this paper, we leave exploring the interesting ramifications of this fact for future work.

366 The next notion we need is that of language restriction w.r.t. a compiler input or output. Specifi-
 367 cally, given a language L and a compiler $\llbracket \cdot \rrbracket_{\mathbf{T}}^S$, $L|_{\llbracket \cdot \rrbracket_{\mathbf{T}}^S}$ indicates the language obtained by considering
 368 only the programs in the input of $\llbracket \cdot \rrbracket_{\mathbf{T}}^S$, i.e., those programs for which the partial compiler is defined
 369 $\llbracket P \rrbracket_{\mathbf{T}}^S \rightarrow \mathbf{P}$. Dually, $L|_{\llbracket \cdot \rrbracket_{\mathbf{T}}^O}$ indicates the language obtained by considering only the programs in the
 370 output of $\llbracket \cdot \rrbracket_{\mathbf{T}}^S$.

371 The final element of technical machinery we need is a way to compare what we call “robust”
 372 language expressiveness. After all, the results we are setting up span different languages: the source and
 373 target languages of the compiler and UCLang. Therefore, it is unsurprising that we need to
 374 compare language expressiveness to enforce that all languages can represent the same programs.
 375 Formally, we say that UCLang can express all the behaviour of a compiler $\llbracket \cdot \rrbracket_{\mathbf{T}}^S$ if UCLang has
 376 a natural translation with the language the compiler takes in input, and with the language the
 377 compiler produces as output.

Definition 8 (UCLang Related Expressiveness).

$$\text{UCLang} \models \llbracket \cdot \rrbracket_{\mathbf{T}}^S \stackrel{\text{def}}{=} S|_{\llbracket \cdot \rrbracket_{\mathbf{T}}^S} \lesssim \text{UCLang} \text{ and } \mathbf{T}|_{\llbracket \cdot \rrbracket_{\mathbf{T}}^O} \lesssim \text{UCLang}$$

378 Finally, we need a way to relate *RHP* compilers between UCLang and between arbitrary languages
 379 S and T . If a program and its compilation are natural translations of their UCLang counterparts,
 380 then the compiler mapping the program to its compilation is *RHP* if and only if the compiler
 381 mapping their UCLang counterparts is also *RHP*.

Theorem 2 (*RHP* for Arbitrary Languages).

$$\begin{aligned} & \text{if } P \sim P_1 \text{ and } \mathbf{P} \sim P_2 \text{ and } \llbracket P \rrbracket_{\mathbf{T}}^S \rightarrow \mathbf{P} \text{ and } \llbracket P_1 \rrbracket_{\text{UCLang}}^{\text{UCLang}} \rightarrow P_2 \\ & \text{then } \vdash \llbracket \cdot \rrbracket_{\text{UCLang}}^{\text{UCLang}} : \text{RHP} \text{ iff } \vdash \llbracket \cdot \rrbracket_{\mathbf{T}}^S : \text{RHP} \end{aligned}$$

382 **PROOF.** (Proven in Isabelle/HOL: Theorem `implementUC`.) Follows from Theorem 1 (UC and
 383 *RHP* Coincide for UCLang) and from the fact that when programs have natural translations as per
 384 Definition 6, they produce the same behaviour. \square

385 **3.2.4 Our Main Result, Formally.** Joining all these results together yields the formal proof of our
 386 main result: UC and *RHP* coincide. We first provide a specialised version of our result (Theorem 3)
 387 before showing the generalisation (Theorem 4) and wrapping up with a discussion of said results.

388 Both versions rely on program natural translations in order to bridge the gap between a program
 389 (resp. a compiled program) and a functionality (resp. a protocol). The specialised version uses what
 390 may seem like a degenerate compiler (in the programming languages sense). Alas, we believe this
 391 ‘degeneration’ may offer insights into protocol language design, which we discuss right after the
 392 theorem. The general version, instead, relies on Definition 8 to draw the more general result about
 393 *RHP* and UC.

Theorem 3 (UC and *RHP* Coincide).

$$\begin{aligned} & \text{if } P \sim \mathbf{F} \text{ and } \pi \sim \mathbf{P} \text{ and } \llbracket \cdot \rrbracket_{\mathbf{T}}^S \stackrel{\text{def}}{=} \llbracket P \rrbracket_{\mathbf{T}}^S \rightarrow \mathbf{P} \\ & \text{then } \vdash \llbracket \cdot \rrbracket_{\mathbf{T}}^S : \text{RHP} \text{ iff } \pi \vdash_{\text{UC}} \mathbf{F} \end{aligned}$$

394 **PROOF.** (Proven in Isabelle/HOL: Theorem `UC_impl_SingletonRC` and Theorem `RC_impl_UC`.)
 395 The co-implication of this theorem is split into the two-coimplications depicted in Figure 2. From left
 396 to right, we obtain the result from Theorem 2 (*RHP* for Arbitrary Languages) and Theorem 1 (UC and
 397 *RHP* Coincide for UCLang). From right to left, we proceed by contradiction. Given an unsimulatable

398 UC attacker against π , we use the natural translation $P \sim F$ and $\pi \sim P$ to construct a counterexample
 399 against the assumption that $\vdash \llbracket \cdot \rrbracket_{\mathbf{T}}^{\mathbf{S}} : RHP$. \square

400 *A Seemingly-Degenerate Compiler.* The compiler considered in Theorem 3 is not defined induc-
 401 tively on the grammar of source programs, as often done in programming languages work [19,
 402 64, 72, 83]. Instead, it works as an input-output match: it translates those input programs for
 403 which it is defined, into the associated output programs. Ideally, we see such a compiler to have an
 404 input-output pair for all known UC results. Then, since protocols are devised as the combination of
 405 smaller protocols, we believe the kind of compiler we consider in Theorem 3 can be used to compile
 406 larger protocols by composing the compilation of the smaller ones. In essence, this kind of compiler
 407 should consider a special source language that is used to express protocols and their composition;
 408 since the details of such a language escape the scope of this paper, they are left for future work.

409 Our results, however, do not just work for such seemingly-degenerate compilers, as demonstrated
 410 by Theorem 4. Indeed, the UC-RC connection we present works for any compiler that is *RHP*,
 411 provided that the compiler translates the program related to the ideal functionality to the program
 412 related to the protocol.

Theorem 4 (UC and *RHP* Coincide, Generalised).

$$\begin{aligned} & \text{if } \text{UCLang} \models \llbracket \cdot \rrbracket_{\mathbf{T}}^{\mathbf{S}} \\ & \text{then } \vdash \llbracket \cdot \rrbracket_{\mathbf{T}}^{\mathbf{S}} : RHP \text{ iff } \left(\forall P, P', F, \pi. \text{ if } \llbracket P \rrbracket_{\mathbf{T}}^{\mathbf{S}} \rightarrow P' \text{ and } P \sim F \text{ and } \llbracket P \rrbracket_{\mathbf{T}}^{\mathbf{S}} \sim \pi \text{ then } \pi \vdash_{\text{UC}} F \right) \end{aligned}$$

413 **PROOF.** (Proven in Isabelle/HOL: Theorem UC_imp_L_RC and Theorem RC_imp_L_UC.)

414 Half of the proof (\Leftarrow) is identical to Theorem 3 and does not rely on the language expressiveness
 415 assumption. The other half uses very similar logic to the proof of Theorem 3. \square

416 *Discussion of the Results.* A way to summarise our result is thus:

417 \quad A security proof between ITMs is a UC proof,
 418 \quad a security proof between arbitrary languages is a *RHP* one.

419 However, oftentimes (mainly for simplicity) we are interested in setting up the connection for the
 420 same source and target languages, which we now refer to as just S (i.e., $\mathbf{T} = S$). For example, this is
 421 what we do in Section 6. This specialisation requires S and UCLang to have natural translations
 422 (as per Definition 7), so the easiest candidates for S are those that are shown to behave like ITMs.
 423 Since ILC has been proven to model ITMs faithfully, that is why we will choose it and extend it in
 424 order to showcase an instance of our connection.

425 Overall, we believe the connection between either S or T and UCLang is not interesting: cryp-
 426 tographers have been modelling real-world code as ITMs for decades. Thus, we believe the most
 427 important interpretation of Theorem 3 is that one should really focus on the S to T translation and
 428 its security proof.

429 At this point, however, one may wonder whether our connection is as precise as it gets, or
 430 whether there are other *RC* notions (or programming languages research notions, as conjectured
 431 by Hicks [58]) that connect with UC. We will prove that our connection is the most precise one
 432 later on, in Section 8.2.

433 *Using our Connection.* We now have a good intuition that the connection holds, and of what it
 434 connects. What are we to make of this connection then, and how can we reap its benefits?

435 In this paper we reap the benefits for UC, attaining rigorous mechanised UC proofs from *RC*
 436 ones. We leave investigating the benefits for *RC* coming from the UC connection for future work.

437 In a nutshell, the first step now is to instantiate abstract languages S and T . The second step is
438 devising a compiler from some ideal functionality written in S to some concrete protocol written in
439 T , and prove that compiler is *RHP*.

440 In this paper, we set S and T to be the same language: RILC (Section 5). This is the simplest
441 solution that lets us showcase the value of our connection by proving UC of a commitment protocol
442 as a *RHP* proof (Section 6). Building on this language also lets us mechanise this proof (Section 7).

443 A different approach would be to define a specification language (for S) and model the language
444 in which some protocol is implemented (i.e., T) and prove *RHP* at the level of the compiler from
445 S to T . While mechanisation is more complex in this case – we conjecture such proofs could be
446 mechanised in Coq⁶ – its applicability to real-world protocols is clearer. We leave exploring this
447 avenue of results for future work.

448 The procedure we describe above is not novel, however, as some work in computer-aided
449 cryptography use a similar approach in their UC proofs [27, 56]. However, no existing work
450 provides a formal proof that allows them to bridge the gap between their language and the world
451 of UC. Thus we believe our result provides a formal justification for the proof approach of these
452 works.

453 Before we showcase our connection, however, we need to unravel a few more assumptions on
454 the semantics of programming languages. In fact, for now we have imposed very few constraints
455 on the semantics of programming languages involved in the connection. Thus, one may wonder
456 whether, in practice, the connection can be instantiated with realistic (formal) languages. This is
457 what we investigate next in Section 4.

458 4 GENERALISING UC COROLLARIES IN RC

459 In this section, we import two important corollaries from UC that are essential to its use for
460 composing secure protocols: the composition theorem and the dummy attacker one. By proving
461 these corollaries in the RC setting, we identify additional conditions for languages that can be
462 used with our connection. Thus this section first proves the composition theorem in the RC setting
463 (Section 4.1) and formulates the dummy attacker theorem (Section 4.2), providing us with concrete
464 requirements for secure composition in programming languages.

465 4.1 The Composition Theorem in RC

466 As discussed in Section 2.1, UC results always talk about the same language: ITMs. As such,
467 the different types of composition (e.g., *protocol composition* between two protocols and *protocol*
468 *execution* composing adversary, protocol and environment to an executable system [37]) are always
469 between ITMs. Conversely, RC results talk about different languages, so when one tries to derive
470 a composition result, one has to deal with composition of programs in different languages. Thus,
471 porting the composition theorem in RC lets us identify additional conditions that the composition
472 of programs in different languages must fulfil.

473 We state these conditions as a set of axioms that regulate the behaviour of RC composition
474 in order to be able to port a key UC result in RC . Specifically, the UC framework comes with a
475 key free theorem called the *Composition Theorem* (Theorem 5), meaning that any protocol that
476 UC-realises a functionality also enjoys Theorem 5. Intuitively, the composition theorem lets one
477 replace a sub-protocol that is part of a larger protocol with its ideal functionality, which is typically
478 much simpler to reason about. Thus, the composition theorem is the basis for the modular analysis
479 of protocols.

⁶We remark that mechanising proofs of RC (and more in general of secure compilation) is not an easy task, and at the time of writing, very few such examples exist [11, 50].

480 Theorem 5 contains the formal definition of the composition theorem. There, and in the following,
 481 we indicate a larger protocol π_{large} composed with a sub-protocol π_{small} (which could also be an
 482 ideal functionality) as $\pi_{\text{large}}^{\pi_{\text{small}}}$.

Theorem 5 (Composition Theorem in UC [37]).

$$\text{if } \pi_s \vdash_{\text{UC}} F_s \text{ then } \pi_1^{\pi_s} \vdash_{\text{UC}} \pi_1^{F_s}$$

483 The theorem states that if a protocol π_s UC-realises a functionality F_s and π_s is used within a
 484 larger protocol π_1 , then protocol π_1 with π_s UC-realises π_1 with F_s in place of π_{small} .

485 Why Theorem 5 is a corollary of UC-realisation can be explained by reasoning about the role
 486 of the environment. Ideally, the environment represents all possible larger protocols π_1 , so UC-
 487 emulation guarantees that no π_1 can distinguish between its sub-part being π_s or F_s .

488 The key notion for this result is the composition operator, so we first discuss this operator in
 489 UC and in *RC* (Section 4.1.1). Then, we present the set of axioms that regulate the behaviour of
 490 composition (Section 4.1.2). As we demonstrate later, providing a concrete instantiation of the *RC*
 491 composition is fairly simple, and there exists plenty of such instances in the real world; moreover,
 492 equivalently easy is showing that these instances respect the axioms. Finally, we state (and prove)
 493 the analog of Theorem 5 in *RC* (Section 4.1.3).

494 4.1.1 Composition Operators.

495 *UC Composition.* Depending on the UC framework, composition is expressed as a program
 496 transformation [37, 60], where two protocols are inlined into a wrapper/sandbox ITM, or as a
 497 rebinding of input and output channels [36]. In both cases, the composition of π_1 with π_s into a
 498 single protocol $\pi_1^{\pi_s}$ restricts the communication interfaces of both protocols. For example, the
 499 sub-protocol π_s cannot communicate directly with the environment, but any such communication
 500 is routed through π_1 . Additionally, π_s and π_1 have different interfaces to the adversary, that is,
 501 there is an adversary for π_s and one for π_1 . Finally, adversaries for π_s and π_1 do not communicate
 502 directly, but only via the environment.

503 *RC Composition.* When defining *RC* composition we rely on canonical programming language
 504 semantics notions and do not model the communication restriction mentioned above for UC. We do
 505 this for the sake of simplicity and generality, and in order to reuse much of the existing semantics
 506 theory developed for *RC*. We leave investigating the semantics of more intricate composition
 507 mechanisms for future work. Essentially, given the visibility qualifiers of a function defined in
 508 a program (e.g., public, private, etc), composition does not change those qualifiers. In UC,
 509 composition can change these qualifiers, and composing $\pi_1^{\pi_s}$ can turn public interfaces of π_s
 510 into private ones.

511 Recall from Section 2.2 that our languages have a notion of programs P and of attackers A that
 512 can be linked together ($A \rightsquigarrow P$) into a complete program. Since the way we define for *RC* can talk
 513 about multiple languages, we define multiple composition operators in *RC*, each differing in its
 514 signature, as we present below.

515 **linking:** \rightsquigarrow is the already-presented linking. It models intra-language composition and its signature
 516 is $A \rightarrow P \rightarrow W$. In UC, linking corresponds to the combination of a protocol with an
 517 adversary.

518 **program FFI:** \otimes_T^S models cross-language composition of programs, so its signature is $P \rightarrow P \rightarrow$
 519 P . As a notational convention, cross-language composition symbols are annotated with
 520 languages: the top-most is the language of the first argument and of the result, while the
 521 bottom-most is the second argument's.

Intuitively, program FFI composition yields a S program made of a T and an S sub-part (resp. P and P), that can communicate with each other. When real-world languages let programs of different languages S and T interoperate with each other, they rely on an actual foreign function interface between S and T . As such, any concrete implementation of this operator must include one such foreign function interface that marshals and unmarshals (i.e., converts) values of one language into the other and vice-versa. In language semantics, this can be modelled as a multilanguage system [20, 74, 87]. We need not use this model in our instantiation of the composition operators (later in this paper, in Sections 5 and 6), because for simplicity, we resort to using a single language for S and T .

This composition operator also exists in UC, where it is dubbed the protocol composition operator, and it has the same signature. In fact, program FFI composition combines two programs into one program while protocol composition combines two protocols into one protocol. Alternatively, certain UC results are formulated in *hybrid* models [38]. There, both functionality and protocol are composed with a functionality that models a setup assumption. For example, if π emulates F in the F_{CRS} -hybrid model, then π contains the common-reference string functionality to realize F . The composition operator presented here is the one allowing such containment to happen.

attacker FFI: \oplus_T^S models a different cross-language parallel composition, since its signature is $A \rightarrow A \rightarrow A$.

Although many works let attackers and programs have the same structure [4, 37, 95], this is not always the case [14, 21, 22, 55, 83]. For this reason, we need a composition operator for attackers that may differ from the program one. This operator is not always given a name in UC, but appears as a construction in the composition proof (e.g., [37, Fig. 8]).

complete FFI: \otimes_T^S models a cross-language composition for complete programs, so its signature is $W \rightarrow W \rightarrow W$.

In order to link any syntactic category of our abstract languages, program and attacker FFI are not sufficient. Therefore we define the complete FFI, too. It defines the composition of systems (pairs of protocol plus adversaries), where the former becomes the environment to the latter.

In UC, the environment is an ITM that provides inputs to both the adversary and protocol. It ‘represents whatever is external to the current protocol execution. This includes other protocol executions and their adversaries, human users, etc.’ [37]. UC’s environment (this ITM) thus serves as a proxy for the higher-level system and the *protocol execution*, composing environment and a pair of adversary and protocol, can be seen as a system composition.

4.1.2 Conditions for RC Composition. To understand what can be a valid or an invalid instance of these composition operators, we define axioms that regulate their behaviour. These axioms are eventually used to derive *RHP*, and thus they are stated in terms of behavioural equivalence (\simeq from Section 2.2).⁷

Axiom 5 requires that any attacker communicating with a program composed of two sub-programs can be decomposed into two attackers, each in the language of the respective program. The reason we need the two programs is to know into what kind of languages to split the attacker. Had we only a single program, we would not know the language of the second attacker.

⁷We conjecture that proving this connection for computational UC may result in a connection with a different *RC* criterion that may not be *RHP*. As such, we believe the following axioms would need to replace \simeq with whatever is appropriate in that connection. For the sake of generality, our Isabelle formalisation is stated in terms of an arbitrary transitive relation.

564 **Axiom 5** (Attacker Decomposition). $\forall A, P, \mathbf{P}. \exists A', A'.$

$$A \bowtie \left(P \otimes_T^S \mathbf{P} \right) \simeq \left(A' \oplus_T^S A' \right) \bowtie \left(P \otimes_T^S \mathbf{P} \right)$$

565 (Formalized in `composition.par_decomp`.)

566 Axiom 6 states that linking an attacker and a program that are themselves composed using their
567 FFI compositions is equivalent to linking each attacker and program and then using complete FFI
568 composition for the result.

Axiom 6 (FFI Recombination).

$$\left(A \oplus_T^S A \right) \bowtie \left(P \otimes_T^S \mathbf{P} \right) \simeq (A \bowtie P) \odot_T^S (A \bowtie \mathbf{P})$$

569 (Formalized in `composition.inter_comp` and `composition.inter_decomp`.)

570 Finally, we formalise the relation between the behaviour of a program and the composition of
571 this program with a new one. Essentially, we require that any two programs that have the same
572 behaviour can be FFI-composed with *the same* program, and still have the same behaviour.

Axiom 7 (Constant Addition).

$$\text{if } P \simeq \mathbf{P} \text{ then } P \odot_S^O P \simeq P \odot_T^O \mathbf{P}$$

573 (Formalized in `composition.const_elim`.)

574 Technically, Axiom 7 can be proven as a co-implication instead of as an implication, since our
575 RC composition operators do not change the visibility qualifiers within programs and attackers.
576 Since the implication is the required direction for the general composition theorem, we state the
577 weaker axiom.

578 **4.1.3 Deriving the Composition Theorem.** We can now state the RC analogue to Theorem 5 as
579 Theorem 6. Here, in the conclusion, we identify the simulator (or, source-level attacker) with S in
580 order to differentiate it from the adversary (or target-level attacker) A . Apart from this cosmetic
581 change, the theorem states that if a compiler is RHP when linking its target and source with a target
582 attacker and a source simulator, then it is still RHP when linking against an external program (P)
583 and then against an external attacker A and simulator S .

Theorem 6 (Composition in RC).

$$\text{if } \forall A. \exists A. A \bowtie \llbracket P \rrbracket \simeq A \bowtie P \text{ then } \forall A. \exists S. A \bowtie P \otimes_T^O \llbracket P \rrbracket \simeq S \bowtie P \otimes_S^O P$$

584 **PROOF.** (Proven in Isabelle/HOL: Theorem `composition`.) Figure 3 provides a visual depiction
585 of this proof and how each of the presented axioms contribute to proving it (via transitivity of \simeq).
586 Below, we indicate assumption $A \bowtie \llbracket P \rrbracket \simeq A \bowtie P$ with HP .

587 □

588 The proof to this theorem is remarkably compact (ten lines of Isabelle/HOL) which makes it easier
589 to comprehend and separate the framework-specific insights (which are compartmentalised in the
590 operators and their axioms) from the structure of the proof. In Section 5.1, we will demonstrate
591 that instantiating these composition operators and proving these axioms for a formal language is
592 easy. Moreover, it simplifies the proof of the composition theorem, which is usually considered the
593 main result of any work presenting a UC-like framework.




Fig. 3. Visual depiction of the proof of Theorem 6.

594 4.2 The Dummy Attacker in RC

595 The second key UC result we want in *RC* is the dummy adversary theorem (Theorem 7). Intuitively,
 596 Theorem 7 states that instead of considering all possible adversaries, one should consider only
 597 a specific one, the *dummy adversary*. This dummy adversary only forwards messages from the
 598 environment to the protocol (which perceives them as coming from the adversary), and relays
 599 messages from the protocol to the adversary to the environment. The common UC frameworks
 600 define the dummy adversary explicitly, e.g., as a stateless ITM that implements a ‘transparent
 601 channel’ [37, 60] or as set of channels between environment and protocol [36]. In the context of
 602 *RC*, we cannot explicitly define the dummy adversary, as the concrete formulation is dependent on
 603 the programming language. Instead, we will axiomatically define how it interacts with the other
 604 part of the system.

605 **Definition 9** (Dummy Adversary (informal)). In UC, the dummy adversary (A_d) is a proxy that
 606 relays all messages from the environment to the protocol and vice versa.

607 With the dummy adversary, we can state its theorem:

Theorem 7 (Dummy Adversary Theorem in UC [37]).

$$\forall A. \exists S. \forall Z. \text{EXEC}(Z, A, \pi) \approx \text{EXEC}(Z, S, F) \text{ iff } \exists S. \forall Z. \text{EXEC}(Z, A_d, \pi) \approx \text{EXEC}(Z, S, F)$$

608 From a formal perspective it is simple to understand why this theorem holds, the \Rightarrow direction of
 609 the co-implication being trivial and the \Leftarrow direction being the interesting one. In fact, the distinction
 610 that the environment provides honest interactions and the adversary provides malicious ones only
 611 exists in theory, but not in any concrete formalisation of the ITMs semantics. So, since both the
 612 adversary and the environment are ITMs that are \forall -quantified, it is sufficient to keep one entity to
 613 provide inputs and observe outputs (in this case the environment) and replace the other (i.e., the
 614 adversary) with a proxy.

615 The de-facto standard for UC proofs is to prove UC-emulation w.r.t. the dummy adversary and
 616 then apply Theorem 7 to obtain UC-emulation in the sense of Definition 2. Importing this theorem
 617 into the *RC* framework provides one key advantage for the proof of *RHP*. Once the simulator is
 618 defined, proving *RHP* amounts to proving trace equivalence of the compiled program and of the
 619 source program linked with the simulator. The key advantage is that these pieces of code are all
 620 defined, and so there is no need to perform any complicated induction and one can just reason
 621 about the semantics of the source and target programs. Even better, proving trace equivalence can
 622 be mechanised by using existing program analysis tools (as we do later with Deepsec in Section 7)
 623 that verify precisely that property.

624 As before, we first need to identify key axioms that regulate the behaviour (and the composition)
 625 of the dummy attacker in *RC* (Section 4.2.1) before proving Theorem 7 in *RC* (Section 4.2.2).

626 4.2.1 *Additional Conditions for Dummy Attacker in RC.* Borrowing from the proof to the dummy
 627 adversary theorem in UC, we rely on a *dummy program*.

628 **Definition 10** (Dummy Program (informal)). The dummy program is a proxy that relays all
 629 messages from the environment to the program it is program-FFI composed with, and vice versa.

630 Intuitively, the dummy program is the dual of the dummy attacker from Definition 11. However,
 631 the dummy program only plays a role in the proof of the dummy attacker theorem and it does not
 632 appear anywhere else in any other formal statement. Likewise, in UC, the analogue notion appears
 633 only in the dummy adversary theorem, as part of the construction of distinguishing environment,
 634 but is never given a name.⁸

635 From a proof perspective, the dummy program appears as part of the attacker, as captured by
 636 Axiom 8. In fact, the Axiom essentially states that any attacker linked with a program ($A \bowtie P$) can be
 637 separated into an equivalent system with the same program, an attacker, and the dummy protocol
 638 ($A' \bowtie (P_d \otimes_T^T P)$). The second statement is essentially the same, but it allows the re-composition of
 639 an attacker (A) and the dummy program (P_d) into a single attacker (A').

640 **Axiom 8** (Dummy Program Decomposition). $\forall A, P, P'. \exists A', A'$.

$$A \bowtie P \simeq A' \bowtie (P_d \otimes_T^T P) \text{ and } A \bowtie (P_d \otimes_S^T P) \simeq A' \bowtie P$$

641 (Formalized in `dummy.prog_split` and `dummy.prog_unsplit`.)

642 Once the adversary part that is the dummy program has been identified, we can also carve out
 643 the part of the adversary that is the dummy attacker, as in Axiom 9.

644 **Definition 11** (Dummy Attacker (informal)). In *RC*, the dummy attacker (A_d) is a proxy that relays
 645 all messages from the environment to the attacker, and vice versa.

646 **Axiom 9** (Dummy Attacker Decomposition). $\forall A, P. \exists A'$.

$$A \bowtie (P_d \otimes_T^T P) \simeq (A' \circledast_T^T A_d) \bowtie (P_d \otimes_T^T P)$$

647 (Formalized in `dummy.adv_split`.)

648 Essentially, these two axioms state that it is possible to isolate the dummy program and the
 649 dummy attacker from an initial adversary, without altering the behaviour of the system.

650 4.2.2 *Deriving the Dummy Attacker Theorem.* Theorem 8 presents the dummy attacker theorem in
 651 *RC*.

Theorem 8 (Dummy Attacker in *RC*).

$$\forall A. \exists A'. A \bowtie [P] \simeq A' \bowtie P \text{ iff } \exists A'. A_d \bowtie [P] \simeq A' \bowtie P$$

652 **PROOF.** (Proven in Isabelle/HOL: Theorem `dummy`) Figure 4 depicts the non-trivial direction of
 653 this proof (\Leftarrow) and how each of the presented axioms contribute to proving it (via transitivity of
 654 \simeq). Below, we indicate assumption $A_d \bowtie [P] \simeq A' \bowtie P$ with *HP*. One interesting bit is that this proof
 655 relies not just on the axioms presented in this section, but on two of the axioms presented in the
 656 previous one: Axiom 6 (FFI Recombination) and Axiom 7 (Constant Addition). \square

657 We have now identified the set of axioms that allow a *RC* result to enjoy key UC properties
 658 such as the composition theorem and the dummy adversary one. Thus, the next section defines a
 659 language that fulfills those axioms (Section 5), in order to later on derive UC results as *RC* ones
 660 (Section 6).

⁸Not to be confused with UC's *dummy parties*, which have an entirely different purpose.

$$\begin{array}{c}
\begin{array}{c} A \bowtie \llbracket P \rrbracket \\ \downarrow \\ A' \bowtie (P_d \otimes_T^T \llbracket P \rrbracket) \end{array} \left. \vphantom{\begin{array}{c} A \bowtie \llbracket P \rrbracket \\ \downarrow \\ A' \bowtie (P_d \otimes_T^T \llbracket P \rrbracket) \end{array}} \right\} \text{Axiom 8} \\
\hline
\begin{array}{c} A' \bowtie (P_d \otimes_T^T \llbracket P \rrbracket) \end{array} \left. \vphantom{\begin{array}{c} A' \bowtie (P_d \otimes_T^T \llbracket P \rrbracket) \end{array}} \right\} \text{Axiom 9} \\
\begin{array}{c} (A'' \oplus_T^T A_d) \bowtie (P_d \otimes_T^T \llbracket P \rrbracket) \\ \downarrow \\ (A'' \bowtie P_d) \otimes_T^T (A_d \bowtie \llbracket P \rrbracket) \end{array} \left. \vphantom{\begin{array}{c} (A'' \oplus_T^T A_d) \bowtie (P_d \otimes_T^T \llbracket P \rrbracket) \\ \downarrow \\ (A'' \bowtie P_d) \otimes_T^T (A_d \bowtie \llbracket P \rrbracket) \end{array}} \right\} \text{Axiom 6} \\
\begin{array}{c} (A'' \bowtie P_d) \otimes_S^T (A' \bowtie P) \\ \downarrow \\ (A'' \oplus_S^T A') \bowtie (P_d \otimes_S^T P) \end{array} \left. \vphantom{\begin{array}{c} (A'' \bowtie P_d) \otimes_S^T (A' \bowtie P) \\ \downarrow \\ (A'' \oplus_S^T A') \bowtie (P_d \otimes_S^T P) \end{array}} \right\} \text{Axiom 6} \\
\begin{array}{c} (A'' \oplus_S^T A') \bowtie (P_d \otimes_S^T P) \end{array} \left. \vphantom{\begin{array}{c} (A'' \oplus_S^T A') \bowtie (P_d \otimes_S^T P) \end{array}} \right\} \text{Axiom 7 (with HP)} \\
\hline
(A'' \oplus_S^T A') \bowtie (P_d \otimes_S^T P) \left. \vphantom{(A'' \oplus_S^T A') \bowtie (P_d \otimes_S^T P)} \right\} \text{Axiom 8} \\
\hline
A \bowtie P
\end{array}$$

Fig. 4. Visual depiction of the proof of Theorem 8.

661 We believe both the composition theorem and the dummy attacker ones are crucial, and thus
662 one should prove all of the presented axioms in whatever language one chooses for RC . However,
663 note that these axioms are not necessary for the connection to hold, they merely provide additional
664 useful results. In particular, the dummy attacker one simplifies the proofs to such a degree that
665 mechanising them becomes possible using existing tools, as we show in Section 7.

666 5 THE REACTIVE, INTERACTIVE λ -CALCULUS (RILC)

667 This section presents the Reactive, Interactive λ -Calculus (RILC), a reactive language that we use
668 to carry out UC proofs as RC ones. RILC extends the Interactive λ -calculus (ILC), so this section
669 first presents the syntax, typing and operational semantics of ILC, alongside the key properties of
670 the language (Section 5.1). Then this section presents the RILC-proper extensions: a trace model
671 as required by RC and a module system to understand where each logical entity (e.g., protocol,
672 attacker) ends (Section 5.2). Finally, this section showcases the properties of RILC, including the
673 fact that it satisfies the axioms of Section 4 and thus it is possible to use it for our connection
674 (Section 5.3).

675 As the name suggests, RILC is a *reactive* language. Reactive languages differ from non-reactive
676 ones in one key element: they do not have a notion of whole program. In non-reactive languages, a
677 program can run only if it is whole: all of its dependencies are resolved and there are no unknown
678 function symbols. Instead, reactive languages have a built-in notion of a black-box environment
679 that a program can both ‘call’ and ‘be called’ (or send a message to and receive a message from).
680 The black box is therefore a catch-all element for all those symbols that are not resolved within the
681 program.

682 Reactive languages are widely used in formal models for cryptographic protocol verification [18],
683 and in our case they are the easiest kind of program for satisfying the axioms of Section 4. This is
684 why we use a reactive language to instantiate our connection. However, we believe it is possible to
685 instantiate the connection with non-reactive languages too, as we discuss in Section 8.3.

686 5.1 The Interactive λ -Calculus (ILC)

687 In a nutshell, ILC [73] is a process calculus that adapts ITMs to a subset of the π -calculus [91]
688 through an affine typing discipline. To ensure that only one process is active (i.e., it can write) at
689 any given time, processes implicitly pass around an affine “write token”. When a process A writes
690 to another process B , process A must have the write token, and by writing it passes the token to
691 process B , which now has the write token. Moreover, to maintain that the order of activations is

692 fully determined, the read endpoints of channels are (non-duplicable) affine resources, and so each
 693 write operation corresponds to a single, unique read operation. Together, these give ILC its central
 694 metatheoretic property of confluence.

695 This section recaps the elements of ILC we borrow from Liao et al. [73], namely its syntax
 696 (Section 5.1.1), its typing (Section 5.1.2) and its operational semantics (Section 5.1.3). Then, in order
 697 to familiarise the reader with ILC processes, it presents some shorthands for ILC notation that we
 698 use throughout this paper as well as an example ILC process (Section 5.1.4).

699 *5.1.1 Syntax.* The syntax of ILC is best described quoting from Liao et al. [73], starting from the
 700 types of the language.

All Types	$U, V ::= A \mid X$
Sendable Types	$S, T ::= I \mid S \times S \mid S + S$
Unrestricted Types	$A, B ::= S \mid A \times A \mid A + A \mid Wr\ S \mid A \rightarrow_{\infty} U \mid A \rightarrow_w U$
Affine Types	$X, Y ::= !A \mid Rd\ S \mid X \otimes X \mid X \oplus X \mid X \rightarrow_1 U$

701 Types (written U, V) are bifurcated into unrestricted types (written A, B) and affine types (written
 702 X, Y). A subset of the unrestricted types are sendable types (written S, T), i.e., the types of values
 703 that can be sent over channels. This restriction ensures that channels model network channels,
 704 which send only data. The sendable types include unit, products, and sums. The unrestricted types
 705 include the sendable types, products, sums, write endpoint types ($Wr\ S$), arrows, and write arrows
 706 ($A \rightarrow_w U$). Write arrows specify unrestricted abstractions for which the write token can be moved
 707 into the affine context of the abstraction body during reduction. The affine types include bang
 708 types, read endpoint types ($Rd\ S$), products, sums, and arrows.

Labels	$\ell ::= \pi \mid w$	Multiplicity Labels	$\pi ::= I \mid \infty$
Channel Endpoint	$c ::= Read\ (d) \mid Write\ (d)$	Channel Names	$d \in \mathcal{D}$
Values	$v ::= unit \mid \lambda_{\ell} x : U. e \mid \langle v, v \rangle_{\ell} \mid inl_{\ell} v \mid inr_{\ell} v \mid c \mid !v$		
Expressions	$e ::= x \mid v \mid \langle e, e \rangle_{\ell} \mid inl_{\ell} e \mid inr_{\ell} e \mid (e)_{\ell} \mid fix_{\ell} (x.e) \mid let_{\pi} x = e \text{ in } e$ $\mid split_{\ell} (e, x_1.x_2.e) \mid case_{\ell} e \text{ of } inl\ x_1 \mapsto e \mid inr\ x_2 \mapsto e$ $\mid !e \mid ie \mid v(x_1, x_2).e \mid wr(e, e) \mid rd(e, x.e) \mid ch(e, x.e, e, x.e) \mid e(\parallel e)$		

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

715 Values in ILC include unit, lambda expressions, pairs, sums, channel endpoints (written c), and
 716 banged values. We distinguish between the names of channel endpoints ($Read\ (d)$ and $Write\ (d)$)
 717 and the channel d itself that binds them. ILC supports a fairly standard feature set of expressions:
 718 pairs construction, left and right tagging, application, fixpoints, let-bindings, pair destruction, tag
 719 destruction, replication and un-replication, channel creation, writing and reading to channels,
 720 choice over multiple channel reads, and forking of a process.

5.1.2 Typing.

Typing Environment	$\Gamma ::= \emptyset \mid \Gamma, (x : A)$
--------------------	---

Affine Environment $\Delta ::= \emptyset \mid \Delta, (x : X) \mid \Delta, \omega$
 Channel Environment $\Psi ::= \emptyset \mid \Psi, (d : S)$

721 Typing relies on three environments: a classical typing environment that binds variables to unre-
 722 stricted types, an affine environment that binds variables to affine types, and a channel environment
 723 binding channel names to sendable types. Notice that the write token ω lives in the affine context,
 724 though it cannot be bound to any variable. Instead, it flows around implicitly by virtue of where
 725 read and write effects are performed, as captured by the typing rules.

726 The typing judgement

$$\Psi; \Delta; \Gamma \vdash e : U$$

727 reads: “under channel environment Ψ , affine environment Δ , and typing environment Γ , expression e
 728 has type U ”.

729 Most typing rules are standard for process calculi and affine types, we report them below and
 730 refer the reader to Liao et al. [73] for an in-depth discussion of them.

(ILC-Type-rdend) $\Psi(d) = S$	(ILC-Type-wrend) $\Psi(d) = S$	(ILC-Type-unit)
$\Psi; \Delta; \Gamma \vdash \text{Read}(d) : Rd \ S$	$\Psi; \Delta; \Gamma \vdash \text{Write}(d) : Wr \ S$	$\Psi; \Delta; \Gamma \vdash \text{unit} : 1$
(ILC-Type-var-u) $\Gamma(x) = A$	(ILC-Type-var-a) $\Delta(x) = X$	(ILC-Type-pair-u) $\Psi; \Delta_1; \Gamma \vdash e_1 : A_1 \quad \Psi; \Delta_2; \Gamma \vdash e_2 : A_2$
$\Psi; \Delta; \Gamma \vdash x : A$	$\Psi; \Delta; \Gamma \vdash x : X$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \langle e_1, e_2 \rangle_\infty : A_1 \times A_2$
(ILC-Type-pair-s) $\Psi; \Delta_1; \Gamma \vdash e_1 : S_1 \quad \Psi; \Delta_2; \Gamma \vdash e_2 : S_2$	(ILC-Type-pair-a) $\Psi; \Delta_1; \Gamma \vdash e_1 : X_1 \quad \Psi; \Delta_2; \Gamma \vdash e_2 : X_2$	(ILC-Type-pair-a) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \langle e_1, e_2 \rangle_1 : X_1 \otimes X_2$
$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \langle e_1, e_2 \rangle_w : S_1 \times S_2$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \langle e_1, e_2 \rangle_1 : X_1 \otimes X_2$	(ILC-Type-inl-u) $\Psi; \Delta; \Gamma \vdash e : A_1$
(ILC-Type-inl-u) $\Psi; \Delta; \Gamma \vdash e : A_1$	(ILC-Type-inl-s) $\Psi; \Delta; \Gamma \vdash e : S_1$	$\Psi; \Delta; \Gamma \vdash \text{inl}_\infty e : A_1 + A_2$
$\Psi; \Delta; \Gamma \vdash \text{inl}_\infty e : A_1 + A_2$	$\Psi; \Delta; \Gamma \vdash \text{inl}_w e : S_1 + S_2$	(ILC-Type-inl-a) $\Psi; \Delta; \Gamma \vdash e : X_1$
(ILC-Type-inl-a) $\Psi; \Delta; \Gamma \vdash e : X_1$	(ILC-Type-inr-u) $\Psi; \Delta; \Gamma \vdash e : A_2$	$\Psi; \Delta; \Gamma \vdash \text{inl}_1 e : X_1 \oplus X_2$
$\Psi; \Delta; \Gamma \vdash \text{inl}_1 e : X_1 \oplus X_2$	$\Psi; \Delta; \Gamma \vdash \text{inr}_\infty e : A_1 + A_2$	(ILC-Type-inr-s) $\Psi; \Delta; \Gamma \vdash e : S_2$
(ILC-Type-inr-s) $\Psi; \Delta; \Gamma \vdash e : S_2$	(ILC-Type-inr-a) $\Psi; \Delta; \Gamma \vdash e : X_2$	$\Psi; \Delta; \Gamma \vdash \text{inr}_w e : S_1 + S_2$
$\Psi; \Delta; \Gamma \vdash \text{inr}_w e : S_1 + S_2$	$\Psi; \Delta; \Gamma \vdash \text{inr}_1 e : X_1 \oplus X_2$	(ILC-Type-split-u) $\Psi; \Delta_1; \Gamma \vdash e : A_1 \times A_2 \quad \Psi; \Delta_2; \Gamma, x_1 : A_1, x_2 : A_2 \vdash e' : U$
(ILC-Type-split-u) $\Psi; \Delta_1; \Gamma \vdash e : A_1 \times A_2 \quad \Psi; \Delta_2; \Gamma, x_1 : A_1, x_2 : A_2 \vdash e' : U$	(ILC-Type-split-s) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_\infty(e, x_1, x_2, e') : U$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_\infty(e, x_1, x_2, e') : U$
$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_\infty(e, x_1, x_2, e') : U$	$\Psi; \Delta_1; \Gamma \vdash e : S_1 \times S_2 \quad \Psi; \Delta_2; \Gamma, x_1 : S_1, x_2 : S_2 \vdash e' : U$	(ILC-Type-split-s) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_w(e, x_1, x_2, e') : U$
(ILC-Type-split-s) $\Psi; \Delta_1; \Gamma \vdash e : S_1 \times S_2 \quad \Psi; \Delta_2; \Gamma, x_1 : S_1, x_2 : S_2 \vdash e' : U$	(ILC-Type-split-a) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_w(e, x_1, x_2, e') : U$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_w(e, x_1, x_2, e') : U$
$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_w(e, x_1, x_2, e') : U$	$\Psi; \Delta_1; \Gamma \vdash e : X_1 \times X_2 \quad \Psi; \Delta_2, x_1 : X_1, x_2 : X_2; \Gamma \vdash e' : U$	(ILC-Type-split-a) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_1(e, x_1, x_2, e') : U$
(ILC-Type-split-a) $\Psi; \Delta_1; \Gamma \vdash e : X_1 \times X_2 \quad \Psi; \Delta_2, x_1 : X_1, x_2 : X_2; \Gamma \vdash e' : U$	(ILC-Type-case-u) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_\ell e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_1(e, x_1, x_2, e') : U$
$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{split}_1(e, x_1, x_2, e') : U$	$\Psi; \Delta_1; \Gamma \vdash e : A_1 + A_2 \quad \Psi; \Delta_2; \Gamma, x_1 : A_1 \vdash e_1 : U \quad \Psi; \Delta_2; \Gamma, x_2 : A_2 \vdash e_2 : U$	(ILC-Type-case-u) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_\ell e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$
(ILC-Type-case-u) $\Psi; \Delta_1; \Gamma \vdash e : A_1 + A_2 \quad \Psi; \Delta_2; \Gamma, x_1 : A_1 \vdash e_1 : U \quad \Psi; \Delta_2; \Gamma, x_2 : A_2 \vdash e_2 : U$	(ILC-Type-case-s) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_\ell e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$
$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$	$\Psi; \Delta_1; \Gamma \vdash e : S_1 + S_2 \quad \Psi; \Delta_2; \Gamma, x_1 : S_1 \vdash e_1 : U \quad \Psi; \Delta_2; \Gamma, x_2 : S_2 \vdash e_2 : U$	(ILC-Type-case-s) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$
(ILC-Type-case-s) $\Psi; \Delta_1; \Gamma \vdash e : S_1 + S_2 \quad \Psi; \Delta_2; \Gamma, x_1 : S_1 \vdash e_1 : U \quad \Psi; \Delta_2; \Gamma, x_2 : S_2 \vdash e_2 : U$	(ILC-Type-case-s) $\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$
$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$	$\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_w e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U$

$$\begin{array}{c}
 \text{742} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : X_1 \oplus X_2 \quad \Psi; \Delta_2, x_1 : X_1; \Gamma \vdash e_1 : U \quad \Psi; \Delta_2, x_2 : X_2; \Gamma \vdash e_2 : U}{\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{case}_1 e \text{ of } \text{inl } x_1 \mapsto e_1 \mid \text{inr } x_2 \mapsto e_2 : U} \text{(ILC-Type-case-a)} \\
 \text{743} \quad \frac{\Psi; \emptyset; \Gamma, x : A \vdash e : U}{\Psi; \Delta; \Gamma \vdash \lambda_{\infty} x : A. e : A \rightarrow_{\infty} U} \text{(ILC-Type-lam-u)} \quad \frac{\Psi; \emptyset; Wr; \Gamma, x : A \vdash e : U}{\Psi; \Delta; \Gamma \vdash \lambda_w x : A. e : A \rightarrow_w U} \text{(ILC-Type-lam-s)} \\
 \text{744} \quad \frac{\Psi; \Delta, x : X; \Gamma \vdash e : U}{\Psi; \Delta; \Gamma \vdash \lambda_1 x : X. e : X \rightarrow_1 U} \text{(ILC-Type-lam-a)} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : A \rightarrow_{\infty} U \quad \Psi; \Delta_2; \Gamma \vdash e' : A}{\Psi; \Delta_1, \Delta_2; \Gamma \vdash (e e')_{\infty} : U} \text{(ILC-Type-app-u)} \\
 \text{745} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : A \rightarrow_w U \quad \Psi; \Delta_2; \Gamma \vdash e' : A}{\Psi; \Delta_1, \Delta_2, Wr; \Gamma \vdash (e e')_w : U} \text{(ILC-Type-app-s)} \\
 \text{746} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : X \rightarrow_1 U \quad \Psi; \Delta_2; \Gamma \vdash e' : X}{\Psi; \Delta_1, \Delta_2; \Gamma \vdash (e e')_1 : U} \text{(ILC-Type-app-a)} \quad \frac{\Psi; \emptyset; \Gamma, x : A \rightarrow_{\infty} U \vdash e : A \rightarrow_{\infty} U}{\Psi; \Delta; \Gamma \vdash \text{fix}_{\infty} (x.e) : A \rightarrow_{\infty} U} \text{(ILC-Type-fix-u)} \\
 \text{747} \quad \frac{\Psi; \emptyset; \Gamma, x : A \rightarrow_w U \vdash e : A \rightarrow_w U}{\Psi; \Delta; \Gamma \vdash \text{fix}_w (x.e) : A \rightarrow_w U} \text{(ILC-Type-fix-s)} \quad \frac{\Psi; \emptyset, x : X \rightarrow_1 U; \Gamma \vdash e : X \rightarrow_1 U}{\Psi; \Delta; \Gamma \vdash \text{fix}_1 (x.e) : X \rightarrow_1 U} \text{(ILC-Type-fix-a)} \\
 \text{748} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : A \quad \Psi; \Delta_2; \Gamma, x : A \vdash e : U}{\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{let}_{\infty} x = e \text{ in } e' : U} \text{(ILC-Type-let-u)} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : X \quad \Psi; \Delta_2, x : A; \Gamma \vdash e : U}{\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{let}_1 x = e \text{ in } e' : U} \text{(ILC-Type-let-a)} \\
 \text{749} \quad \frac{\Psi; \Delta; \Gamma \vdash e : A}{\Psi; \Delta; \Gamma \vdash !e : !A} \text{(ILC-Type-bang)} \quad \frac{\Psi; \Delta; \Gamma \vdash e : !A}{\Psi; \Delta; \Gamma \vdash ie : A} \text{(ILC-Type-gnab)} \quad \frac{\Psi; \Delta, x_1 : Rd S; \Gamma, x_2 : Wr S \vdash e : U}{\Psi; \Delta; \Gamma \vdash v(x_1, x_2).e : U} \text{(ILC-Type-nu)} \\
 \text{750} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : S \quad \Psi; \Delta_2; \Gamma \vdash e' : Wr S}{\Psi; \Delta_1, \Delta_2, \omega; \Gamma \vdash \text{wr}(e, e') : 1} \text{(ILC-Type-wr)} \\
 \text{751} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e : Rd S \quad \Psi; \Delta_2, \omega, x : !S \otimes Rd S; \Gamma \vdash e' : U \quad \omega \notin \Delta_2}{\Psi; \Delta_1, \Delta_2; \Gamma \vdash \text{rd}(e, x.e') : U} \text{(ILC-Type-rd)} \\
 \text{752} \quad \frac{\Psi; \Delta_1; \Gamma \vdash e_1 : Rd S \quad \Psi; \Delta_2; \Gamma \vdash e_2 : Rd T \quad \Psi; \Delta_3, \omega, x : !S \otimes Rd S \otimes Rd T; \Gamma \vdash e'_1 : U}{\Psi; \Delta_3, \omega, x : !T \otimes Rd S \otimes Rd T; \Gamma \vdash e'_2 : U \quad \omega \notin \Delta_3} \text{(ILC-Type-choice)} \\
 \text{753} \quad \frac{\Psi; \Delta_1, \Delta_2, \Delta_3; \Gamma \vdash \text{ch}(e_1, x.e'_1, e_2, x.e'_2) : U}{\Psi; \Delta_1; \Gamma \vdash e : U \quad \Psi; \Delta_2; \Gamma \vdash e' : V} \text{(ILC-Type-fork)} \\
 \Psi; \Delta_1, \Delta_2; \Gamma \vdash e(\parallel e') : U
 \end{array}$$

5.1.3 Operational Semantics.

Process Names	$p, q \in \mathcal{P}$	Names Sets	$\Sigma ::= \emptyset \mid \Sigma, d \mid \Sigma, p$
Process Pool	$\Pi ::= \emptyset \mid \Pi, p : e$	Configurations	$C ::= \langle \Sigma; \Pi \rangle$
Evaluation Ctx.	$E ::= [\cdot] \mid \langle E, e \rangle_{\ell} \mid \langle v, E \rangle_{\ell} \mid \text{inl}_{\ell} E \mid \text{inr}_{\ell} E \mid (E e)_{\ell} \mid (v E)_{\ell} \mid \text{let}_{\pi} x = E \text{ in } e$ $\mid \text{split}_{\ell}(E, x_1.x_2.e) \mid \text{case}_{\ell} E \text{ of } \text{inl } x_1 \mapsto e \mid \text{inr } x_2 \mapsto e \mid !E \mid iE$ $\mid \text{wr}(E, e) \mid \text{wr}(v, E) \mid \text{rd}(E, x.e) \mid \text{ch}(E, x.e, e, x.e) \mid \text{ch}(e, x.e, E, x.e)$		

The operational semantics of ILC relies on the notion of configuration C , which is a tuple of (existing) channel and process names Σ , and a pool of running and terminated processes Π . The operational

semantics of ILC is a small-step concurrent semantics that relies on a contextual semantics for the reduction of expressions. The following judgements comprise the operational semantics of ILC:

$C \equiv C'$	congruence rules between configurations
$c \rightsquigarrow c'$	channel endpoint communication: c communicates to c'
$C \hookrightarrow C'$	configuration reduction: C takes a step and becomes C'
$e \hookrightarrow_0 e'$	expression primitive reduction: e takes a primitive step and becomes e'

754 The rules of the operational semantics are standard for process calculi; we report them below
755 and refer the reader to Liao et al. [73] for an in-depth discussion of them.

$$\begin{array}{c}
\text{(ILC-eq-conf)} \\
\frac{\Pi \equiv_{\text{perm}} \Pi'}{\langle \Sigma; \Pi \rangle \equiv \langle \Sigma; \Pi' \rangle} \\
\text{(ILC-connect)} \\
\frac{\text{Wr } d \rightsquigarrow \text{Rd } d}{e \hookrightarrow_0 e'} \\
\text{(ILC-Sem-local)} \\
\frac{\langle \Sigma; \Pi, p : E[e] \rangle \hookrightarrow \langle \Sigma; \Pi, p : E[e'] \rangle}{q \notin \Sigma} \\
\text{(ILC-Sem-fork)} \\
\frac{\langle \Sigma; \Pi, p : E[e \parallel e'] \rangle \hookrightarrow \langle \Sigma, q; \Pi, p : E[e], q : e' \rangle}{d \notin \Sigma} \\
\text{(ILC-Sem-nu)} \\
\frac{\langle \Sigma; \Pi, p : E[v(x_1, x_2).e] \rangle \hookrightarrow \langle \Sigma, d; \Pi, p : E[e[\text{Read}(d) / x_1][\text{Write}(d) / x_2]] \rangle}{c_2 \rightsquigarrow c_1} \\
\text{(ILC-Sem-read-write)} \\
\frac{\langle \Sigma; \Pi, p : E[\text{rd}(c_1, x.e)], q : E'[\text{wr}(v, c_2)] \rangle \hookrightarrow \langle \Sigma; \Pi, p : E[e[\langle !v, c_1 \rangle_1 / x]], q : E'[\text{unit}] \rangle}{c \rightsquigarrow c_i \text{ for } i \in 1..2} \\
\text{(ILC-Sem-choice)} \\
\frac{\langle \Sigma; \Pi, p : E[\text{ch}(c_1, x_1.e_1, c_2, x_2.e_2)], q : E'[\text{wr}(v, c)] \rangle \hookrightarrow \langle \Sigma; \Pi, p : E[e_i[\langle !v, c_1, c_2 \rangle_1 / x_i]], q : E'[\text{unit}] \rangle}{C_1 \equiv C'_1 \quad C'_1 \hookrightarrow C'_2 \quad C'_2 \equiv C_2} \\
\text{(ILC-Sem-cong)} \\
\frac{C_1 \hookrightarrow C'_1 \quad C'_1 \hookrightarrow C_2 \quad C_2 \equiv C_2}{C_1 \hookrightarrow C_2} \\
\text{(ILC-Sem-ex-letin)} \quad \text{(ILC-Sem-ex-beta)} \\
\frac{\text{let}_\pi x = v \text{ in } e \hookrightarrow_0 e[v / x] \quad (\lambda_\ell x. e)_\ell \hookrightarrow_0 e[v / x]}{\text{split}_\ell(\langle \langle v_1, v_2 \rangle_\ell, x_1.x_2.e \rangle \hookrightarrow_0 e[v_1 / x_1][v_2 / x_2])} \\
\text{(ILC-Sem-ex-split)} \\
\frac{\text{split}_\ell(\langle \langle v_1, v_2 \rangle_\ell, x_1.x_2.e \rangle \hookrightarrow_0 e[v_1 / x_1][v_2 / x_2])}{\text{case}_\ell \text{ inl}_\ell v \text{ of inl } x_1 \mapsto e \mid \text{inr } x_2 \mapsto e' \hookrightarrow_0 e[v / x_1]} \\
\text{(ILC-Sem-ex-case-l)}
\end{array}$$

$$\begin{array}{c}
 \text{(ILC-Sem-ex-case-r)} \\
 \frac{\text{case}_\ell \text{ inr}_\ell v \text{ of inl } x_1 \mapsto e \mid \text{inr } x_2 \mapsto e' \hookrightarrow_0 e' [v / x_2]}{\text{(ILC-Sem-ex-fix)} \quad \text{(ILC-Sem-ex-gnabang)}} \\
 \frac{\text{fix}_\ell (x.e) \hookrightarrow_0 e[\text{fix}_\ell (x.e) / x]}{i!v \hookrightarrow_0 v}
 \end{array}$$

769 **5.1.4 Shorthands and ILC Example.** The syntax of ILC lacks many useful real-world shorthands. In
 770 this section we present the semantics for the following such useful shorthands:

- 771 • $\text{fwd}(\text{Read}(d), \text{Write}(d'))$ defines a proxy that forwards all messages read on channel d to
 772 channel d' (Rule ILC-eq-fwd);
- 773 • \otimes models the usual binary operations (+, -, =, etc.) on values (Rule ILC-Sem-ex-bop);
- 774 • Rules ILC-Sem-ex-ifte-t and ILC-Sem-ex-ifte-f show if-then-else reductions that can be ob-
 775 tained from case destruction of tagged values;
- 776 • $\text{loop}(e)$ models an unbound loop (Rule ILC-Sem-ex-loop);
- 777 • $\text{loop}(n)(e)$ models a bounded loop that runs for n iterations (Rules ILC-Sem-ex-loopN
 778 and ILC-Sem-ex-loop1).

$$\begin{array}{c}
 \text{(ILC-eq-fwd)} \\
 \frac{\langle \Sigma; \Pi, p : E [\text{fwd}(\text{Read}(d), \text{Write}(d'))] \rangle \equiv \langle \Sigma; \Pi, p : E [\text{rd}(\text{Read}(d), x.\text{wr}(x, \text{Write}(d')))] \rangle}{\text{(ILC-Sem-ex-bop)} \quad \text{(ILC-Sem-ex-ifte-t)} \quad \text{(ILC-Sem-ex-ifte-f)}} \\
 \frac{v'' = \otimes(v, v')}{v \otimes v' \hookrightarrow_0 v''} \quad \frac{\text{if true then } e \text{ else } e' \hookrightarrow_0 e}{\text{(ILC-Sem-ex-loopN)}} \quad \frac{\text{if false then } e \text{ else } e' \hookrightarrow_0 e'}{\text{(ILC-Sem-ex-loop1)}} \\
 \frac{\text{loop}(e) \hookrightarrow_0 e; \text{loop}(e)}{\text{(ILC-Sem-ex-loop)}} \quad \frac{n > 1}{\text{loop}(n)(e) \hookrightarrow_0 e; \text{loop}(n-1)(e)} \quad \frac{\text{loop}(1)(e) \hookrightarrow_0 e}{}
 \end{array}$$

783 We defer presenting an example of ILC code (and its semantics) to Section 6. Note that, in
 784 the examples, we often massage the ILC syntax to enhance readability. For example, we will
 785 write code in an imperative form (more similar to other similar languages used by tools that
 786 perform protocol verification). Moreover, we will write $\text{wr}(\text{Msg } v)$ from $\text{Write}(PQ)$ instead of
 787 writing $\text{wr}(\text{Msg } v, \text{Write}(PQ))$, assuming that messages carry a defining header (in this case, Msg).
 788 Finally, we will write $\text{rd}(\text{Msg } x)$ from $\text{Read}(PQ)$ instead of writing a more complex read expression
 789 that reads on channel end $\text{Read}(PQ)$ and then checks that the read value matches some header
 790 Msg and binds the read value in x .

791 5.2 RILC: Modules, Traces and Cryptographic Additions to ILC

792 This section presents our addition to ILC that define RILC. The first such addition is a series of
 793 cryptographic primitives that are used to describe the hash-based commitment protocol of Section 6
 794 (Section 5.2.1).

795 The second one is a module system that is used to clearly identify the different modules (read,
 796 protocols and sub-protocols) that comprise a RILC program (Section 5.2.2).

797 The final addition is a reactive trace semantics whose purpose is twofold (Section 5.2.3). First,
 798 the reactive semantics lifts the language into a reactive setting, where RILC programs (now called
 799 modules) interact with an unspecified environment. Second, the traces capture any interaction that
 800 happens on the interface between modules and the environment.

801 **5.2.1 Cryptographic Additions.** Vanilla ILC does not contain cryptographic primitives, but the
 802 authors provide some additions to ILC [73, Appendix] in order to model a commitment protocol
 803 (the same protocol we also model later). In this version of RILC, we add cryptographic primitives

804 for symbolic pseudorandom generation with a trapdoor and xor-ing of symbolic values. Our
 805 cryptographic additions are handled in a separate part of the program state, so that they can be
 806 easily extended and replaced with minimal effort for ensuring their additions do not violate any
 807 metatheoretical property.

Security Parameter	$\lambda \in \mathbb{N}$	Random Seed	$n^{rd} ::= b_1 \cdots b_j$
Seals	$\sigma \in \mathcal{S}$	Values	$v ::= \cdots \mid \sigma$
Expressions	$e ::= \cdots \mid \text{takernd} \mid \text{keygen}(e) \mid \text{prg}(e, e) \mid \text{invert}(e, e) \mid \text{xors}(e, e)$		
Evaluation Contexts	$E ::= \cdots \mid \text{keygen}(E) \mid \text{prg}(E, e) \mid \text{prg}(v, E) \mid \text{invert}(E, e)$ $\mid \text{invert}(v, E) \mid \text{xors}(E, e) \mid \text{xors}(v, E)$		

808 To model the output of a symbolic prg, we introduce the notion of seals σ , which are values and
 809 opaque tokens (which have been used to model other cryptographic primitives such as additions in
 810 other work [9, 93, 94]). There are a number of expressions that deal with cryptographic primitives:
 811 taking a λ -long random bitstring, generating a key, prg-hashing a value with a key, inverting the
 812 result of a hash via a trapdoor, and xor-ing values. In general, we assume a security parameter λ
 813 which is an integer and a randomness parameter n^{rd} , which is a string of bits whose length is a
 814 polynomial function of λ .

815 *Typing of Cryptographic Additions.* Typing cryptographic additions is straightforward (and
 816 therefore omitted), since they do not pass the write token around. The only meaningful change to
 817 the typing is that *takernd* needs to know the length of the security parameter in order to tell the
 818 type of the bitstring it returns. Thus, typing rules need to be extended to carry around λ . Since this
 819 is a minor, tedious addition that mostly clutters the notation, we elide that annotation in subsequent
 820 judgements.

Operational Semantics of Cryptographic Additions.

Cryptographic Stores	$G ::= \emptyset \mid G, \text{xor}(\sigma, v, v) \mid G, \langle v, v', \langle \sigma, \perp, \perp \rangle \rangle \mid G, \langle v, v', \langle \sigma, v'', \sigma' \rangle \rangle$
Leaked Knowledge	$H ::= \emptyset \mid H, v$
Cryptographic Configuration	$K ::= \langle H; G; \lambda; n^{rd} \rangle$

821 In order to present the semantics of the cryptographic expressions, we need to define additional
 822 elements of the runtime state that the operational semantics rules rely on.

823 Cryptographic stores G are a store needed to collect any information required by the crypto-
 824 graphic additions. In this work, the store contains: the result of *xors* operations, bindings for a
 825 random number v and a fresh key v' bound to a trapdoor σ that has not been used (thus the two
 826 \perp), and bindings for a key v used to create the prg σ' for a value v'' . Leaked knowledge represents
 827 all values that have been leaked, the specific way in which leakage takes place (i.e., how modules
 828 can leak data to the attacker) is presented later, in Section 5.2.3. Cryptographic configurations K
 829 are the part of a runtime state that contains all cryptography-related additions.

The operational semantics that handles cryptographic additions follows these judgements

$K \triangleright e \hookrightarrow_0 K' \triangleright e'$	Primitive reduction relying on the cryptographic configuration
$\langle K, C \rangle \hookrightarrow \langle K', C' \rangle$	Cryptographic configuration takes a step

830 Below are the semantics rules of the cryptographic additions.

$$\begin{array}{c}
 \text{(RILC-Sem-take)} \\
 831 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad n^{rnd} = b_0 \cdots b_{\lambda-1} \cdots b_j \quad \lambda - 1 \leq j}{n' = b_0 \cdots b_{\lambda-1} \quad n_c^{rnd} = b_\lambda \cdots b_j \quad K' = \langle H; G; \lambda; n_c^{rnd} \rangle} \\
 \hline
 K \triangleright \text{takernd} \hookrightarrow_0 K' \triangleright n' \\
 \text{(RILC-Sem-keygen)} \\
 832 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad G' = G, \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, \perp, \perp \rangle \rangle \quad \text{fresh}(v_{rand}, G)}{\text{fresh}(v_{pubkey}, G) \quad \text{fresh}(\sigma_{trap}, G) \quad K' = \langle H; G'; \lambda; n^{rnd} \rangle} \\
 \hline
 K \triangleright \text{keygen}(v_{rand}) \hookrightarrow_0 K' \triangleright \langle v_{pubkey}, \sigma_{trap} \rangle_\infty \\
 \text{(RILC-Sem-keygen)} \\
 833 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, \perp, \perp \rangle \rangle \in G}{K \triangleright \text{keygen}(v_{rand}) \hookrightarrow_0 K \triangleright \langle v_{pubkey}, \sigma_{trap} \rangle_\infty} \\
 \text{(RILC-Sem-trapdoor)} \\
 834 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad \text{fresh}(\sigma, G) \quad \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, \perp, \perp \rangle \rangle \in G}{G' = G, \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, v_{in}, \sigma \rangle \rangle \setminus \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, \perp, \perp \rangle \rangle \quad K' = \langle H; G'; \lambda; n^{rnd} \rangle} \\
 \hline
 K \triangleright \text{prg}(v_{pubkey}, v_{in}) \hookrightarrow_0 K' \triangleright \sigma \\
 \text{(RILC-Sem-trapdoor-exists)} \\
 835 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, v_{in}, \sigma \rangle \rangle \in G}{K \triangleright \text{prg}(v_{pubkey}, v_{in}) \hookrightarrow_0 K \triangleright \sigma} \\
 \text{(RILC-Sem-invert-true)} \\
 836 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad \exists v_{in}. \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, v_{in}, \sigma \rangle \rangle \in G}{K \triangleright \text{invert}(\langle v_{pubkey}, \sigma_{trap} \rangle_\infty, \sigma) \hookrightarrow_0 K \triangleright \text{true}} \\
 \text{(RILC-Sem-invert-false)} \\
 837 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad \nexists v_{in}. \langle v_{rand}, v_{pubkey}, \langle \sigma_{trap}, v_{in}, \sigma \rangle \rangle \in G}{K \triangleright \text{invert}(\langle v_{pubkey}, \sigma_{trap} \rangle_\infty, \sigma) \hookrightarrow_0 K \triangleright \text{false}} \\
 \text{(RILC-Sem-xor-seal)} \\
 838 \quad \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad K' = \langle H; G'; \lambda; n^{rnd} \rangle}{\text{if } xor(v, _ _) \notin G \text{ then } G' = G, xor(v, v', \sigma) \text{ and } \text{fresh}(\sigma, G)} \\
 \text{if } xor(v, v', \sigma) \in G \text{ then } G' = G \\
 \text{if } xor(v, v'', \sigma'') \in G \text{ and } v' \neq v'' \text{ then } G' = G, xor(v, v', \sigma) \text{ and } \text{fresh}(\sigma, G)} \\
 \hline
 K \triangleright \text{xors}(v, v') \hookrightarrow_0 K' \triangleright \sigma \\
 \text{(RILC-Sem-crypto)} \qquad \text{(RILC-Sem-local)} \\
 839 \quad \frac{K \triangleright e \hookrightarrow_0 K' \triangleright e' \qquad \langle C \rangle \hookrightarrow \langle C \rangle}{\langle K; \Sigma; \Pi, p : E[e] \rangle \hookrightarrow \langle K'; \Sigma; \Pi, p : E[e'] \rangle \qquad \langle K; C \rangle \hookrightarrow \langle K; C \rangle}
 \end{array}$$

840 Rule RILC-Sem-take takes the first λ -long bitstring from the randomness bits n^{rnd} and updates
 841 that accordingly. Rule RILC-Sem-keygen takes a random value v_{rand} and generates a fresh binding
 842 for a key v_{pubkey} with a known trapdoor σ_{trap} . Rule RILC-Sem-keygen returns the existing binding
 843 in case the key generation is called with a known input. Rule RILC-Sem-trapdoor calculates the
 844 fresh prg σ of a value v_{in} with some key v_{pubkey} . Rule RILC-Sem-trapdoor-exists is used in case the
 845 prg of some value and some key has already been calculated, in which case, the result is fetched
 846 from the cryptographic store G . Rule RILC-Sem-invert-true acknowledges a correct inversion of a
 847 prg σ given its key v_{pubkey} and the trapdoor σ_{trap} , while Rule RILC-Sem-invert-false is used when

848 the inversion is used on a prg with either the wrong key or trapdoor. Rule RILC-Sem-xor-seal
849 models a simple, one-directional exclusive or.

850 Rule RILC-Sem-crypto lifts the cryptographic primitive reductions to the cryptographic state
851 while Rule RILC-Sem-local lifts any previous reduction from Section 5.1.3 that does not use crypto-
852 graphic primitives to the cryptographic state.

In the following, we define the initial cryptographic state via function $K_0(\cdot)$, which takes a security parameter and a random bitstring in input and returns the empty cryptographic state.

$$K_0(\lambda, n^{rnd}) \stackrel{\text{def}}{=} \langle \emptyset, \emptyset, \lambda, n^{rnd} \rangle$$

Cryptography for the Environment. As we have hinted (but will only show in Section 5.2.3), RILC is a reactive language, whose model contains a black-box environment that can send messages to the reactive program being run. The cryptographic additions should also let the semantics calculate what is the environment allowed to compute, and this is what the cryptographic environment semantics describes. The judgement for this semantics is:

$$\vdash_{env} K \rightsquigarrow K' \triangleright v \quad \text{From knowledge } K, \text{ the environment produces value } v \\ \text{and augments its knowledge to } K'$$

853 Below are the rules for this judgement, which we keep to a minimum for simplicity.

$$\begin{array}{c} \text{(RILC-Sem-EC-base)} \\ \frac{v \neq \sigma}{\vdash_{env} K \rightsquigarrow K \triangleright v} \\ \text{(RILC-Sem-EC-history)} \\ \frac{K = \langle H; G; \lambda; n^{rnd} \rangle \quad v \in H}{\vdash_{env} K \rightsquigarrow K \triangleright v} \\ \text{(RILC-Sem-EC-take)} \\ \frac{K \triangleright \text{takernd} \hookrightarrow_0 K' \triangleright n'}{\vdash_{env} K \rightsquigarrow K' \triangleright n'} \\ \text{(RILC-Sem-EC-xor)} \\ \frac{\vdash_{env} K \rightsquigarrow K' \triangleright v \quad \vdash_{env} K'' \rightsquigarrow K''' \triangleright v' \quad K'' \triangleright \text{xors}(v, v') \hookrightarrow_0 K''' \triangleright \sigma}{\vdash_{env} K \rightsquigarrow K_f \triangleright \sigma} \\ \begin{array}{c} K''' = \langle H; G; \lambda; n^{rnd} \rangle \\ K_f = \langle H, \sigma; G; \lambda; n^{rnd} \rangle \end{array} \\ \text{(RILC-Sem-EC-keygen)} \\ \frac{\vdash_{env} K \rightsquigarrow K' \triangleright v \quad K' \triangleright \text{keygen}(v) \hookrightarrow_0 K'' \triangleright \langle v_{pubkey}, \sigma_{trap} \rangle_\infty}{\vdash_{env} K \rightsquigarrow K_f \triangleright \langle v_{pubkey}, \sigma_{trap} \rangle_\infty} \\ \begin{array}{c} K'' = \langle H; G; \lambda; n^{rnd} \rangle \\ K_f = \langle H, \sigma_{trap}, v_{pubkey}; G; \lambda; n^{rnd} \rangle \end{array} \\ \text{(RILC-Sem-EC-prg)} \\ \frac{\vdash_{env} K \rightsquigarrow K' \triangleright v \quad \vdash_{env} K'' \rightsquigarrow K''' \triangleright v' \quad K'' \triangleright \text{prg}(v, v') \hookrightarrow_0 K''' \triangleright \sigma}{\vdash_{env} K \rightsquigarrow K_f \triangleright \sigma} \\ \begin{array}{c} K''' = \langle H; G; \lambda; n^{rnd} \rangle \\ K_f = \langle H, \sigma; G; \lambda; n^{rnd} \rangle \end{array} \end{array}$$

858 An environment can generate any value that is a ground value (Rule RILC-Sem-EC-base), any
859 value that has been sent to it in the past (Rule RILC-Sem-EC-history), or any random value from
860 the randomness bitstring (Rule RILC-Sem-EC-take). An environment can also compute the xor of
861 any value it could generate (Rule RILC-Sem-EC-xor) and generate a fresh key (Rule RILC-Sem-EC-
862 keygen), an act that leaks the key trapdoor to the environment knowledge. Finally, an environment
863 can create a fresh prg, so long as the values used in its computation were known to it (Rule RILC-
864 Sem-EC-prg). Given that our treatment of cryptographic elements is symbolic, we do not concern
865 ourselves with modelling a semantics with a computational bound for the environment. We believe
866 lifting this simplification can be done by relying on existing work, e.g., [77].

5.2.2 A Module System.

Modules

$M ::= (D; \Pi; I; X)$

Endpoint Types	$\tau ::= Rd\ S \mid Wr\ S$
Imports, Exports, Definitions	$I, X, D ::= \emptyset \mid I, c : \tau$
Processes	$\Pi ::= \emptyset \mid \Pi, p : rd(Read(d), x.e')$

867 The key element of a module system is the definition of modules (M), which are a tuple consisting
 868 of defined channel names, process definitions, imports and exports. First, any top-level process starts
 869 by reading from a channel, and thus we change the top-level definition of processes accordingly.
 870 For well-typedness reasons, that top-level read channel, alongside any other channel used in the
 871 process of the body form the definitions D of a module. Imports define the dependencies of the
 872 module: those channel names are not defined by the module and thus external modules (or the
 873 environment) will provide an implementation for them. Exports define what a module supplies for
 874 external modules to use, essentially any channel name that is defined but not exported is local to
 875 the module. Any channel end c that appears in definitions, imports or exports carries its endpoint
 876 type, meaning, it describes what kind of channel end is and how it is supposed to be used.

877 In the following, for simplicity, with a slight abuse of notation, we assume that these types
 878 mention sessions, i.e., the full sequence of the type of messages that are to be exchanged on the
 879 channel. Note that it would be possible to simply use a different channel after each message is
 880 communicated, but this would make many an example hard to follow. We leave a full formalisation
 881 of session types for RILC for future work.

882 The operation that happens on modules is linking, denoted with $M_1 \bowtie M_2$ (Rule RILC-Linking).

$$\begin{array}{c}
 \text{(RILC-Linking)} \\
 \frac{M_1 = D_1; \Pi_1; I_1; X_1 \quad M_2 = D_2; \Pi_2; I_2; X_2}{I = I_1 \setminus X_2 \cup I_2 \setminus X_1 \quad X = X_1 \setminus I_2 \cup X_2 \setminus I_1 \quad D = D_1 \cup D_2 \quad \Pi = \Pi_1 \cup \Pi_2} \\
 M_1 \bowtie M_2 = D; \Pi; I; X
 \end{array}$$

884 Module linking does what one expects, it generates another module whose processes and definitions
 885 are the union of the processes and definitions of its sub-modules. Then, the resulting module has
 886 all the imports and exports of its sub-modules, minus those that are fulfilled by the other module.
 887 That is, the resulting module has all imports of the two sub-modules, minus their exports, and all
 888 the exports of the two sub-modules, minus their imports. When performing the set difference (as
 889 in $I_1 \setminus X_2$), we assume that subtracting a write (resp. read) end for a channel remove the respective
 890 read (resp. write) end from the set, i.e., removing $Write(d)$ from $\{Read(d), Read(e)\}$ results in
 891 $\{Read(e)\}$.

892 **Note on Notation:** From the type of linking ($\bowtie : M \times M \rightarrow M$), it is evident that the top-level
 893 notion of partial programs that is of interest is indeed modules. In the previous part of the paper,
 894 we used metavariable P to range over such partial programs. In the case of RILC specifically, we
 895 use M to indicate what was previously denoted with P . We do not rebind the metavariable P for
 896 RILC, which can be then thought as an alternative to M .

Typing of Modules. Modules are typed according to these new judgements.

$\vdash M$	Well-typed module
$D, I \vdash \Pi$	Well-typed process with its defined channels and imports

897 The typing of modules is straightforward (Rule RILC-Type-Module). A module is well-typed if its
 898 processes are, according to the process typing (Rule RILC-Type-Process), which in turns relies on
 899 the expression typing of Section 5.1.2. Note that module and process typing do not tamper with the
 900 write token, nor with the affinity of resources, and thus their addition easily preserve the properties
 901 of ILC.

$$\begin{array}{c}
\text{(RILC-Type-Module)} \\
M = (D; \Pi; I; X) \quad \Pi = \Pi_1, \dots, \Pi_n \quad \bigcap_{i \in 1..n} \Pi_i.p = \emptyset \\
\hline
D = \bigcup_{i \in 1..n} D_i \quad I = \bigcup_{i \in 1..n} I_i \quad X \subseteq D \quad \forall i \in 1..n. D; I \vdash \Pi_i \\
\hline
\vdash M \\
\text{(RILC-Type-Process)} \\
\Pi = p : rd(\text{Read}(d), x.e') \quad \text{Read}(d) : Rd \ S \in D \\
D; \emptyset; \emptyset \vdash \text{Read}(d) : Rd \ S \quad D, I; \emptyset; \emptyset \vdash rd(\text{Read}(d), x.e) : U \\
\hline
D; I \vdash \Pi
\end{array}$$

904 With a slight abuse of notation we write D and I to mean their channel and types that are used to
905 create the Ψ used for typing expressions in the premise of Rule RILC-Type-Process.

906 **5.2.3 A Reactive Trace Semantics.** The reactive trace semantics that makes RILC a reactive language
907 relies on two key notions: reactive programs and traces. A reactive program contains some internal
908 state, which is not directly observable from an external observer and reacts to a sequence of inputs
909 from an environment by producing a sequence of observable outputs. After each input, the program
910 may update its internal state, allowing all past inputs to influence an output. By definition, a reactive
911 program is really a component of a larger system that provides it inputs. In RILC, the reactive
912 program is a module (and we use the two words interchangeably in this section) and the larger
913 system is an unspecified black-box that represents the classical UC environment. The input and
914 output sequences of messages exchanged between the reactive program and the environment are
915 what constitute an interaction trace (from Section 2).

Reactive Configurations	$\Omega ::= \langle \eta; K; \Xi; C \rangle$
Token Tag	$\eta ::= \varnothing \mid \perp$
External Names	$\Xi ::= \emptyset \mid \Xi, c : S$

916 From an operational semantics perspective, dealing with reactive programs changes the notion
917 of configuration, which now becomes a reactive configuration Ω . A reactive configuration keeps
918 track of the write token in the token tag η , in order to know whether the reactive program (\perp) or
919 the environment (\varnothing) has the right to write. Then, it also keeps track of the cryptographic store
920 K , as presented in Section 5.2.1. The reactive configuration then contains a list of external names
921 Ξ , that is, a list of those names that the reactive program relies on and that the environment is
922 supposed to handle the other end of. Finally, the reactive configuration contains the elements of an
923 ILC configuration from Section 5.1.3, namely the existing names Σ and the process pool Π .

Actions	$\tau ::= \text{send } v \text{ on } c? \mid \text{send } v \text{ on } c! \mid \epsilon$	Interaction Traces	$\bar{\tau} ::= \emptyset \mid \tau \cdot \bar{\tau}$
Probabilities	$\rho \in 0..1$	Traces	$\bar{t} ::= (\rho, \bar{\tau})$

924 Traces \bar{t} are pairs of a probability ρ and an interaction trace $\bar{\tau}$. An interaction trace is a sequence
925 of actions τ . Actions include: the environment sending a message v on channel c to the reactive
926 program ($\text{send } v \text{ on } c?$), the reactive program sending a message v on channel c to the environment
927 ($\text{send } v \text{ on } c!$), and a silent action ϵ . The actions use $?$ and $!$ decorations borrowed from the process
928 calculi literature [91] to indicate the “direction” of the message. Note that the semantics will insist
929 that interaction traces start with a $?$ action and then alternate between $!$ and $?$ ones.

Operational Reactive Trace Semantics. The semantics relies on the following judgements, where we recap all previous judgements too, for the sake of completeness.

$\Omega \equiv \Omega'$	Configurations are equivalent, as in Section 5.1.3
$c \rightsquigarrow c'$	c writes where c' reads, as in Section 5.1.3
$e \hookrightarrow_0 e'$	Expression reduction, as in Section 5.1.3
$C \hookrightarrow C'$	Configuration takes a step, as in Section 5.1.3
$\langle K, C \rangle \hookrightarrow \langle K, C' \rangle$	Crypto configuration takes a step, as in Section 5.2.1
$\Omega \xrightarrow{\tau} \Omega'$	Reactive configuration takes a step with action τ
$\Omega \xrightarrow{\bar{\tau}} \Omega'$	Reactive configuration big-steps with interaction trace $\bar{\tau}$

930 Many of the judgements of the semantics we rely on are untouched from previous definitions.
 931 Below are those rules of the semantics that deal with the robustness and the traces.

$$\begin{array}{c}
 \text{(RILC-Sem-Write)} \\
 \frac{c \in \Xi \quad K = \langle H; G; \lambda; n^{md} \rangle \quad K' = \langle H, v; G; \lambda; n^{md} \rangle}{\langle \perp; K; \Xi; \Sigma; \Pi, p : E [wr(v, c)] \rangle \xrightarrow{\text{send } v \text{ on } c!} \langle \omega; K'; \Xi; \Sigma; \Pi, p : E [unit] \rangle} \\
 \text{(RILC-Sem-Read)} \\
 \frac{c \in \Xi \quad \vdash_{env} K \rightsquigarrow K' \triangleright v}{\langle \omega; K; \Xi; \Sigma; \Pi, p : E [rd(c, x.e)] \rangle \xrightarrow{\text{send } v \text{ on } c?} \langle \perp; K'; \Xi; \Sigma; \Pi, p : E [e[\langle v, c \rangle / x]] \rangle} \\
 \text{(RILC-Sem-choice)} \\
 \frac{c_i \in \Xi \quad \vdash_{env} K \rightsquigarrow K' \triangleright v}{\langle \omega; K; \Xi; \Sigma; \Pi, p : E [ch(c_1, x_1.e_1, c_2, x_2.e_2)] \rangle \hookrightarrow \langle \perp; K'; \Xi; \Sigma; \Pi, p : E [e_i[\langle !v, c_1, c_2 \rangle_1 / x_i]] \rangle} \\
 \text{(RILC-Sem-Other)} \\
 \frac{\langle K; \Sigma; \Pi \rangle \hookrightarrow \langle K'; \Sigma; \Pi' \rangle}{\langle \perp; K; \Xi; \Sigma; \Pi \rangle \xrightarrow{\epsilon} \langle \perp; K'; \Xi; \Sigma; \Pi' \rangle} \\
 \text{(RILC-Sem-Error)} \\
 \frac{}{\langle \perp; K; \Xi; \Sigma; \Pi, p : E [error] \rangle \xrightarrow{\bar{\tau}} \langle \omega; K; \Xi; \Sigma; \emptyset \rangle}
 \end{array}$$

937 Rule RILC-Sem-Write states that when the reactive program has the write token ($\eta = \perp$), and it is
 938 writing on a channel c that belongs to the environment ($c \in \Xi$), then the environment can read the
 939 value that is being written, and that value is added to the environment knowledge K' . This is the
 940 way leakage (as left hanging from Section 5.2.1) is generated: values that are written into channels
 941 that belong to external names are effectively leaked to the environment (that is, to the attacker).

942 Dually, Rule RILC-Sem-Read states that when the environment has the write token ($\eta = \omega$),
 943 and the reactive program is reading on a channel c that belongs to the environment, then the
 944 environment can send a value v that the program reads. The value is generated according to
 945 what the environment can compute according to the environment cryptographic reductions of
 946 Section 5.2.1 ($\vdash_{env} K \rightsquigarrow K' \triangleright v$).

947 Rule RILC-Sem-Other relies on the previously defined semantics to execute a reduction within
 948 the reactive program (note that this rule can only be triggered when the program has the write
 949 token).

950 Finally, Rule RILC-Sem-Error halts the computation whenever the *error* expression is executed.

$$\begin{array}{c}
\text{(Trace-Refl)} \\
\frac{}{\Omega \xRightarrow{\epsilon} \Omega} \\
951 \\
\text{(Trace-Action)} \\
\frac{\Omega \xrightarrow{\tau} \Omega'' \quad \Omega'' \xRightarrow{\bar{\tau}} \Omega'}{\Omega \xRightarrow{\tau \cdot \bar{\tau}} \Omega'} \\
\text{(Trace-No-Action)} \\
\frac{\Omega \xrightarrow{\epsilon} \Omega'' \quad \Omega'' \xRightarrow{\bar{\tau}} \Omega'}{\Omega \xRightarrow{\bar{\tau}} \Omega'}
\end{array}$$

Reactive Program Traces. The following judgements define when does a reactive program generate a (set of) traces.

$$\begin{array}{ll}
W; \lambda; n^{rnd} \rightsquigarrow \{\bar{t}\} & \text{Reactive program } W \text{ produces a set of traces } \{\bar{t}\} \\
& \text{with security parameter } \lambda \text{ and randomness } n^{rnd} \\
Behav(W, \lambda) = \{\bar{t}\} & \text{Behaviours (i.e., set of traces } \bar{t}) \text{ of a reactive program } W \\
& \text{with a security parameter } \lambda
\end{array}$$

Before we explain each judgement and their rules, we describe some auxiliary functions.

$$NR(\lambda) = \{n \mid n : [Bits] \text{ and } \|[Bits]\| = f(\lambda) \text{ and } f : \mathbb{N} \rightarrow \mathbb{N} \text{ and } \vdash \text{poly}(f)\}$$

952 $NR(\lambda)$ generates all the random bitstrings $\{n\}$ that a reactive program can use, and the length of
953 these bitstrings is a polynomial function ($\vdash \text{poly}(f)$) of the security parameter λ . In the following,
954 we use $\|\cdot\|$ to indicate the length of a list, or the cardinality of a set.

$$\begin{array}{c}
\text{(RILC-Initial State)} \\
\frac{W = D; \Pi; I; X \quad \vdash W \quad \Xi = I \cup X \quad \Sigma = \text{dom}(\Pi) \cup \text{dom}(\Xi)}{\Omega_0(W; \lambda; n^{rnd}) = \langle \omega; K_0(\lambda, n^{rnd}); \Xi; \Sigma; \Pi \rangle} \\
955 \\
\text{(RILC-Terminal State)} \\
\frac{\nexists \Omega', \tau. \Omega \xrightarrow{\tau} \Omega'}{\vdash \Omega : \text{term}}
\end{array}$$

956 Given a complete (as per the definition of “complete” given in Section 2.2) reactive program W (and
957 the security parameter, and a randomness bitstring n^{rnd}), Rule RILC-Initial State defines the starting
958 state for that program. Dually, Rule RILC-Terminal State tells that a configuration is terminal, i.e.,
959 it cannot step any more.

(RILC-Interaction Traces)

$$\frac{}{W; \lambda; n^{rnd} \rightsquigarrow \left\{ (\bar{\tau}, \rho) \mid \exists \Omega. \Omega_0(W; \lambda; n^{rnd}) \xRightarrow{\bar{\tau}} \Omega \text{ and } \vdash \Omega : \text{term} \text{ and } \rho = 1/(\|NR(\lambda)\|) \right\}}$$

963 Rule RILC-Interaction Traces calculates the set of traces of a complete reactive program. The
964 probability of each trace is calculated as the probability of obtaining the trace over all randomnesses
965 that exist.

(RILC-Interaction Traces)

$$\frac{}{Behav(W, \lambda) = \left\{ (\bar{\tau}, \rho) \mid W; \lambda; n^{rnd} \rightsquigarrow (\bar{\tau}, \rho_0) \text{ and } \rho = \sum_{\rho_0} \text{ and } n^{rnd} \in NR(\lambda) \right\}}$$

967 Rule RILC-Interaction Traces calculates the behaviour of a complete reactive program. The prob-
968 ability ρ of each trace is calculated as the sum of all the probabilities ρ_0 to generate that trace
969 calculated with any possible randomness n^{rnd} taken from the space of randomnesses $NR(\lambda)$.

970 5.3 Properties of RILC

971 RILC enjoys a number of properties, first that its additions from Section 5.2 still preserve well-
972 typedness and confluence from the original ILC work (Section 5.3.1). Then, RILC fulfils the axioms
973 of Section 4, so it can be used in our connection and it enjoys the composition theorem and the
974 dummy adversary one (Section 5.3.2).

975 **5.3.1 Language Properties.** In order to state confluence for RILC we define two projection functions
 976 on traces that extract the inputs ($\cdot|_I$) and outputs ($\cdot|_O$) actions of a trace:

$$\begin{array}{c}
 \text{(Inputs - empty)} \\
 \hline
 \emptyset|_I = \epsilon \\
 \text{(Outputs - empty)} \\
 \hline
 \emptyset|_O = \epsilon
 \end{array}
 \qquad
 \begin{array}{c}
 \text{(Inputs - ?)} \\
 \tau = \text{send } v \text{ on } c? \\
 \hline
 \tau \cdot \bar{\tau}|_I = \tau \cdot \bar{\tau}|_I \\
 \text{(Outputs - ?)} \\
 \tau = \text{send } v \text{ on } c? \\
 \hline
 \tau \cdot \bar{\tau}|_O = \bar{\tau}|_O
 \end{array}
 \qquad
 \begin{array}{c}
 \text{(Inputs - !)} \\
 \tau = \text{send } v \text{ on } c! \\
 \hline
 \tau \cdot \bar{\tau}|_I = \bar{\tau}|_I \\
 \text{(Outputs - !)} \\
 \tau = \text{send } v \text{ on } c! \\
 \hline
 \tau \cdot \bar{\tau}|_O = \tau \cdot \bar{\tau}|_O
 \end{array}$$

979 Additionally, we formalise the fact that during two executions the environment changes the
 980 cryptographic state in the same manner. We say that two such executions are consistent wrt
 981 cryptographic inputs, denote this fact as $\vdash \cdot \parallel \cdot : CIC$ and formalise it as follows:

$$\begin{array}{c}
 \text{(Cryptographic input Consistency - Base)} \\
 \hline
 \vdash \Omega_1 \xrightarrow{\epsilon} \Omega_1 \parallel \Omega_2 \xrightarrow{\epsilon} \Omega_2 : CIC \\
 \text{(Cryptographic input Consistency - Inductive)} \\
 \vdash \Omega'_1 \xrightarrow{\bar{i}} \Omega'_1 \parallel \Omega'_2 \xrightarrow{\bar{i}} \Omega'_2 : CIC \quad \vdash \Omega_1 \xrightarrow{\tau} \Omega'_1 \parallel \Omega_2 \xrightarrow{\tau} \Omega'_2 : CIC \\
 \hline
 \vdash \Omega_1 \xrightarrow{\tau} \Omega'_1 \xrightarrow{\bar{i}} \Omega'_1 \parallel \Omega_2 \xrightarrow{\tau} \Omega'_2 \xrightarrow{\bar{i}} \Omega'_2 : CIC \\
 \text{(Cryptographic input Consistency - Single-nop)} \qquad \text{(Cryptographic input Consistency - Single-in)} \\
 \tau = \epsilon \vee \tau = \tau! \qquad \Omega'_1.K = \Omega'_2.K \\
 \hline
 \vdash \Omega_1 \xrightarrow{\tau} \Omega'_1 \parallel \Omega_2 \xrightarrow{\tau} \Omega'_2 : CIC \qquad \vdash \Omega_1 \xrightarrow{\tau?} \Omega'_1 \parallel \Omega_2 \xrightarrow{\tau?} \Omega'_2 : CIC
 \end{array}$$

985 RILC has a full confluence property (Theorem 9) telling that if a program state Ω performs two
 986 different traces up to two terminal states Ω_1 and Ω_2 with the same environment-supplied inputs,
 987 and with the same changes to the cryptographic state from the environment, then the terminal
 988 states are renamings of each other (with renaming function $g(\cdot)$), and the outputs of the traces are
 989 also the same.

Theorem 9 (Full Confluence).

$$\begin{array}{l}
 \text{if } \Omega \xrightarrow{\bar{i}} \Omega_1 \text{ and } \Omega \xrightarrow{\bar{i}'} \Omega_2 \text{ and } \vdash \Omega_1 : \text{term} \text{ and } \vdash \Omega_2 : \text{term} \\
 \text{and } \bar{i}|_I = \bar{i}'|_I \text{ and } \vdash \Omega \xrightarrow{\bar{i}} \Omega_1 \parallel \Omega \xrightarrow{\bar{i}'} \Omega_2 : CIC \\
 \text{then } \Omega_1 = g(\Omega_2) \text{ and } \vdash g(\Omega_2) : \text{term} \text{ and } \bar{i}|_O = \bar{i}'|_O
 \end{array}$$

990 The proof of Theorem 9 is an unsurprising induction over the generated traces, with the silent
 991 actions case being inherited from Liao et al. [73]. The only interesting case is the single visible
 992 action lemma (Lemma 1).

Lemma 1 (Generalised Confluence Single).

$$\begin{array}{l}
 \text{if } \Omega \xrightarrow{\tau} \Omega_1 \text{ and } f(\Omega) \xrightarrow{\tau'} \Omega_2 \text{ and } \tau|_I = \tau'|_I \text{ and } \Omega_1.K = \Omega_2.K \\
 \text{then } \Omega_1 = g(\Omega_2) \text{ and } \tau|_O = \tau'|_O
 \end{array}$$

993 In the proof of Lemma 1, we have to reason about the rules generating the single actions, which
 994 are the key additions to RILC. Thus, this proof essentially shows that our additions to RILC still
 995 keep the language confluent, assuming that the environment provides the same inputs and the
 996 same changes to the cryptographic state.

997 5.3.2 *Fulfilling the Axioms of Section 4.* The Axioms of Section 4 are split in two groups, those
 998 needed for the composition theorem and those needed for the dummy adversary one.

999 *Axioms for Composition.* In order to state the Axioms of Section 4.1 in RILC, we need to instan-
 1000 tiate the different abstract composition operators (\bowtie , \otimes_T^S , $\oplus_T^S \otimes_T^S$) in RILC. Since RILC has a single
 1001 composition operator (\bowtie from Rule RILC-Linking), we instantiate all the operators with \bowtie . We
 1002 believe in other languages where the distinction between attacker and program needs to be carried
 1003 around (e.g., as in the work of El-Korashy et al. [51]), we could not instantiate all the operators
 1004 with a single one, but we would have to define different ones.

1005 The key property that we rely on for proving the axioms is linking commutativity (Lemma 2),
 1006 whose proof is straightforward and thus left in the technical report [86].

1007 **Lemma 2** (Linking Commutativity). $A \bowtie M \simeq M \bowtie A$

1008 Below are the statements of the axioms instantiated for RILC. As their proofs are tedious and
 1009 not insightful, we leave them in the technical report [86]. First we present the ones related to
 1010 composition:

1011 **Axiom 5** $\forall A_0, M. \exists A_1, A_2. A_0 \bowtie M \simeq (A_1 \bowtie A_2) \bowtie M$

1012 **Axiom 6** $(A_1 \bowtie A_2) \bowtie (M_1 \bowtie M_2) \simeq (A_1 \bowtie M_1) \bowtie (A_2 \bowtie M_2)$

1013 **Axiom 7** if $M_1 \simeq M_2$ then $M_3 \bowtie M_1 \simeq M_3 \bowtie M_2$

1014 From these Axioms, we can derive Corollary 1 below, which is Theorem 6 (Composition in *RC*)
 1015 specified for RILC:

Corollary 1 (Composition for RILC).

$$\text{if } \forall A. \exists A'. A \bowtie [P] \simeq A' \bowtie P \text{ then } \forall A''. \exists S. A'' \bowtie P_S \bowtie [P] \simeq S \bowtie P_S \bowtie P$$

1016 *Axioms for the Dummy Attacker.* The axioms related to the dummy attacker (Section 4.2) cannot
 1017 possibly be stated (let alone proved) in a general fashion. In fact, different modules (which model
 1018 different protocols) provide different signatures, and thus there is not really a single dummy attacker
 1019 that fits all protocols.⁹ However, we notice that once the signature of a module is fixed, and the
 1020 meaning of each imported and exported channels is known, it is possible to devise a dummy attacker
 1021 for that protocol. Even more, we conjecture that for every protocol it is possible to devise a dummy
 1022 attacker (though we leave stating and proving such a conjecture for future work). Thus, we do not
 1023 provide a general proof of Axiom 8 and Axiom 9, but leave them to be proved on a per-protocol
 1024 basis.

1025 Specifically, we will report on the proofs of those Axioms for the specific protocol we consider
 1026 in this paper in Section 6.2.4.

1027 With our candidate formal language for instantiating the connection, we now turn to reaping its
 1028 benefits and show how to derive UC proofs as *RC* ones.

1029 6 RIGOROUS, SCALABLE UC PROOFS AS *RC* PROOFS

1030 This section first provides some background on UC and *RC* proofs and it clarifies how to use our
 1031 connection to provide UC proofs via *RC* ones (Section 6.1). Then it provides two instances of UC
 1032 proofs done via *RC* ones. The first one is for a known result: UC of a single-bit commitment in the
 1033 static corruption setting [41, 73] (Section 6.2). The second one is instead for a simple but (to our
 1034 knowledge) novel result: UC of the same protocol, but with adaptive corruptions (Section 6.3).

⁹We note that there may be some simplifying assumptions that fix the channels in a signature (e.g., one channel per protocol with a unit that is responsible to dispatch the messages to the protocol parties), but that is beyond the scope of this paper.

1035 6.1 Background on UC and RC Proofs

1036 We now recount UC (Section 6.1.1) and RC proofs (Section 6.1.2), before drawing conclusions on
1037 the benefits to reap from each approach (Section 6.1.3).

1038 *6.1.1 UC Proofs.* As mentioned earlier, the vast majority of UC proofs exploits the dummy-
1039 adversary theorem (Theorem 7) and thus consist of two steps. First, authors establish the existence
1040 of the simulator by stating it explicitly. There is usually little explanation about why it is defined
1041 as it is, since the simulator is the result of a fair amount of trial and error. Second, the authors
1042 prove the indistinguishability property at the heart of UC: $\text{EXEC}(Z, A_d, \pi) \approx \text{EXEC}(Z, S, F)$. This
1043 is usually an induction over the length of the trace from the perspective of the environment, or
1044 more precisely, over the number of outgoing messages from the environment, which mark the
1045 beginning of what is often called an *epoch*. The epoch ends when the environment regains control
1046 by receiving a message from the attacker or protocol. For each epoch, these proofs show (via case
1047 distinction over the message) that some state invariant holds; that will ensure that the message the
1048 environment receives at the end of the epoch is indistinguishable in the ideal and the real world.
1049 Consequently, the state invariant links the state of the parties in the ideal world (S, F) and in the
1050 real world (A_d, π) to each other. A recurring pattern in UC proofs is that the simulator S tracks a
1051 large part of the state of the protocol π to correctly simulate behaviour that is independent from F .

1052 The main work of the proof is hence in the aforementioned case distinction: no matter what
1053 the environment sends to the protocol/ideal functionality, the state invariant holds. The majority
1054 of these cases is about the simulator correctly tracking the protocol state. For one, this is tedious:
1055 state invariants can become very large, since interactions in the ideal world may involve multiple
1056 messages between S and F . For two, and somewhat tragically, this is mechanical but not easily
1057 mechanisable (as admitted by the authors of one mechanised UC proof [43], see Section 6.1). The
1058 mathematician is essentially performing a symbolic execution in a deterministic system. If the
1059 language used to describe protocol, simulator and functionality were to have a formal semantics,
1060 this tedious task could be automated. In many pen-and-paper proofs, this language has no formal
1061 semantics and it is thus not very surprising, albeit disappointing, that these laborious steps are
1062 skipped entirely. The proof concentrates on the few cases where a cryptographic argument is
1063 necessary, establishing a reduction of this form: if, given the state invariant, the environment can
1064 compute a message, then the same environment can be transformed into an attacker for some
1065 cryptographic game

1066 Summarizing, there is a lot of potential for automation in these proofs, should they be conducted
1067 in a formal language. We believe some of the low-hanging fruits here are (i) the symbolic execution
1068 to prove state equivalence in the majority of cases and (ii) the automated synthesis of the state
1069 equivalence relation. A more challenging task is the partial synthesis of the simulator. Given a
1070 partial simulator and a partial state equivalence that fixes the cryptographically interesting parts of
1071 the interaction, provide a simulator and state equivalence relation that produces equivalent traces
1072 for the interactions that are not defined by the partial simulator.

1073 While the informal description above captures the majority of UC proofs, we note that some
1074 follow a more structured approach. The game-playing technique [31] consists of structuring
1075 cryptographic proofs as a series of small syntactic refinement steps in an informal pseudocode
1076 language. The technique is also sometimes used in UC proofs (e.g., [16, 42]) and it has been
1077 mechanised in EasyUC but, quoting Canetti et al. [43]: “despite the relative simplicity [of their case
1078 studies, proving UC] took an immense amount of work” [43]. The authors point out the need for a
1079 domain specific language that is coroutine-based, as opposed to the underlying EasyCrypt [28]
1080 language they used, as well as the lack of symbolic evaluation. Moreover, proving composition
1081 inside EasyCrypt was only possible for specific instances, but not in general. The authors state that

1082 a general proof must be performed at the level of EasyCrypt’s metatheory, i.e., without automation.
 1083 Our composition axiom can help this task. Later work in EasyCrypt [27], heavily improved on the
 1084 proof effort and automation by aligning the message passing semantics more closely to EasyCrypt’s
 1085 procedure-based communication model. As this enforces a hierarchical communication model (in
 1086 contrast to UC, where the environment can directly access lower level functionalities), the authors
 1087 suggest to think of their work as a backend for [43] and leave a sound translation for future work.
 1088 The main concern of their work [27] is mechanising the complexity analysis of programs – which
 1089 is a precondition for the validity of composition in the computational setting.

1090 *6.1.2 RC Proofs.* Most *RC* proofs follow a proof technique inherited from secure compilation
 1091 works [80] called *Backtranslation* [11, 12, 14, 83, 85]. The goal of a backtranslation is to define
 1092 how to construct any source attacker (the existentially-quantified A in Definition 4) starting from
 1093 whatever target information available (i.e., what precedes said existentially-quantified A). If the
 1094 used information is the target attacker A , then the proof is called *Context-based Backtranslation*, if
 1095 the used information is the target trace, then it is called *Trace-based Backtranslation*.

1096 Once the backtranslation is defined, the *RC* proofs essentially proceed by induction over the
 1097 target trace that needs to be replicated in the source. Most times those traces follow this inductive
 1098 principle: they are either empty or the concatenation of a trace with a list of attacker-generated
 1099 actions ending with a control-transferring action to the program ($\bar{\tau} \cdot \tau?$), followed by a sequence of
 1100 program-generated actions ending with a control-transferring action to the context ($\bar{\tau}' \cdot \tau!$).

1101 The proof sets up a simulation argument between the target context performing $\bar{\tau} \cdot \tau?$ and its
 1102 backtranslated counterpart, which needs to be proven to produce $\bar{\tau} \cdot \tau?$ in the source. Following
 1103 established compilation-style proofs [72], this is done by setting up a cross-language state relation
 1104 ($\Omega \sim \Omega$) between the states of the source attacker (Ω) and those of its backtranslation-generated
 1105 counterpart (Ω). Then, after the $?$ -decorated action, control is transferred to the compiled program.
 1106 The proof then proceeds by a classical compiler correctness argument which ensures that the
 1107 program performs $\bar{\tau}' \cdot \tau!$ given that its compilation perform $\bar{\tau}' \cdot \tau!$. This proof is carried out with a
 1108 different cross-language state relation relating states of the source program and its compilation
 1109 ($\Omega \approx \Omega$). Notably, the relation when control is in the attacker code (\sim) tends to enforce a stronger
 1110 invariant than the other one (\approx), and thus there is need to weaken and strengthen the relation. The
 1111 weakening and strengthening often happens when the compiler performs some security-related
 1112 checks that are placed as a wrapper around the program. It is in fact common for secure compilers
 1113 to act as “security wrappers” around correctly-compiled code [6, 34, 49, 82].

1114 *RC* works that use the backtranslation proof technique typically do not use reactive languages,
 1115 and thus they cannot write the analogous of the dummy attacker theorem. Thus, backtranslation
 1116 proofs require quite some effort, with research being active in simplifying said effort [51]. This
 1117 effort is especially great when these proofs are mechanised. At the time of writing, the only such
 1118 mechanised effort required 20K lines of Coq just for a part of the backtranslation correctness
 1119 proof [50].

1120 *6.1.3 The Best of Both Worlds.* In the following, we will use our connection (Theorem 3) to provide
 1121 rigorous UC proofs, relying on the benefits of *RC* proofs. Specifically, we will consider a protocol
 1122 and its ideal functionality, coded in RILC. By using the RILC semantics, we will perform the
 1123 reductions in the real world (for now manually) and compute the real-world traces. Then, we will
 1124 perform the ideal world reductions and show that the ideal traces they yield are the same (trace
 1125 and probability-wise) as the real-world ones. This will let us conclude that the compiler from the
 1126 ideal functionality to the protocol is *RHP*, and thus the latter UC-realises the former.

1127 6.2 UC Commitments as RC Proofs for the Static Corruption Model

1128 In order to perform this proof we need to define the following RILC processes:

- 1129 (1) the ideal functionality F (Section 6.2.1);
- 1130 (2) the protocol P, Q that UC-realises F (Section 6.2.2);
- 1131 (3) the simulator S used in the proof (Section 6.2.3).

1132 Finally, the proper proof is presented in Section 6.2.4, followed by a discussion Section 6.2.5. This
 1133 proof rests on a few assumptions, which we inherit from the original UC proofs from Canetti and
 1134 Fischlin [41], and that we present below.

1135 *Assumptions for the Static Corruption Commitment.* Below are the assumptions that hold for the
 1136 protocol of this section. Note that these assumptions are not imposed by us, but they are inherited
 1137 by the work that defined the protocol in the first place. We simply recap them here for simplicity.

- 1138 (1) the protocol assumes authenticated channels, i.e., if a party A receives a message m , it knows
 1139 the identity of the party B that sent m .
- 1140 (2) the only party that can be corrupted is the committer (P), since corrupting the receiver
 1141 does not let the attacker break any of the commitment properties. This is borrowed from
 1142 the simulator of Canetti and Fischlin [41], which only considers the effect of corruption on
 1143 commitments created by the corrupted party—in other words, corruption is only relevant for
 1144 a party when that party acts as a sender.

1145 *6.2.1 Ideal Functionality.* The ideal functionality is presented in Listing 1, with comments explain-
 1146 ing the code. In all our code snippets when we write a channel XY , it is a channel used from process
 X to send a message read by process Y .

```

1 process F =
2   rd Commit b from Read(PF)      // receive a commitment message for a bit 'b'
3   wr Receipt to Write(FS)        // acknowledge that the commitment has been initiated
4   rd Open from Read(PF)          // receive a message to open the commitment
5   wr Opened b to Write(FS)       // open the commitment to the right bit 'b'

```

Listing 1. Ideal Functionality for the single-bit commitment.

1147 An interesting bit of this code is its *module signature*, which is in Listing 2. The reason we deem
 1148 such signatures interesting is that it describes the channels that will be used.

```

1 imports If = Write(FS)
2 definitions Df = Read(PF), Write(FS)
3 exports Xf = Read(PF)

```

Listing 2. Module signature for ideal functionality for the single-bit commitment.

1149
 1150 *6.2.2 Protocol.* The protocol consists of two parties (a committer P and a receiver Q) whose
 1151 code is in Listing 3 and of a common reference string whose code is in Listing 4. In the static
 1152 corruption model, corruption is defined at the beginning of the computation with a message from
 1153 the environment. That is what happens on Line 2 and that is the reason why the code of the
 1154 committer P is split in two parts: the non-corrupted behaviour (Line 3 to Line 20) and the corrupted
 1155 one (Line 22 to Line 30). Both parts contain effectively the same preamble (Line 4 to Line 9). Those
 1156 are additional messages that we inserted *before* the proper protocol to fix the common reference
 1157 string (CRS)-related interleavings, since this simplifies the proof. The CRS (presented in Listing 4
 1158 and described below) is a piece of code that P, Q , and Z all interact with. To both simplify the proof

```

1 process P =
2   rd M from Read(ZP) // receive the static corruption message
3   if M == CrptNO then // in case no party is corrupted
4     wr StartCrs to Write(PCrs) // start the CRS functionality
5     rd Started from Read(CrsP) // ack the CRS started
6     wr Ok to Write(PZ) // return control to the environment,
7     // which will call the CRS
8     rd WaitCrs from Read(CrsP) // synchronise with the environment
9     wr Waited to Write (PCrs)
10
11 // here starts the commitment protocol,
12 // as taken from the literature
13 rd Commit b from Read(ZP) // receive a commitment message for a bit 'b'
14 wr GetCRS to Write(PCrs) // request the CRS
15 rd PublicStrings s pk0 pk1 from Read(CrsP) // receive the CRS
16 let r = takernd
17 let y = (if b == 0 then prg pk0 r else xors(prg(pk1, r), s)) // set y to
18 // the commitment of bit 'b' using the 'prg'
19 wr Commit' y to Write(PQ) // send the commitment 'y' to the receiver
20 rd Open from Read(ZP) // receive a message to open the commitment
21 wr Open' b r to Write(PQ) // forward that request to the receiver
22 else // if the corruption message is not 'CrptNO',
23 // the committer is corrupted
24 // these synchronisation messages are as above
25 wr StartCrs to Write(PCrs)
26 rd Started from Read(CrsP)
27 wr Ok to Write(PZ)
28 rd WaitCrs from Read(CrsP)
29 wr Waited to Write (PCrs)
30
31 // below is the corrupted committer: a proxy
32 // from the environment to the receiver
33 Loop(2)(fwd(Read(ZP))(Write(PQ))) // we expect 2 iterations, thus the bound
34
35 process Q =
36 rd Commit' y from Read(PQ) // receive the commitment 'y' from the committer
37 wr GetCRS to Write(QCrs) // request the CRS
38 rd PublicStrings s pk0 pk1 from Read(CrsQ) // receive the CRS
39 wr Receipt to Write(QZ) // acknowledge the commitment has been initiated
40 rd Open' b r from Read(PQ) // receive original commitment bit 'b'
41 // and randomness 'r'
42 // below: calculate if the communication
43 // of 'b', 'r' and 'y' was done properly
44 if (b == 0 && y == prg(pk0, r)) ∨ (b == 1 && y == xors(prg(pk1, r), s))
45 wr Opened b to Write(QZ)
46 else error // some parameter does not match: abort

```

Listing 3. Single-bit commitment protocol.

1159 and to establish a security argument (as explained in Section 6.2.4), we fix the interaction to be first
 1160 Z, then P, then Q.

1161 Rather than describing the exchange in text, we provide a diagrammatic representation of the
 1162 expected message exchanges in Figure 5. Apart from the additional code that explicitates the
 1163 corruption and from the code that simplifies the rigorous reasoning in the proofs, the essence of
 1164 the protocol code is the same as in its previous definition by Canetti and Fischlin [41].

1165 Intuitively, the CRS functionality lets the parties read from a bit string that is sampled from a
 1166 specified distribution during a one-time setup phase. For the purposes of the protocol, this setup
 1167 phase is assumed to be run honestly. The protocol's use of a CRS sidesteps the impossibility of
 1168 constructing a simulator for secure commitments in the *plain model* (a setting that makes no
 1169 assumptions about an honest setup phase); this impossibility result was proved by Canetti and

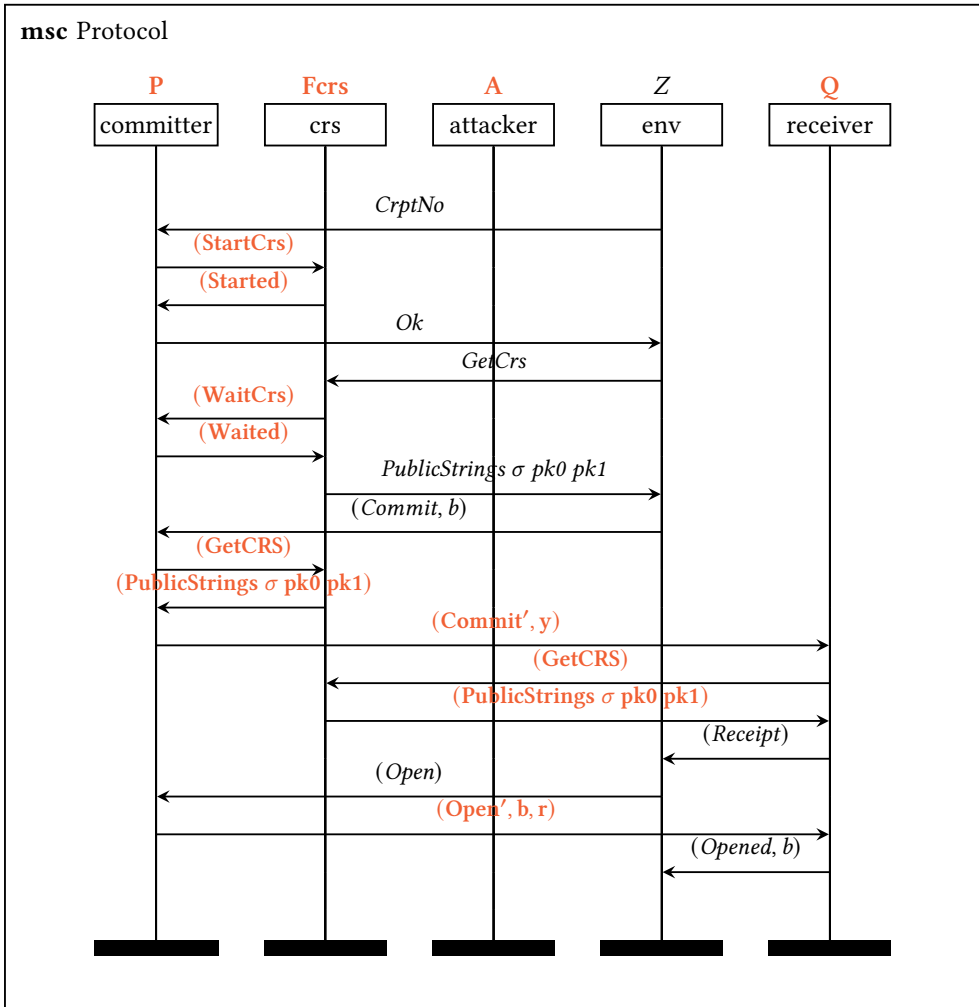


Fig. 5. Diagram representation of the protocol interaction in the no-corruption case. Red actions are *internal* steps while black ones are those that constitute a trace.

1170 Fischlin [41, §3]. As for the protocol, the core intuition behind this code and its original formulation
 1171 is unchanged from the original paper [41].

1172 The signature for the protocol (in Listing 5) highlights an interesting difference with the signature
 1173 of the ideal functionality presented in Listing 2. In fact, this signature is *different* from the one
 1174 of the ideal functionality: *and rightly so!* In fact, from a module perspective, at the protocol level,
 1175 this code is all we consider, while at the functionality level, we do not just consider F , but also the
 1176 simulator (which we discuss next). Thus, the signature of the protocol module should be the same
 1177 as the signature of the module obtained from linking the functionality and the simulator. As we
 1178 show below, this is the case, but in order to argue this point, we need to present the simulator first.

1179 **6.2.3 Simulator.** The simulator for this proof contains a number of code snippets: the simulator,
 1180 the fake common reference string, and the simulated parties. The simulator is presented in Listing 6.

```

1 process Crs =
2   rd StartCrs from Read(PCrs)
3   wr Started to Write(CrsP)
4   rd GetCrs from Read(ZCrs)
5   let s = takernd
6   let r0 = takernd
7   let r1 = takernd
8   let pk0, _ = keygen(r0)
9   let pk1, _ = keygen(r1)
10  wr WaitCrs to Write(CrsP)
11  rd Waited from Read(PCrs)
12  wr PublicStrings s pk0 pk1 to Write(CrsZ)
13  rd Getcrs from Read(PCrs)
14  wr PublicStrings s pk0 pk1 to Write(CrsP)
15  rd GetCrs from Read(QCrs)
16  wr PublicStrings s pk0 pk1 to Write(CrsQ)

```

Listing 4. Protocol CRS

```

1 imports Ip = Write(PCrs), Write(PZ), Write(PQ)
2 definitions Dp = Write(PCrs), Write(PZ), Write(PQ), Read(ZP), Read(CrsP)
3 exports Xp = Read(ZP), Read(CrsP)
4 imports Iq = Write(QCrs), Write(QZ)
5 definitions Dq = Write(QCrs), Write(QZ), Read(PQ), Read(CrsQ)
6 exports Xq = Read(PQ), Read(CrsQ)
7 imports Icrs = Write(CrsP), Write(CrsZ), Write(CrsQ)
8 definitions Dcrs = Write(CrsP), Write(CrsZ), Write(CrsQ), Read(PCrs), Read(ZCrs), Read(
9   QCrs)
9 exports Xcrs = Read(PCrs), Read(ZCrs), Read(QCrs)
10
11 // the resulting signature after linking P, Q, and Crs is
12 imports I = Write(PZ), Write(QZ), Write(CrsZ)
13 definitions D = Dp ∪ Dq ∪ Dcrs
14 exports X = Read(ZP), Read(ZCrs)

```

Listing 5. Protocol Interface

1181 The simulator receives the static corruption message (Line 2) and behaves differently in case
 1182 P is corrupted (Line 12) or not (Line 31). In either case, the simulator starts the fake crs and
 1183 receives the fake parameters. If P is not corrupted, the simulator acts as a proxy for the messages it
 1184 receives Lines 8 and 11. If P is corrupted, the simulator receives the intermediate *Commmitt' y'* and
 1185 *Open'* messages (Line 18 and Line 32 respectively). Then, it uses its knowledge of the fake CRS to
 1186 calculate whether the received commitment bit y' is in the image of the 0 or 1 key. Once it learns
 1187 this information, it sends the correct bit (0 or 1) to the ideal functionality (Line 26 and Line 28
 1188 respectively). From here on, the simulator acts as a proxy, and in the end it checks whether the
 1189 opened bit was correct (Line 33).

1190 The fake common reference string is responsible for setting up a common reference string whose
 1191 backdoor is known to the simulator (Listing 7). The fake CRS is a process we inherit from the
 1192 original proof of UC for this protocol [41]. The reason we need this in the ideal world is dictated by
 1193 some reductions that happen in the real world, which we now discuss. Essentially, in the case P is
 1194 corrupted, it could behave maliciously or honestly. In order to detect the honest case, the simulator
 1195 must know whether its messages are conformant to the data communicated from the CRS (i.e.,
 1196 whether it used the string s , and the keys $pk0$ and $pk1$). Thus, the first thing the simulator does
 1197 is starting the fake CRS process, which computes a known common reference string s and that
 1198 communicates the string as well as both the keys and the trapdoors to the simulator (Line 8). With

```

1 process S =
2   rd !M from Read(ZP)
3   if M == CrptNO then
4     wr CrptNO to Write(SP)
5     rd !Ack from Read(PS)
6     wr FakeSetup to Write(SCrs)           // starts the fake CRS
7     rd !Fake pk0 pk1 td0 td1 s from Read(CrsS) // receives the fake CRS parameters
8     wr Ok to Write(SP)
9     rd !Receipt from Read(FS)
10    wr Receipt to Write(SQ)
11    fwd(Read(FS))(Write(SQ))
12  else
13    wr CrptP to Write(SP)
14    rd !Ack from Read(PS)
15    wr FakeSetup to Write(SCrs)           // starts the fake CRS
16    rd !Fake pk0 pk1 td0 td1 s from Read(CrsS) // receives the fake CRS parameters
17    wr Ok to Write(SP)
18    rd !Commit' y from Read(PS)           // receive the Commit' message from P
19    let g0 = (invert(<pk0, td0>, y))       // use the trapdoor to calculate if the
20                                           // committed bit 'y' in the image of
21                                           // the key for 0
22    let g1 = (invert(<pk1, td1>, xors(y, s))) // use the trapdoor to calculate if the
23                                           // committed bit 'y' in the image of
24                                           // the key for 1
25    if g0 then
26      wr Commit 0 to Write(SP)           // instruct P (and then the F) to commit to 0
27    elif g1 then
28      wr Commit 1 to Write(SP)           // instruct P (and then the F) to commit to 1
29    else
30      wr Commit 0 to Write(SP)           // fake commitment, just to resume other processes
31      fwd(Read(FS))(Write(SQ))
32      rd !Open' b r from Read(PS)         // receive the Open' message from a (malicious?) P
33      if (b == 0 && g0) ∨ (b == 1 && g1) then
34        wr Open to Write SP              // instruct P (and then the F) to Open
35        fwd(Read(FS))(Write(SQ))         // forward the Opened message to Q
36      else error

```

Listing 6. Simulator for the single-bit commitment.

```

1 process FAKECRS =
2   rd !FakeSetup from Read(SCrs)
3   let pk0, td0 = keygen([1..1]) // lambda-long ones
4   let pk1, td1 = keygen([1..1]) // lambda-long ones
5   let r0 = takernd
6   let r1 = takernd
7   let s = xors(prg(pk0, r0), prg(pk1, r1))
8   wr Fake pk0 pk1 td0 td1 s to Write(CrsS)
9   rd GetCrs from Read(ZCrs)
10  wr WaitCrs to Write(CrsP)
11  rd Waited from Read(PCrs)
12  wr PublicStrings s pk0 pk1 to Write(CrsZ)

```

Listing 7. Fake CRS for the single bit commitment.

1199 the trapdoor, the simulator will be able to understand if the message received from P (Line 18) is
1200 honest or not.

1201 The dummy parties are elements that are present in the literature too, and here we prove that
1202 they are needed in order to make the signatures of the modules match (Listing 8). Essentially, the


```

1 process P =
2   rd !M from Read(ZP)
3   wr M to Write(PS)
4   rd !M from Read(SP)
5   if M == CrptNO then
6     wr Ack to Write(PS)
7     rd Ok from Read(SP)
8     wr Ok to Write(PZ)
9     rd WaitCrs from Read(CrsP)
10    wr Waited to Write(PCrs)
11    loop(2)(
12      fwd(Read(ZP))(Write(PF))
13    )
14  else // M == CrptP
15    wr Ack to Write(PS)
16    rd WaitCrs from Read(CrsP)
17    wr Waited to Write(PCrs)
18    loop(2)(
19      fwd(Read(ZP))(Write(PS))
20      fwd(Read(SP))(Write(PF))
21    )
22
23 process Q =
24   fwd(Read(SQ))(Write(QZ))

```

Listing 8. Simulator Parties

1203 dummy parties act as proxies: they forward environment messages to the functionality (Lines 3
 1204 and 19), and simulator messages to the environment (Lines 8 and 24) (or to the functionality Line 12).
 The signature of the simulator is presented in Listing 9.

```

1 imports Is = Write(SP), Write(SCrs), Write(SQ)
2 definitions Ds = Write(SP), Write(SCrs), Write(SQ), Read(ZP), Read(PS), Read(CrsS), Read
3   (FS)
4 exports Xs = Read(ZP), Read(PS), Read(CrsS), Read(FS)
5 imports Ifake = Write(CrsS), Write(CrsP), Write(CrsZ)
6 definitions Dfake = Write(CrsS), Write(CrsP), Write(CrsZ), Read(SCrs), Read(ZCrs), Read(
7   PCrs)
8 exports Ifake = Read(SCrs), Read(ZCrs), Read(PCrs)
9 imports Idp = Write(PS), Write(PZ), Write(PCrs), Write(PF)
10 definitions Ddp = Write(PS), Write(PZ), Write(PCrs), Write(PF), Read(ZP), Read(SP), Read
11   (CrsP)
12 exports Idp = Read(ZP), Read(SP), Read(CrsP)
13 imports Idq = Write(QZ)
14 definitions Ddq = Write(QZ), Read(SQ)
15 exports Idq = Read(SQ)
16
17 // the resulting signature after linking P, Q, FAKECRS, and S is
18 imports I = Write(CrsZ), Write(PZ), Write(PF), Write(QZ)
19 definitions D = Ds ∪ Dfake ∪ Ddp ∪ Ddq
20 exports X = Read(ZP), Read(FS), Read(ZCrs), Read(ZP)
21
22 // the resulting signature after linking the result above with F is
23 imports I = Write(CrsZ), Write(PZ), Write(QZ)
24 definitions D = Ds ∪ Dfake ∪ Ddp ∪ Ddq ∪ Df
25 exports X = Read(ZP), Read(ZCrs)

```

Listing 9. Simulator Signature.

1206 When linking the simulator with the ideal functionality, we obtain a single process whose
 1207 signature – in terms of imports and export – is the same of the protocol signature from Listing 5.
 1208 This is necessary, since such a (static) difference would already make any UC (or *RHP*) result
 1209 impossible.

1210 *6.2.4 RC Proof.* We can now define a compiler that translates Listing 1 into Listing 3 (Definition 12)
 1211 and prove it is *RHP* (Theorem 10).

Definition 12 (Compiler for the Static Corruption Commitment).

$$\llbracket \cdot \rrbracket \stackrel{\text{def}}{=} \llbracket P_{\text{comm}} \rrbracket \rightarrow P_{\text{comm}} \quad \text{where } P_{\text{comm}} \text{ is Listing 1}$$

$$P_{\text{comm}} \text{ is Listing 3}$$

Theorem 10 (The $\llbracket \cdot \rrbracket$ Compiler is *RHP*).

$$\vdash \llbracket \cdot \rrbracket : RHP$$

1212 Like most proofs, the one of Theorem 10 is tedious, so we now report the main insights we
 1213 gained while doing it and leave the full proof in the technical report [86]. We believe these steps
 1214 may be helpful for those that want to have a paper proof before attempting a mechanisation such
 1215 as the one presented in the next section (as we did).

1216 **General structure:** First, we devise a dummy attacker for the commitment and prove Axiom 9
 1217 and Axiom 8, which yield Theorem 8 (Dummy Attacker in *RC*) (again, these uninformative
 1218 proofs are left in the technical report). This lets us get rid of the target attacker and, in turn,
 1219 ensures that we have all the code (protocol, ideal functionality, simulator) whose behaviour
 1220 needs to be analysed for the proof.

1221 Then, the proof proceeds by specialising the trace-based backtranslation proof technique used
 1222 in secure compilation work [11, 79, 83]. Intuitively, the proof consists of analysing the target-
 1223 level traces and generate the simulator in such a way that it replicates all $?$ -decorated actions
 1224 in the trace. The $!$ -decorated actions are then generated by the source program, provided it is
 1225 called with the correct arguments. Ensuring the source program is called with the correct
 1226 arguments can be complex, for example the proofs for the one-bit commitment presented
 1227 here rely on a *FAKECRS* functionality and on the invertibility of the *prg* function. Since
 1228 the application of the backtranslation proof technique is more interesting in the adaptive
 1229 corruption case, we provide more details there (see Section 6.3.4).

1230 Once the simulator is devised, the proof proceeds by matching the reductions in the source
 1231 (P_{comm} plus simulator) with those in the target (P_{comm}), while keeping track of what actions
 1232 generate traces. While seemingly dull, this exercise alone is useful in getting rid of inconsis-
 1233 tencies such as the ideal functionality sending a message with some data, and the protocol
 1234 sending a message with more data.

1235 **Interleavings:** One complication when doing this proof is the amount of interleavings that arise
 1236 due to the choices that the environment can do. For example, the environment can choose to
 1237 (1) set up the CRS and then talk to the committer, or (2) talk to the committer and then set up
 1238 the CRS. For simplicity, we streamline these interleavings, which, in turn, leaves fewer cases
 1239 to reason about in the proof.

1240 This is the reason why, for example, we added the code from Line 4 to Line 9 to the protocol
 1241 of Listing 3. Of course, changes in the target behaviour propagate to the source, and thus our
 1242 simulator in Listing 6 also contains code that reduces interleavings.

1243 **Z talks to CRS:** One of the interleavings that we fixed is the environment talking to the CRS,
 1244 which is an event that must happen prior to the protocol itself. The reason this must happen

has to deal with one of the cases of the proof. When the sender party is corrupted, it forwards all messages from the environment. Thus, we have two cases: either the environment behaves like an honest sender, or it does not. In the former case, we need to know that the environment has communicated with the CRS in order to derive the goodness of the values that are being sent. Essentially, if we do not force this communication, then we have more cases to consider for the environment to not behave like a honest sender. Thus, to simplify the proof, we fixed the communication to happen before the protocol.

Security argument: One simplifying assumption of our model is the symbolic crypto model that we have added into RILC (in Section 5.2.1). This common assumption [8, 9, 88, 93] ensures that the security argument we need to consider is solely about the randomness distributions in traces.

Specifically, when unravelling the reductions of the protocol, we obtain the following traces. Note that for simplicity, we present them in symbolic form, that is, instead of listing all the single-bit combinations, we represent their values abstractly, with a symbol: Also, we elide the trace of events between the environment and the `CrS` of Listing 4 with \dots

- $((\text{CrptNo})? \cdot (\text{Ok})! \cdot \dots \cdot (\text{Commit } b)? \cdot (\text{Receipt})! \cdot (\text{Open})? \cdot (\text{Opened } b)!, 1/2)$
- $((\text{CrptP})? \cdot (\text{Ok})! \cdot \dots \cdot (\text{Commit } y)? \cdot (\text{Receipt})! \cdot (\text{Open } b \ r)? \cdot (\text{Opened } b)!, 1/2 \times 1/\varphi)$
- $((\text{CrptP})? \cdot (\text{Ok})! \cdot \dots \cdot (\text{Commit } y)? \cdot (\text{Receipt})! \cdot (\text{Open } b \ r)? \cdot (\text{Error})!, 1/2 \times 1/(1 - \varphi))$

As we can see, the total of these traces amount to 1, so they capture all the behaviour of the protocol. In these traces, b and y are single-bit values (i.e., either 0 or 1), while r is a λ -long sequence of bits. The probabilities are the most interesting bit: the first $1/2$ comes from the environment choice on the corruption of P , i.e., that is the probability of the environment sending the first message as a `CrptNo` or as a `CrptP`. Probability $1/\varphi$ is the probability of the environment guessing both b and r .

Following the steps mentioned above, we follow the source reductions and see that the source modules (simulator, ideal functionality, dummy parties and fake CRS) perform the same traces, at least wrt the emitted actions. We then have to argue that the distribution of all target traces matches the source one. Wrt the choice of the initial message, they are chosen by the environment using the same message distribution, so the first $1/2$ probability exists in source traces too. Wrt the probability of the environment guessing b and r , we notice that the environment has the same probability of generating those values, since source and target values come from the same set, and since the environment does not know more in the source than in the target.

Thus we can conclude that the source does exactly the same traces presented above.

6.2.5 Excess Code in Protocol. After seeing the amount of code that is required in order to apply our approach, we see the protocol code having too much boilerplate code as a possible criticism. We agree with this point, and we believe to be justified in having added said boilerplate, given the simplification they provide. Moreover, it is worth noting that the additional boilerplate surrounds the original protocol code and thus can be morally separated from it.

Some of the boilerplate code has been added as part of the attacker model (e.g., Line 2 in Listing 3 where the corruption message is received) and that could have been handled differently. For example, given the static corruption setting, we may want to say that in order to prove that a protocol UC-realises a functionality, one has to provide different implementations: one for the protocol and one for each of the corruption cases considered. That would simplify the presentation of the protocol, and protocol implementors would benefit from that.

1291 The additional synchronisation messages provide another part of the boilerplate, but we believe
 1292 that they simplify the pen-and-paper approach significantly, since there would many interleavings
 1293 to consider. An alternative would be to fix a scheduling of messages, and prove the protocol UC-
 1294 realises a functionality given a certain scheduling. Finally, as we move to using tools to mechanise
 1295 these proofs, that boilerplate may not be necessary. As we report in Section 7, our encoding of this
 1296 proof in DEEPSEC lets us strip the protocol, the CRS, and the simulator of the synchronisation
 1297 messages.

1298 6.3 UC Commitments as RC Proofs for the Adaptive Corruption Model

1299 For the adaptive corruption setting, we need to adapt both the ideal functionality and the protocol.
 1300 The adaptive setting is a difficult one for commitments: it is known that adaptive security requires
 1301 either an exotic assumption (e.g., securely erasable memory) or weaker security properties (Hirt et al.
 1302 [59, §1.2] discuss in recent work that sidesteps this issue). For simplicity we take the latter approach,
 1303 demonstrating a scheme that sacrifices the hiding property. Thus, we start this section with the
 1304 new ideal functionality (Section 6.3.1). Then we present the protocol (Section 6.3.2), followed by
 1305 the simulator (Section 6.3.3) and the proof itself (Section 6.3.4). As before, this proof has also be
 1306 mechanised in DEEPSEC, which we report in the next section.

1307 *6.3.1 Ideal Functionality.* The only difference between the ideal functionality for the adaptive
 1308 corruption case (Listing 10) and the one for the static corruption case (Listing 1) are the additional
 1309 messages on Lines 3 and 4. The functionality immediately leaks the commitment bit (Line 3) and
 1310 then waits for input to resume the commitment (Line 4).

```

1 fob =
2   rd Commit b from Read(PF)      // receive the commitment bit
3   wr Leak b to Write(FS)        // immediately leak it
4   rd DoCommit from Read(PF)     // resume the functionality as before
5   wr Receipt to Write(FS)
6   rd Open from Read(PF)
7   wr Opened b to Write(FS)

```

Listing 10. Ideal functionality for an only-binding commitment.

1311 *6.3.2 Protocol (and CRS).* As before, we enrich all real-world parties (i.e., the committer, the receiver,
 1312 and the CRS) with additional synchronisation messages, to limit the amount of interleavings we
 1313 have to deal with in the proof. The first party that starts is the **CRS**, which awaits a message from
 1314 the environment and then generates the CRS as before in Listing 3. Then it sends the public strings
 1315 message to the environment first, then to the committer, and finally to the receiver. The CRS acts as
 1316 the synchronisation party between committer and receiver, with the additional messages **WaitCrs**,
 1317 **Waitcd**, **SynchOpen**, **Synched**, **SynchOpen'**, and **Synched'**.

1318 The committer (**P**) and the receiver (**Q**) protocols are very similar to their static-corruption
 1319 counterparts save for one key difference. There is no communication that happens on authenticated
 1320 channels directly between **P** and **Q**: all protocol-related messages are transmitted on an attacker
 1321 interface. In our convention, that is visible in the channel names used in the communication: they
 1322 are of the form **AP**, **PA**, **AQ** or **QA**. This indicates that the message sent or received by the protocol
 1323 party goes through the attacker (and thus possibly the environment) before reaching the other
 1324 party.

1325 In the protocol description we comment the additional lines that have been added for synchroni-
 1326 sation, as well as those lines that have been added to give up the hiding property.

```

1 CRS =
2   rd GetCrS from Read(ZCrS)
3   let s = takernd
4   let r0 = takernd
5   let r1 = takernd
6   let pk0, _ = keygen(r0)
7   let pk1, _ = keygen(r1)
8   wr WaitCrS to Write(CrSP) // start the committer
9   rd Waited from Read(PCrS) // ack the committer started
10  wr PublicStrings s pk0 pk1 to Write(CrSZ) // send values to environment
11  rd GetCrS from Read(PCrS) // ack the committer wants the CRS
12  wr WaitCrS to Write(CrSQ) // start the receiver
13  rd Waited from Read(QCrS) // ack the receiver started
14  wr PublicStrings s pk0 pk1 to Write(CrSP) // send values to the committer
15  rd GetCrS from Read(QCrS) // ack the receiver wants the CRS
16  wr SyncOpen to Write(CrSP) // tell the committer it can open
17  // the commitment
18  rd Synched from Read(PCrS) // ack the committer can open
19  wr PublicStrings s pk0 pk1 to Write(CrSQ) // send values to the receiver
20  rd SynchOpen' from Read(PCrS) // ack the committe can forward the open
21  // to the receiver
22  wr SynchOpen to Write(CrSQ) // tell the receiver it can open
23  // the commitment
24  rd Synched from Read(QCrS) // ack the receiver can open
25  wr Synched' to Write(CrSP) // tell the committer to send
26  // the open message

```

Listing 11. CRS for adaptive corruption.

1327 **6.3.3 Simulator.** The simulator follows the same structure of the previous one, save for one detail.
 1328 Since the value of the committed bit is known, the simulator does not have to rely on inverting the
 1329 *prg* used by the (*fake*) CRS in order to know what the commitment bit is.

1330 The fake CRS (Listing 14) is essentially the same as before save for the different synchronisation
 1331 messages, which we indicate with a comment.

1332 Since there are more messages being exchanged between the environment and the functionality,
 1333 the dummy parties (Listing 15) contain additional proxy messages. For simplicity, some of these
 1334 messages are forwarded to the simulator, as in Lines 7 and 9. Those are messages that the environ-
 1335 ment should send to the simulator directly. To avoid having the simulator listen on the environment
 1336 channel, and then synchronise with the dummy party, we let the dummy parties be the only entities
 1337 that talk to the environment. The dummy parties then forward to the simulator those messages
 1338 that are not meant for the simulator (in this case, the *Commit'* and the *Open'* messages).

1339 As before, the interfaces of the linked source parties (dummy *P*, dummy *Q*, *S*, *FAKECRS* and *fob*)
 1340 are the same as the interfaces of the linked target ones (*P*, *Q*, and *CRS*).

1341 6.3.4 RC Proof.

Definition 13 (Compiler for the Adaptive Corruption Commitment).

$$\llbracket \cdot \rrbracket_{comm}^{adap} \stackrel{\text{def}}{=} \llbracket P_{comm} \rrbracket \rightarrow P_{comm} \quad \text{where } P_{comm} \text{ is Listing 10}$$

$$P_{comm} \text{ is Listing 12}$$

Theorem 11 (The $\llbracket \cdot \rrbracket_{comm}^{adap}$ Compiler is RHP).

$$\vdash \llbracket \cdot \rrbracket_{comm}^{adap} : RHP$$

1342 The proof follows the same structure of the one presented in Section 6.2.4, so we leave its details
 1343 for the technical report [86]. The traces generated in the adaptive corruption case contain more

```

1 processes P =
2   rd WaitCrs from Read(CrsP)           // added synchronisation
3   wr Waited to Write(PCrs)            // added synchronisation
4   rd Commit b from Read(ZP)
5   wr Leak b to Write(PA)               // added for no hiding
6   rd DoCommit from Read(AP)           // added for no hiding
7   wr GetCRS to Write(PCrs)
8   rd PublicStrings s pk0 pk1 from Read(CrsP)
9   let r = takernd
10  let y = (if b == 0 then prg pk0 r else xors(prg(pk1, r), s))
11  wr Commit' y to Write(PA)
12  rd SyncOpen from Read(CrsP)          // added synchronisation
13  wr Synched to Write(PCrs)           // added synchronisation
14  rd Open from Read(AP)
15  wr SyncOpen' to Write(PCrs)          // added synchronisation
16  rd Synched' from Read(CrsP)         // added synchronisation
17  wr Open' b r to Write(PA)
18
19 Q =
20  rd WaitCrs from Read (CrsQ)          // added synchronisation
21  wr Waited to Write(QCrs)            // added synchronisation
22  rd Commit' y from Read(AQ)
23  wr GetCRS to Write(QCrs)
24  rd PublicStrings s pk0 pk1 from Read(CrsQ)
25  wr Receipt to Write(QA)
26  rd SynchOpen from Read(CrsQ)        // added synchronisation
27  wr Synched to Write(QCrs)          // added synchronisation
28  rd Open' b r from Read(AQ)
29  if (b == 0 && y == prg(pk0, r)) \\/ (b == 1 && y == xors(prg(pk1, r), s))
30    wr Opened b to Write(QZ)
31  else
32    error

```

Listing 12. Only-binding commitment protocol for adaptive corruption.

```

1 S =
2   rd Starts from Read(CrsS)
3   wr FakeSetup from Write(SCrs)
4   rd Fake pk0 pk1 td0 td1 s from Read(CrsS)
5   wr Ok to Write(SCrs)
6   rd Leak b from Read(FS)
7   wr Leak b to Write(SQ)
8   rd Receipt from Read(FS)
9   let r = takernd
10  let y = if b == 0 then prg pk0 r else xors(prg(pk1, r), s)
11  wr Commit' y to Write(SQ)
12  rd Commit' y' from Read(PS)
13  wr Receipt to Write(SQ)
14  rd Opened b from Read(FS)
15  wr Open' b r to Write(SQ)
16  rd Open' b' r' from Read(PS)
17  if (y == y') && (b == b') && (r == r')
18    wr Opened b to Write(SQ)
19  else
20    error

```

Listing 13. Simulator for adaptive corruption.

1344 actions than the previous ones, as presented below. For simplicity, we only show the action trace,
 1345 without the probabilities, and abstract the values symbolically. We also elide the error trace where

```

1 FAKECRS =
2   rd GetCrs from Read(ZCrs)           // added synchronisation
3   wr Starts to Write(CrsS)           // added synchronisation
4   rd !FakeSetup from Read(SCrs)
5   let pk0, td0 = keygen([1..1])
6   let pk1, td1 = keygen([1..1])
7   let r0 = takernd
8   let r1 = takernd
9   let s = xors(prg(pk0, r0), prg(pk1, r1))
10  wr Fake pk0 pk1 td0 td1 s to Write(CrsS)
11  rd Ok from Read(SCrs)
12  wr WaitCrs to Write(CrsP)           // added synchronisation
13  rd Waited from Read(PCrs)          // added synchronisation
14  wr PublicStrings s pk0 pk1 to Write(CrsZ)

```

Listing 14. Fake CRS for adaptive corruption.

```

1 P =
2   rd WaitCrs from Read(CrsP)
3   wr Waited to Write(PCrs)
4   Loop(2)(
5     fwd(Read(ZP))(Write(PF)) // forward from environment to functionality
6   )
7   fwd(Read(ZP))(Write(PS)) // forward from environment to simulator
8   fwd(Read(ZP))(Write(PF)) // forward from environment to functionality
9   fwd(Read(ZP))(Write(PS)) // forward from environment to simulator
10
11 Q =
12   Loop(5)(
13     fwd(Read(SQ))(Write(QZ))
14   )

```

Listing 15. Simulator Parties for adaptive corruption.

1346 the last action is replaced by an error message. This case happens when the environment does not
 1347 act as a forwarder for the *Commit'* and the *Open'* messages, but in the case it acts maliciously and
 1348 tampers with the communicated y , b and r values.

- 1349 • *GetCrs?* · *PublicStrings s pk0 pk1!* · *Commit b?* · *Receipt b!* · *DoCommit?* · *Commit' y!* · *Commit' y'?* ·
- 1350 *Receipt!* · *Open?* · *Open' b r!* · *Open' b' r'?* · *Opened b!*

1351 Applying the backtranslation proof technique to this trace yields the following elements of the
 1352 simulator S , of the *FAKECRS*, and of the proxies P and Q :

- 1353 • *GetCrs?* gets backtranslated to Line 2 of *FAKECRS*
- 1354 • *PublicStrings s pk0 pk1!* is Line 14 of *FAKECRS*
- 1355 • *Commit b?* is Line 5 in P (iteration 1 of the loop)
- 1356 • *Receipt b!* is Line 13 in Q (iteration 1 of the loop)
- 1357 • *DoCommit?* is Line 5 in P (iteration 2 of the loop)
- 1358 • *Commit' y!* is Line 13 in Q (iteration 2 of the loop)
- 1359 • *Commit' y'?* is Line 7 in P
- 1360 • *Receipt!* is Line 13 in Q (iteration 3 of the loop)
- 1361 • *Open?* is Line 8 in P
- 1362 • *Open' b r!* is Line 13 in Q (iteration 4 of the loop)
- 1363 • *Open' b' r'?* is Line 9 in P
- 1364 • *Opened b!* is Line 13 in Q (iteration 5 of the loop)

1365 This initial part of the backtranslation essentially builds all of the proxies, as well as the entry and
 1366 the exit point of the *FAKECRS*.

1367 At this point, since we know that there must only be one party computing at each time, we
 1368 can follow the trail of messages in order to start filling out the rest of the missing code. For
 1369 example, fixing the interleavings between the parties and the *FAKECRS* creates Lines 2, 3, 4 in
 1370 the *FAKECRS* and the related Lines 2, 3 in *S*. The bulk of the *FAKECRS* (Lines 5 to 9) cannot be
 1371 automatically generated by the backtranslation, but Lines 10 to 13 follow directly from having fixed
 1372 the interleavings. Thus, once Lines 10 to 13 in *FAKECRS* are crafted, the related lines (Line 4, 5 in *S*,
 1373 and Lines 2, 3 in *P*) are also a straightforward addition. As another example of straightforward code
 1374 generation, notice that message **Commit b?** is forwarded from the proxy of *P* to the functionality,
 1375 and then it causes a *Leak* message to be sent into the simulator. This generates Lines 6 and 7 in *S*,
 1376 and one iteration of the loop in *Q*.

1377 With this procedure, this specialised trace-based backtranslation proof technique has built all
 1378 of the code in *S*, *P*, *Q* and *FAKECRS* except for Lines 9 and 10 in *S* and Lines 5 to 9 in *FAKECRS*.
 1379 This is expected, and in secure compilation work that use this proof technique, a bit of creativity is
 1380 also needed in order to make the proof go through.¹⁰ We envision that, with the proof automation
 1381 described in Section 7, one can employ program synthesis techniques to automate the search space
 1382 for the missing creative bits, but we leave investigating this for future work.

1383 One interesting insight we encountered while doing said proof was the inability to complete it
 1384 while trying to retain both the binding and hiding properties. In fact, there comes a time (Line 11)
 1385 when performing the reductions that the simulator must create the commitment value (*y* in the
 1386 code). If the protocol is trying to retain the hiding property, the simulator has no data to conjure *y*
 1387 from, so the proof gets stuck. On the other hand, if the protocol gives up hiding, it knows enough
 1388 data to create such *y*, as done in Line 10.

1389 The proofs provided in this section already showcase how to make good use of our connection
 1390 in order to obtain rigorous, formal UC proofs. However, as we point out, many of the steps of these
 1391 proofs can be automated, and this insight is what we expand upon in the next section, where we
 1392 mechanise the proofs presented here.

1393 7 MECHANISING THESE PROOFS

1394 This section reports on how to mechanise the proofs of the previous section. First, this section
 1395 justifies the choice of the tool for this mechanisation: DEEPSEC (Section 7.1). Then it details how
 1396 to model protocols in DEEPSEC, i.e., how to go from the pen-and-paper protocol definitions of
 1397 the previous section to the internal representation of DEEPSEC (Section 7.2). Finally, this section
 1398 reports the lessons learned while doing this mechanisation effort (Section 7.3).

1399 7.1 Picking the Tool

1400 To explore the potential for mechanisation of the previous proofs, we investigated the applicability
 1401 of off-the-shelf tools for the analysis of the previous case study. Our requirements were twofold.
 1402 First, the tool must provide a high-degree of automation, since the proofs required to manually
 1403 check the semantics of each program reduction. Second, the tool must rely on a language whose
 1404 communication model is similar to the one of RILC. We could not use existing ILC-based tools since
 1405 they are currently limited to type-checking processes, with no symbolic execution of processes.
 1406 Since we are interested in knowing what is *possible* from a mechanisation perspective, we accept a
 1407 (small) semantic gap between RILC and the native language of the tool.

¹⁰For example, code that operates type conversion needs to be crafted and added by the backtranslation for the proof of RSC of a compiler between typed and untyped code [83].

1408 Thus, our tool choice fell on DEEPSEC [44], which is a decision procedure for process equivalence
 1409 in the applied- π calculus from the perspective of a Dolev-Yao attacker (or, environment, in RILC
 1410 terms). The Dolev-Yao attacker is an unusual choice in the context of UC, but our focus is on the
 1411 possibility of fully automating the trace-based analysis. Künnemann, Patrignani, and Cecchetti [68]
 1412 perform a similar analysis in the computational model and employ CryptoVerif as a backend. Even
 1413 though CryptoVerif is known for its comparatively high degree of automation, their case study
 1414 requires manual guidance. Of course, CryptoVerif has a probabilistic semantics, which is harder to
 1415 reason about than DEEPSEC's non-deterministic, non-probabilistic execution model. In this work,
 1416 however, we argue that automation is possible, once suitable abstractions have been established.

1417 DEEPSEC represents protocols in the applied- π calculus, whose semantics is fairly close to the
 1418 one of RILC and thus omitted. The semantics of applied π has three notable limitations. First,
 1419 applied π does not have the cryptographic primitives required by the protocol (i.e., no *xors*, no *prg*,
 1420 etc.). Second, applied π lacks an explicit module system like the one of RILC. Finally, applied π
 1421 processes are not necessarily confluent. We need to account for these while modelling protocols
 1422 and we explain in detail how we do this accounting in the next section.

1423 Another limitation is that the decision procedure of DEEPSEC is undecidable [3], so the tool
 1424 needs to bound the number of processes running in parallel by some fixed number. Fortunately, the
 1425 commit protocol we chose as fixed in Section 6.2 only covers a single session, so our verification is
 1426 not affected by this limitation.

1427 7.2 Modelling Protocols in DEEPSEC

1428 We now describe the choices behind modelling channels, primitives, and the various entities of the
 1429 protocol.

1430 *Modelling Channels.* DEEPSEC channels are one-directional, untyped channels, akin to RILC
 1431 channel ends, which can be labelled as public or private. Communication over public channels
 1432 ends up in DEEPSEC traces while communication over private ones does not. Thus, we set up one
 1433 DEEPSEC channel for each channel end in Section 6.2, that is, one channel between all pairs of
 1434 protocol, subprotocol, environment and attacker that talk to each other. The channels that interact
 1435 with the environment are public, while those that do not are private.

1436 In general, any channel whose read and write ends are both specified in the protocol should be
 1437 private, to ensure communication on that channel does not generate any trace. Dually, if at least
 1438 one of the read and write ends is not specified in the protocol, that channel should be public, since
 1439 it models communication with the environment.

1440 *Modelling Primitives.* The first limitation of DEEPSEC is the lack of built-in cryptographic prim-
 1441 itives that the commitment protocol relies on. This is a design choice of DEEPSEC, and the tool
 1442 lets programmers define arbitrary primitives (via keyword *fun*) and their reductions (via keyword
 1443 *reduc*). Thus we add reductions for the *prg* and its inversion, as well as for *xors* (i.e., for those rules
 1444 added in Section 5.2.1).

```

1445 1 fun prg/2.
1446 2 fun keygen/1.
1448 3 reduc invert(prg(keygen(trapdoor), seed), trapdoor) → seed.
1449 4
1450 5 fun xor/2.
1451 6 reduc dexor1(a, xor(a, b)) → b.
1453 7 reduc dexor2(b, xor(a, b)) → a.
  
```

1454 Observe that the XOR theory is incomplete. It is missing the associativity, commutativity, the
 1455 neutrality of 1 and elimination property of 0. As Deepsec cannot handle equational theories, we
 1456 had to simplify XOR in this regard.

1457 *Modelling Entities as Processes.* The second limitation of applied π is its lack of a module system.
 1458 To overcome this, we identify applied π processes to have the same role as RILC modules. Thus we
 1459 use process composition (\parallel) as module composition (\rightsquigarrow), since they behave in the same way.

1460 We code a process for each of the entity in Section 6.2, in the real world we have `committer` (for
 1461 P), `receiver` (for Q), and `CRS`, while in the ideal world we have the ideal functionality `COM` (for F),
 1462 and then `dummy-committer` (for P), `dummy-receiver` (for Q), `fakeCRS` (for $FAKECRS$), and `simulator`
 1463 (for S). For the sake of space, we do not show the code of these processes here, we provide a list of
 1464 all processes at the end of this section, together with a link to the code.

1465 Below we present the code `COM`, which we use to highlight how DEEPSEC processes are essentially
 1466 a complete transliteration (modulo syntactic difference) of their RILC counterparts.

```
1467
1468 1 let COM =
1469 2   in(p2f,x22); let (=Commit,b) = x22 in
1470 3   out(f2a,Receipt);
1471 4   in(p2f,x23); let =Open = x23 in
1472 5   out(f2a,(Opened,b)).
```

1474 The statement `in(p2f,x22)` reads from a channel from the dummy party (p) to the functionality
 1475 (f) and binds the result into variable `x22`. Then `let (=Commit,b) = x22 in` checks that the message
 1476 in `x22` has the form `(=Commit, b)`, for some value `b`. This is just like `rd Commit b from Read(PF)` in
 1477 Listing 1. Dually, `out(f2a,Receipt)` writes to the channel from functionality (f) to simulator (a),
 1478 just like `wr Receipt to Write(FS)` in RILC (cfr Listing 1).

1479 Unfortunately, the previous list of processes is not sufficient to successfully model the commit-
 1480 ment protocol in DEEPSEC because of the last limitation of the tool: its lack of confluence. As they
 1481 stand, there are interleavings of the execution that complicate the reasoning of trace equivalence.

1482 Thus, we introduce another process, essentially an environment wrapper between the environ-
 1483 ment in DEEPSEC and the processes we have defined. This environment wrapper fixes the flow
 1484 of messages from the environment, essentially forcing the scheduler policy, making the wrapped
 1485 processes confluent. This, in turn, eliminates the need to synchronise different entities (as done in
 1486 Sections 6.2 and 6.3), since the environment wrapper provides a synchronisation entity already.
 1487 The environment wrapper has the following code, which is an input-output alternation (we use the
 1488 variable names to indicate what message is expected to be passed each time):

```
1489
1490 1 let env =
1491 2   in(z2e,x30corrupt);
1492 3   out(e2p,x30corrupt);
1493 4   in(p2e,x31ok);
1494 5   out(e2z,x31ok);
1495 6   in(z2e,x32getcrs);
1496 7   out(e2s,x32getcrs);
1497 8   in(s2e,x33pubstrings);
1498 9   out(e2z,x33pubstrings);
149910  in(z2e,x34commit);
150011  out(e2p,x34commit);
150112  in(q2e,x35receipt);
150213  out(e2z,x35receipt);
150314  in(z2e,x36open);
150415  out(e2p,x36open);
150516  in(q2e,x37opened);
150617  out(e2z,x37opened);
150718  0.
```

1509 All wrapped processes talk to the environment via channel ends that end with `e`, for example
 1510 an input channel from the environment to the simulator is called `e2s`. The environment wrapper
 1511 provides a binding for all those channels and forwards them through `e2z` or `z2e` channels, where `z`
 1512 is used to identify the DEEPSEC environment. The channels between the environment wrapper
 1513 and all wrapped processes are set to private.

1514 We compose the real and ideal world from these processes as follows:

```
1515 let realw = env | committer | receiver | CRS.  

  1516 1 let idealw = env | dummy_committer | dummy_receiver | COM | sim | fakecrs.  

  1517 2  

  1518 3
```

1519 In order to test that `realw` and `idealw` are trace equivalent, we issue the following DEEPSEC
 1520 command, which returns true, attesting to the equivalence of the two systems.

```
1521 query trace_equiv(idealw, realw).  

  1522 1  

  1523 2
```

1524 The DEEPSEC files can be found at the project page, <https://uc-is-sc.github.io/>, specifically at:

1525 <https://uc-is-sc.github.io/deepsec/deepsec-commitment.zip>

1526 and they include the following files. Note that for simplicity, we split the proof of the corruption
 1527 case in two files, one where the no-corruption case, and one for the case of the corruption of the
 1528 committer:

1529 **com-nocor.dps**: contains the proof for the static corruption case when no party is corrupted;
 1530 **com-corP.dps**: contains the proof for the static corruption case when the committer is corrupted;
 1531 **com-adaptive.dps**: contains the proof for the adaptive corruption case.

1532 7.3 Lessons Learned

1533 One of the benefits of adding the `env` process is that we can debug the trace equivalence proof
 1534 step-by-step. Essentially, the `env` process is a sequence of input-output pairs that identifies what
 1535 the environment is sending and what it is reading from whatever world (be it real or ideal) it is
 1536 communicating with. By considering such input-output pairs incrementally (i.e., first the first pair,
 1537 then the first two pairs and so forth), we can incrementally check that trace equivalence holds. And
 1538 when it does not hold, we know precisely at which step did the equivalence fail. We found “attacks”
 1539 at almost every step, about half of them due simple typos, and the rest of them pointing out issues
 1540 with the setup or the simulator related to the scheduling messages meant to achieve confluence.

1541 The graphical user interface of DEEPSEC allowed us to step through counter-examples in both
 1542 real and ideal world, identifying the source of these issues. Analysis took a few seconds, so we
 1543 could often fix them by trial-and-error.

1544 In summary, the symbolic analysis of DEEPSEC is an excellent starting point for mechanisation,
 1545 and the automation obtained due to the Dolev-Yao model is very useful. Even though a cryptographic
 1546 analysis would give stronger result, we observe that the majority of messages the environment
 1547 receives is literally the same. Hence a hybrid approach that combines cryptographic reduction with
 1548 symbolic execution could be beneficial.

1549 As previously mentioned, the confluence requirement of RILC results in unnecessarily compli-
 1550 cated and yields unnatural models of protocols. This is a drawback that RILC inherits from the
 1551 classical UC frameworks. While confluence simplifies the analysis (there are fewer traces to match
 1552 between the real and ideal world, even though they are longer), DEEPSEC was designed specifically
 1553 to handle asynchronous communication, heavily exploiting parallel computing. Comparing our
 1554 model with standard DEEPSEC models, we see no reason why DEEPSEC could not handle similar
 1555 models with real asynchrony. We are thus confident that future verification tools could easily
 1556 draw from the algorithm of DEEPSEC to handle asynchronous, non-deterministic communication.

1557 Compared to RILC, this would improve the generality and readability of the protocol models we
1558 analyse.

1559 8 DISCUSSION

1560 This section discusses the relation of our results to other notions in UC and RC. First, we briefly
1561 discuss other UC notions and how they respect the axioms of Section 3 (Section 8.1). Then we
1562 discuss why our connection is the most precise, by showing that no other RC notion connects with
1563 UC, and that an existing connecting conjecture is false (Section 8.2). Finally, we argue the different
1564 choice of programming language style (reactive VS non-reactive, Section 8.3).

1565 8.1 Different UC Notions

1566 Several closely related notions of composable, simulation-based security exist in the literature,
1567 including GNUC [60], IITM [67], EasyUC [43], ILC [73], iUC [36], among others. In this section
1568 we argue that the four axioms of Section 3, which that section discusses in the context of UC as
1569 defined by Canetti [37], also hold in all of the previously mentioned models (and likely hold in
1570 other models from the literature). As in the rest of this paper, we leave the many issues raised by
1571 *computational* notions of UC to future work.

1572 To begin, we note that Axiom 1 and Axiom 3 hold by inspection for all of these systems: these
1573 two axioms merely require defining a trace model, EXEC_T, and EXEC commensurate with the
1574 computational model of the underlying notion of UC.

1575 Axiom 4—the fact that non-probabilistic environments are complete—holds in all of these systems
1576 for the same reason it does in the UC notion of Canetti [37]: in all cases, the definition of simulation
1577 universally quantifies the environment *after* binding the protocol and the adversary. If there exists
1578 a probabilistic distinguishing environment, then there exists at least one set of random choices
1579 for which the environment succeeds, which we can hard-code into a deterministic environment.
1580 For GNUC, see [60, Def. 7]; for IITM, see [67, Def. 1]; for ILC, see [73, §6.3]; for iUC, see [36,
1581 p. 8]. EasyUC mirrors Canetti’s UC definition in this regard, and Canetti explicitly states that
1582 non-probabilistic environments are complete [37, p. 47].

1583 Axiom 2—the fact that finite traces contain the final bit—also holds in all of these systems. In
1584 GNUC, the environment’s output is an arbitrary string, which is post-processed by a distinguisher
1585 that returns a bit [60, Defs. 6 and 7]; this distinguisher is exactly the extraction function β of
1586 Axiom 4. In IITM, either the environment writes a single bit to a decision tape or execution halts
1587 without writing such a bit, in which case the result is 0 by definition; the same holds in iUC [36, p. 7].
1588 In ILC, execUC outputs a single bit [73, §6.3]. In both EasyUC and Canetti’s UC, the environment
1589 terminates the trace with a final bit.

1590 More generally, the axioms of Section 3 appear to hold because of the *structure* of composable
1591 security—the execution model, quantifier ordering, and distinguishing procedure. Thus, we expect
1592 these axioms to hold for essentially all related notions of composable security.

1593 8.2 Wrong Connections

1594 *RHP* is the only criterion we could have chosen to instantiate our connection. Other candidates
1595 include other RC criteria from the work of Abate et al. [14], as well as a conjecture from Hicks
1596 [58]: fully-abstract compilation (*FAC*) [1]. This section first debunks the connection with other RC
1597 criteria (Section 8.2.1) before disproving the *FAC* conjecture (Section 8.2.2).

1598 *8.2.1 RHP and Other RC Criteria.* Abate et al. [14] provide RC criteria for preserving all known
1599 classes of hyperproperties, and arrange them in a partial order on their expressiveness. In order to

1600 show that *RHP*, and only *RHP*, coincides with UC, we prove that all neighbouring elements to *RHP*
 1601 do not coincide with UC. These elements are:

- 1602 • *RSCHP*, the preservation of subset-closed hyperproperties (*SCHP*);
- 1603 • *RTP*, the preservation of arbitrary trace properties, which has been proven to be strictly
 1604 weaker than *RHP*.

1605 *RSCHP Also Works.* *SCHP* intuitively formalises the notion of refinement, and it has been proven
 1606 to be what correct compilers preserve through compilation (in absence of attackers though) [12].

1607 Arbitrary hyperproperties are defined as sets of sets of traces, a hyperproperty *H* is *SCHP* if for
 1608 any set of set of traces *b* in *H*, then *H* also contains all subsets of *b*. Formally:

$$SCHP \stackrel{\text{def}}{=} \{H \mid \forall b' \subseteq b. \text{ if } b \in H \text{ then } b' \in H\}$$

1609 There is a single difference between *RHP* and *RSCHP* (below): while the former is a refinement
 1610 notion, the latter is not. Instead, it is a coincidence, as the target and source behaviours are connected
 1611 by a co-implication.

Definition 14 (Robust Subset-Closed Hyperproperty Preservation).

$$[\cdot] \vdash RSCHP \stackrel{\text{def}}{=} \forall P, A. \exists A. \text{Behav}(A \bowtie [P]) \subseteq \text{Behav}(A \bowtie P)$$

$$\text{or equivalently : } \forall P, A. \exists A. \forall \bar{t}. \text{ if } A \bowtie [P] \rightsquigarrow \bar{t} \text{ then } A \bowtie P \rightsquigarrow \bar{t}$$

1612 Interestingly, this notion also coincides with UC (proven in Isabelle/HOL in Theorem *RSCHP*toUC
 1613 and Theorem UCto*RSCHP*). This suggests that with the UC trace model, arbitrary hyperproperties
 1614 and subset-closed ones collapse. In fact, traces in UC contain probabilities, and given a set of traces,
 1615 the sum of their probabilities must be exactly 1. *RSCHP* mandates that the traces of one system
 1616 ($A \bowtie [P]$) are contained in the other ($A \bowtie P$), so the latter cannot contain any additional trace with
 1617 non-zero probability. Given that there cannot be any trace with zero probability (they cannot be
 1618 emitted by any program, ever), the two sets must coincide, and thus the implication in *RSCHP* is
 1619 equivalent to the co-implication of *RHP*.

1620 Similar collapses of classes of hyperproperties are not novel, and there can be language operators
 1621 that cause it (as discussed in Abate et al. [14]). Since they deter from the main point of the paper,
 1622 we leave investigating additional collapses and the ramifications of this collapse for future work.

1623 *Other notions: Preserving Trace Properties, Hypersafety Properties and Relaxed Hypersafety Prop-*
 1624 *erties.* Trace properties, formally defined as sets of traces [2, 71, 92], capture many functional
 1625 properties of programs. Intuitively a compiler attains *RTP* if it preserves arbitrary trace properties
 1626 through compilation [14].

1627 The difference between *RHP* and *RTP* (below) is in the order of quantifiers: in the former the
 1628 traces are quantified last, while in the latter, they are quantified *before* the source context.

Definition 15 (Robust Trace-Preserving Compilation).

$$[\cdot] \vdash RTP \stackrel{\text{def}}{=} \forall P, A, \bar{t}. \exists A. \text{ if } A \bowtie [P] \rightsquigarrow \bar{t} \text{ then } A \bowtie P \rightsquigarrow \bar{t}$$

1629 However, as reported by Datta et al. [46], there exists a UC-like notion whose quantifiers are in
 1630 the same order as *RTP*: Reactive Simulatability (RS) [23]. Indeed, at the formal level, the only thing
 1631 that changes between RS and UC is the quantifiers ordering:

Definition 16 (Reactive Simulatability).

$$\pi \vdash_{RS} F \stackrel{\text{def}}{=} \forall A, Z. \exists S. \text{EXEC}[Z, A, \pi] \approx \text{EXEC}[Z, S, F]$$

1632 Kuesters et al. [67] state that RS and UC coincide for IITM but not for the notion of UC of
 1633 Canetti [37]. This would imply that *RTP* and *RHP* also coincide, given additional setting-specific
 1634 assumptions. We believe the connection between RS and *RTP* can be formally proven, though we
 1635 leave investigating such specific assumptions for future work.

1636 Additional collapses whose investigation we also leave for future work include arbitrary hyper-
 1637 safety properties and relaxed hypersafety properties. Arbitrary *HSP* are hypersafety properties,
 1638 which include meaningful security properties such as many flavours of non-interference [45]. While
 1639 safety properties are identified by a set of bad prefixes (that the safety property does not extend),
 1640 hypersafety properties are identified by finite sets of sets of bad prefixes (which the hypersafety
 1641 property does not extend). Canonically, *HSP* are defined as finite sets of sets of prefixes, and in the
 1642 finiteness of the first set lies the difference with *SCHP*. Assume we relax the notion of hypersafety
 1643 to encompass *infinite* sets of finite prefixes and call this set *relaxed Hypersafety (XHSP)*

1644 We conjecture that in the specific trace model of UC, *HSP* and *XHSP* coincide, since communica-
 1645 tion is bounded, value size is bounded and thus there cannot be an infinite set of traces. We believe
 1646 that *XHSP* and *SCHP* collapse with the additional common assumptions on the trace model used in
 1647 UC. This would just serve the purpose of illustrating what class does UC morally correspond to, so
 1648 we leave this for future work.

1649 **8.2.2 Disproving The FAC Conjecture.** Fully-abstract compilation (*FAC*) has been the de-facto
 1650 standard for secure compilation until the recent proposal of *RC* [80]. Unlike *RC*, *FAC* does not
 1651 rely on a notion of traces, but on the notion of contextual equivalence [89] in order to specify
 1652 security properties. Two programs P_1 and P_2 are contextually equivalent (indicated as $P_1 \approx_{\text{ctx}} P_2$)
 1653 if they co-terminate no matter what attackers (program contexts) they link against. Formally, if
 1654 we indicate single reduction steps as \hookrightarrow and n such reductions as \hookrightarrow^n , termination of a (whole)
 1655 program is defined as follows:

$$W \Downarrow \stackrel{\text{def}}{=} \exists n. W \hookrightarrow^n \dots \hookrightarrow^n W' \not\hookrightarrow$$

1656 which leads to the following definition of contextual equivalence:

$$P_1 \approx_{\text{ctx}} P_2 \stackrel{\text{def}}{=} \forall A. A \bowtie P_1 \Downarrow \text{ iff } A \bowtie P_2 \Downarrow$$

1657 A compiler is fully abstract (indicated as $\vdash [\cdot] : \text{FAC}$) if it preserves (and reflects) contextual
 1658 equivalence of source programs in their compiled counterparts. Formally

$$\vdash [\cdot] : \text{FAC} \stackrel{\text{def}}{=} \forall P_1, P_2. P_1 \approx_{\text{ctx}} P_2 \text{ iff } [P_1] \approx_{\text{ctx}} [P_2]$$

1659 If we unfold the definitions of source and target contextual equivalence, we obtain the following
 1660 statement:

$$\forall A. (A \bowtie P_1 \Downarrow \text{ iff } A \bowtie P_2 \Downarrow) \text{ iff } \forall A. (A \bowtie [P_1] \Downarrow \text{ iff } A \bowtie [P_2] \Downarrow)$$

1661 Thus, like *RC*, *FAC* considers a robust notion of security, since it universally quantifies over all
 1662 target program contexts.

1663 The simplest argument against the conjecture that *FAC* is UC is counting arguments. While UC
 1664 is *propositional*, in that it talks about a single entity (i.e., the protocol or the ideal functionality) in
 1665 each domain (concrete and abstract), *FAC* is *relational*, in that it talks about pairs of source and
 1666 target programs. As such, when trying to derive an equivalence between the two notions, it is not
 1667 possible to connect all of the elements: if P_1 is F and $[P_1]$ is π , what do P_2 and $[P_2]$ correspond to?

1668 More generally, *FAC* talks about *arbitrary* program equivalences, since P_1 and P_2 can be chosen
 1669 arbitrarily. Instead, UC talks about preserving the expected behaviour of an ideal functionality
 1670 encoded in a protocol which interacts with real-world attackers.

1671 *What about Contextual Equivalence?* If we model the ‘expected behaviour’ as simply terminating
 1672 or diverging, then we can see a similarity between the notion of contextual equivalence and UC
 1673 (despite its lack of an existentially-quantified simulator). There, it is clearer what the two programs
 1674 correspond to: P_1 is the protocol while P_2 is the ideal functionality linked with the simulator.

1675 Using contextual equivalence for proving UC-like notions has been recently investigated by
 1676 Morrisett et al. [77]. We believe this approach is valid, and it can be seen as a specialisation of our
 1677 result that uses termination as the environment final bit instead of traces. In fact they rely on the
 1678 dummy attacker theorem whose application we justify in Section 4.2. Additionally, since contextual
 1679 equivalence only talks about a single language, it can only be applied to the setting where the
 1680 language of the protocol and of the ideal functionality coincide. While we have done the same here,
 1681 our result is more general, and using *RHP* lets one use different languages for the protocol and for
 1682 the ideal functionality.

1683 8.3 Reactive VS Non-Reactive Languages

1684 We have seen how the connection can be set up with a reactive language (RILC), but as we mentioned
 1685 before, we believe the connection holds also for non-reactive languages.

1686 Recall that the main difference between reactive and non-reactive languages is that reactive ones
 1687 have no notion of whole program, while non-reactive ones do. A whole program has a well-defined
 1688 entry point for the `main` and it cannot be effectively extended further. Since all the dependencies
 1689 are resolved, and the `main` is defined, any code that is added to a whole program is essentially dead
 1690 code.

1691 In reactive languages (such as RILC), there is an explicit entity in the semantics that models the
 1692 environment (a configuration with the write token: $\langle \omega; K; \Xi; C \rangle$), and any communication with
 1693 it yields a trace action. With non-reactive programs, there is no such environment entity in the
 1694 semantics, so one should ask what is a trace action in this setting.

1695 To answer this question, notice that trace actions are what gets communicated on the environment
 1696 interface, so we must identify the environment in the non-reactive setting. We believe that in this
 1697 setting the environment effectively gets merged with the attacker. After all, as for the dummy
 1698 attacker theorem, there is no need to have two universally-quantified entities (environment and
 1699 attacker), and since there is no environment in this setting, both roles are covered by the attacker.
 1700 Thus, any communication on the interface between the attacker and the program needs to generate
 1701 a trace action.

1702 Notice, however, that the definition of traces is central to our development. By changing the
 1703 notion of traces and how they are defined, some of the presented axioms need changing. We do
 1704 believe, however, that none of the existing results break, and we now briefly discuss why.

1705 *Theorem 3 (UC and RHP Coincide).* This result is untouched by the language setting.

1706 *Theorem 6 (Composition in RC).* In order to derive this result we rely on a number of composition
 1707 operators. First, \sphericalangle needs not return a whole program, lest its output be not composable with any
 1708 other program. Then \odot_T^S must also not return a whole program, just a complete one, the axioms that
 1709 use it need not change (Axiom 6 (FFI Recombination) and Axiom 7 (Constant Addition)). Finally,
 1710 some technical machinery may be needed in order for Axiom 7 (Constant Addition) to make sense.
 1711 Recall that its statement is: $P \simeq \mathbf{P}$ then $P \odot_S^Q P \simeq P \odot_T^Q \mathbf{P}$. Here, the programs of the premise would
 1712 be whole programs, i.e., they define a `main` and they are complete. It is therefore unclear what
 1713 happens when something else links against them, so long as the new program defines the `main`, it
 1714 can be called, otherwise it is dead code. We expect the notion of the main entry point needs to be
 1715 tweaked in order for this composition to make sense and leave studying this for future work.

1716 *Theorem 8 (Dummy Attacker in RC).* Without a notion of environment, the dummy attacker
1717 theorem makes little sense: we cannot replace the only attacker entity with a dummy! However,
1718 we can borrow existing results from work that used a trace model similar to those required here in
1719 order to obtain similar results to the dummy attacker theorem.

1720 Recall that in this case we need traces to be those actions that are generated across the attacker-
1721 program interface. Such trace semantics have been studied aplenty in the context of fully-abstract
1722 trace semantics [10, 65, 81], i.e., a trace semantics that is as precise as contextual equivalence. A
1723 key element of those semantics is that they simplify reasoning by eliding the notion of attackers,
1724 which gets abstracted away. Now suppose T is equipped with a trace semantics that yields just
1725 those traces. If the trace semantics of T is fully-abstract with respect to behavioural equivalence, we
1726 can use the trace semantics and eliminate the target attacker from the theorem statement. This is
1727 essentially the same result we get in reactive languages when using the dummy attacker theorem.

1728 With one such trace semantics for T , we conjecture that we can follow a proof technique akin to
1729 the one used in reactive languages. Essentially, from the target traces we can build the source-level
1730 attacker (or, simulator) A . Then, we will have all source elements (program and simulator) and we
1731 can calculate whether they are trace-equivalent with the target compiled program. This would be
1732 a simplification of an existing proof technique called trace-based backtranslation, which is often
1733 used in secure compilation work [11, 79, 83, 85]. We leave investigating this novel proof technique
1734 for future work.

1735 9 RELATED WORK

1736 *Programming Languages Meets Cryptography.* Researchers have tried to connect programming
1737 languages and cryptography at length, as attested by the Dagstuhl seminar of 2014 [58]. Existing
1738 results span language models as well as tools that bring the two worlds closer together.

1739 From the language side, a number of languages provide more rigorous formalisation of those
1740 languages used by cryptographers. ILC is one such language, which the authors used as a way to
1741 encode the semantics of Interactive Turing Machines [73]. As shown, ILC is a good candidate to
1742 showcase our connection.

1743 IPDL is another such language, which is equipped with an equational logic that can be used to
1744 reason about contextual equivalence of IPDL programs [77]. IPDL is also a good candidate for our
1745 connection, and its equational logic can likely be extended to reason about trace equivalence of
1746 IPDL programs (once they are lifted to the reactive setting). In the case where the source and target
1747 languages differ, however, the IPDL logic would not be useful, since it is bound to a single language.

1748 To show the high-level security of some low-level cryptographic primitives researchers have
1749 come up with high-level calculi [3, 5, 7, 8, 52] or cryptographic-aware compilers [17] – a large body
1750 of work recently surveyed by Hastings et al. [57]. Interestingly, Acay et al. [17] also discovered one
1751 half of our connection (RC implies UC), but were not interested in the other half¹¹, as we showcase
1752 in this work.

1753 When it comes to tools, there are a number of verification tools for expressing cryptographic
1754 protocols and checking some of their properties. These tools include DEEPSEC [44] EasyCrypt [28]
1755 CryptoVerif [32] and Squirrel [25]. All of them have a reactive formal language as the semantic
1756 foundation for their protocol specification language, but none is really built to prove UC . With the
1757 results of this work, we demonstrate, in general, that tools that are built to prove (trace) equivalences
1758 can be used to provide UC mechanised proofs.

1759 *UC Works.* Universal composition was introduced in the seminal work of Canetti [37], but the
1760 idea of describing security via an ideal functionality that describes all non-attacking network

¹¹Personal communication.

1761 traces via simulation goes back much further [54]. We already discussed UC and some of its
1762 descendants [36, 73]. Our work is a generalisation of this concept that removes the need to specify
1763 the communication down to the machine model. Our main result relates robust compilation to
1764 perfect emulation. For a detailed discussion of the challenges concerning *computational* emulation,
1765 we refer to Hofheinz et al. [63].

1766 Concepts from universal composability were transferred from ITMs to other languages, but
1767 also embeddings between different UC-like frameworks have been shown (e.g., UC [37] into
1768 IITM [90]). One line of research considers the intricacies of UC in the applied-pi calculus [33, 48].
1769 This line of research formulates emulation using process equivalences in the applied-pi calculus
1770 and composition using a rewriting of the channels used inside a process. This requires to track
1771 which channels are used for network communication, which for the environment and which
1772 are external. This requires several well-formedness conditions on the processes and complicates
1773 the definitions, while essentially describing the interfaces at the meta level rather than inside
1774 of the language. By contrast, the composition operator for RILC (Section 5.2.2) benefits from a
1775 module system, combining defined interfaces in a straight-forward way. In the end, both types of
1776 composition behave very similarly. The former approach is directly compatible with existing tools
1777 such as ProVerif and DEEPSEC (which do not support a module system), albeit the well-formedness
1778 conditions need to be checked manually. Our approach is conceptually simpler and separates more
1779 clearly between language semantics, module system and meta theory. In terms of programming
1780 languages, UC has been adopted to a fragment of Java [70] and (concepts like ideal functionality
1781 and emulation) to a fragment of F# [53]. Each of these works considers UC (or related concepts)
1782 within their respective target language. Viewing UC as a form of robust compilation stands to
1783 simplify deriving similar results for other languages and help relate these results across languages.

1784 The results just mentioned come with entirely different verification methods, the former using
1785 program dependency graphs to obtain emulation from non-interference [70], the latter using
1786 refinement types to obtain assert-rely-style compositionality by type-checking [53]. At the protocol
1787 level, the UC variants related to the applied-pi calculus [33, 48] rely on off-the-shelf checkers for
1788 protocol equivalences. Delaune and Hirschi [47] provided a survey about these in 2007; a more
1789 recent survey with a more general scope also contains a section on equivalence properties [26].
1790 By and large, protocol equivalence checkers can handle protocols with an unbounded number of
1791 sessions if the ‘two worlds’ are structurally similar or if syntactic conditions can ensure that results
1792 for a small number of sessions translate to the unbounded model – sometimes called *small-model*
1793 *results*. For a bounded number of sessions, and in the Dolev-Yao model there are multiple decision
1794 procedures, providing convenience and automation, with DEEPSEC being the most recent and most
1795 advanced. This points to an interesting research question: can we structure the composition operator
1796 so that small-model results apply and we can use these decision procedures (in the Dolev-Yao
1797 model?).

1798 In terms of direct mechanisations of UC, we already discussed two formalisations in the cryptography-
1799 centric theorem-prover EasyCrypt [27, 43] in Section 6.1.1. Constructive Cryptography, a UC-like
1800 framework, was formalised in Isabelle/HOL [30]. These two are the most advanced mechanisation
1801 efforts for UC-like models in the computational model. Both require an enormous amount of effort
1802 and expertise for modelling and proving, but provide a very high degree of assurance. There is hope
1803 though: there are verification tools in the computational model that provide more automation [32]
1804 by using program transformations that are computationally justified. A related line of research
1805 justifies deduction systems that perform symbolic reasoning in the computational model [25, 29],
1806 likewise promising to increase the amount of automation. Some of these [25, 32] model protocols in
1807 a computational variant of the applied-pi calculus, suggesting that the choice between the Dolev-Yao
1808 model and the computational model is not such a fundamental questions. The translation and the

1809 composition operators might not be all that different and the most automated reasoning tools in
1810 the computational model take inspiration from those in the Dolev-Yao model.

1811 Datta et al. [46] present a number of UC-like notions and compare their expressiveness. This looks
1812 like a good starting point to further connect programming language notions with cryptographic
1813 ones, as hinted in Section 8.2.1.

1814 The idea of providing high-level abstractions for cryptographic protocols in programming
1815 languages predates UC, of course. Abadi et al. [8], for example, translate a high-level language with
1816 an abstraction for channels to a low-level language that implements this channel via encrypted
1817 and channel communication. On the first glance, this provides *less* flexibility: the abstraction
1818 result corresponds not to UC as a concept, but to a singular emulation result. On the other hand,
1819 this approach allows, in principle, a tighter integration into the language. Consider, for example,
1820 communication protocols, which are better abstracted as a channel, versus cryptographic primitives
1821 like encryption, which are better abstracted to symbolic terms. Exploring whether UC results can
1822 be related to *RC* this way – which is competing to the approach presented in this work – remains
1823 an open problem.

1824 *RC Works*. The theory of robust compilation was recently devised by Abate et al. [14] and later
1825 expanded to account for differences between the source and target trace models [13]. We believe
1826 our connection does not need to account for source and target trace difference since UC deals with
1827 the same language for protocols and functionalities.

1828 A novel criterion for secure compilation, *RC* has been compared to existing secure compilation
1829 notions such as fully-abstract compilation [15]. We reported that fully abstract compilation is not a
1830 good candidate for our connection in Section 8.2.2.

1831 *RC* has been used to reason about the preservation of arbitrary safety properties through compila-
1832 tion [83, 84] and through translation validation [35] even with mechanised proofs [51]. Additionally,
1833 it has been used to reason about the absence of leaks due to speculation attacks such as Spectre [85].

1834 10 CONCLUSION

1835 This paper presented a striking connection between the cryptographic framework of Universal
1836 Composability (UC) and the secure compilation criterion called Robust Hyperproperty Preservation
1837 (*RHP*). This paper first formalised (and mechanised in Isabelle) this connection and then identified
1838 a the requirements for lifting this connection to an *RHP*-compiler between arbitrary languages.
1839 Then, the paper presented a formal language that fulfills these requirements, encoded a single-
1840 commitment bit protocol in that language, and proved that the protocol is UC by proving that the
1841 compiler generating that protocol attains *RHP*. Finally, the paper mechanised the *RHP* proof in
1842 Deepsec for both the static and the dynamic corruption cases, providing a scalable, mechanised
1843 proof of UC via our connection.

1844 ACKNOWLEDGMENTS

1845 The authors would like to thank Carmine Abate, Amal Ahmed, Dan Boneh, Deepak Garg, John
1846 Mitchell, Jeremy Thibault for useful feedback and discussions. This work was partially supported
1847 by the Office of Naval Research for support through grant N00014-18-1-2620, Accountable Protocol
1848 Customization; the German Federal Ministry of Education and Research (BMBF) through funding for
1849 the CISPA-Stanford Center for Cybersecurity (FKZ: 13N1S0762); the Italian Ministry of Education
1850 through funding for the Rita Levi Montalcini grant (call of 2019).

REFERENCES

- 1851
- 1852 [1] Martín Abadi. 1998. Protection in Programming-Language Translations. In *ICALP'98*. 868–883.
- 1853 [2] Martín Abadi, Bowen Alpern, Krzysztof R. Apt, Nissim Francez, Shmuel Katz, Leslie Lamport, and Fred B. Schneider.
- 1854 1991. Preserving Liveness: Comments on "Safety and Liveness from a Methodological Point of View". *Inf. Process. Lett.*
- 1855 40, 3 (1991), 141–142. DOI : [http://dx.doi.org/10.1016/0020-0190\(91\)90168-H](http://dx.doi.org/10.1016/0020-0190(91)90168-H)
- 1856 [3] Martín Abadi and Véronique Cortier. 2006. Deciding Knowledge in Security Protocols under Equational Theories.
- 1857 *Theoretical Computer Science* 367, 1-2 (Nov. 2006), 2–32. DOI : <http://dx.doi.org/10.1016/j.tcs.2006.08.032>
- 1858 [4] Martín Abadi and Cédric Fournet. 2001. Mobile Values, New Names, and Secure Communication. In *POPL*. ACM,
- 1859 104–115.
- 1860 [5] Martín Abadi and Cédric Fournet. 2001. Mobile Values, New Names, and Secure Communication. *SIGPLAN Not.* 36, 3
- 1861 (Jan. 2001), 104–115. DOI : <http://dx.doi.org/10.1145/373243.360213>
- 1862 [6] Martín Abadi, Cédric Fournet, and Georges Gonthier. 1998. Secure Implementation of Channel Abstractions. In *IEEE*
- 1863 *Symposium on Logic in Computer Science*. 105–116.
- 1864 [7] Martín Abadi, Cédric Fournet, and Georges Gonthier. 2000. Authentication Primitives and their Compilation. In
- 1865 *Principles of Programming Languages*. ACM, 302–315.
- 1866 [8] Martín Abadi, Cédric Fournet, and Georges Gonthier. 2002. Secure Implementation of Channel Abstractions. *Information*
- 1867 *and Computation* 174, 1 (2002), 37–83. DOI : <http://dx.doi.org/10.1006/inco.2002.3086>
- 1868 [9] Martín Abadi and Andrew D. Gordon. 1999. A Calculus for Cryptographic Protocols: The spi Calculus. *Inf. Comput.*
- 1869 148, 1 (1999), 1–70.
- 1870 [10] Martín Abadi and Gordon D. Plotkin. 2012. On Protection by Layout Randomization. *ACM Transactions on Information*
- 1871 *and System Security* 15, Article 8 (July 2012), 8:1–8:29 pages.
- 1872 [11] Carmine Abate, Arthur Azevedo de Amorim, Roberto Blanco, Ana Nora Evans, Guglielmo Fachini, Catalin Hritcu,
- 1873 Théo Laurent, Benjamin C. Pierce, Marco Stronati, and Andrew Tolmach. 2018. When Good Components Go Bad:
- 1874 Formally Secure Compilation Despite Dynamic Compromise (*CCS '18*).
- 1875 [12] Carmine Abate, Roberto Blanco, Ștefan Ciobăcă, Adrien Durier, Deepak Garg, Catalin Hritcu, Marco Patrignani, Éric
- 1876 Tanter, and Jeremy Jérémy. 2020. Trace-Relating Compiler Correctness and Secure Compilation. In *Programming*
- 1877 *Languages and Systems - 29th European Symposium on Programming, ESOP 2020 (ESOP)*. 1–28. DOI : [http://dx.doi.org/10.](http://dx.doi.org/10.1007/978-3-030-44914-8_1)
- 1878 [1007/978-3-030-44914-8_1](http://dx.doi.org/10.1007/978-3-030-44914-8_1)
- 1879 [13] Carmine Abate, Roberto Blanco, Ștefan Ciobăcă, Adrien Durier, Deepak Garg, Cătălin Hrițcu, Marco Patrignani,
- 1880 Éric Tanter, and Jérémy Thibault. 2021. An Extended Account of Trace-Relating Compiler Correctness and Secure
- 1881 Compilation. *ACM Trans. Program. Lang. Syst.* 43, 4, Article 14 (nov 2021), 48 pages. DOI : [http://dx.doi.org/10.1145/](http://dx.doi.org/10.1145/3460860)
- 1882 [3460860](http://dx.doi.org/10.1145/3460860)
- 1883 [14] Carmine Abate, Roberto Blanco, Deepak Garg, Cătălin Hrițcu, Marco Patrignani, and Jérémy Thibault. 2019. Journey
- 1884 Beyond Full Abstraction: Exploring Robust Property Preservation for Secure Compilation. In *2019 IEEE 32th Computer*
- 1885 *Security Foundations Symposium (CSF 2019)*.
- 1886 [15] Carmine Abate, Matteo Busi, and Stelios Tsampas. 2021. Fully Abstract and Robust Compilation: And How to Reconcile
- 1887 the Two, Abstractly. In *Programming Languages and Systems - 19th Asian Symposium, APLAS 2021, Chicago, IL, USA,*
- 1888 *October 17-18, 2021, Proceedings (Lecture Notes in Computer Science)*, Hakjoo Oh (Ed.), Vol. 13008. Springer, 83–101.
- 1889 DOI : http://dx.doi.org/10.1007/978-3-030-89051-3_6
- 1890 [16] Michel Abdalla, Manuel Barbosa, Tatiana Bradley, Stanisław Jarecki, Jonathan Katz, and Jiayu Xu. 2020. Universally
- 1891 Composable Relaxed Password Authenticated Key Exchange. In *Advances in Cryptology – CRYPTO 2020 (Lecture Notes*
- 1892 *in Computer Science)*, Daniele Micciancio and Thomas Ristenpart (Eds.). Springer International Publishing, Cham,
- 1893 278–307. DOI : http://dx.doi.org/10.1007/978-3-030-56784-2_10
- 1894 [17] Coşku Acay, Rolph Recto, Joshua Gancher, Andrew C. Myers, and Elaine Shi. 2021. Viaduct: An Extensible, Optimizing
- 1895 Compiler for Secure Distributed Programs. In *Proceedings of the 42nd ACM SIGPLAN International Conference on*
- 1896 *Programming Language Design and Implementation (PLDI 2021)*. Association for Computing Machinery, New York, NY,
- 1897 USA, 740–755. DOI : <http://dx.doi.org/10.1145/3453483.3454074>
- 1898 [18] Luca Aceto, Anna Ingólfssdóttir, Kim Guldstrand Larsen, and Jiri Srba. 2007. *Reactive Systems: Modelling, Specification*
- 1899 *and Verification*. Cambridge University Press.
- 1900 [19] Amal Ahmed and Matthias Blume. 2008. Typed Closure Conversion Preserves Observational Equivalence. In *International*
- 1901 *Conference on Functional Programming*. ACM, 157–168.
- 1902 [20] Amal Ahmed and Matthias Blume. 2011. An Equivalence-Preserving CPS Translation via Multi-Language Semantics. In
- 1903 *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (ICFP '11)*. ACM, 431–444.
- 1904 [21] Mihhail Aizatulin, Andrew D. Gordon, and Jan Jürjens. 2011. Extracting and Verifying Cryptographic Models from C
- 1905 Protocol Code by Symbolic Execution. In *Proceedings of the 18th ACM Conference on Computer and Communications*
- 1906 *Security - CCS '11*. ACM Press, Chicago, Illinois, USA, 331. DOI : <http://dx.doi.org/10.1145/2046707.2046745>

- 1907 [22] Michael Backes, Robert Künnemann, and Esfandiar Mohammadi. 2016. Computational Soundness for Dalvik Bytecode.
 1908 In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*. Association
 1909 for Computing Machinery, New York, NY, USA, 717–730. DOI: <http://dx.doi.org/10.1145/2976749.2978418>
- 1910 [23] Michael Backes, Birgit Pfitzmann, and Michael Waidner. 2004. Secure Asynchronous Reactive Systems. (04 2004).
 1911 Technical report 082.
- 1912 [24] Michael Backes, Birgit Pfitzmann, and Michael Waidner. 2007. The reactive simulatability (RSIM) framework for
 1913 asynchronous systems. *Inf. Comput.* 205, 12 (2007), 1685–1720. DOI: <http://dx.doi.org/10.1016/j.ic.2007.05.002>
- 1914 [25] David Baelde, Stéphanie Delaune, Charlie Jacomme, Adrien Koutsos, and Solène Moreau. 2021. An Interactive Prover
 1915 for Protocol Verification in the Computational Model. In *SP 2021 - 42nd IEEE Symposium on Security and Privacy*
 1916 (*Proceedings of the 42nd IEEE Symposium on Security and Privacy (s & P'21)*). San Fransisco / Virtual, United States.
- 1917 [26] Manuel Barbosa, Gilles Barthe, Karthik Bhargavan, Bruno Blanchet, Cas Cremers, Kevin Liao, and Bryan Parno.
 1918 2021. SoK: Computer-Aided Cryptography. In *2021 IEEE Symposium on Security and Privacy (SP)*. 777–795. DOI:
 1919 <http://dx.doi.org/10.1109/SP40001.2021.00008>
- 1920 [27] Manuel Barbosa, Gilles Barthe, Benjamin Grégoire, Adrien Koutsos, and Pierre-Yves Strub. 2021. Mechanized Proofs
 1921 of Adversarial Complexity and Application to Universal Composability. In *Proceedings of the 2021 ACM SIGSAC*
 1922 *Conference on Computer and Communications Security*. ACM, Virtual Event Republic of Korea, 2541–2563. DOI:
 1923 <http://dx.doi.org/10.1145/3460120.3484548>
- 1924 [28] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. 2011. Computer-Aided Security
 1925 Proofs for the Working Cryptographer. In *Advances in Cryptology – CRYPTO 2011*, Phillip Rogaway (Ed.). Springer
 1926 Berlin Heidelberg, Berlin, Heidelberg, 71–90.
- 1927 [29] Gilles Barthe, Benjamin Grégoire, Charlie Jacomme, Steve Kremer, and Pierre-Yves Strub. 2019. Symbolic Methods in
 1928 Computational Cryptography Proofs. In *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*. 136–13615.
 1929 DOI: <http://dx.doi.org/10.1109/CSF.2019.00017>
- 1930 [30] David Basin, Andreas Lochbihler, Ueli Maurer, and S. Reza Sefidgar. 2021. Abstract Modeling of System Communication
 1931 in Constructive Cryptography Using CryptHOL. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*.
 1932 1–16. DOI: <http://dx.doi.org/10.1109/CSF51468.2021.00047>
- 1933 [31] Mihir Bellare and Phillip Rogaway. 2004. *The Game-Playing Technique*. Technical Report.
- 1934 [32] B. Blanchet. 2008. A Computationally Sound Mechanized Prover for Security Protocols. *IEEE Trans. Dependable and*
 1935 *Secure Comput.* 5, 4 (Oct. 2008), 193–207. DOI: <http://dx.doi.org/10.1109/TDSC.2007.1005>
- 1936 [33] Florian Böhl and Dominique Unruh. 2013. Symbolic Universal Composability. In *2013 IEEE 26th Computer Security*
 1937 *Foundations Symposium*. 257–271. DOI: <http://dx.doi.org/10.1109/CSF.2013.24>
- 1938 [34] William J. Bowman and Amal Ahmed. 2015. Noninterference for free. In *ICFP*. ACM.
- 1939 [35] Matteo Busi, Pierpaolo Degano, and Letterio Galletta. 2022. Towards Effective Preservation of Robust Safety Properties.
 1940 In *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing (SAC '22)*. Association for Computing
 1941 Machinery, New York, NY, USA, 1674–1683. DOI: <http://dx.doi.org/10.1145/3477314.3507084>
- 1942 [36] Jan Camenisch, Stephan Krenn, Ralf Küsters, and Daniel Rausch. 2019. iUC: Flexible Universal Composability Made
 1943 Simple. In *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of*
 1944 *Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part III (Lecture Notes in Computer*
 1945 *Science)*, Vol. 11923. Springer, 191–221.
- 1946 [37] R. Canetti. 2001. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *Proceedings*
 1947 *42nd IEEE Symposium on Foundations of Computer Science*. IEEE, Newport Beach, CA, USA, 136–145. DOI: <http://dx.doi.org/10.1109/SFCS.2001.959888>
- 1948 [38] Ran Canetti. 2001. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *Proceedings of*
 1949 *the 42Nd IEEE Symposium on Foundations of Computer Science (FOCS '01)*. IEEE Computer Society, Washington, DC,
 1950 USA, 136–. <http://dl.acm.org/citation.cfm?id=874063.875553>
- 1951 [39] Ran Canetti. 2020. Universally Composable Security. *J. ACM* 67, 5 (Sept. 2020), 28:1–28:94. DOI: <http://dx.doi.org/10.1145/3402457>
- 1952 [40] Ran Canetti, Asaf Cohen, and Yehuda Lindell. 2014. A Simpler Variant of Universally Composable Security for Standard
 1953 Multiparty Computation. Cryptology ePrint Archive, Report 2014/553. (2014). <https://eprint.iacr.org/2014/553>.
- 1954 [41] Ran Canetti and Marc Fischlin. 2001. Universally Composable Commitments. In *Advances in Cryptology – CRYPTO*
 1955 *2001*, Joe Kilian (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 19–40.
- 1956 [42] Ran Canetti, Palak Jain, Marika Swanberg, and Mayank Varia. Universally Composable End-to-End Secure Messaging:
 1957 A Modular Analysis. In *CRYPTO'2022*. 78.
- 1958 [43] Ran Canetti, Alley Stoughton, and Mayank Varia. 2019. EasyUC: Using EasyCrypt to Mechanize Proofs of Universally
 1959 Composable Security. In *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*. 167–16716. DOI: <http://dx.doi.org/10.1109/CSF.2019.00019>
- 1960
 1961
 1962

- 1963 [44] Vincent Cheval, Steve Kremer, and Itsaka Rakotonirina. 2018. DEEPSEC: Deciding Equivalence Properties in Security
 1964 Protocols Theory and Practice. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, 529–546.
 1965 DOI : <http://dx.doi.org/10.1109/SP.2018.00033>
- 1966 [45] Michael R. Clarkson and Fred B. Schneider. 2010. Hyperproperties. *Journal of Computer Security* 18, 6 (2010), 1157–1210.
 1967 DOI : <http://dx.doi.org/10.3233/JCS-2009-0393>
- 1968 [46] Anupam Datta, Ralf Küsters, John C. Mitchell, and Ajith Ramanathan. 2005. On the Relationships Between Notions of
 1969 Simulation-based Security. In *Proceedings of the Second International Conference on Theory of Cryptography (TCC'05)*.
 1970 Springer-Verlag, 476–494. DOI : http://dx.doi.org/10.1007/978-3-540-30576-7_26
- 1971 [47] Stéphanie Delaune and Lucca Hirschi. 2017. A Survey of Symbolic Methods for Establishing Equivalence-Based
 1972 Properties in Cryptographic Protocols. *Journal of Logical and Algebraic Methods in Programming* 87 (Feb. 2017),
 1973 127–144. DOI : <http://dx.doi.org/10.1016/j.jlamp.2016.10.005>
- 1974 [48] Stéphanie Delaune, Steve Kremer, and Olivier Pereira. 2009. Simulation Based Security in the Applied Pi Calculus.
 1975 *Cryptology ePrint Archive* (2009).
- 1976 [49] Dominique Devriese, Marco Patrignani, Frank Piessens, and Steven Keuchel. 2017. Modular, Fully-abstract Compilation
 1977 by Approximate Back-translation. *Logical Methods in Computer Science* Volume 13, Issue 4 (Oct. 2017).
- 1978 [50] Akram El-Korashy, Roberto Blanco, Jérémy Thibault, Adrien Durier, Deepak Garg, and Catalin Hritcu. 2022. Se-
 1979 curePtrs: Proving Secure Compilation with Data-Flow Back-Translation and Turn-Taking Simulation. (2022).
 1980 arXiv:cs.PL/2110.01439
- 1981 [51] Akram El-Korashy, Stelios Tsampas, Marco Patrignani, Dominique Devriese, Deepak Garg, and Frank Piessens. 2021.
 1982 CapablePtrs: Securely Compiling Partial Programs Using the Pointers-as-Capabilities Principle. In *34th IEEE Computer
 1983 Security Foundations Symposium, CSF 2021, Dubrovnik, Croatia, June 21-25, 2021*. 1–16. DOI : [http://dx.doi.org/10.1109/
 1985 CSF51468.2021.00036](http://dx.doi.org/10.1109/

 1984 CSF51468.2021.00036)
- 1985 [52] Cédric Fournet, Markulf Kohlweiss, and Pierre-Yves Strub. 2011. Modular Code-Based Cryptographic Verification. In
 1986 *Proceedings of the 18th ACM Conference on Computer and Communications Security - CCS '11*. ACM Press, Chicago,
 1987 Illinois, USA, 341. DOI : <http://dx.doi.org/10.1145/2046707.2046746>
- 1988 [53] Cédric Fournet, Markulf Kohlweiss, and Pierre-Yves Strub. 2011. Modular Code-Based Cryptographic Verification. In
 1989 *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS '11)*. Association for Computing
 1990 Machinery, New York, NY, USA, 341–350. DOI : <http://dx.doi.org/10.1145/2046707.2046746>
- 1991 [54] Oded Goldreich. 1990. A Note on Computational Indistinguishability. *Inf. Process. Lett.* 34, 6 (1990), 277–281. DOI :
 1992 [http://dx.doi.org/10.1016/0020-0190\(90\)90010-U](http://dx.doi.org/10.1016/0020-0190(90)90010-U)
- 1993 [55] Roberto Guanciale, Musard Balliu, and Mads Dam. 2020. InSpectre: Breaking and Fixing Microarchitectural Vulnerabil-
 1994 ities by Formal Analysis. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*.
 1995 ACM, Virtual Event USA, 1853–1869. DOI : <http://dx.doi.org/10.1145/3372297.3417246>
- 1996 [56] H. Haagh, A. Karbyshev, S. Oechsner, B. Spitters, and P. Strub. 2018. Computer-Aided Proofs for Multiparty Computation
 1997 with Active Security. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, Los
 1998 Alamitos, CA, USA, 119–131. DOI : <http://dx.doi.org/10.1109/CSF.2018.00016>
- 1999 [57] Marcella Hastings, Brett Hemenway, Daniel Noble, and Steve Zdancewic. 2019. SoK: General Purpose Compilers
 2000 for Secure Multi-Party Computation. In *2019 IEEE Symposium on Security and Privacy (SP)*. 1220–1237. DOI : [http://
 2002 dx.doi.org/10.1109/SP.2019.00028](http://

 2001 dx.doi.org/10.1109/SP.2019.00028)
- 2002 [58] Michael Hicks. 2014. Formal Reasoning in PL and Crypto. (2014). Available at: [http://www.pl-
 2004 enthusiast.net/2014/12/23/formal-reasoning-pl-crypto/](http://www.pl-

 2003 enthusiast.net/2014/12/23/formal-reasoning-pl-crypto/).
- 2004 [59] Martin Hirt, Chen-Da Liu-Zhang, and Ueli Maurer. 2021. Adaptive Security of Multi-Party Protocols, Revisited. In
 2005 *Theory of Cryptography Conference*.
- 2006 [60] Dennis Hofheinz and Victor Shoup. 2011. GNUC: A New Universal Composability Framework. *Cryptology ePrint
 2007 Archive*. (2011). <http://eprint.iacr.org/>
- 2008 [61] Dennis Hofheinz and Dominique Unruh. 2006. Simulatable security and polynomially bounded concurrent compos-
 2009 ability. In *IEEE Symposium on Security and Privacy*.
- 2010 [62] Dennis Hofheinz, Dominique Unruh, and Jörn Müller-Quade. 2013. Polynomial runtime and composability. *Journal of
 2011 Cryptology* 26, 3 (2013), 375–441.
- 2012 [63] Dennis Hofheinz, Dominique Unruh, and Jörn Müller-Quade. 2013. Polynomial Runtime and Composability. *J Cryptol*
 2013 26, 3 (July 2013), 375–441. DOI : <http://dx.doi.org/10.1007/s00145-012-9127-4>
- 2014 [64] Chung-Kil Hur and Derek Dreyer. 2011. A Kripke Logical Relation Between ML and Assembly. In *Principles of
 2015 Programming Languages*. 133–146.
- 2016 [65] Alan Jeffrey and Julian Rathke. 2005. Java Jr.: Fully abstract trace semantics for a core Java language. In *ESOP'05
 2017 (LNCS)*, Vol. 3444. Springer, 423–438.
- 2018 [66] Matthias Kruse and Marco Patrignani. 2022. Composing Secure Compilers.

- [67] Ralf Kuesters, Max Tuengerthal, and Daniel Rausch. 2013. *The IITM Model: A Simple and Expressive Model for Universal Composability*. Technical Report 025.
- [68] Robert Künnemann, Marco Patrignani, and Ethan Cecchetti. Computationally Bounded Robust Compilation and Universally Composable Security. In *CSF 2024*.
- [69] Ralf Küsters. 2006. Simulation-Based Security with Inexhaustible Interactive Turing Machines. In *Computer Security Foundations Workshop*. IEEE Computer Society, 309–320.
- [70] Ralf Kusters, Tomasz Truderung, Bernhard Beckert, Daniel Bruns, Michael Kirsten, and Martin Mohr. 2015. A Hybrid Approach for Proving Noninterference of Java Programs. In *2015 IEEE 28th Computer Security Foundations Symposium*. IEEE, Verona, 305–319. DOI : <http://dx.doi.org/10.1109/CSF.2015.28>
- [71] Leslie Lamport and Fred B. Schneider. 1984. Formal Foundation for Specification and Verification. In *Distributed Systems: Methods and Tools for Specification, An Advanced Course*. 203–285. DOI : http://dx.doi.org/10.1007/3-540-15216-4_15
- [72] Xavier Leroy. 2009. A Formally Verified Compiler Back-end. *Journal of Automated Reasoning* 43, 4 (2009), 363–446. <http://dx.doi.org/10.1007/s10817-009-9155-4>
- [73] Kevin Liao, Matthew A. Hammer, and Andrew Miller. 2019. ILC: A Calculus for Composable, Computational Cryptography. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2019)*. ACM, 640–654. DOI : <http://dx.doi.org/10.1145/3314221.3314607>
- [74] Jacob Matthews and Robert Bruce Findler. 2009. Operational Semantics for Multi-language Programs. *ACM Trans. Program. Lang. Syst.* 31, 3 (April 2009), 12:1–12:44.
- [75] Ueli Maurer. 2011. Constructive Cryptography - A New Paradigm for Security Definitions and Proofs. In *Theory of Security and Applications - Joint Workshop, TOSCA 2011, Saarbrücken, Germany, March 31 - April 1, 2011, Revised Selected Papers (Lecture Notes in Computer Science)*, Sebastian Mödersheim and Catuscia Palamidessi (Eds.), Vol. 6993. Springer, 33–56. DOI : http://dx.doi.org/10.1007/978-3-642-27375-9_3
- [76] John C. Mitchell, Rahul Sharma, Deian Stefan, and Joe Zimmerman. 2012. Information-Flow Control for Programming on Encrypted Data. In *2012 IEEE 25th Computer Security Foundations Symposium*. 45–60. DOI : <http://dx.doi.org/10.1109/CSF.2012.30>
- [77] Greg Morrisett, Elaine Shi, Kristina Sojakova, Xiong Fan, and Joshua Gancher. 2021. IPDL: A Simple Framework for Formally Verifying Distributed Cryptographic Protocols. Cryptology ePrint Archive, Paper 2021/147. (2021). <https://eprint.iacr.org/2021/147> <https://eprint.iacr.org/2021/147>.
- [78] Marco Patrignani. 2020. Why Should Anyone use Colours? or, Syntax Highlighting Beyond Code Snippets. (2020). arXiv:cs.SE/2001.11334
- [79] Marco Patrignani, Pieter Agten, Raoul Strackx, Bart Jacobs, Dave Clarke, and Frank Piessens. 2015. Secure Compilation to Protected Module Architectures. *ACM Trans. Program. Lang. Syst.* 37, Article 6 (April 2015), 6:1–6:50 pages.
- [80] Marco Patrignani, Amal Ahmed, and Dave Clarke. 2019. Formal Approaches to Secure Compilation A Survey of Fully Abstract Compilation and Related Work. *ACM Comput. Surv.* 51, 6, Article 125 (Jan. 2019), 36 pages.
- [81] Marco Patrignani and Dave Clarke. 2015. Fully abstract trace semantics for protected module architectures. *Elsevier COMLAN* 42, 0 (2015), 22 – 45. DOI : <http://dx.doi.org/10.1016/j.cl.2015.03.002>
- [82] Marco Patrignani, Dominique Devriese, and Frank Piessens. 2016. On Modular and Fully Abstract Compilation. In *Computer Security Foundations Symposium*.
- [83] Marco Patrignani and Deepak Garg. 2019. Robustly Safe Compilation. In *Programming Languages and Systems - 28th European Symposium on Programming, ESOP 2019 (ESOP'19)*.
- [84] Marco Patrignani and Deepak Garg. 2021. Robustly Safe Compilation, an Efficient Form of Secure Compilation. *ACM Trans. Program. Lang. Syst.* 43, 1, Article 1 (feb 2021), 41 pages. DOI : <http://dx.doi.org/10.1145/3436809>
- [85] Marco Patrignani and Marco Guarnieri. 2021. Exorcising Spectres with Secure Compilers. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21)*. Association for Computing Machinery, New York, NY, USA, 445–461. DOI : <http://dx.doi.org/10.1145/3460120.3484534>
- [86] Marco Patrignani, Riad S. Wahby, and Robert Künnemann. 2019. Universal Composability is Secure Compilation. (2019). arXiv:1910.08634
- [87] James T. Perconti and Amal Ahmed. 2014. Verifying an Open Compiler Using Multi-language Semantics. In *ESOP (LNCS)*, Vol. 8410. 128–148.
- [88] Benjamin Pierce and Eijiro Sumii. 2000. Relating Cryptography and Polymorphism. (2000). <http://www.kb.ecei.tohoku.ac.jp/~sumii/pub/infohide.pdf> manuscript.
- [89] Gordon D. Plotkin. 1977. LCF Considered as a Programming Language. *Theoretical Computer Science* 5 (1977), 223–255.
- [90] Daniel Rausch, Ralf Küsters, and Céline Chevalier. 2022. Embedding the UC Model into the IITM Model. In *Advances in Cryptology – EUROCRYPT 2022 (Lecture Notes in Computer Science)*, Orr Dunkelman and Stefan Dziembowski (Eds.). Springer International Publishing, Cham, 242–272. DOI : http://dx.doi.org/10.1007/978-3-031-07085-3_9
- [91] Davide Sangiorgi and David Walker. 2001. *PI-Calculus: A Theory of Mobile Processes*. Cambridge University Press, New York, NY, USA.

- 2076 [92] Fred B. Schneider. 2000. Enforceable security policies. *ACM Transactions of Information Systems Security* 3, 1 (2000),
2077 30–50. <http://www.cs.cornell.edu/fbs/publications/EnfSecPols.pdf>
- 2078 [93] Eijiro Sumii and Benjamin C. Pierce. 2003. Logical Relations for Encryption. *J. Comput. Secur.* 11, 4 (July 2003), 521–554.
- 2079 [94] Eijiro Sumii and Benjamin C. Pierce. 2004. A Bisimulation for Dynamic Sealing. In *Principles of Programming Languages*.
2080 161–172.
- 2081 [95] Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan
2082 Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-
2083 Béguélin. 2016. Dependent Types and Multi-Monadic Effects in F*. In *Proceedings of the 43rd Annual ACM SIGPLAN-*
2084 *SIGACT Symposium on Principles of Programming Languages (POPL '16)*. Association for Computing Machinery, New
2085 York, NY, USA, 256–270. DOI:<http://dx.doi.org/10.1145/2837614.2837655>