

Abstract Modeling of Systems Communication in Constructive Cryptography using CryptHOL

David Basin
Dept. of Computer Science
ETH Zürich, Switzerland
basin@inf.ethz.ch

Andreas Lochbihler
Digital Asset
Zurich, Switzerland
mail@andreas-lochbihler.de

Ueli Maurer
Dept. of Computer Science
ETH Zürich, Switzerland
maurer@inf.ethz.ch

S. Reza Sefidgar
Dept. of Computer Science
ETH Zürich, Switzerland
reza.sefidgar@inf.ethz.ch

Abstract—Proofs in simulation-based frameworks have the greatest rigor when they are machine checked. But the level of details in these proofs surpasses what the formal-methods community can handle with existing tools. Existing formal results consider streamlined versions of simulation-based frameworks to cope with their complexity. Hence, a central question is how to abstract details from composability results and enable their formal verification.

In this paper, we focus on the modeling of systems communication in composable security statements. The existing formal models consider fixed communication patterns to reduce the complexity of their proofs. However, as we will show, this can affect the reusability of security statements. We propose an abstract approach to modeling systems communication in Constructive Cryptography that avoids this problem. Our approach is suitable for mechanized verification and we use CryptHOL, a framework for developing mechanized cryptography proofs, to implement it in the Isabelle/HOL theorem prover. As a case study, we formalize the construction of a secure channel using the Diffie-Hellman key exchange and a one-time-pad.

Note to reviewers: All results presented in this paper are formalized using Isabelle/HOL. We provide this formalization as a supplementary material, which is available online [1], and explain its structure in Appendix A.

I. INTRODUCTION

A. Problem context

Simulation-based frameworks [2], [7], [11], [13], also called composability frameworks, pave the way for modular reasoning about cryptographic constructs independently of their application context. They follow the *Ideal-world—Real-world* [9] paradigm to security proofs, where a distributed protocol’s security is defined by comparing its execution, i.e., the real world, to an idealized specification that intrinsically satisfies the desired properties.

Devising rigorous security statements in these frameworks is very challenging. The way that notions of algorithms, runtime, corruptions, etc. are baked into composability theorems in a framework such as Universal Composability (UC) [7] introduces considerable details into composable security statements. The resulting complexity is beyond what the current scientific review processes can afford. The formal-methods community have not been able to alleviate the situation either. The level of details in simulation-based frameworks surpasses what the existing formal-methods tools can reasonably handle.

Simulation-based frameworks’ approach to modeling systems communication is one of the main reasons for the complexity of their resulting proofs. In these frameworks, components communicate via network tapes and a central scheduler called the adversary Turing machine. This concrete approach increases the complexity of ideal specifications and makes the security arguments intricate [6]. As such, the existing formal results such as EasyUC [8] consider a simplified version of the simulation-based frameworks with restricted communication capabilities between components. However, as we will show in Section III-A, such simplifications affect the reusability of security statements.

Constructive Cryptography (CC) [16], [17], [18] proposes a fundamental shift in how security statements are made and proved. It introduces an abstract approach to composable security arguments that allows one to focus on a particular aspect of security proofs without being distracted by other details. This makes composable security statements manageable for protocol designers. However, the existing CC results [12], [15] do not delve into the details of systems communication.

B. Our solution

We propose an abstract approach to modeling systems communication in composable security statements by means of generic ideal functionalities. We follow CC’s methodology whereby specifications (and implementations) are *sets of* ideal (or real respectively) functionalities. This is in contrast to UC-style frameworks that define such functionalities as single objects and enables us to model systems communication without introducing a centralized adversary machine.

First, we propose a semantic domain called *Fused Resource Templates* (FRT) that abstracts over the systems communication patterns. FRTs are suitable for mechanizing composable security proofs, and we formalize all the results presented in this paper using CryptHOL [4], [14], a framework for mechanizing cryptographic proofs in the proof assistant Isabelle/HOL [19]. Second, we show how FRTs constitute an instantiation of the CC framework. We formalize operators that enable us to compose FRTs and thereby define complex FRTs in terms of simpler ones. We then introduce the notion of secure constructions using FRTs and provide composition theorems for such constructions. Finally, we demonstrate the practicality of our approach in a case study. We formalize the construction

of a secure channel by composing the Diffie-Hellman key exchange with one-time-pad encryption.

Our result is based on Lochbihler et al. [15]’s formalization of CC. To keep this paper self contained and preserve its presentation flow, we recap this work in various parts of the paper. We provide a detailed comparison in Section VI

C. Contributions

The results presented in this paper contribute to both the cryptography and the formal-methods communities. Our approach to capturing communication patterns increases the abstraction of ideal functionalities and makes composable security arguments amenable to mechanized verification. In detail:

- We propose the first CC instantiation that explicates the details of systems communication. The role of systems communication models has not been studied in the CC literature, although we show that reusability and modularity of CC proofs depends on it (cf. Section III-A). We propose an abstract approach to capturing systems communication patterns in CC proofs.
- We formalize our approach in a theorem prover and thereby enable the mechanized verification of composable security arguments. By abstracting over systems communication patterns, we allow protocol designers to focus on the main security concerns while ensuring the reusability and modularity of security statements. We provide proof rules that enable modular reasoning about concrete (and asymptotic) security statements. We use these rules to formalize the Diffie-Hellman key exchange and one-time-pad that are comparatively simpler and shorter than any existing formalization.

D. Structure

We start with a short recap of simulation-based frameworks and CC in Section II. In Section III, we explain the importance of a precise systems communication modeling for the modularity and reusability of security proofs. We then present our formalization of FRTs and their operators. In Section IV, we define concrete security in terms of FRT constructions and provide theorems for security proof compositions. In Section V, we present our case study to demonstrate our approach’s applicability. Finally, in Sections VI and VII, we compare with related work and draw conclusions.

II. PRELIMINARIES

This section introduces background that is used in the rest of the paper. We review the composable security proofs and how it is treated in simulation-based frameworks and CC. We focus on system communication modeling in composable security arguments. In each subsection, we use the notation that is common among the research community for the topic under discussion.

A. Composable security arguments

The ideal-world—real-world [9] approach to security enables modular reasoning about security statements. In this approach, a protocol’s security is defined by comparing its execution, i.e., the real world, to an idealized specification that satisfies the desired properties by definition: a protocol π *realizes* (or implements) an idealized functionality \mathcal{F} if there exists a *simulator* that can simulate π ’s behavior (in an adversarial environment) by interacting with \mathcal{F} . The so called *Composability Theorems* extend the above idea to an arbitrary application context. Consider a protocol $\rho|\mathcal{F}$, where the $|$ operator denotes protocol composition, that realizes an ideal functionality \mathcal{G} ; then the protocol $\rho|\pi$ also realizes \mathcal{G} if the protocol π realizes the ideal functionality \mathcal{F} .

The composability of security proofs depends on how generic their communication modeling is. To enable security proofs being reused, ideal functionalities must be independent from concrete setups. For instance, suppose that F can be realized by the protocols π , $\alpha|\beta$, and $\gamma|\delta|\mathcal{H}$. Then, F should abstractly represent all of these protocols by capturing their essential properties. In particular, F must be independent from the number of sub-protocols and their interactions, i.e., the semantics of the $|$ operator, in its realizations. Tying an ideal functionality to a particular execution flow essentially prevents all the realizations that do not follow this flow.

B. UC-style frameworks

In simulation-based frameworks [2], [7], [11], [13], (sub-) protocols are modeled using interactive Turing machines with network tapes and unique identifiers. At each point of a protocol’s execution, only one of these machines is active and the others wait for new inputs. A special Turing machine, called the adversary, schedules the activation of protocol components: it is activated after each non-adversary machine halts and determines the next Turing machine to activate. The adversary plays the role of a mediator too. When two machines want to communicate, the sender informs the adversary about the message¹ content and destination, i.e., the receiving Turing machine’s identifier, and the adversary forwards the message according to its underlying corruption model.

The generality of the communication model in simulation-based frameworks stems from the central adversary Turing machine that is universally quantified in security definitions. Consider the compound protocol $\rho|\mathcal{F}$, and interpret $|$ according to the execution model explained above. ρ ’s subroutine calls to \mathcal{F} result in a transfer of the execution flow that is essential for the completion of ρ ’s execution. The subroutine call is carried out by two message broadcasts: the first message carries the “request” details while the second one conveys the “response” information. These messages need not be passed consecutively; therefore, each of ρ ’s subroutine calls may correspond to an arbitrary sequence of Turing machine activations starting with

¹A message can be any information that is transferred between protocol components, including meta-level information, subroutine calls, or protocol-specific data.

ρ 's "request" message and ending with \mathcal{F} 's (or any of its realization's) "response" message. As such, \mathcal{F} 's semantics abstractly captures the execution of an arbitrary number (and order) of Turing machines realizing it.

The simulation-based frameworks' concrete approach to modeling can affect the rigor of their resulting proofs. As Camenisch et al. put it [6], protocol descriptions in simulation-based frameworks include meta-level and model specific information that makes them unnecessarily complicated; but, ignoring such details can lead to ill-defined specifications, sketchy proofs, and flawed results. The existing formal methods tools have not been able to alleviate the situation either. The level of details in simulation-based frameworks surpasses what the formal-methods community can reasonably handle with existing techniques. As such, the existing formalization results consider stripped-down versions of simulation-based frameworks by restricting the communication and corruption models [8]. Hence, they are not applicable in a wide range of scenarios.

C. Constructive Cryptography

Constructive Cryptography (CC) [16], [17], [18] is an abstract paradigm for developing a theory of cryptography. Instead of focusing on concrete systems and proving their properties, CC studies system classes, i.e., the shared behavior of similar systems, and their transformations. As such, the notions of runtime, algorithms, and complexity are not intrinsic parts of every definition and proof in CC; they just represent classification strategies.

Modular reasoning about system classifications, which are called *specifications* in CC, is based on the concepts of *Resources* and *Converters*. Every aspect of individual systems, including the adversary's capabilities and information leakage is made explicit as a resource with named input/output interfaces. A specification is a set of resources. Parties can change a resource's behavior by attaching converters to their designated interfaces on that resource. It is possible to access multiple resources at the same time and attach many converters. For example, $x^a y^b \triangleright [R, S]$ denotes the attachment of converters x and y to resources R and S , where a and b are injective mappings between converter and resource interfaces and $[_, _]$ denotes the parallel access to two resources.² One can combine converters and interface attachments into a single *Protocol*³ π and use the notation $\pi[R, S]$ instead. Let $\pi\mathcal{R} = \{\pi R \mid R \in \mathcal{R}\}$ denote the lifting of a protocol's attachment to specifications and define $[\mathcal{R}, \mathcal{S}] = \{[R, S] \mid R \in \mathcal{R} \wedge S \in \mathcal{S}\}$. According to CC's terminology, π *constructs* the specification \mathcal{S} from \mathcal{R} if and only if $\pi\mathcal{R} \subseteq \mathcal{S}$. In the ideal-world—real-world terminology, $\pi\mathcal{R}$ and \mathcal{S} play the role of the real-world implementation and ideal functionality respectively. The following *composability properties* hold for any protocols π and π' and arbitrary specifications \mathcal{R} , \mathcal{S} , and \mathcal{T} :

- 1) If $\pi\mathcal{R} \subseteq \mathcal{S}$ and $\pi'\mathcal{S} \subseteq \mathcal{T}$, then $\pi'\pi\mathcal{R} \subseteq \mathcal{T}$.

²The order of resources is unimportant in this notation; however, one can reorder converters only if they attach to disjoint sets of interfaces.

³The meaning of the term "protocol" here is different from UC-style frameworks, where it refers to one or more Turing machines working together.

- 2) If $\pi\mathcal{R} \subseteq \mathcal{S}$, then $\pi[\mathcal{R}, \mathcal{T}] \subseteq [\mathcal{S}, \mathcal{T}]$.

Common security notions are expressed as particular forms of specifications. For example, consider the simulator-based notion of information-theoretic security. Using a simulator σ corresponds to specifications of the form $\sigma\mathcal{S}$, where the simulator is defined in terms of a converter that only attaches to the adversary interfaces. Indistinguishability is analysed by *relaxing* specifications as $\mathcal{R}^\epsilon = \{S \mid R \in \mathcal{R} \wedge \mathfrak{d}(R, S) \leq \epsilon\}$, where $\mathfrak{d}(R, S)$ is the least upper bound on the advantages of all distinguishers in distinguishing the resources R and S . To define information-theoretic security, we combine these two forms: in a two-party setting with parties A and B and the adversary E, the construction of a specification \mathcal{S} from \mathcal{R} using the protocol $\pi = x^A y^B$ is information-theoretically secure iff there exist a simulator $\sigma = z^E$ such that $\pi\mathcal{R} \subseteq (\sigma\mathcal{S})^\epsilon$. Computational security is defined similarly using computational indistinguishability relaxation $\square^{<\epsilon}$, which is defined like \square^ϵ except that it considers computationally bounded distinguishers. In the above definition, note that we are using party names as a placeholder for interface mappings. Henceforth we drop the distinction between simulators and the converter that they attach and refer to both σ and z as the simulator.

To better understand the above concepts, consider the following example. A secure communication channel between two parties Alice and Bob can be established using a shared key and an authenticated channel: the parties use the key to encrypt their message and transmit the ciphertext via the authenticated channel. Figure 1a is a pictorial model of such scenario in CC. Consider the protocol π that consists of the encryption converter *Enc* and the decryption converter *Dec*. Alice and Bob attach *Enc* and *Dec* to their sides respectively. The adversary Eve controls the communication network but does not have access to the pre-shared keys. This is modeled using the adversary interface on resources: the authenticated channel resource *Auth* allows the adversary to look at (but not edit) the channel's content and drop or delay messages; however, the key resource *Key* does not leak any information through its adversary interface.

Security is expressed using specifications. Let $\mathcal{R}_{auth-key}$ denote the specification that only contains the parallel composition of the *Auth* and *Key* resources and let \mathcal{S}_{sec} denote the singleton set that contains the ideal secure channel *Sec*. Here, *Sec* is a resource that is similar to *Auth* except that it leaks the length of the channel's content. The encryption scheme used by the converters *Enc* and *Dec* is information-theoretically secure if there exists a simulator σ , represented as the *Sim* converter in the diagram, for which $\pi\mathcal{R}_{auth-key} \subseteq (\sigma\mathcal{S}_{sec})^\epsilon$. So, to prove the encryption scheme secure, it suffices to prove the indistinguishability of the two composed resources (the grey rectangles) in Figures 1a and 1b.

III. A NEW CC SEMANTIC DOMAIN

We propose Fused Resource Templates (FRT) as a class of CC specifications that allow us to abstract over the communication patterns in security proofs and is suitable for mechanized

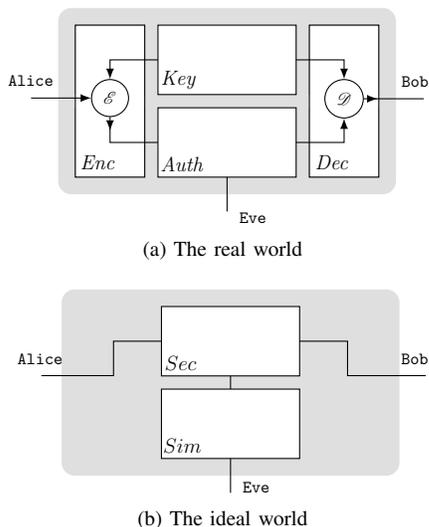


Fig. 1: Alice, Bob, and Eve are the interface names. \mathcal{E} and \mathcal{D} denote the encryption and decryption functions, respectively. Each gray rectangle represents the resource built from the resources and converters in its interior.

verification. To motivate our proposal, we show that compositional security statements should not be tied to a fixed communication pattern, as is the case in existing CC case studies.

A. Motivation

It is common to simplify the presentation of security proofs by omitting details that are considered somehow trivial; however, leaving out such details sometimes does more harm than good. Consider the construction example in Section II-C. There, we limit the extent of our proofs' reusability by defining the specifications $\mathcal{R}_{auth-key}$ and \mathcal{R}_{Sec} as singleton sets. That is, the security of the protocol π , representing the *Enc* and *Dec* converters, is only guaranteed in contexts that use the concrete resources *Auth*, *Key*, and *Sec*. But is it possible to ignore this limitation and interpret the aforementioned resources as placeholders for arbitrary keys and channels? Such simplifications are omnipresent in existing CC case studies: proofs are carried out in a concrete (and usually underspecified) context and readers are left to generalize them. In this section, we will show that singleton specifications do not suffice for composable security proofs.

We start by considering the impact of singleton specifications on the modularity and reusability of security proofs. The point of composability is that the realization of a specification can be decomposed into small and easy-to-understand steps, which can then be combined modularly. For example, in Figure 1, we may combine π with various key-exchange protocols to obtain different realizations of a secure channel; however, we want to prove π 's security only once. That is, we do want to abstract away from the details of the key-exchange protocol in π 's proof.

This cannot be achieved if we start from a specification with just a single resource like *Key*, because different key exchange protocols have different activation patterns for the

parties. This is how the problem with generic communication modeling, which we explained in Section II-A, arises in CC proofs. For example, in a three-round Diffie-Hellman protocol π_{3dh} , Alice asks for Bob's half-key and sends her half-key back after receiving Bob's answer. So before Alice puts her message on *Auth*, the execution of π_{3dh} would lead to the sequence Alice \rightarrow Bob \rightarrow Alice of party activations. Hence, *Key* would have to already model this activation sequence so that π_{3dh} can securely construct *Key* and be composed with π . Clearly, the security proof for π should not have to deal with this particular sequence of activations (and all the possible failed executions of the key exchange protocol). Even worse, another key-exchange protocol may lead to a different activation sequence. For example, a two-round Diffie-Hellman key exchange π_{2dh} yields either the activation sequences Alice \rightarrow Bob \rightarrow Alice or Bob \rightarrow Alice before Alice sends her message on the authenticated channel. Those activation sequences cannot be added to the single *Key* resource as each key exchange protocol would then have to support *all* these activation sequences. So, the security argument for the secure channel construction is entangled with a specific key exchange protocol. This hinders the modular composition of protocols.

Varying party activations is not the only problem of singleton specifications. The information that the resource *Key* should leak to Eve also depends on the key exchange protocol. For example, the two-phase Diffie-Hellman protocol π_{2dh} uses two authenticated channels to transport the half-keys from Alice to Bob and vice versa. Those channels expose Eve interfaces; queries on those interfaces will be rejected until a message has been entered into the channel. The simulator in π_{2dh} 's security proof must account for this behavior. Yet this is not possible when Eve's interface of the ideal resource *Key* leaks no information to the simulator. What must be leaked depends on the particular key exchange protocol, e.g. adapting the key specification's leakage to π_{2dh} becomes problematic in π_{3dh} 's security proof that utilizes three authenticated channels.

These two problems can be avoided if specifications contain many resources. For example, the specification for keys would contain one resource for each type of key-exchange protocols: one that is suitable for constructions with two authenticated channels, another one for constructions with three authenticated channel, and so forth. However, the security proofs must treat all resources in such a specification uniformly; otherwise we would again need a separate proof for each protocol. So, we need a semantic domain that defines each specification by describing the common aspects of the resources that it contains.

B. Fused Resource Templates

We define FRTs in terms of Lochbihler et al. [15]'s formalization of resources. Each FRT describes a parametrized resource: all instances share a common behavior while each instance exhibits a particular party-activation pattern.

1) *A Short Recap of Resources:* In CC [12], probabilistic reactive systems are called resources. Consider a bit string that is sampled uniformly from the set of bit strings of length n . One can model such a random bit string as a probabilistic

reactive system that gets an integer n as input and outputs n random bits. In general, each system output may depend on its previous inputs and outputs, e.g. a random oracle that answers repeated queries with the same output. It is therefore convenient to model resources by a transition function and an initial state: for a state s and an input x , the transition function $tr(s, x)$ returns a probability distribution over the successor state and the output.

Lochbihler et al. [15] hide the internal state of resources to simplify the proofs and definitions. The resource type existentially quantifies over the state type in the pair of the transition function and the initial state:

$$\mathbb{R}(\alpha, \beta) \equiv \exists \sigma. (\sigma \Rightarrow \alpha \Rightarrow \mathbb{D}(\beta \times \sigma)) \times \sigma,$$

where α represents the type of inputs, β the type of outputs, \times the product type, \Rightarrow the function space, and $\mathbb{D}(\gamma)$ probability distributions over γ . The Isabelle/HOL formalization uses co-inductive datatypes, in short co-datatypes, instead of existential types. Let $e : \chi$ denote that expression e has type χ ; the resource type is the final co-algebra for the transition function $run : \mathbb{R}(\alpha, \beta) \Rightarrow \alpha \Rightarrow \mathbb{D}(\beta \times \mathbb{R}(\alpha, \beta))$ that returns a distribution on the output and the successor resource for a given input. Finality means two resources are equal if their input-output behaviors are bisimilar. The following co-datatype formalizes this:

$$\text{codatatype } \mathbb{R}(\alpha, \beta) = \text{Resource } (\alpha \Rightarrow \mathbb{D}(\beta \times \mathbb{R}(\alpha, \beta))).$$

The existing formalizations [4], [15] introduce two more components that facilitate the definition of resources with multiple interfaces. First, CryptHOL’s oracle composition operator $+_{\circ}$ interleaves two transition functions $tr_i : \sigma \Rightarrow \alpha_i \Rightarrow \mathbb{D}(\beta_i \times \sigma)$ (for $i = 1, 2$) operating on a shared state of type σ into one transition function $tr_1 +_{\circ} tr_2 : \sigma \Rightarrow \alpha_1 + \alpha_2 \Rightarrow \mathbb{D}((\beta_1 + \beta_2) \times \sigma)$, where $\alpha + \beta$ denotes the disjoint union of the types α and β . Following the CryptHOL terminology, we use the terms “oracle” and “probabilistic transition function” interchangeably in this paper. Second, the function RES converts an oracle into a resource, i.e., it introduces the existential type quantifier and hides the state. $RES(tr, s_0)$ is the co-inductive counterpart of a probabilistic transition system where the state is explicit, represented using the transition function tr and the initial state s_0 .

2) *Fused Resources*: FRT instances correspond to particular members of a resource family, but how do we know if an instance belongs to a family? Remember that family members share a common behavior while each of them will exhibit a particular behavior too. FRTs only focus on the common behaviors and treat the particular ones abstractly as a parameter in their description. Nevertheless, we must ensure that every instantiation of such parameter will not affect the “common behavior”. For example, consider the specification for secure channels. The common behavior of every resource in the family of secure channels is that they would *at most* leak the length of the transmitted messages. So, the FRT describing this family must prohibit any instantiation that can leak more information.

Each FRT consists of a *Core* part, describing the common behavior, and a set of *Rest* parts. When an FRT is instantiated, the Core part is *Fused* with one of the Rest parts and results

in a special kind of resource, which we call a *Fused Resource*, that enforces one-way information flow from its Rest to Core. In what follows, we explain how a Fused Resource can be defined in terms of a *Fusing* function.

The Core part is implemented as a record $core \equiv \langle \langle cinit := \dots, cpoke := \dots, cfunc := \dots \rangle \rangle$ with three fields: an initial Core state $cinit : \sigma$, a probabilistic event handler $cpoke : \sigma \Rightarrow \varepsilon \Rightarrow \mathbb{D}(\sigma)$, and a probabilistic transition function $cfunc : \sigma \Rightarrow \alpha \Rightarrow \mathbb{D}(\beta \times \sigma)$. The $cpoke$ field is an event handler that defines a notification mechanism: given the current Core state and an event, it defines a probability distribution on the successor state. The $cfunc$ field describes the input-output behavior of the interfaces: given the current Core state and an input, it defines a probability distribution on the pair of the output and the successor state.

The Rest part communicates with the Core ones via poke events. They are defined in terms of a record $rest \equiv \langle \langle rinit := \dots, rfunc := \dots \rangle \rangle$ that consists of an initial Rest state $rinit : \tau$ and a probabilistic event-augmented transition function $rfunc : \tau \Rightarrow \chi \Rightarrow \mathbb{D}((\delta \times \mathbb{L}(\varepsilon)) \times \tau)$, where $\mathbb{L}(\varepsilon)$ denotes the type of events list. The $cfunc$ field defines the input-output behavior of Rest interfaces as well as the information they leak to the Core part: given the current Rest state and an input, it defines a probability distribution on the output, the successor rest state, and a list of events that the Fused Resource’s core part will be notified about.

Given a $rest \equiv \langle \langle rinit := ri_{def}, rfunc := rf_{def} \rangle \rangle$ and a $core \equiv \langle \langle cinit := ci_{def}, cpoke := cp_{def}, cfunc := cf_{def} \rangle \rangle$ with the types defined above, the Fusing function $fuse(core, rest)$ outputs a probabilistic transition function with type $(\sigma \times \tau) \Rightarrow \alpha + \chi \Rightarrow \mathbb{D}((\beta + \delta) \times (\sigma \times \tau))$ that enforces a one-way information flow from *rest* to *core* by means of events. Consider the value constructors $Left_$ and $Right_$ for the disjoint union $+$ type, and let tr_{fuse} and (s, t) denote the resulting oracle and its internal state respectively. For a given input x , the output of $tr_{fuse}((s, t), x)$ is a probability distribution on the pair of the output y and the successor state (s', t') that is determined in one of two ways. First, if x is of the form $Left\ x'$, then y and s' are determined by $cf_{def}(s, x')$ and $t' = t$. Second, if x is of the form $Right\ x'$, then $rf_{def}(t, x')$ determines y, t' , and a list of events es from which we calculate $s' = foldl(cp_{def}, s, es)$.

The Fusing function provides a failure mechanism too. Every probability distribution in this paper considers a sample space that provides a distinguished bottom element \perp for modeling failures. By combining the probability and option monads, we define the $fuse$ function such that any failures in calls to cp_{def}, cf_{def} , and rf_{def} will fail the whole Fused Resource, i.e., all subsequent queries to tr_{fuse} are answered with \perp independently of which interface they are sent to.

We can combine all the aforementioned building blocks into a one-liner using the technique explained in Section III-B1. Let \bullet denote the operator for accessing record fields. For Core and Rest records $core$ and $rest$ mentioned above, $FUSE(core, rest) \equiv RES(fuse(core, rest), (core \bullet cinit, rest \bullet rinit))$ is a Fused Resource of type $\mathbb{R}(\alpha + \chi, \beta + \delta)$. We denote the type of such Fused Resource as $\mathbb{F}(\alpha, \chi, \beta, \delta)$ to distinguish

it from typical resources of the same type.

Core records suffice to describe the behavior that is common to every instantiation of an FRT. Let \spadesuit and \clubsuit respectively denote the types of *core* and *rest* records mentioned above. Given a core $core : \spadesuit$, an FRT $\{\{core\}\} : \clubsuit \Rightarrow \mathbb{F}(\alpha, \chi, \beta, \delta)$ is a function from Rest records to Fused Resources such that:

$$\{\{core\}\}(rest) \equiv FUSE(core, rest). \quad (1)$$

When essential, we specify the type of suitable Rest records using a subscript $\{\{core\}\}_{\clubsuit}$.

3) *A concrete example:* We now explain how FRTs enable the modeling of specifications that include resources with arbitrary interfaces and activation patterns. The main idea is to define Core such that it returns \perp when an unwanted sequence of poke events has been received.

As an example, we formalize the ideal specification for keys as the FRT $\{\{key\}\}$. It represents all realizations in which two parties Alice and Bob execute a protocol to share a symmetric key. We define the fields of the core record $key \equiv (c_{init} := ci_{key}, c_{poke} := cp_{key}, c_{func} := cf_{key})$ in the given order. The state consists of a pair $(kernel, shell)$. Here, $kernel$ is a value of the form $Void$ or $Hold k$ that indicates whether the key has already been generated; and $shell \subseteq \{Alice, Bob\}$ keeps track of the enabled interfaces, where "disabled" means the interface would return \perp if queried. The initial state ci_{key} is $(Void, \{\})$, which indicates that the key has not been generated yet and both parties' interfaces are disabled.

The probabilistic event handler cp_{key} keeps track of two classes of events: the event *Init* updates the state to $Hold k$ by uniformly sampling k from the key's domain and the events of the form *Open party* will update the state by inserting *party* in $shell$. Repeated events are not allowed and immediately invoke the failure mechanism.

The probabilistic transition function $cf_{key} \equiv tr_{Eve-k} \oplus tr_{Alice-k} \oplus tr_{Bob-k}$ describes the interfaces for Alice, Bob, and the adversary, also called Eve, that are composed using the oracle composition operator. The adversary interface's functionality tr_{Eve-k} leaks no information to Eve; it is a dummy oracle that does not change the state and answers every input with a unit value \otimes . The interface functionality for Alice is defined using $tr_{Alice-k}$ which outputs a key k upon a unit input \otimes if the state is a pair $(Hold k, shell)$ such that $Alice \in shell$. Queries to $tr_{Alice-k}$ will invoke the failure mechanism if the state does not satisfy the aforementioned conditions. Bob's interface functionality tr_{Bob-k} is the same as Alice's interface except that it checks for $Bob \in shell$ upon receiving queries.

Figure 2 is a pictorial representation of the ideal specification for keys. $\{\{key\}\}$ represents all Fused Resources that accept queries on Alice and Bob interfaces after one or more queries to the i_1, \dots, i_n interfaces, which trigger the key generation event and enable (at least) one of the parties' interfaces. Either Alice or Bob may attempt to receive the key first; it is only important that their corresponding interface has been enabled during one of the event triggers.

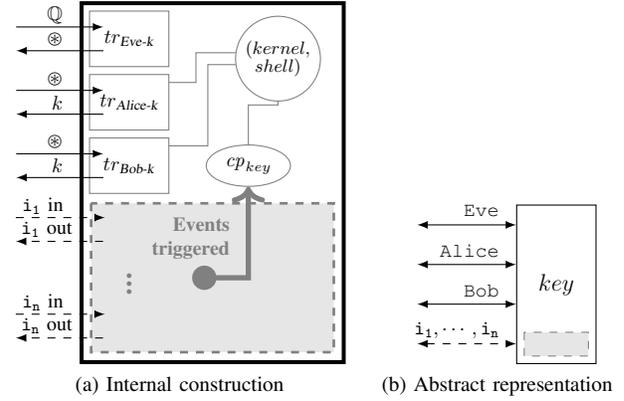


Fig. 2: Formalizing the ideal specification for keys. The rest part is shown as a gray area that accepts various instantiations. Rest interfaces are shown using dashed arrows. To simplify the diagrams, we may merge multiple interfaces into a single arrow like i_1, \dots, i_n above.

C. Constructions using FRTs

Defining complex systems in terms of simpler ones is the essence of composable security statements. CC captures this using the concepts of simultaneously-available resources and converter attachments, which are lifted to specifications as explained in Section II-C. That is, compound resources describing complex probabilistic systems can be defined as the attachment of simpler building blocks, i.e., converters and simpler resources. We show how these concepts are expressed using FRTs. Our formalization reuses the converters and resources from Lochbihler et al.'s [15].

1) *A short recap of compound resources:* Converters are probabilistic reactive systems that internally interact with other reactive systems. A converter has two types (or classes) of interfaces. First, the *external interfaces* are used to receive inputs and send outputs. Second, computing an output may invoke many interactions on the *internal interfaces*. A converter appears like a resource when one of its external interfaces is queried, however, such a query may invoke other converters and resources prior to its response. A converter's internal interface is formalized as a CryptHOL [4] *Generative Probabilistic Value* (GPV) $\mathbb{G}(\xi, \gamma, \rho)$: a GPV produces results of type ξ by interacting with arbitrary reactive systems that answer queries of type γ with responses of type ρ . Formally,

$$\text{codatatype } \mathbb{G}(\xi, \gamma, \rho) = \text{GPV } \mathbb{D}(\xi + (\gamma \times (\rho \Rightarrow \mathbb{G}(\xi, \gamma, \rho)))).$$

A converter then is modeled using an initial state $s_0 : \sigma$ and a transition function $tr : \sigma \Rightarrow \alpha \Rightarrow \mathbb{G}(\beta \times \sigma, \gamma, \rho)$ that describes a GPV on the output and the new state of the converter given its current state and an input. Similar to resources, Lochbihler et al. [15] provide an operator $CNV(tr, s_0)$ that can hide a converter's internal state and represent it as a value of the following co-datatype:

$$\text{codatatype } \mathbb{C}(\alpha, \beta, \gamma, \rho) = \text{Converter } (\alpha \Rightarrow \mathbb{G}(\beta \times \mathbb{C}(\alpha, \beta, \gamma, \rho), \gamma, \rho)).$$

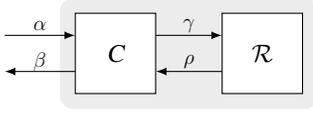


Fig. 3: Attaching a converter to a resource.

As shown in Figure 3, attaching $C : \mathbb{C}(\alpha, \beta, \gamma, \rho)$ to a resource $R : \mathbb{R}(\gamma, \rho)$, which means R responds to C 's internal queries, creates a new resource $C \triangleright R : \mathbb{R}(\alpha, \beta)$, which is represented as the surrounding gray rectangle.

Lochbihler et al. model multiple interfaces i_1, \dots, i_n , which answer queries of types $\alpha_1, \dots, \alpha_n$ with responses of types β_1, \dots, β_n , using disjoint unions as a single interface with inputs of type $\alpha_1 + \dots + \alpha_n$ and outputs of type $\beta_1 + \dots + \beta_n$. So interfaces are identified by their position in the sum type. The attachment of (multiple) converters to a subset of (simultaneously available) resources' interfaces are defined in terms of the following operators:

- Let $R_1 : \mathbb{R}(\alpha_1, \beta_1)$ and $R_2 : \mathbb{R}(\alpha_2, \beta_2)$ be resources. Their parallel composition $R_1 \parallel R_2 : \mathbb{R}(\alpha_1 + \alpha_2, \beta_1 + \beta_2)$ directs queries α_1 to R_1 and α_2 to R_2 and forwards the responses accordingly.
- Let $C_1 : \mathbb{C}(\alpha_1, \beta_1, \gamma_1, \eta_1)$ and $C_2 : \mathbb{C}(\alpha_2, \beta_2, \gamma_2, \eta_2)$ be two converters. Their parallel composition $C_1 \mid C_2 : \mathbb{C}(\alpha_1 + \alpha_2, \beta_1 + \beta_2, \gamma_1 + \gamma_2, \eta_1 + \eta_2)$ makes both converters available at the same time, analogous to parallel resource composition.
- Let $C_1 : \mathbb{C}(\alpha, \beta, \gamma, \eta)$ and $C_2 : \mathbb{C}(\gamma, \eta, \chi, \delta)$ be two converters. Their sequential composition $C_1 \odot C_2 : \mathbb{C}(\alpha, \beta, \chi, \delta)$ uses C_2 to answer C_1 's queries on C_1 's internal interface.
- The identity converter $\mathbb{1} : \mathbb{C}(\alpha, \beta, \alpha, \beta)$ simply forwards all queries and responses from the external to the internal interface and vice versa.

We provide an example to clarify how the above operators are used. Consider resource R_1 with one interface x , resource R_2 with two interfaces y and z , and a converter C with an external interface e and two internal interfaces a and b . The compound resource $((\mathbb{1} \mid C) \odot C_W) \triangleright (R_1 \parallel R_2)$ results from attaching a and b to z and x respectively, where C_W is defined in terms of special converters called *Wirings* [15] that transpose $x + (y + z)$ to $y + (z + x)$ ⁴. To avoid boilerplate notation and operators, we often describe converter attachments using interface names similar to Section II-C. Moreover, we will drop the distinction between the terms "converter" and "protocol" from now on since multiple converter and resource attachments are encoded as a single attachment written in terms of composition operators.

2) *Compound Fused Resources and FRTs*: The aforementioned operators do not preserve the structure of Fused Resources. For instance, given the Fused Resource $F : \mathbb{F}(\alpha, \chi, \beta, \delta)$ and the converter $C : \mathbb{C}(\chi + \alpha, \delta + \beta, \alpha + \chi, \beta + \delta)$ that swaps interface positions, then $C \triangleright F$ may be inexpressible

⁴The overloaded $+$ operator on interface names is only used to indicate that disjoint unions are not associative and Wirings enable interface reordering.

in terms of *FUSE* due to its inverted information flow. We now show typical cases that do preserve the structure of Fused Resources, and thus of FRTs. As we explain in Section IV-A, these cases suffice for the step in the security proof that allows us to abstract over the communication behavior and focus on the common properties. In the following four lemmas, consider the Core records $cr_i : \spadesuit_i$, the Rest records $rs_i : \clubsuit_i$, and the Fused Resources $F_i \equiv \text{FUSE}(cr_i, rs_i)$ for $i \in \{1, 2\}$. The first lemma explains the simultaneous access to Fused Resources.

Lemma 1. For $F_i : \mathbb{F}(\alpha_i, \chi_i, \beta_i, \delta_i)$, consider the Fused Resource $F : \mathbb{F}(\alpha_1 + \alpha_2, \chi_1 + \chi_2, \beta_1 + \beta_2, \delta_1 + \delta_2)$ defined as $F \equiv \text{FUSE}(cr_1 \parallel_c cr_2, rs_1 \parallel_r rs_2)$, where \parallel_c and \parallel_r stand for the parallel composition of Core and Rest records respectively, analogous to $+_{\oplus}$ for oracles. Then $F = C_{F-PL} \triangleright (F_1 \parallel F_2)$, where C_{F-PL} is the wiring that groups the Core and Rest interfaces of F_1 and F_2 together.

As a corollary of the following lemma, converters may attach to just the Core (or similarly the Rest) interfaces of a Fused Resource since either of the parallel converters could be replaced with the identity converter $\mathbb{1}$.

Lemma 2. For $F_1 : \mathbb{F}(\alpha, \chi, \beta, \delta)$ and $F_2 : \mathbb{F}(\pi, \mu, \theta, \nu)$, consider the converters $C_c : \mathbb{C}(\pi, \theta, \alpha, \beta)$ and $C_r : \mathbb{C}(\mu, \nu, \chi, \delta)$. Then $F_2 = (C_c \mid C_r) \triangleright F_1$ if $cr_2 = C_c \triangleright_c cr_1$ and $rs_2 = C_r \triangleright_r rs_1$, where \triangleright_c attaches a converter to the cfunc oracle of a Core and \triangleright_r does so for the rfunc oracle of a Rest part.

The next two lemmas determine if a converter's simultaneous attachment to Core and Rest interfaces results in a new Fused Resource. Using them, one may rewrite an arbitrary converter C 's attachment to F_2 , i.e., $C \triangleright F_2$, into $(C \odot _) \triangleright F_1$, where $_$ is filled with some wiring, and then (if possible) apply the previous lemma. These lemmas stem from the one-way information flow from Rest to Core interfaces of Fused Resources.

Lemma 3. For $F_1 : \mathbb{F}(\pi + \alpha, \chi, \theta + \beta, \delta)$, let I refer to the left-most (w.l.o.g.) Core interface. Assume that I 's θ responses do not depend on α queries, and let cr_2 and rs_2 be the records that result from removing I 's functionality from cr_1 and adding it to rs_1 . Then $(F_2 : \mathbb{F}(\alpha, \pi + \chi, \beta, \theta + \delta)) = C_{F-CR} \triangleright F_1$, where C_{F-CR} is the wiring converter that reorders the interface positions accordingly.

Lemma 4. For $F_1 : \mathbb{F}(\alpha, \pi + \chi, \beta, \theta + \delta)$, let I refer to the left-most (w.l.o.g.) Rest interface. Assume that δ responses do not depend on I 's π queries, and let rs_2 and cr_2 be the records that result from removing I 's functionality from rs_1 and adding it to cr_1 . Then $(F_2 : \mathbb{F}(\pi + \alpha, \chi, \theta + \beta, \delta)) = C_{F-RC} \triangleright F_1$, where C_{F-RC} is the wiring converter that reorders the interface positions accordingly.

While these lemmas refer to individual Fused resources, they naturally lift pointwise to FRTs. So these cases also preserve the FRT structure. We write $C \{ \! \! \! \clubsuit \! \! \! \}$ for the FRT that consists of the Fused resources $C \triangleright \{ \! \! \! \clubsuit \! \! \! \}$ (rs) for $rs \in \clubsuit$. This notation is analogous to Section II-C where $\pi \mathcal{R}$ represented the attachment of the protocol π to the specification \mathcal{R} .

For Lemma 1, the resulting FRT $\{\{cr_1\|_c cr_2\}\}_{\clubsuit_1 \times \clubsuit_2}$ is indexed over $\clubsuit_1 \times \clubsuit_2 = \{rs_1\|_r rs_2 \mid rs_1 \in \clubsuit_1 \wedge rs_2 \in \clubsuit_2\}$. This set does not include all Rest records that could be fused to the parallel composition of the Core records cr_1 and cr_2 . This is because the internal state of Rest records in $\clubsuit_1 \times \clubsuit_2$ can be split into two independent components, one for each of the fixed Rest records rs_1 and rs_2 ; however such a decomposition is not possible in general.

IV. SECURE CONSTRUCTIONS WITH FRTs

As explained in Section II-C, simulation-based security for specifications \mathcal{R} and \mathcal{S} and a protocol π demands that there is a simulator $\sigma = z^E$, attaching a converter z to the adversary interface E (short form of Eve), such that

$$\forall R \in \mathcal{R}. \exists S \in \mathcal{S}. \mathfrak{d}(\pi R, \sigma S) \leq \epsilon. \quad (2)$$

Note that the same simulator σ must work for all $R \in \mathcal{R}$. Accordingly, a formal proof of simulation-based security must also work for all $R \in \mathcal{R}$. To that end, \mathcal{R} must be expressed in such a way that the properties relevant for the security proof are explicit and the irrelevant details can be abstracted away. In this section, we show that FRTs are well-suited for composable simulation-based security: the Core record captures the relevant properties and the Rest record hides the details of the party activation patterns.

First, we specialize (2) to the case where \mathcal{R} and \mathcal{S} are FRTs. We only define the information-theoretic security here; the computational security notion is defined analogously.

Definition 1 (Information-theoretic concrete security). *Let $\{\{real\}\}_{\clubsuit_R}$ and $\{\{ideal\}\}_{\clubsuit_I}$ be FRTs and let π be the protocol, i.e., a converter that attaches to the Core user interface of $\{\{real\}\}$. Then $\pi\{\{real\}\}$ is an information-theoretically ϵ -secure realization of the ideal specification $\{\{ideal\}\}$ if there exist a simulator σ that attaches to the adversary interface of Core and Rest, and a Rest embedding function $f : \clubsuit_R \Rightarrow \clubsuit_I$ such that*

$$\mathfrak{d}(\pi \triangleright \{\{real\}\}(rs), \sigma \triangleright \{\{ideal\}\}(f(rs))) \leq \epsilon \quad (3)$$

for all $rs \in \clubsuit_R$.

Note that the Rest embedding f in (3) skolemizes the existential $\exists S \in \mathcal{S}$ in (2). This works thanks to the shape of FRTs, namely both \mathcal{R} and \mathcal{S} are parametrized by the rest records.

We extend this notion to asymptotic security by introducing a security parameter η and requiring that ϵ is negligible in η . As is customary in CryptHOL [4], our formalization uses Isabelle/HOL's module system for that. Accordingly, we obtain an asymptotic security statement for every concrete security statement essentially for free.

A. Proving a Protocol Secure

Proving a protocol π secure boils down to establishing a bound on the advantage of a distinguisher. This is typically done as a sequence of game transformations using the existing tool set from CryptHOL [4], [14], similar to Lochbihler et al.'s formalization of CC [15]. We now explain the common

pattern for FRTs. We also augment this tool set with a new up-to proof rule for trace equivalence of resources.

Lochbihler et al. [15] have already formalized the distinguishability bound $\mathfrak{d}(R, S)$ on two concrete resources R and S using explicit distinguishers. Formally, a distinguisher $D : \mathbb{G}(\mathbb{B}, \alpha, \beta)$ is a GPV that returns a Boolean \mathbb{B} after having interacted with a resource through queries α and responses β . We write $D \blacktriangleright R : \mathbb{D}(\mathbb{B})$ for the resulting probability distribution over Booleans \mathbb{B} . D 's advantage $adv(D, R, S)$ of distinguishing R and S is then given by

$$adv(D, R, S) = |(D \blacktriangleright R) ! True - (D \blacktriangleright S) ! True|,$$

where $p ! x$ denotes the probability that the distribution p assigns to the elementary event x . Then, $\mathfrak{d}(R, S) \leq \epsilon$ iff $adv(D, R, S) \leq \epsilon$ for all distinguishers D . Here, D ranges over all distinguishers in the information-theoretic setting and over computationally bounded distinguishers in the computational setting.

The proof of (3) must work for all $rs \in \clubsuit_R$. It should therefore concentrate on the Cores *real* and *ideal* of the Fused Resources $\{\{real\}\}(rs)$ and $\{\{ideal\}\}(f(rs))$. The first step is therefore to pretend that the fusing is part of the distinguisher. For comparison, this step corresponds to applying the dummy adversary lemma in UC.

On the real side $\pi \triangleright \{\{real(rs)\}\}$, Lemma 2 shows that the converter π can be attached directly to *real* rather than the Fused Resource, because π attaches only to the Core interface. This transformed core *real'* yields a transformed FRT $\{\{real'\}\}_{\clubsuit_R}$.

On the ideal side, however, we cannot directly use Lemma 2 to push the simulator σ through *fuse* because σ attaches to the adversary Core and Rest interfaces. This way, the simulator can create a flow of information from the core part to the rest part, which cannot happen in a Fused Resource. We now exploit that the simulator must work for all rs in the same way, in particular for the Rest record that does not provide any interfaces at all. Accordingly, σ uses only to the interfaces that f adds on top of rs 's. By a similar argument, $f(rs)$ answers queries on these additional interfaces without consulting rs itself. Moreover, f may translate poke events from rs to poke events for the ideal Core and this translation is independent of rs 's state.⁵ So we can move these additional interfaces to *ideal* by Lemma 4. Then, σ only attaches to Core interfaces and can thus be integrated into the Core by Lemma 2. This new core *ideal'* yields a new FRT $\{\{ideal'\}\}_{\clubsuit_R}$.

In summary, it suffices to prove

$$adv(D, \{\{real'\}\}(rs), \{\{ideal'\}\}(rs)) \leq \epsilon$$

⁵This argument can be made precise in the computational setting: The skolemization f in (3) of the existential $\exists S \in \mathcal{S}$ in (2) ensures that f chooses the Rest records in $\{\{ideal\}\}$ uniformly. Formally, this translates to f being relationally parametric in the queries and responses of rs . So we can decompose $f(rs)$ into two parts with disjoint state: (i) rs and (ii) additional independent interfaces and a translator for poke events coming from rs .

In the information-theoretic setting, this uniformity argument does not apply and the argument therefore need not hold in pathological cases. We have not yet encountered such a case in practice though.

for all D . Now, the Fused Resources $\{\!\{real'\}\!\}(rs)$ and $\{\!\{ideal'\}\!\}(rs)$ use the same Rest record rs . Since the Core and Rest in a Fused Resource operate on separate states, we can understand a Core record as a resource of its own, with the poke events as an additional interface that returns only dummy responses \otimes . Accordingly, we have

$$adv(D, \{\!\{real'\}\!\}(rs), \{\!\{ideal'\}\!\}(rs)) = adv(D', real', ideal')$$

for some D' . This key proof step abstracts over the party activation patterns: from now on, the security proof can focus on the cores $real'$ and $ideal'$. Since these are like resources, this proof can be approached as described in the literature [15].

Frequently, such a proof must transform the resource into a particular shape, e.g., to prepare for the application of a reduction argument. Lochbihler et al. [15] found that these steps can often be justified by trace equivalence. They gave the following bisimulation-style proof rule for establishing trace equivalence of resources, where $dirac(x)$ is the one-point distribution on x , and $\overline{run}(p, a)$ denotes the weighted combination of running a resource from the distribution p of resources with query a , and $p|_x$ conditions a distribution p over pairs $\alpha \times \beta$ on the elementary event $x : \alpha$.

Theorem 1 ([15, Thm. 1]). *Two resources R and S are trace equivalent iff there exists a relation X between distributions of resources such that*

- 1) X relates $dirac(R)$ to $dirac(S)$.
- 2) Whenever $(p, q) \in X$, then for all queries a , $\overline{run}(p, a)$ and $\overline{run}(q, a)$ have the same marginal distribution on responses and $(\overline{run}(p, a)|_b, \overline{run}(q, a)|_b) \in X$ for each possible response b to the query a .

The distributions in the relation X capture the so-far unobserved probabilistic choices of the resources. For example, suppose that one resource R eagerly samples a random value at the start whereas the other resource S delays the probabilistic choice until it is actually needed. So there is a time interval when R has already sampled the value, but not S , i.e., the chosen value has not yet been observed through the interactions with R . During this time, the relation X conserves this unobserved randomness in a distribution on R 's side. When S finally samples the value and uses it to compute a response, this conserved randomness is used in the proof to justify the indistinguishability of the responses.

In practice, the randomness may be observed only in parts or not at all in some runs. We frequently experienced such situations in our case studies, especially when the adversary interferes or a failure event happens. Yet, the closure condition on X in Thm. 1 requires that X contains all reachable combinations of distributions in such cases, even though we know that there is no point in conserving the randomness in this branch of the proof. This adds considerable bloat to the definition of X and to the trace equivalence proof because condition (2) must hold for every pair of distributions in X , which may require even further pairs in X , and so on. In some examples, these ‘‘unnecessary’’ cases significantly outnumbered the relevant cases.

To counter this case explosion, we introduce a closure operator $[X]$ on X and generalize the above theorem to an up-to proof rule. The closure $[X]$ of a relation X of distributions is defined inductively as follows, where $p \ggg f$ denotes the p -weighted combination of a family f of distributions, i.e., $(p \ggg f)!x = \sum_y p!y \cdot f(y)!x$.

- 1) Whenever $(p, q) \in X$, then $(p, q) \in [X]$.
- 2) Let f, g be two families of distributions over the same countable index set I and p be a distribution over I . If $(f(i), g(i)) \in [X]$ for all i in p 's support, then $(p \ggg f, p \ggg g) \in [X]$.

We can now show the up-to version of the trace equivalence proof rule. It differs from Thm. 1 only in the closure $[X]$ instead of X in condition 2.

Theorem 2 (Trace equivalence upto). *Two resources R and S are trace equivalent iff there exists a relation X between distributions of resources such that*

- 1) X relates $dirac(R)$ to $dirac(S)$.
- 2) Whenever $(p, q) \in X$, then for all queries a , $\overline{run}(p, a)$ and $\overline{run}(q, a)$ have the same marginal distribution on responses and $(\overline{run}(p, a)|_b, \overline{run}(q, a)|_b) \in [X]$ for each possible response b to the query a .

This theorem can significantly reduce the number of cases in a trace equivalence proof. In the above example with the unnecessary randomness, we discard the randomness with two constant families $f(i) = dirac(R)$ and $g(i) = dirac(S)$. Not only does this step reduce the number of proof cases, it also simplifies the remaining cases because these cases start from the single resource in the one-point distribution rather than a distribution of resources.

B. Composability

We now show that our security definition 1 yields composable security statements. In CC, a protocol π typically uses multiple resources R_1, \dots, R_n to create a resource $\pi \triangleright (R_1 \parallel \dots \parallel R_n)$, which should be hard to distinguish from a simulated ideal resource $\sigma \triangleright S$. Similarly, another protocol ρ may use S and possibly other resources to achieve another ideal resource T . To abstract over the party activation patterns of these resources, we consider FRTs rather than individual resources. This ensures that the security proof for ρ is independent of the patterns that we need for R_i during the composition.

We first give composability theorems for concrete security statements of individual resources and then use them to establish composability for FRTs. As before, we focus on the information-theoretic setting; the computational setting is analogous by restricting the class of distinguishers to computationally bounded ones.

Definition 2 (Concrete security for individual resources). *A resource R ϵ -securely realizes a resource S if there is a converter σ that attaches to S 's adversary interfaces such that*

$$adv(D, R, \sigma \triangleright S) \leq \epsilon$$

for all distinguishers D .

Clearly, if $\pi\{\text{real}\}_{\clubsuit_R}$ ϵ -securely realizes the FRT $\{\text{ideal}\}_{\clubsuit_I}$ (Def. 1), then $\pi \triangleright \{\text{real}\}(rs)$ ϵ -securely realizes the Fused Resource $\{\text{ideal}\}(f(rs))$ in the sense of Def. 2 for all $rs \in \clubsuit_R$ as witnessed by the simulator σ .

Secure realization between individual instances is compositional. The theorem below generalizes the asymptotic security statements from [15] to the concrete setting.

Theorem 3 (Composability for individual resources).

- 1) Every resource R 0-securely realizes itself.
- 2) If R_1 ϵ_1 -securely realizes S_1 as witnessed by σ_1 and R_2 ϵ_2 -securely realizes S_2 as witnessed by σ_2 , then $R_1 \parallel R_2$ $(\epsilon_1 + \epsilon_2)$ -securely realizes $S_1 \parallel S_2$ as witnessed by $\sigma_1 \sigma_2$.
- 3) If R ϵ -securely realizes S as witnessed by σ and S ϵ' -securely realizes T as witnessed by τ , then R $(\epsilon + \epsilon')$ -securely realizes T as witnessed by $\tau \odot \sigma$.
- 4) If R ϵ -securely realizes S as witnessed by σ and C is a converter that attaches to R 's user interface, then $C \triangleright R$ ϵ -securely realizes $C \triangleright S$ as witnessed by σ .

These composability results suffice to show that secure realization between FRTs is also compositional. As a FRT $\{\text{cr}\}_{\clubsuit}$ is a family of Fused Resources $\{\text{cr}\}(rs)$ for $rs \in \clubsuit$, we obtain composability by lifting Thm. 3 pointwise to specifications. Note that the simulator and the rest embeddings remain independent of the chosen Fused Resources, as required by Def. 1.

Theorem 4 (Composability for FRTs).

- 1) The identity protocol $\mathbb{1}$ applied to a FRT $\{\text{cr}\}_{\clubsuit}$ 0-securely realizes the FRT $\{\text{cr}\}_{\clubsuit'}$ for $\clubsuit' \supseteq \clubsuit$ with the simulator being the identity converter and the rest embedding being the identity function.
- 2) Let $\pi_i\{\text{real}_i\}_{\clubsuit_{R_i}}$ ϵ_i -securely realizes the FRT $\{\text{ideal}_i\}_{\clubsuit_{I_i}}$ with simulator σ_i and rest embedding f_i for $i = 1, 2$. The parallel composition

$$(\pi_1 \mid \pi_2)\{\text{real}_1 \parallel_c \text{real}_2\}_{\clubsuit_{R_1} \times \clubsuit_{R_2}}$$

$(\epsilon_1 + \epsilon_2)$ -securely realizes

$$\{\text{ideal}_1 \parallel_c \text{ideal}_2\}_{\clubsuit_{I_1} \times \clubsuit_{I_2}}$$

with simulator $\sigma_1 \sigma_2$ and rest embedding $f(rs_1, rs_2) = (f_1(rs_1), f_2(rs_2))$.

- 3) Let $\pi\{\text{real}\}_{\clubsuit_R}$ ϵ -securely realize the FRT $\{\text{middle}\}_{\clubsuit_M}$ with simulator σ and rest embedding f and let $\rho\{\text{middle}\}_{\clubsuit'_M}$ ϵ' -securely realize $\{\text{ideal}\}_{\clubsuit_I}$ with simulator τ and rest embedding g such that $\clubsuit_M \subseteq \clubsuit'_M$. Then, $\pi\{\text{real}\}_{\clubsuit_R}$ $(\epsilon + \epsilon')$ -securely realizes $\{\text{ideal}\}_{\clubsuit_I}$ with simulator $\tau \odot \sigma$ and rest embedding $g \circ f$.
- 4) If $\pi\{\text{real}\}_{\clubsuit_R}$ ϵ -securely realizes the FRT $\{\text{ideal}\}_{\clubsuit_I}$ and C is a converter that attaches to the Core user interface, then $C\pi\{\text{real}\}_{\clubsuit_R}$ ϵ -securely realize $C\{\text{ideal}\}_{\clubsuit_I}$ with the same simulator and rest embedding.

V. CASE STUDIES

We now use our framework to formalize the construction of a secure channel using three authenticated channels. For brevity, we only provide a high-level overview and use diagrams to present the main steps of our formalization [1].

We apply the following conventions in diagrams. Fused Resources are depicted on the right; we use subscripted \mathcal{R} as short names for them, e.g. $\mathcal{R}_{\text{aut1}}$ corresponds to the FRT $\{\text{aut1}\}$ that is instantiated with the Rest record rs_1 in Figure 1a. Converters' attachment to Fused Resources is represented using arrows, where solid and dashed arrows represent the attachment to Core and Rest interfaces respectively. Such an attachment results in new Fused Resources that are presented using rounded rectangles with different colors or borders. Fused Resources or converters that are in the same column are composed in parallel using \parallel or \mid respectively. The left-most interfaces of every diagram belong to Eve, Alice, Bob, and "Rest" from top to bottom. For simplicity, we merge multiple Rest or adversary interfaces into a single arrow and mark their meeting point using a black circle. The interfaces are preserved among all the diagrams, so we name them only in some of the subfigures.

Our example construction combines two cryptographic primitives in Figure 4a. First, the Diffie-Hellman key exchange protocol with the converters C_{dhA} and C_{dhB} is used to construct a key \mathcal{R}_{key} from two of the authenticated channels, i.e., $\mathcal{R}_{\text{aut1}}$ and $\mathcal{R}_{\text{aut2}}$ that are in the opposite directions. Second, we use one-time-pad in the converters C_{enc} and C_{dec} to construct a secure channel \mathcal{R}_{sec} from the key \mathcal{R}_{r1} , which is the result of Diffie-Hellman key exchange, and the last authenticated channel $\mathcal{R}_{\text{aut3}}$.

The aforementioned components behave as their name suggests. The key and channel resources are defined as we have explained in Section III-B3. The C_{enc} and C_{dec} converters are stateless and each provide a single external interface for Alice and Bob respectively. When the C_{enc} converter is queried with a message m , it fetches a key k via its internal interface that is attached to C_{dhA} and forwards $m \oplus k$, where \oplus is the xor operator for bit-strings, to its internal interface that is connected to the authenticated channel $\mathcal{R}_{\text{aut3}}$. The C_{dec} converter works analogously to decrypt the cipher-text that it fetches from $\mathcal{R}_{\text{aut3}}$.

The C_{dhA} and C_{dhB} are more involved. We describe C_{dhA} as a representative of these converters since they are like duals. C_{dhA} is a stateful converter with two external interfaces. The first external interface, which is placed at the top and redirected to \mathcal{R}_{r1} 's rest part, stores a half-key x in its state upon receiving a unit query \otimes and puts g^x into the authenticated channel $\mathcal{R}_{\text{aut1}}$, where g is the cyclic group's generator. The second external interface, which is attached to C_{dec} , can only get queried after the query to the first interface; otherwise the failure mechanism is triggered. It fetches the half-key g^y , which is sent by the other party, from the authenticated channel $\mathcal{R}_{\text{aut2}}$ and stores $(g^y)^x$ in its state.

Figure 4 depicts the central steps of the security proof. First, corresponding to steps from Figure 4a to 4e, we prove that the specification \mathcal{R}_{r1} securely realizes the key specification \mathcal{R}_{key} . Second, corresponding to the dotted area in Figures 4e and 4f,

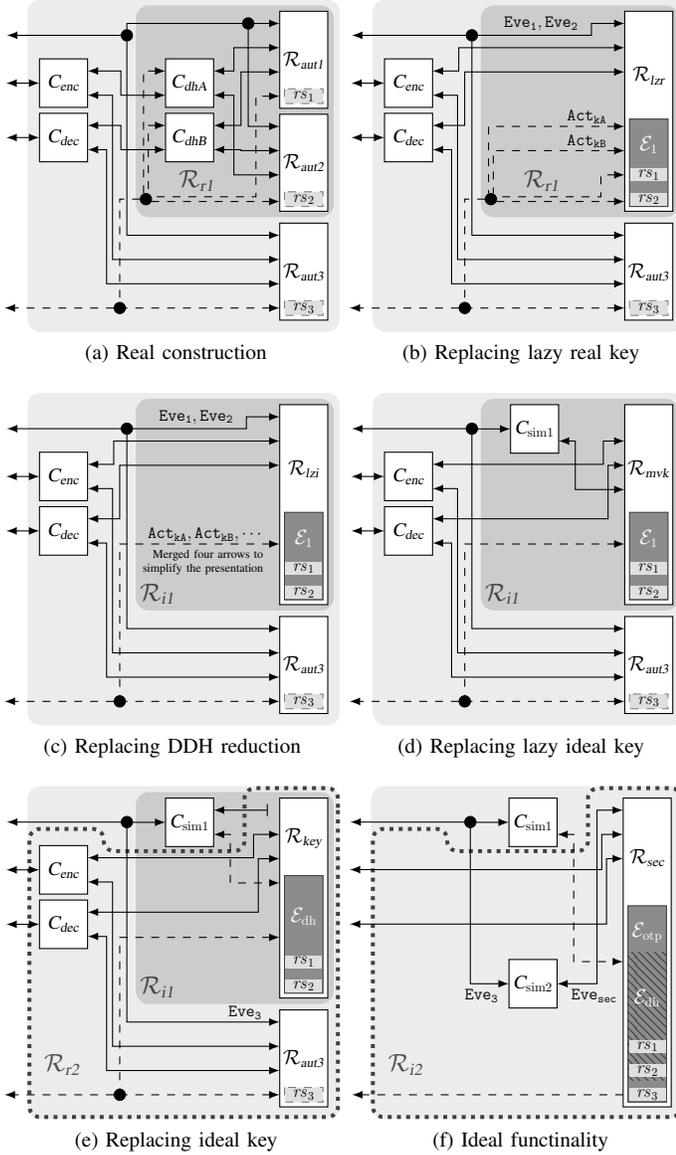


Fig. 4: Composing two security statements.

we show that the specification \mathcal{R}_{r2} , i.e., the Fused Resource that is presented using dashed borders and results from attaching C_{enc} and C_{dec} converters to \mathcal{R}_{key} and \mathcal{R}_{aut3} , securely realizes the secure channel specification \mathcal{R}_{sec} . Finally, justifying the steps from Figure 4a to 4f, the two constructions are combined using the composition theorems.

The secure realization of \mathcal{R}_{key} using \mathcal{R}_{r1} requires a proof with three major steps. First, going from Figure 4a to Figure 4b, we make our key-exchange protocol lazy. That is, we prove that \mathcal{R}_{r1} is trace equivalent to a fused resource \mathcal{R}_{lzr} that postpones the half-key samplings until one of the parties requires the key. Note that the aforementioned Fused Resources provide the same interfaces to the outside world. Let's call C_{dhA} and C_{dhB} 's top interface Act_{kA} and Act_{kB} respectively; furthermore, let's call \mathcal{R}_{aut1} and \mathcal{R}_{aut2} 's adversary interfaces Eve_1 and Eve_2 .

In the Fused Resource \mathcal{R}_{r1} , the Diffie-Hellman half-keys x and y are sampled separately upon the queries to Act_{key-a} and Act_{key-b} respectively. However, \mathcal{R}_{lzr} samples the pair (x, y) when it receives a query on any of Act_{kA} , Act_{kB} , Eve_1 , or Eve_2 . This proof step utilizes Lemma 2 to reason about the attachment of Diffie-Hellman converters to the authenticated channels and Lemma 3 to move the protocol initiation queries, i.e., Act_{kA} and Act_{kB} , to \mathcal{R}_{lzr} 's rest part.

Second, corresponding to the backwards steps from Figure 4e to Figure 4c, we apply a similar procedure to make the ideal specification \mathcal{R}_{key} lazy. That is, \mathcal{R}_{i1} is trace equivalent to a FRT \mathcal{R}_{lzi} that samples the outputs of the adversary and user interfaces at the same time. Therefore, instead of sampling x , y , and z separately upon the queries to Act_{kA} , Act_{kB} , and either of Eve_1 or Eve_2 respectively, \mathcal{R}_{lzi} samples the triple (x, y, z) whenever any of the aforementioned interfaces are queried. Furthermore, \mathcal{R}_{lzi} resembles the combined behavior of C_{sim1} and \mathcal{E}_{dh} using the rest embedding \mathcal{E}_1 : it keeps track of Eve_1 and Eve_2 's forward queries and the events received on the rest interfaces, e.g. Act_{kA} and Act_{kB} , to trigger appropriate events on its core part. The proof step from Figure 4e to Figure 4d uses a corollary of Lemma 4 to move \mathcal{E}_{dh} to \mathcal{R}_{mvk} 's core part; and the proof step to Figure 4c uses Lemma 2 to reason about the attachment of the resulting specification to C_{sim1} . Note that C_{sim1} is not connected to \mathcal{R}_{key} 's core adversary interface that answers all queries with \otimes . We use an unconnected arrow to signify the fact that C_{sim1} 's behavior does not depend on the input on that interface.

Third, corresponding to the step between Figures 4b and 4c, we prove an arbitrary distinguisher D 's advantage in distinguishing \mathcal{R}_{lzr} and \mathcal{R}_{lzi} is bounded by D 's advantage in the decisional Diffie-Hellman game, in short *ddh*. This proof step is based on a reduction in which \mathcal{R}_{lzr} and \mathcal{R}_{lzi} query an external oracle to receive the triples $(g^x, g^y, g^{(x*y)})$ and (g^x, g^y, g^z) respectively, which are used to answer all the user and adversary queries.

The secure realization of \mathcal{R}_{sec} using \mathcal{R}_{r2} requires less effort. This corresponds to the step from Figure 4e to Figure 4f. We prove that \mathcal{R}_{r2} is trace-equivalent to the fused resource resulting from attaching C_{sim2} to \mathcal{R}_{sec} . The simulator C_{sim2} does not need to communicate with \mathcal{R}_{sec} 's rest part. Upon receiving a look query on Eve_3 , it queries Eve_{sec} with a look query to get the length of \mathcal{R}_{sec} 's content and outputs a uniformly sampled bit string with the same length; the output is stored in C_{sim2} 's state to answer future look queries. The rest of the adversary queries to C_{sim2} , e.g. forward or drop queries, are simply forwarded to \mathcal{R}_{sec} . The rest embedding \mathcal{E}_{otp} keeps track of the events that are received on \mathcal{R}_{key} and \mathcal{R}_{aut3} 's rest interfaces to trigger the corresponding events on \mathcal{R}_{sec} 's core. As mentioned earlier, the composition theorem combines this construction step with the previous ones.

In summary, we obtain the following security results. Let ϵ denote the distinguishing advantage for the decisional Diffie-Hellman game and consider the protocol $\pi \equiv \mathbb{1} | ((C_{enc} \odot C_{dhA}) | (C_{dec} \odot C_{dhB}))$. Then, $\pi \{ \{ aut1 |_{c} aut2 \} |_{c} aut3 \}$ ϵ -securely realizes the specification $\{ \{ sec \} \}$ with the simulator $(C_{sim1} | C_{sim2}) | \mathbb{1}$

and the rest embedding $\mathcal{E}_{\text{otp}}(\mathcal{E}_{\text{dh}}(rs_1 \parallel_r rs_2) \parallel_r rs_3)$.

VI. DISCUSSION AND RELATED WORK

In this section, we discuss how our formalization relates to the existing approaches to composable security arguments and compare it with relevant formalization results.

Our result is an instance of CC. As mentioned in Section II-C, CC is a theory that enables the modular reasoning about system classes, i.e., the shared behavior of similar systems. FRTs introduce a formal approach to abstract over the systems communication patterns in a family of systems that exhibit similar behavior.

Our work is the first that formulates requirements for system communication modeling and provides a rigorous solution as an instance of CC. In Section III-A, we explained the importance of such a solution for the reusability of security proofs. The pen-and-paper presentations of existing CC results [12], [16], [17], [18] do not necessitate providing details on the communication model and, by eliding them, they must be understood as just a high-level overview of security proofs and their composition. However, in our approach, security proofs are mathematical objects and leaving out any details would invalidate them.

Comparing to the simulation-based frameworks [2], [7], [11], [13], we consider a less intricate type of system communication. Nevertheless, there are scenarios in which our approach excels. In simulation-based frameworks, the adversary Turing machine controls the communication between all protocol components. The composition theorems in these frameworks fix the corruption model but consider arbitrary adversary machines besides that. This resembles the role that the FRTs' Rest part plays in our definition of secure constructions and their composition. We provide composition theorems that consider arbitrary Rests; however, we only consider a one-way information flow from the Rest to the Core part, which corresponds to the non-adversary components in the simulation-based frameworks. But note that we allow multiple Rests in our model, e.g. when two FRTs are composed in parallel. This is convenient for modeling scenarios where parties are not corrupted by the same adversary, which would be challenging in simulation-based framework since they use a central adversary machine.

We build on an existing formalization of CC in CryptHOL [15], to which we refer as CCHOL in this section. We reuse large parts of CCHOL's formalization of resources, converters, distinguishers, and various composition operators that were briefly recapped in the paper. Overall, our formalization required 9.7K lines of lemmas and definitions, where the case study constitutes 3.5k lines and the rest form a library of reusable resource definitions and proof rules that can be used for modelling and reasoning about other security protocols.

We compare our result with CCHOL in four respects. First, the limitations mentioned in Section III-A apply to CCHOL proofs since it focuses on individual resources rather than specifications. For example, CCHOL's one-time-pad case study cannot be composed with a Diffie-Hellman key exchange like in Section V because CCHOL's notion of ideal key has a fixed communication pattern and does not allow for the

activations that are needed to exchange the half-keys. Second, trace-equivalence proofs in our result are shorter and less complicated than those in CCHOL using the up-to proof rule for trace equivalence (Thm. 2). Third, CCHOL comes only with an asymptotic security definition whereas we provide both concrete and asymptotic statements and composition theorems. Finally, CCHOL's asymptotic security notion is too strong: in addition to $\mathfrak{d}(\pi \triangleright R, \sigma \triangleright S)$ being negligible in the security parameter, CCHOL requires a notion of functional correctness, which guarantees that the real resource provides the distinguisher with at least the same capabilities as the ideal resource. This hinders composability even further. For example, a secure channel does not securely realize an authenticated channel because the adversary cannot learn the contents of the channel. So one cannot use a secure channel when only an authenticated channel is needed.

There are few results on formalizing simulation-based proofs. Most of them [5], [10] focus on individual (or class of) protocols and hence they do not offer composition theorems that can be applied in arbitrary context. EasyUC [8], which is the most similar to our work, does support the composition of security proofs. It uses EasyCrypt [3] to formalize a simplified version of the Universal Composability framework with a restricted adversary machine. Such simplifications are essential since the formalization of Universal Composability in its most general form is challenging, even with a state-of-the-art formal-methods tool like EasyCrypt.

We compare EasyUC with our result in two respects. First, EasyUC's restriction on the adversary machine affects the reusability of security proofs, as explained in Section III-A, since it leads to a fixed communication pattern among the non-adversary machines. Second, formalizing security statements in EasyUC is more difficult. This is a benefit that stems from CC's abstract approach to cryptography; we don't need to delve into concrete details that are intrinsic part of every security statement in UC-style frameworks. For instance, the EasyUC formalization that constructs a secure channel using Diffie-Hellman key exchange and a one-time-pad constitutes 18K lines of proofs and definitions, where 12K lines are devoted to composing the concrete security statement alone.

VII. CONCLUSION AND FUTURE WORK

We have presented an abstract approach to communication modeling in Constructive Cryptography that is suitable for mechanized verification. We highlight the importance of systems communication patterns on the reusability of security proofs and offer a rigorous approach that allows protocol designers to abstractly capture it. We explain the limitations of existing formalized composability results that do not model systems communications to a full extent. By lifting such limitations, we support the modular formalization of a wider range of scenarios than existing formal-methods tools support. Carrying out further case studies and enhancing automation support is left as future work.

REFERENCES

- [1] Anonymous. Formalization source and HTML output. https://www.dropbox.com/s/eixl861jcaqk0xs/CSF_CC-CryptHOL.zip?dl=0.
- [2] M. Backes, B. Pfizmann, and M. Waidner. A general composition theorem for secure reactive systems. In *Theory of Cryptography (TCC 2003)*, *Proceedings*, volume 2950 of *LNCS*, pages 335–354. Springer, 2003.
- [3] G. Barthe, B. Grégoire, S. Héraud, and S. Z. Béguelin. Computer-aided security proofs for the working cryptographer. In *Advances in Cryptology (CRYPTO 2011)*, *Proceedings*, volume 6841 of *LNCS*, pages 71–90. Springer, 2011.
- [4] D. A. Basin, A. Lochbihler, and S. R. Sefidgar. CryptHOL: Game-based proofs in higher-order logic. *Journal of Cryptology*, 33(2):494–566, 2020.
- [5] D. Butler, D. Aspinall, and A. Gascón. How to simulate it in isabelle: Towards formal proof for secure multi-party computation. In *Interactive Theorem Proving (ITP 2017)*, *Proceedings*, volume 10499 of *LNCS*, pages 114–130. Springer, 2017.
- [6] J. Camenisch, R. R. Enderlein, S. Krenn, R. Küsters, and D. Rausch. Universal composition with responsive environments. In *Advances in Cryptology (ASIACRYPT 2016)*, *Proceedings, Part II*, volume 10032 of *LNCS*, pages 807–840, 2016.
- [7] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Foundations of Computer Science (FOCS 2001)*, *Proceedings*, pages 136–145. IEEE Computer Society, 2001.
- [8] R. Canetti, A. Stoughton, and M. Varia. EasyUC: Using EasyCrypt to mechanize proofs of universally composable security. In *Computer Security Foundations Symposium (CSF 2019)*, *Proceedings*, pages 167–183. IEEE, 2019.
- [9] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [10] H. Haagh, A. Karbyshev, S. Oechsner, B. Spitters, and P.-Y. Strub. Computer-aided proofs for multiparty computation with active security. In *Computer Security Foundations (CSF 2018)*, *Proceedings*, pages 119–131. IEEE Computer Society, 2018.
- [11] D. Hofheinz and V. Shoup. GNUC: A new universal composability framework. *Journal of Cryptology*, 28(3):423–508, 2015.
- [12] D. Jost and U. Maurer. Overcoming impossibility results in composable security using interval-wise guarantees. In *Advances in Cryptology (CRYPTO 2020)*, *Proceedings, Part I*, volume 12170 of *LNCS*, pages 33–62. Springer, 2020.
- [13] R. Küsters and M. Tuengerthal. The IITM model: A simple and expressive model for universal composability. *IACR Cryptology ePrint Archive*, 2013:25, 2013.
- [14] A. Lochbihler. Probabilistic functions and cryptographic oracles in higher order logic. In *European Symposium on Programming (ESOP 2016)*, *Proceedings*, volume 9632 of *LNCS*, pages 503–531. Springer, 2016.
- [15] A. Lochbihler, S. R. Sefidgar, D. Basin, and U. Maurer. Formalizing constructive cryptography using cryptHOL. In *Computer Security Foundations Symposium (CSF 2019)*, *Proceedings*, pages 152–166. IEEE, 2019.
- [16] U. Maurer. Constructive cryptography - A new paradigm for security definitions and proofs. In *Theory of Security and Applications - Joint Workshop (TOSCA 2011)*, *Revised Selected Papers*, volume 6993 of *LNCS*, pages 33–56. Springer, 2011.
- [17] U. Maurer and R. Renner. Abstract cryptography. In *Innovations in Computer Science (ICS 2010)*, *Proceedings*, pages 1–21. Tsinghua University Press, 2011.
- [18] U. Maurer and R. Renner. From indifferentiability to constructive cryptography (and back). In *Theory of Cryptography Conference (TCC 2016)*, *Proceedings, Part I*, volume 9985 of *LNCS*, pages 3–24. Springer, 2016.
- [19] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-order Logic*, volume 2283 of *LNCS*. Springer, 2002.

APPENDIX

A. Guide to the source theory files

All the definitions and lemmas in this paper are formalized and verified in the Isabelle/HOL proof assistant. In what follows, we provide a guide for the reader to navigate the source theories, which are available online [1].

The root directory consists of many theory files, the **Specifications** folder that contains the ideal specifications for keys and channels, and the **Constructions** folder that stores our case study’s formalization. The theory files in the root directory contain the lemmas and definitions that correspond to their names:

- **Fold_Spmf.thy** and **Goodies.thy** formalize the probabilistic fold function and a series of helper lemmas that are used in the other theory files.
- **More_CC.thy** and **State_Isomorphism.thy** formalize our extensions to the theory of resources and converters. The lemmas and definitions in this theory file are not specific to Fused Resources only.
- **Observe_Failure.thy** formalizes the indistinguishability of resources when a bottom element \perp is added to every distribution’s sample space to model failures.
- **Fused_Resource.thy** formalizes Fused Resources, their trace equivalence, and various operators on Core and Rest records. In particular, the proposition `trace'_eq_simI_upto` proves the Theorem 2 for Fused Resources.
- **Construction_Utility.thy** formalizes common building blocks for defining compound Fused Resources. In particular, the propositions `parallel_oracle_fuse` and `attach_parallel_fuse'` prove the Lemmas 1 and 2; and the propositions `fuse_ishift_core_to_rest` and `move_simulator_interface` prove the Lemmas 3 and 4 respectively.
- **Concrete_Security.thy** formalizes the notion of information-theoretic concrete security. In particular, the propositions `constructive_security_obsf_trivial`, `parallel_constructive_security_obsf`, `constructive_security_obsf_composability`, and `constructive_security_obsf_lifting_usr` prove the four claims of Theorem 3.
- **Asymptotic_Security.thy** extends the above to asymptotic notion of security.

The **Specifications** folder provides the ideal specifications for keys and channels. We explained the details of **Key.thy** in Section III-B3.

Our case study has three main theory files, which are stored in the folder **Constructions**.

- **One_Time_Pad.thy** formalizes the construction of a secure channel from a key and an authenticated channel using one-time-pad encryption.
- **Diffie_Hellman.thy** constitutes the formal construction of a key from two authenticated channels using the Diffie-Hellman key exchange.
- **DH_OTP.thy** stores the final lemma that states the security of the aforementioned constructions’ composition.