

A tutorial introduction to CryptHOL

Andreas Lochbihler¹ and S. Reza Sefidgar²

¹Digital Asset (Switzerland) GmbH, Zurich, Switzerland,
mail@andreas-lochbihler.de

²Institute of Information Security, Department of Computer Science,
ETH Zurich, Zurich, Switzerland,
reza.sefidgar@inf.ethz.ch

Abstract

This tutorial demonstrates how cryptographic security notions, constructions, and game-based security proofs can be formalized using the CryptHOL framework. As a running example, we formalize a variant of the hash-based ElGamal encryption scheme and its IND-CPA security in the random oracle model. This tutorial assumes familiarity with Isabelle/HOL basics and standard cryptographic terminology.

1 Introduction

CryptHOL [2, 9] is a framework for constructing rigorous game-based proofs using the proof assistant Isabelle/HOL [13]. Games are expressed as probabilistic functional programs that are shallowly embedded in higher-order logic (HOL) using CryptHOL’s combinators. The security statements, both concrete and asymptotic, are expressed as Isabelle/HOL theorem statements, and their proofs are written declaratively in Isabelle’s proof language Isar [18]. This way, Isabelle mechanically checks that all definitions and statements are type-correct and each proof step is a valid logical inference in HOL. This ensures that the resulting theorems are valid in higher-order logic.

This tutorial explains the CryptHOL essentials using a simple security proof. Our running example is a variant of the hashed ElGamal encryption scheme [5]. We formalize the scheme, the indistinguishability under chosen plaintext (IND-CPA) security property, the computational Diffie-Hellman (CDH) hardness assumption [3], and the security proof in the random oracle model. This illustrates how the following aspects of a cryptographic security proof are formalized using CryptHOL:

- Game-based security definitions (CDH in §2.1 and IND-CPA in §2.4)
- Oracles (a random oracle in §2.2)
- Cryptographic schemes, both generic (the concept of an encryption scheme) and a particular instance (the hashed Elgamal scheme in §2.5)
- Security statements (concrete and asymptotic, §3.2 and §4.2)
- Reductions (from IND-CPA to CDH for hashed Elgamal in §3.1)

- Different kinds of proof steps (§§3.3–3.8):
 - Using intermediate games
 - Defining failure events and applying indistinguishability-up-to lemmas
 - Equivalence transformations on games

This tutorial assumes that the reader knows the basics of Isabelle/HOL and game-based cryptography and wants to get hands-on with CryptHOL. The semantics behind CryptHOL’s embedding in higher-order logic and its soundness are not discussed; we refer the reader to the scientific articles for that [2, 9]. Shoup’s tutorial [16] provides a good introduction to game-based proofs. The following Isabelle features are frequently used in CryptHOL formalizations; the tutorials are available from the Documentation panel in Isabelle/jEdit.

- Function definitions (tutorials `prog-prove` and `functions`, [8])
- Locales (tutorial `locales`, [1])
- The Transfer package [7]

This document is generated from a corresponding Isabelle theory file available online [11].¹ It contains this text and all examples, including the security definitions and proofs. We encourage all readers to download the latest version of the tutorial and follow the proofs and examples interactively in Isabelle/HOL. In particular, a Ctrl-click on a formal entity (function, constant, theorem name, ...) jumps to the definition of the entity.

We split the tutorial into a series of recipes for common formalization tasks. In each section, we cover a familiar cryptography concept and show how it is formalized in CryptHOL. Simultaneously, we explain the Isabelle/HOL and functional programming topics that are essential for formalizing game-based proofs.

1.1 Getting started

CryptHOL is available as part of the Archive of Formal Proofs [10]. Cryptography formalizations based on CryptHOL are arranged in Isabelle theory files that import the relevant libraries.

```
theory CryptHOL-Tutorial imports
  CryptHOL.CryptHOL
begin
```

The file `CryptHOL.CryptHOL` is the canonical entry point into CryptHOL. For the hashed Elgamal example in this tutorial, the CryptHOL library contains everything that is needed. Additional Isabelle libraries can be imported if necessary.

¹The tutorial has been added to the Archive of Formal Proofs after the release of Isabelle2018. Until the subsequent Isabelle release, the tutorial is only available in the development version at https://devel.isa-afp.org/entries/Game_Based_Crypto.html. The version for Isabelle2018 is available at http://www.andreas-lochbihler.de/pub/crypthol_tutorial.zip.

2 Modelling cryptography using CryptHOL

This section demonstrates how the following cryptographic concepts are modelled in CryptHOL.

- A security property without oracles (§2.1)
- An oracle (§2.2)
- A cryptographic concept (§2.3)
- A security property with an oracle (§2.4)
- A concrete cryptographic scheme (§2.5)

2.1 Security notions without oracles: the CDH assumption

In game-based cryptography, a security property is specified using a game between a challenger and an adversary. The probability of an adversary to win the game against the challenger is called its advantage. A cryptographic construction satisfies a security property if the advantage for any “feasible” adversary is “negligible”. A typical security proof reduces the security of a construction to the assumed security of its building blocks. In a concrete security proof, it is therefore not necessary to formally define “feasibility” and “negligibility”, as the security statement establishes a concrete relation between the advantages of specific adversaries.² We return to asymptotic security statements in §4.

A formalization of a security property must therefore specify all of the following:

- The operations of the scheme (e.g., an algebraic group, an encryption scheme)
- The type of adversary
- The game with the challenger
- The advantage of the adversary as a function of the winning probability

For hashed Elgamal, the cyclic group must satisfy the computational Diffie-Hellman assumption. To keep the proof simple, we formalize the equivalent list version of CDH.

Definition (The list computational Diffie-Hellman game). Let \mathcal{G} be a group of order q with generator \mathbf{g} . The List Computational Diffie-Hellman (LCDH) assumption holds for \mathcal{G} if any “feasible” adversary has “negligible” probability in winning the following **LCDH game** against a challenger:

1. The challenger picks x and y randomly (and independently) from $\{0, \dots, q-1\}$.
2. It passes \mathbf{g}^x and \mathbf{g}^y to the adversary. The adversary generates a set L of guesses about the value of \mathbf{g}^{xy} .

²The cryptographic literature sometimes abstracts over the adversary and defines the advantage to be the advantage of the best “feasible” adversary against a game. Such abstraction would require a formalization of feasibility, for which CryptHOL currently does not offer any support. We therefore always consider the advantage of a specific adversary.

3. The adversary wins the game if $\mathbf{g}^{xy} \in L$.

The scheme for LCDH uses only a cyclic group. To make the LCDH formalisation reusable, we formalize the LCDH game for an arbitrary cyclic group \mathcal{G} using Isabelle’s module system based on locales. The locale *list-cdh* fixes \mathcal{G} as a finite cyclic group that has elements of type *'grp*. Basic facts about finite groups are formalized in the CryptHOL theory *CryptHOL.Cyclic-Group*.³

```
locale list-cdh = cyclic-group  $\mathcal{G}$ 
for  $\mathcal{G}$  :: 'grp cyclic-group (structure)
begin
```

The LCDH game does not need oracles. The adversary is therefore just a probabilistic function from two group elements to a set of guesses, which are group elements again. In CryptHOL, the probabilistic nature is expressed by the adversary returning a discrete subprobability distribution over sets of guesses, as expressed by the type constructor *spmf*. We define the following abbreviation as a shorthand for the type of LCDH adversaries.⁴

```
type-synonym 'grp' adversary = 'grp'  $\Rightarrow$  'grp'  $\Rightarrow$  'grp' set spmf
```

The LCDH game itself is expressed as a function from the adversary \mathcal{A} to the subprobability distribution of the adversary winning. CryptHOL provides operators to express these distributions as probabilistic programs and reason about them using program logics:

- The *do* notation desugars to monadic sequencing in the monad of subprobabilities [17]. Intuitively, every line $x \leftarrow p$; samples an element x from the distribution p . The sampling is independent unless the distribution p depends on previously sampled variables. At the end of the block, the *return-spmf* _ returns whether the adversary has won the game.
- *sample-uniform* n denotes the uniform distribution over the set $\{0, \dots, n - 1\}$.
- *order* \mathcal{G} denotes the order of \mathcal{G} and $([\wedge]) :: 'grp \Rightarrow nat \Rightarrow 'grp$ is the group exponentiation operator.

The LCDH game formalizes the challenger’s behavior against an adversary \mathcal{A} . In the following definition, the challenger randomly (and independently) picks two natural numbers x and y that are between 0 and \mathcal{G} ’s order and passes them to the adversary. The adversary then returns a set *zs* of guesses for $\mathbf{g}^{x * y}$, where \mathbf{g} is the generator of \mathcal{G} . The game finally returns a *boolean* that indicates whether the adversary produced a right guess. Formally, *game* \mathcal{A} is a *boolean* random variable.

```
definition game :: 'grp adversary  $\Rightarrow$  bool spmf where
  game  $\mathcal{A}$  = do {
```

³The syntax directive **structure** tells Isabelle that all group operations in the context of the locale refer to the group \mathcal{G} unless stated otherwise.

Isabelle automatically adds the locale parameters and the assumptions on them to all definitions and lemmas inside that locale. Of course, we could have made the group \mathcal{G} an explicit argument of all functions ourselves, but then we would not benefit from Isabelle’s module system, in particular locale instantiation.

⁴Actually, the type of group elements has already been fixed in the locale *list-cdh* to the type variable *'grp*. Unfortunately, such fixed type variables cannot be used in type declarations inside a locale in Isabelle2018. The **type-synonym** *adversary* is therefore parametrized by a different type variable *'grp'*, but it will be used below only with *'grp*.

```

  x ← sample-uniform (order  $\mathcal{G}$ );
  y ← sample-uniform (order  $\mathcal{G}$ );
  zs ←  $\mathcal{A}$  (g [^] x) (g [^] y);
  return-spmf (g [^] (x * y) ∈ zs)
}

```

The advantage of the adversary is equivalent to its probability of winning the LCDH game, which is measured using the *spmf* function.

definition *advantage* :: 'grp adversary \Rightarrow real
where *advantage* \mathcal{A} = *spmf* (game \mathcal{A}) True

end

This completes the formalisation of the LCDH game and we close the locale *list-cdh* with **end**. The above definitions are now accessible under the names *game* and *advantage*. Furthermore, when we later instantiate the locale *list-cdh*, they will be specialized to the given parameters. We will return to this topic in §2.5.

2.2 A Random Oracle

A cryptographic oracle grants an adversary black-box access to a certain information or functionality. In this section, we formalize a random oracle, i.e., an oracle that models a random function with a finite codomain. In the Elgamal security proof, the random oracle represents the hash function: the adversary can query the oracle for a value and the oracle responds with the corresponding “hash”.

Like for the LCDH formalization, we wrap the random oracle in the locale *random-oracle* for modularity. The random oracle will return a *bitstring*, i.e. a list of booleans, of length *len*.

type-synonym *bitstring* = *bool list*

locale *random-oracle* =
fixes *len* :: *nat*
begin

In CryptHOL, oracles are modeled as probabilistic transition systems that given an initial state and an input, return a subprobability distribution over the output and the successor state. The type synonym (*'s, 'a, 'b*) *oracle'* abbreviates *'s \Rightarrow 'a \Rightarrow ('b \times 's) spmf*.

A random oracle accepts queries of type *'a* and generates a random bitstring of length *len*. The state of the random oracle remembers its previous responses in a mapping of type *'a \rightarrow bitstring*. Upon a query *x*, the oracle first checks whether this query was received before. If so, the oracle returns the same answer again. Otherwise, the oracle randomly samples a bitstring of length *len*, stores it in its state, and returns it alongside with the new state.

type-synonym *'a state* = *'a \rightarrow bitstring*

definition *oracle* :: *'a state \Rightarrow 'a \Rightarrow (bitstring \times 'a state) spmf*
where
oracle σ *x* = (case σ *x* of
 None \Rightarrow do {

```

    bs ← spmf-of-set (nlists UNIV len);
    return-spmf (bs, σ(x ↦ bs)) }
| Some bs ⇒ return-spmf (bs, σ)

```

Initially, the state of a random oracle is the empty map *empty*, as no queries have been asked. For readability, we introduce an abbreviation:

abbreviation (*input*) *initial* :: 'a state **where** *initial* ≡ *Map.empty*

This actually completes the formalization of the random oracle. Before we close the locale, we prove two technical lemmas:

1. The lemma *lossless-oracle* states that the distribution over answers and successor states is *lossless*, i.e., a full probability distribution. Many reasoning steps in game-based proofs are only valid for lossless distributions, so it is generally recommended to prove losslessness of all definitions if possible.
2. The lemma *fresh* describes random oracle's behavior when the query is fresh. This lemma makes it possible to automatically unfold the random oracle only when it is known that the query is fresh.

lemma *lossless-oracle* [*simp*]: *lossless-spmf* (*oracle* σ *x*)
by(*simp add: oracle-def split: option.split*)

lemma *fresh*:
oracle σ *x* =
 (do { *bs* ← *spmf-of-set* (*nlists* UNIV *len*);
 return-spmf (*bs*, σ(*x* ↦ *bs*)) })
if σ *x* = *None*
using that **by**(*simp add: oracle-def*)

end

Remark: Independence is the default. Note that - *spmf* represents a discrete probability distribution rather than a random variable. The difference is that every *spmf* is independent of all other *spmf*s. There is no implicit space of elementary events via which information may be passed from one random variable to the other. If such information passing is necessary, this must be made explicit in the program. That is why the random oracle explicitly takes a state of previous responses and returns the updated states. Later, whenever the random oracle is used, the user must pass the state around as needed. This also applies to adversaries that may want to store some information.

2.3 Cryptographic concepts: public-key encryption

A cryptographic concept consists of a set of operations and their functional behaviour. We have already seen two very simple examples: the cyclic group in §2.1 and the random oracle in §2.2. We have formalized both of them as locales; we have not modelled their functional behavior as this is not needed for the proof. In this section, we now present a more realistic example: public-key encryption with oracle access.

A public-key encryption scheme consists of three algorithms: key generation, encryption, and decryption. They are all probabilistic and, in the most general case, they may

access an oracle jointly with the adversary. As before, the operations are modelled as parameters of a locale, *ind-cpa-pk*.

- The key generation algorithm *key-gen* outputs a public-private key pair.
- The encryption operation *encrypt* takes a public key and a plaintext of type *'plain* and outputs a ciphertext of type *'cipher*.
- The decryption operation *decrypt* takes a private key and a ciphertext and outputs a plaintext.
- Additionally, the predicate *valid-plains* tests whether the adversary has chosen a valid pair of plaintexts. This operation is needed only in the IND-CPA game definition in the next section, but we include it already here for convenience.

```

locale ind-cpa-pk =
  fixes key-gen :: ('pubkey × 'privkey, 'query, 'response) gpv
  and encrypt :: 'pubkey ⇒ 'plain ⇒ ('cipher, 'query, 'response) gpv
  and decrypt :: 'privkey ⇒ 'cipher ⇒ ('plain, 'query, 'response) gpv
  and valid-plains :: 'plain ⇒ 'plain ⇒ bool
begin

```

The three actual operations are generative probabilistic values (GPV) of type $(-, 'query, 'response) gpv$. A GPV is a probabilistic algorithm that has not yet been connected to its oracles. The interface to the oracle is abstracted in the two type parameters *'query* for queries and *'response* for responses. As before, we omit the specification of the functional behavior, namely that decrypting an encryption with a key pair returns the plaintext.

2.4 Security notions with oracles: IND-CPA security

In general, there are several security notions for the same cryptographic concept. For encryption schemes, an indistinguishability notion of security [6] is often used. We now formalize the notion indistinguishability under chosen plaintext attacks (IND-CPA) for public-key encryption schemes. Goldwasser et al. [15] showed that IND-CPA is equivalent to semantic security.

Definition (IND-CPA). Let *key-gen*, *encrypt* and *decrypt* denote a public-key encryption scheme. The IND-CPA game is a two-stage game between the *adversary* and a *challenger*:

Stage 1 (find):

1. The challenger generates a public key *pk* using *key-gen* and gives the public key to the adversary.
2. The adversary returns two messages *m*₀ and *m*₁.
3. The challenger checks that the two messages are a valid pair of plaintexts. (For example, both messages must have the same length.)

Stage 2 (guess):

1. The challenger flips a coin *b* (either 0 or 1) and gives *encrypt pk m_b* to the adversary.

2. The adversary returns a bit b' .

The adversary wins the game if his guess b' is the value of b . Let P_{win} denote the winning probability. His advantage is $|P_{win} - 1/2|$

Like for the encryption scheme, we will define the game such that the challenger and the adversary have access to a shared oracle, but the oracle is still unspecified. Consequently, the corresponding CryptHOL game is a GPV, like the operations of the abstract encryption scheme. When we specialize the definitions in the next section to the hashed Elgamal scheme, the GPV will be connected to the random oracle.

The type of adversary is now more complicated: It is a pair of probabilistic functions with oracle access, one for each stage of the game. The first computes the pair of plaintext messages and the second guesses the challenge bit. The additional *'state'* parameter allows the adversary to maintain state between the two stages.

type-synonym (*'pubkey'*, *'plain'*, *'cipher'*, *'query'*, *'response'*, *'state'*) *adversary* =
 (*'pubkey'* \Rightarrow ((*'plain'* \times *'plain'*) \times *'state'*, *'query'*, *'response'*) *gpv*)
 \times (*'cipher'* \Rightarrow *'state'* \Rightarrow (*bool*, *'query'*, *'response'*) *gpv*)

The IND-CPA game formalization below follows the above informal definition. There are three points that need some explanation. First, this game differs from the simpler LCDH game in that it works with GPVs instead of SPMFs. Therefore, probability distributions like coin flips *coin-spmf* must be lifted from SPMFs to GPVs using the coercion *lift-spmf*. Second, the assertion *assert-gpv* (*valid-plains* m_0 m_1) ensures that the pair of messages is valid. Third, the construct *TRY - ELSE -* catches a violated assertion. In that case, the adversary's advantage drops to 0 because the result of the game is a coin flip.

fun *game* :: (*'pubkey'*, *'plain'*, *'cipher'*, *'query'*, *'response'*, *'state'*) *adversary*
 \Rightarrow (*bool*, *'query'*, *'response'*) *gpv*
where
game (\mathcal{A}_1 , \mathcal{A}_2) = *TRY* *do* {
 (pk , sk) \leftarrow *key-gen*;
 (m_0 , m_1), σ \leftarrow \mathcal{A}_1 pk ;
assert-gpv (*valid-plains* m_0 m_1);
 b \leftarrow *lift-spmf* *coin-spmf*;
cipher \leftarrow *encrypt* pk (*if* b *then* m_0 *else* m_1);
 b' \leftarrow \mathcal{A}_2 *cipher* σ ;
Done ($b' = b$)
} *ELSE* *lift-spmf* *coin-spmf*

Figure 1 visualizes this game as a grey box. The dashed boxes represent parameters of the game or the locale, i.e., parts that have not yet been instantiated. The actual probabilistic program is shown on the left half, which uses the dashed boxes as sub-programs. Arrows in the grey box from the left to the right pass the contents of the variables to the sub-program. Those in the other direction bind the result of the sub-program to new variables. The arrows leaving box indicate the query-response interaction with an oracle. The thick arrows emphasize that the adversary's state is passed around explicitly. The double arrow represents the return value of the game. We will use this to define the adversary's advantage.

As the oracle is not specified in the game, the advantage, too, is parametrized by the oracle, given by the transition function *oracle* :: (*'s'*, *'query'*, *'response'*) *oracle'* and the

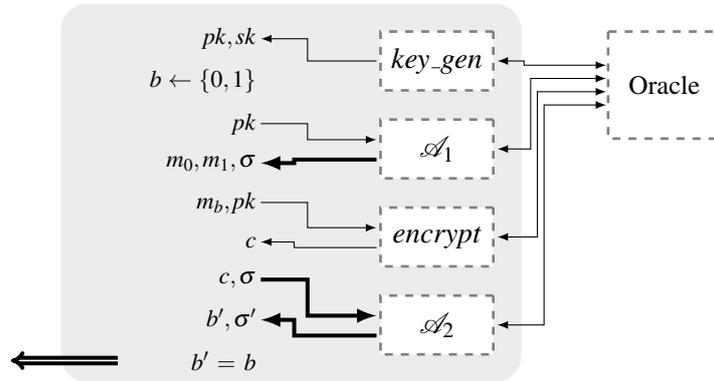


Figure 1: Graphic representation of the generic IND-CPA game.

initial state $\sigma :: 's$ its initial state. The operator *run-gpv* connects the game with the oracle, whereby the GPV becomes an SPMF.

```
fun advantage :: (' $\sigma$ ', 'query', 'response') oracle'  $\times$  ' $\sigma$ '
 $\Rightarrow$  ('pubkey', 'plain', 'cipher', 'query', 'response', 'state') adversary  $\Rightarrow$  real
```

where

```
advantage (oracle,  $\sigma$ )  $\mathcal{A} = |spmf (run-gpv oracle (game \mathcal{A}) \sigma) True - 1/2|$ 
```

end

2.5 Concrete cryptographic constructions: the hashed ElGamal encryption scheme

With all the above modelling definitions in place, we are now ready to explain how concrete cryptographic constructions are expressed in CryptHOL. In general, a cryptographic construction builds a cryptographic concept from possibly several simpler cryptographic concepts. In the running example, the hashed ElGamal cipher [5] constructs a public-key encryption scheme from a finite cyclic group and a hash function. Accordingly, the formalisation consists of three steps:

1. Import the cryptographic concepts on which the construction builds.
2. Define the concrete construction.
3. Instantiate the abstract concepts with the construction.

First, we declare a new locale that imports the two building blocks: the cyclic group from the LCDH game with namespace *lcdh* and the random oracle for the hash function with namespace *ro*. This ensures that the construction can be used for arbitrary cyclic groups. For the message space, it suffices to fix the length *len-plain* of the plaintexts.

```
locale hashed-elgamal =
  lcdh: list-cdh  $\mathcal{G}$  +
  ro: random-oracle len-plain
for  $\mathcal{G} :: 'grp$  cyclic-group (structure)
and len-plain :: nat
begin
```

Second, we formalize the hashed ElGamal encryption scheme. Here is the well-known informal definition.

Definition (Hashed Elgamal encryption scheme). Let G be a cyclic group of order q that has a generator g . Furthermore, let h be a hash function that maps the elements of G to bitstrings, and \oplus be the xor operator on bitstrings. The Hashed-ElGamal encryption scheme is given by the following algorithms:

Key generation Pick an element x randomly from the set $\{0, \dots, q-1\}$ and output the pair (g^x, x) , where g^x is the public key and x is the private key.

Encryption Given the public key pk and the message m , pick y randomly from the set $\{0, \dots, q-1\}$ and output the pair $(g^y, h(pk^y) \oplus m)$. Here \oplus denotes the bitwise exclusive-or of two bitstrings.

Decryption Given the private key sk and the ciphertext (α, β) , output $h(\alpha^{sk}) \oplus \beta$.

As we can see, the public key is a group element, the private key a natural number, a plaintext a bitstring, and a ciphertext a pair of a group element and a bitstring.⁵ For readability, we introduce meaningful abbreviations for these concepts.

type-synonym $'grp'$ *pub-key* = $'grp'$
type-synonym $'grp'$ *priv-key* = nat
type-synonym *plain* = $bitstring$
type-synonym $'grp'$ *cipher* = $'grp' \times bitstring$

We next translate the three algorithms into CryptHOL definitions. The definitions are straightforward except for the hashing. Since we analyze the security in the random oracle model, an application of the hash function H is modelled as a query to the random oracle using the GPV *hash*. Furthermore, we define the plaintext validity predicate to check the length of the adversary's messages produced by the adversary.

abbreviation $hash :: 'grp \Rightarrow (bitstring, 'grp, bitstring) gpv$
where
 $hash\ x \equiv Pause\ x\ Done$

definition $key-gen :: ('grp\ pub-key \times 'grp\ priv-key) spmf$
where
 $key-gen = do \{$
 $\quad x \leftarrow sample-uniform\ (order\ \mathcal{G});$
 $\quad return-spmf\ (\mathbf{g}\ [\wedge] x, x)$
 $\}$

definition $encrypt :: 'grp\ pub-key \Rightarrow plain \Rightarrow ('grp\ cipher, 'grp, bitstring) gpv$
where
 $encrypt\ \alpha\ msg = do \{$
 $\quad y \leftarrow lift-spmf\ (sample-uniform\ (order\ \mathcal{G}));$
 $\quad h \leftarrow hash\ (\alpha\ [\wedge] y);$
 $\quad Done\ (\mathbf{g}\ [\wedge] y, h\ [\oplus] msg)$
 $\}$

definition $decrypt :: 'grp\ priv-key \Rightarrow 'grp\ cipher \Rightarrow (plain, 'grp, bitstring) gpv$

⁵More precisely, the private key ranges between 0 and $q-1$ and the bitstrings are of length $len\ plain$. However, Isabelle/HOL's type system cannot express such properties that depend on locale parameters.

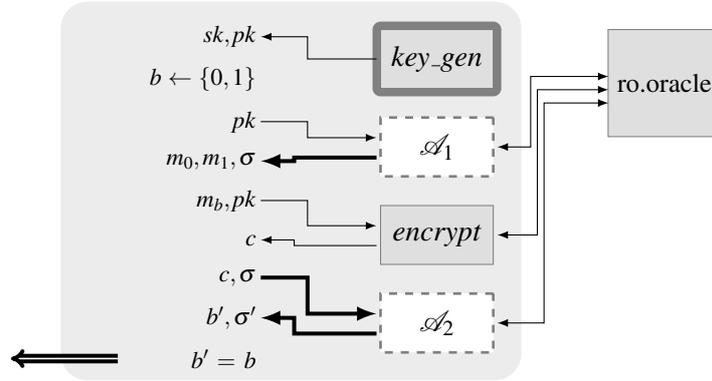


Figure 2: The IND-CPA game instantiated with the Hashed-ElGamal encryption scheme and accessing a random oracle.

where

```

decrypt  $x = (\lambda(\beta, \zeta)$ . do {
   $h \leftarrow \text{hash}(\beta \uparrow x)$ ;
  Done ( $\zeta \oplus h$ )
})

```

definition *valid-plains* :: *plain* \Rightarrow *plain* \Rightarrow *bool*

where

```

valid-plains msg1 msg2  $\longleftrightarrow$  length msg1 = len-plain  $\wedge$  length msg2 = len-plain

```

The third and last step instantiates the interface of the encryption scheme with the hashed Elgamal scheme. This specializes all definition and theorems in the locale *ind-cpa-pk* to our scheme.

sublocale *ind-cpa*: *ind-cpa-pk* (*lift-spmf* *key-gen*) *encrypt* *decrypt* *valid-plains* .

Figure 2 visualizes the instantiation. In comparison to Fig. 1, the boxes for the key generation and the encryption algorithm have been instantiated with the hashed Elgamal definitions from this section. We nevertheless draw the boxes to indicate that the definitions of these algorithms has not yet been inlined in the game definition. The thick grey border around the key generation algorithm denotes the *lift-spmf* operator, which embeds the probabilistic *key-gen* without oracle access into the type of GPVs with oracle access. The oracle has also been instantiated with the random oracle *oracle* imported from *hashed-elgamal*'s parent locale *random-oracle* with prefix *ro*.

3 Cryptographic proofs in CryptHOL

This section explains how cryptographic proofs are expressed in CryptHOL. We will continue our running example by stating and proving the IND-CPA security of the hashed Elgamal encryption scheme under the computational Diffie-Hellman assumption in the random oracle model, using the definitions from the previous section. More precisely, we will formalize a reduction argument (§3.1) and bound the IND-CPA advantage using the CDH advantage. We will *not* formally state the result that CDH hardness in the cyclic group implies IND-CPA security, which quantifies over all fea-

sible adversaries—to that end, we would have to formally define feasibility, for which CryptHOL currently does not offer any support.

The actual proof of the bound consists of several game transformations. We will focus on those steps that illustrate common steps in cryptographic proofs (§3.3–§3.8).

3.1 The reduction

The security proof involves a reduction argument: We will derive a bound on the advantage of an arbitrary adversary in the IND-CPA game *game* for hashed Elgamal that depends on another adversary’s advantage in the LCDH game *game* of the underlying group. The reduction transforms every IND-CPA adversary \mathcal{A} into a LCDH adversary *elgamal-reduction* \mathcal{A} , using \mathcal{A} as a black box. In more detail, it simulates an execution of the IND-CPA game including the random oracle. At the end of the game, the reduction outputs the set of queries that the adversary has sent to the random oracle. In detail, the reduction works as follows given a two part IND-CPA adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ (Figure 3 visualizes the reduction as the dotted box):

1. It receives two group elements α and β from the LCDH challenger.
2. The reduction passes α to the adversary as the public key and runs \mathcal{A}_1 to get messages m_1 and m_2 . The adversary is given access to the random oracle with the initial state *empty*.
3. The assertion checks that the adversary returns two valid plaintexts, i.e., m_1 and m_2 are strings of length *len-plain*.
4. Instead of actually performing an encryption, the reduction generates a random bitstring h of length *len-plain* (*nlists UNIV len-plain* denotes the set of all bitstrings of length *len-plain* and *spmf-of-set* converts the set into a uniform distribution over the set.)
5. The reduction passes (β, h) as the challenge ciphertext to the adversary in the second phase of the IND-CPA game.
6. The actual guess b' of the adversary is ignored; instead the reduction returns the set *dom s'* of all queries that the adversary made to the random oracle as its guess for the CDH game.
7. If any of the steps after the first phase fails, the reduction’s guess is the set *dom s* of oracle queries made during the first phase.

```
fun elgamal-reduction
  :: ('grp pub-key, plain, 'grp cipher, 'grp, bitstring, 'state) ind-cpa.adversary
  ⇒ 'grp lcdh.adversary
```

where

```
elgamal-reduction ( $\mathcal{A}_1, \mathcal{A}_2$ )  $\alpha \beta = do$  {
  (( $m_1, m_2$ ),  $\sigma$ ),  $s$  ← exec-gpv ro.oracle ( $\mathcal{A}_1 \alpha$ ) ro.initial;
  TRY do {
    - :: unit ← assert-spmf (valid-plains  $m_1 m_2$ );
     $h$  ← spmf-of-set (nlists UNIV len-plain);
    ( $b', s'$ ) ← exec-gpv ro.oracle ( $\mathcal{A}_2 (\beta, h) \sigma$ )  $s$ ;
    return-spmf (dom s')
  } ELSE return-spmf (dom s)
}
```

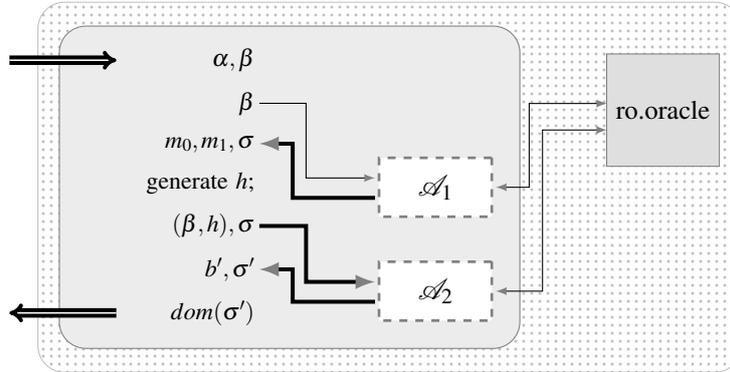


Figure 3: The reduction for the Elgamal security proof.

3.2 Concrete security statement

A concrete security statement in CryptHOL has the form: Subject to some side conditions for the adversary \mathcal{A} , the advantage in one game is bounded by a function of the transformed adversary's advantage in a different game.⁶

theorem *concrete-security*:

assumes *side conditions for \mathcal{A}*

shows $\text{advantage}_1 \mathcal{A} \leq f(\text{advantage}_2(\text{reduction } \mathcal{A}))$

For the hashed Elgamal scheme, the theorem looks as follows, i.e., the function f is the identity function.

theorem *concrete-security-Elgamal*:

assumes *lossless: ind-cpa.lossless \mathcal{A}*

shows $\text{ind-cpa.advantage}(\text{ro.oracle}, \text{ro.initial}) \mathcal{A} \leq \text{lcdh.advantage}(\text{Elgamal-reduction } \mathcal{A})$

Such a statement captures the essence of a concrete security proof. For if there was a feasible adversary \mathcal{A} with non-negligible advantage against the *game*, then *Elgamal-reduction* \mathcal{A} would be an adversary against the *game* with at least the same advantage. This implies the existence of an adversary with non-negligible advantage against the cryptographic primitive that was assumed to be secure. What we cannot state formally is that the transformed adversary *Elgamal-reduction* \mathcal{A} is feasible as we have not formalized the notion of feasibility. The readers of the formalization must convince themselves that the reduction preserves feasibility.

In the case of *Elgamal-reduction*, this should be obvious from the definition (given the theorem's side condition) as the reduction does nothing more than sampling and redirecting data.

Our proof for the concrete security theorem needs the side condition that the adversary is lossless. Losslessness for adversaries is similar to losslessness for subprobability

⁶A security proof often involves several reductions. The bound then depends on several advantages, one for each reduction.

distributions. It ensures that the adversary always terminates and returns an answer to the challenger. For the IND-CPA game, we define losslessness as follows:

definition (in *ind-cpa-pk*) *lossless*
 $:: ('pubkey, 'plain, 'cipher, 'query, 'response, 'state) adversary \Rightarrow bool$
where
 $lossless = (\lambda(\mathcal{A}_1, \mathcal{A}_2). (\forall pk. lossless-gpv \mathcal{I}\text{-full}(\mathcal{A}_1 pk))$
 $\wedge (\forall cipher \sigma. lossless-gpv \mathcal{I}\text{-full}(\mathcal{A}_2 cipher \sigma)))$

So now let's start with the proof.

proof –

As a preparatory step, we split the adversary \mathcal{A} into its two phases \mathcal{A}_1 and \mathcal{A}_2 . We could have made the two phases explicit in the theorem statement, but our form is easier to read and use. We also immediately decompose the losslessness assumption on \mathcal{A} .⁷

obtain $\mathcal{A}_1 \mathcal{A}_2$ **where** \mathcal{A} [simp]: $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ **by** (cases \mathcal{A})
from *lossless* **have** *lossless1* [simp]: $\wedge pk. lossless-gpv \mathcal{I}\text{-full}(\mathcal{A}_1 pk)$
and *lossless2* [simp]: $\wedge \sigma cipher. lossless-gpv \mathcal{I}\text{-full}(\mathcal{A}_2 \sigma cipher)$
by(*auto simp add: ind-cpa.lossless-def*)

3.3 Recording adversary queries

As can be seen in Fig. 2, both the adversary and the encryption of the challenge ciphertext use the random oracle. The reduction, however, returns only the queries that the adversary makes to the oracle (in Fig. 3, h is generated independently of the random oracle). To bridge this gap, we introduce an *interceptor* between the adversary and the oracle that records all adversary's queries.

define *interceptor* $:: 'grp set \Rightarrow 'grp \Rightarrow (bitstring \times 'grp set, -, -) gpv$
where
 $interceptor \sigma x = (do \{$
 $h \leftarrow hash x;$
 $Done (h, insert x \sigma)$
 $\})$ **for** σx

We integrate this interceptor into the *game* using the *inline* function as illustrated in Fig. 4 and name the result *game₀*.

define *game₀* **where**
 $game_0 = TRY do \{$
 $(pk, -) \leftarrow lift\text{-}spmf\ key\text{-}gen;$
 $((m_1, m_2), \sigma), s \leftarrow inline\ interceptor (\mathcal{A}_1 pk) \{ \};$
 $assert\text{-}gpv (valid\text{-}plains m_1 m_2);$
 $b \leftarrow lift\text{-}spmf\ coin\text{-}spmf;$
 $c \leftarrow encrypt\ pk (if\ b\ then\ m_1\ else\ m_2);$
 $(b', s') \leftarrow inline\ interceptor (\mathcal{A}_2 c \sigma) s;$
 $Done (b' = b)$
 $\} ELSE lift\text{-}spmf\ coin\text{-}spmf$

⁷Later in the proof, we will often prove losslessness of the definitions in the proof. We will not show them in this document, but they are in the Isabelle sources from which this document is generated.

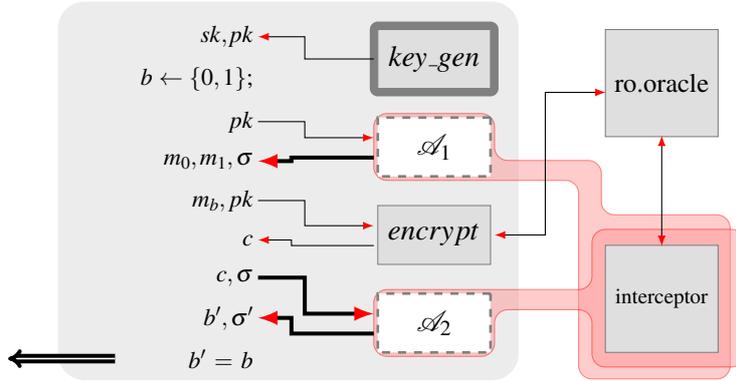


Figure 4: The IND-CPA game after expanding the key generation algorithm’s definition and inlining the query-recording hash oracle. The red boxes represent the inline operator.

We claim that the above modifications do not affect the output of the IND-CPA game at all. This might seem obvious since we are only logging the adversary’s queries without modifying them. However, in a formal proof, this needs to be precisely justified.

More precisely, we have been very careful that the two games *game* \mathcal{A} and *game*₀ have identical structure. They differ only in that *game*₀ uses the adversary $(\lambda pk. inline\ interceptor\ (\mathcal{A}_1\ pk)\ \emptyset, \lambda cipher\ \sigma. inline\ interceptor\ (\mathcal{A}_2\ cipher\ \sigma))$ instead of \mathcal{A} . The formal justification for this replacement happens in two steps:

1. We replace the oracle transformer *interceptor* with *id-oracle*, which merely passes queries and results to the oracle.
2. Inlining the identity oracle transformer *id-oracle* does not change an adversary and can therefore be dropped.

The first step is automated using Isabelle’s Transfer package [7], which is based on Mitchell’s representation independence [12]. The replacement is controlled by so-called transfer rules of the form $R\ x\ y$ which indicates that x shall replace y ; the correspondence relation R captures the kind of replacement. The *transfer* proof method then constructs a constraint system with one constraint for each atom in the proof goal where the correspondence relation and the replacement are unknown. It then tries to solve the constraint system using the rules that have been declared with the attribute $[transfer\ rule]$. Atoms that do not have a suitable transfer rule are not changed and their correspondence relation is instantiated with the identity relation ($=$).

The second step is automated using Isabelle’s simplifier.

In the example, the crucial change happens in the state of the oracle transformer: *interceptor* records all queries in a set whereas *id-oracle* has no state, which is modelled with the singleton type *unit*. To capture the change, we define the correspondence relation cr on the states of the oracle transformers. (As we are in the process of adding this state, this state is irrelevant and cr is therefore always true. We nevertheless have to make an explicit definition such that Isabelle does not automatically beta-reduce terms, which would confuse *transfer*.) We then prove that it relates the initial states and that cr is a bisimulation relation for the two oracle transformers. The bisimulation proof itself is automated, too: A bit of term rewriting (**unfolding**) makes the two

oracle transformers structurally identical except for the state update function. Having proved that the state update function $\lambda\text{-}\sigma$. σ is a correct replacement for *insert* w.r.t. *cr*, the *transfer-prover* then lifts this replacement to the bisimulation rule. Here, *transfer-prover* is similar to *transfer* except that it works only for transfer rules and builds the constraint system only for the term to be replaced.

The theory source of this tutorial contains a step-by-step proof to illustrate how transfer works.

```

{ define cr :: unit  $\Rightarrow$  'grp set  $\Rightarrow$  bool where cr  $\sigma$   $\sigma'$  = True for  $\sigma$   $\sigma'$ 
  have [transfer-rule]: cr () {} by (simp add: cr-def) — initial states
  have [transfer-rule]: ((=)  $\implies$  cr  $\implies$  cr) ( $\lambda\text{-}\sigma$ .  $\sigma$ ) insert — state update
    by (simp add: rel-fun-def cr-def)
  have [transfer-rule]: — cr is a bisimulation for the oracle transformers
    (cr  $\implies$  (=)  $\implies$  rel-gpv (rel-prod (=) cr) (=)) id-oracle interceptor
    unfolding interceptor-def [abs-def] id-oracle-def [abs-def] bind-gpv-Pause bind-rpv-Done
    by transfer-prover
  have ind-cpa.game  $\mathcal{A}$  = game0 unfolding game0-def  $\mathcal{A}$  ind-cpa.game.simps
    by transfer (simp add: bind-map-gpv o-def ind-cpa.game.simps split-def)
}

```

3.4 Equational program transformations

Before we move on, we need to simplify *game₀* and inline a few of the definitions. All these simplifications are equational program transformations, so the Isabelle simplifier can justify them. We combine the *interceptor* with the random oracle *oracle* into a new oracle *oracle'* with which the adversary interacts.

```

define oracle' :: 'grp set  $\times$  ('grp  $\rightarrow$  bitstring)  $\Rightarrow$  'grp  $\Rightarrow$  -
where oracle' = ( $\lambda$ (s,  $\sigma$ ) x. do {
  (h,  $\sigma'$ )  $\leftarrow$  case  $\sigma$  x of
    None  $\Rightarrow$  do {
      bs  $\leftarrow$  spmf-of-set (nlists UNIV len-plain);
      return-spmf (bs,  $\sigma$ (x  $\mapsto$  bs)) }
    | Some bs  $\Rightarrow$  return-spmf (bs,  $\sigma$ );
  return-spmf (h, insert x s,  $\sigma'$ )
})
have *: exec-gpv ro.oracle (inline interceptor  $\mathcal{A}$  s)  $\sigma$  =
  map-spmf ( $\lambda$ (a, b, c). ((a, b), c)) (exec-gpv oracle'  $\mathcal{A}$  (s,  $\sigma$ )) for  $\mathcal{A}$   $\sigma$  s
by (simp add: interceptor-def oracle'-def ro.oracle-def Let-def
  exec-gpv-inline exec-gpv-bind o-def split-def cong del: option.case-cong-weak)

```

We also want to inline the key generation and encryption algorithms, push the *TRY* - *ELSE* - towards the assertion (which is possible because the adversary is lossless by assumption), and rearrange the samplings a bit. The latter is automated using *monad-normalisation* [14].

```

have game0: run-gpv ro.oracle game0 ro.initial = do {
  x  $\leftarrow$  sample-uniform (order  $\mathcal{G}$ );
  y  $\leftarrow$  sample-uniform (order  $\mathcal{G}$ );
  b  $\leftarrow$  coin-spmf;
  (((msg1, msg2),  $\sigma$ ), (s, s-h))  $\leftarrow$ 
    exec-gpv oracle' ( $\mathcal{A}$  ( $\mathbf{g}$  [ $\uparrow$ ] x)) ({} , ro.initial);
  TRY do {
    - :: unit  $\leftarrow$  assert-spmf (valid-plains msg1 msg2);

```

```

(h, s-h') ← ro.oracle s-h (g [^] (x * y));
let cipher = (g [^] y, h [⊕] (if b then msg1 else msg2));
(b', (s', s-h'')) ← exec-gpv oracle' (A2 cipher σ) (s, s-h');
return-spmf (b' = b)
} ELSE do {
  b ← coin-spmf;
  return-spmf b
}
}
including monad-normalisation
by(simp add: game0-def key-gen-def encrypt-def * exec-gpv-bind bind-map-spmf assert-spmf-def
  try-bind-assert-gpv try-gpv-bind-lossless split-def o-def if-distrib lcdh.nat-pow-pow)

```

This call to Isabelle’s simplifier may look complicated at first, but it can be constructed incrementally by adding a few theorems and looking at the resulting goal state and searching for suitable theorems using **find-theorems**. As always in Isabelle, some intuition and knowledge about the library of lemmas is crucial.

- We knew that the definitions *game₀-def*, *key-gen-def*, and *encrypt-def* should be unfolded, so they are added first to the simplifier’s set of rewrite rules.
- The equations *exec-gpv-bind*, *try-bind-assert-gpv*, and *try-gpv-bind-lossless* ensure that the operator *exec-gpv*, which connects the *game₀* with the random oracle, is distributed over the sequencing. Together with ***, this gives the adversary access to *oracle'* instead of the interceptor and the random oracle, and makes the call to the random oracle in the encryption of the chosen message explicit.
- The theorem *lcdh.nat-pow-pow* rewrites the iterated exponentiation $(g [^] x) [^] y$ to $g [^] (x * y)$.
- The other theorems *bind-map-spmf*, *assert-spmf-def*, *split-def*, *o-def*, and *if-distrib* take care of all the boilerplate code that makes all these transformations type-correct.

Note that the state of the oracle *oracle'* is changed between \mathcal{A}_1 and \mathcal{A}_2 . Namely, the random oracle’s part *s-h* may change when the chosen message is encrypted, but the state that records the adversary’s queries *s* is passed on unchanged.

3.5 Capturing a failure event

Suppose that two games behave the same except when a so-called failure event occurs. Then the chance of an adversary distinguishing the two games is bounded by the probability of the failure event. In other words, the simulation of the reduction is allowed to break if the failure event occurs. In the running example, such an argument is a key step to derive the bound on the adversary’s advantage. But to reason about failure events, we must first introduce them into the games we consider. This is because in CryptHOL, the probabilistic programs describe probability distributions over what they return (*return-spmf*). The variables that are used internally in the program are not accessible from the outside, i.e., there is no memory to which these are written. This has the advantage that we never have to worry about the names of the variables, e.g., to avoid clashes. The drawback is that we must explicitly introduce all the events that we are interested in.

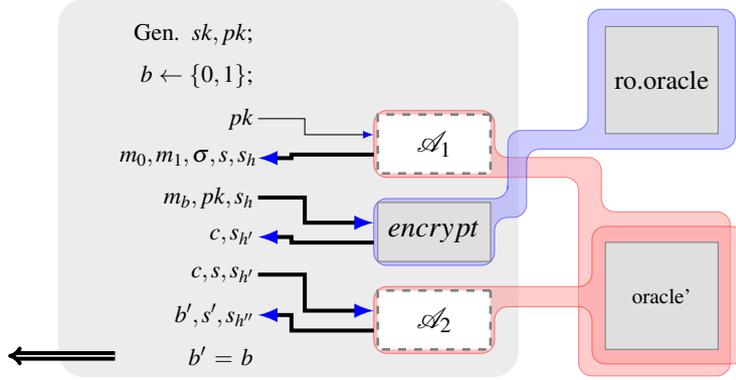


Figure 5: The IND-CPA game after flattening. The blue box around the encryption algorithm and the random oracle represents the expanded definition of them.

Introducing a failure event into a game is straightforward. So far, the games $game$ and $game_0$ simply denoted the probability distribution of whether the adversary has guessed right. For hashed Elgamal, the simulation breaks if the adversary queries the random oracle with the same query $\mathbf{g}[\cdot](x * y)$ that is used for encrypting the chosen message m_b . So we simply change the return type of the game to return whether the adversary guessed right *and* whether the failure event has occurred. The next definition $game_1$ does so. (Recall that $oracle'$ stores in its first state component s the queries by the adversary.) In preparation of the next reasoning step, we also split off the first two samplings, namely of x and y , and make them parameters of $game_1$.

```

define game1 :: nat ⇒ nat ⇒ (bool × bool) spmf
where game1 x y = do {
  b ← coin-spmf;
  ((m1, m2), σ), (s, s-h) ← exec-gpv oracle' (A1 (g [·] x)) ({}, ro.initial);
  TRY do {
    - :: unit ← assert-spmf (valid-plains m1 m2);
    (h, s-h') ← ro.oracle s-h (g [·] (x * y));
    let c = (g [·] y, h [⊕] (if b then m1 else m2));
    (b', (s', s-h'')) ← exec-gpv oracle' (A2 c σ) (s, s-h');
    return-spmf (b' = b, g [·] (x * y) ∈ s')
  } ELSE do {
    b ← coin-spmf;
    return-spmf (b, g [·] (x * y) ∈ s)
  }
} for x y

```

It is easy to prove that $game_0$ combined with the random oracle is a projection of $game_1$ with the sampling added:

```

let ?sample = λf :: nat ⇒ nat ⇒ - spmf. do {
  x ← sample-uniform (order ℒ);
  y ← sample-uniform (order ℒ);
  f x y }
have game0-game1:
  run-gpv ro.oracle game0 ro.initial = map-spmf fst (?sample game1)
by (simp add: game0 game1-def o-def split-def map-try-spmf map-scale-spmf)

```

3.6 Game hop based on a failure event

A game hop based on a failure event changes one game into another such that they behave identical unless the failure event occurs. The *fundamental-lemma* bounds the absolute difference between the two games by the probability of the failure event. In the running example, we would like to avoid querying the random oracle when encrypting the chosen message. The next game $game_2$ is identical except that the call to the random oracle $oracle$ is replaced with sampling a random bitstring.

```

define  $game_2 :: nat \Rightarrow nat \Rightarrow (bool \times bool) \text{ spmf}$ 
where  $game_2\ x\ y = do$  {
   $b \leftarrow coin\text{-}spmf$ ;
   $((m_1, m_2), \sigma), (s, s\text{-}h) \leftarrow exec\text{-}gpv\ oracle' (\mathcal{A}_1 (\mathbf{g} [\wedge] x)) (\{\}, ro.\text{initial})$ ;
  TRY  $do$  {
    -  $:: unit \leftarrow assert\text{-}spmf (valid\text{-}plains\ m_1\ m_2)$ ;
     $h \leftarrow spmf\text{-of}\text{-}set (nlists\ UNIV\ len\text{-}plain)$ ;
    — We do not query the random oracle for  $\mathbf{g} [\wedge] (x * y)$ , but instead sample a random bitstring
     $h$  directly. So the rest differs from  $game_1$  only if the adversary queries  $\mathbf{g} [\wedge] (x * y)$ .
     $let\ cipher = (\mathbf{g} [\wedge] y, h \oplus (if\ b\ then\ m_1\ else\ m_2))$ ;
     $(b', (s', s\text{-}h')) \leftarrow exec\text{-}gpv\ oracle' (\mathcal{A}_2\ cipher\ \sigma) (s, s\text{-}h)$ ;
     $return\text{-}spmf (b' = b, \mathbf{g} [\wedge] (x * y) \in s')$ 
  } ELSE  $do$  {
     $b \leftarrow coin\text{-}spmf$ ;
     $return\text{-}spmf (b, \mathbf{g} [\wedge] (x * y) \in s)$ 
  }
} for  $x\ y$ 

```

To apply the *fundamental-lemma*, we first have to prove that the two games are indeed the same except when the failure event occurs.

```

have  $rel\text{-}spmf (\lambda (win, bad) (win', bad'). bad = bad' \wedge (\neg bad' \longrightarrow win = win')) (game_2\ x\ y)$ 
 $(game_1\ x\ y)$  for  $x\ y$ 

```

proof –

This proof requires two invariants on the state of $oracle'$. First, $s = dom\ s\text{-}h$. Second, s only becomes larger. The next two statements capture the two invariants:

```

interpret  $inv\text{-}oracle' : callee\text{-}invariant\text{-}on\ oracle' (\lambda (s, s\text{-}h). s = dom\ s\text{-}h) \mathcal{I}\text{-full}$ 
by  $unfold\text{-}locales(auto\ simp\ add: oracle'\text{-}def\ split: option.\text{split}\text{-}asm\ if\text{-}split)$ 
interpret  $bad : callee\text{-}invariant\text{-}on\ oracle' (\lambda (s, -). z \in s) \mathcal{I}\text{-full}$  for  $z$ 
by  $unfold\text{-}locales(auto\ simp\ add: oracle'\text{-}def)$ 

```

First, we identify a bisimulation relation $?X$ between the different states of $oracle'$ for the second phase of the game. Namely, the invariant $s = dom\ s\text{-}h$ holds, the set of queries are the same, and the random oracle's state (a map from queries to responses) differs only at the point $\mathbf{g} [\wedge] (x * y)$.

```

let  $?X = \lambda (s, s\text{-}h) (s', s\text{-}h'). s = dom\ s\text{-}h \wedge s' = s \wedge s\text{-}h = s\text{-}h' (\mathbf{g} [\wedge] (x * y) := None)$ 

```

Then, we can prove that $?X$ really is a bisimulation for $oracle'$ except when the failure event occurs. The next statement expresses this.

```

let  $?bad = \lambda (s, s\text{-}h). \mathbf{g} [\wedge] (x * y) \in s$ 
let  $?R = (\lambda (a, s1') (b, s2'). ?bad\ s1' = ?bad\ s2' \wedge (\neg ?bad\ s2' \longrightarrow a = b \wedge ?X\ s1'\ s2'))$ 
have  $bisim : rel\text{-}spmf\ ?R (oracle'\ s1\ plain) (oracle'\ s2\ plain)$ 
if  $?X\ s1\ s2$  for  $s1\ s2$  plain using  $that$ 

```

by(*auto split: prod.splits intro!: rel-spmf-bind-refl simp add: oracle'-def rel-spmf-return-spmf2 fun-upd-twist split: option.split dest!: fun-upd-eqD*)

have *inv: callee-invariant oracle' ?bad*

— Once the failure event has happened, it will not be forgotten any more.

by(*unfold-locales*)(*auto simp add: oracle'-def split: option.split-asm*)

Now we are ready to prove that the two games $game_1$ and $game_2$ are sufficiently similar. The Isar proof now switches into an **apply** script that manipulates the goal state directly. This is sometimes convenient when it would be too cumbersome to spell out every intermediate goal state.

show *?thesis*

unfolding *game₁-def game₂-def*

— Peel off the first phase of the game using the structural decomposition rules *rel-spmf-bind-refl* and *rel-spmf-try-spmf*.

apply(*clarsimp intro!: rel-spmf-bind-refl simp del: bind-spmf-const*)

apply(*rule rel-spmf-try-spmf*)

subgoal *TRY for b m₁ m₂ σ s s-h*

apply(*rule rel-spmf-bind-refl*)

— Exploit that in the first phase of the game, the set s of queried strings and the map of the random oracle $s-h$ are updated in lock step, i.e., $s = \text{dom } s-h$.

apply(*drule inv-oracle'.exec-gpv-invariant; clarsimp*)

— Has the adversary queried the random oracle with $\mathbf{g} [\cdot] (x * y)$ during the first phase?

apply(*cases g [\cdot] (x * y) ∈ s*)

subgoal *True* — Then the failure event has already happened and there is nothing more to do. We just have to prove that the two games on both sides terminate with the same probability.

by(*auto intro!: rel-spmf-bindI1 rel-spmf-bindI2 lossless-exec-gpv[where $\mathcal{I} = \mathcal{I}$ -full] dest!: bad.exec-gpv-invariant*)

subgoal *False* — Then let's see whether the adversary queries $\mathbf{g} [\cdot] (x * y)$ in the second phase. Thanks to *ro.fresh*, the call to the random oracle simplifies to sampling a random bitstring.

apply(*clarsimp iff del: domIff simp add: domIff ro.fresh intro!: rel-spmf-bind-refl*)

apply(*rule rel-spmf-bindI[where $R = ?R$]*)

— The lemma *exec-gpv-oracle-bisim-bad-full* lifts the bisimulation for *oracle'* to the adversary \mathcal{A}_2 interacting with *oracle'*.

apply(*rule exec-gpv-oracle-bisim-bad-full[OF - - bisim inv inv]*)

apply(*auto simp add: fun-upd-idem*)

done

done

subgoal *ELSE by*(*rule rel-spmf-refl*) *clarsimp*

done

qed

Now we can add the sampling of x and y in front of $game_1$ and $game_2$, apply the *fundamental-lemma*.

hence *rel-spmf (λ(win, bad) (win', bad'). (bad ↔ bad') ∧ (¬ bad' → win ↔ win'))*
(*?sample game₂*) (*?sample game₁*)

by(*intro rel-spmf-bind-refl*)

hence *|measure (measure-spmf (?sample game₂)) {(win, -). win} - measure (measure-spmf (?sample game₁)) {(win, -). win}|*

≤ measure (measure-spmf (?sample game₂)) {(-, bad). bad}

unfolding *split-def by*(*rule fundamental-lemma*)

moreover

The *fundamental-lemma* is written in full generality for arbitrary events, i.e., sets of elementary events. But in this formalization, the events of interest (correct guess and

failure) are elementary events. We therefore transform the above statement to measure the probability of elementary events using *spmf*.

```

have measure (measure-spmf (?sample game2)) {(win, -). win} = spmf (map-spmfst (?sample
game2)) True
  and measure (measure-spmf (?sample game1)) {(win, -). win} = spmf (map-spmfst (?sample
game1)) True
  and measure (measure-spmf (?sample game2)) {(-, bad). bad} = spmf (map-spmf snd (?sample
game2)) True
  unfolding spmf-conv-measure-spmf measure-map-spmf by (auto simp add: vimage-def split-def)
ultimately have hop12:
  |spmf (map-spmfst (?sample game2)) True - spmf (map-spmfst (?sample game1)) True|
  ≤ spmf (map-spmf snd (?sample game2)) True
  by simp

```

3.7 Optimistic sampling: the one-time-pad

This step is based on the one-time-pad, which is an instance of optimistic sampling. If two runs of the two games in an optimistic sampling step would use the same random bits, then their results would be different. However, if the adversary's choices are independent of the random bits, we may relate runs that use different random bits, as in the end, only the probabilities have to match. The previous game hop from *game₁* to *game₂* made the oracle's responses in the second phase independent from the encrypted ciphertext. So we can now change the bits used for encrypting the chosen message and thereby make the ciphertext independent of the message.

To that end, we parametrize *game₂* by the part that does the optimistic sampling and call this parametrized version *game₃*.

```

define game3 :: (bool ⇒ bitstring ⇒ bitstring ⇒ bitstring spmf) ⇒ nat ⇒ nat ⇒ (bool × bool)
spmf
where game3 f x y = do {
  b ← coin-spmf;
  ((m1, m2), σ), (s, s-h) ← exec-gpv oracle' (A1 (g [^] x)) ({}, ro.initial);
  TRY do {
    - :: unit ← assert-spmf (valid-plains m1 m2);
    h' ← f b m1 m2;
    let cipher = (g [^] y, h');
    (b', (s', s-h')) ← exec-gpv oracle' (A2 cipher σ) (s, s-h);
    return-spmf (b' = b, g [^] (x * y) ∈ s')
  } ELSE do {
    b ← coin-spmf;
    return-spmf (b, g [^] (x * y) ∈ s)
  }
} for f x y

```

Clearly, if we plug in the appropriate function *?f*, then we get *game₂*:

```

let ?f = λ b m1 m2. map-spmf (λ h. (if b then m1 else m2) [⊕] h) (spmf-of-set (nlists UNIV
len-plain))
have game2-game3: game2 x y = game3 ?f x y for x y
  by (simp add: game2-def game3-def Let-def bind-map-spmf xor-list-commute o-def)

```

The *one-time-pad* lemma now allows us to remove the exclusive or with the chosen message, because the resulting distributions are the same. The proof is slightly non-trivial because the one-time-pad lemma holds only if the xor'ed bitstrings have the right

length. The congruence rules *try-spmf-cong* *bind-spmf-cong* [*OF refl*] *if-cong* [*OF refl*] extract this information from the program of the game.

```

let ?f' =  $\lambda b m_1 m_2$ . spmf-of-set (nlists UNIV len-plain)
have game3: game3 ?f x y = game3 ?f' x y for x y
by(auto intro!: try-spmf-cong bind-spmf-cong[OF refl] if-cong[OF refl]
  simp add: game3-def split-def one-time-pad valid-plains-def
  simp del: map-spmf-of-set-inj-on bind-spmf-const split: if-split)

```

The rest of the proof consists of simplifying *game3* ?f'. The steps are similar to what we have shown before, so we do not explain them in detail. The interested reader can look at them in the theory file from which this document was generated. At a high level, we see that there is no need to track the adversary's queries in *game2* or *game3* any more because this information is already stored in the random oracle's state. So we change the *oracle'* back into *oracle* using the Transfer package. With a bit of rewriting, the result is then the *game* for the adversary *elgamal-reduction* \mathcal{A} . Moreover, the guess *b'* of the adversary is independent of *b* in *game3* ?f, so the first boolean returned by *game3* ?f' is just a coin flip.

```

have game3-bad: map-spmf-snd (?sample (game3 ?f')) = lcdh.game (elgamal-reduction  $\mathcal{A}$ )
have game3-guess: map-spmf-fst (game3 ?f' x y) = coin-spmf for x y

```

3.8 Combining several game hops

Finally, we combine all the (in)equalities of the previous steps to obtain the desired bound using the lemmas for reasoning about reals from Isabelle's library.

```

have ind-cpa.advantage (ro.oracle, ro.initial)  $\mathcal{A}$  = |spmf (map-spmf-fst (?sample game1)) True
  - 1 / 2|
  using ind-cpa-game-eq-game0 by(simp add: game0-game1 o-def)
also have ... = |1 / 2 - spmf (map-spmf-fst (?sample game1)) True|
  by(simp add: abs-minus-commute)
also have 1 / 2 = spmf (map-spmf-fst (?sample game2)) True
  by(simp add: game2-game3 game3 o-def game3-guess spmf-of-set)
also have |... - spmf (map-spmf-fst (?sample game1)) True|  $\leq$  spmf (map-spmf-snd (?sample
  game2)) True
  by(rule hop12)
also have ... = lcdh.advantage (elgamal-reduction  $\mathcal{A}$ )
  by(simp add: game2-game3 game3 game3-bad lcdh.advantage-def o-def del: map-bind-spmf)
finally show ?thesis .

```

This completes the concrete proof and we can end the locale *hashed-elgamal*.

qed

end

4 Asymptotic security

An asymptotic security statement can be easily derived from a concrete security theorem. This is done in two steps: First, we have to introduce a security parameter η to everything. Only then can we state asymptotic security. The proof is easy given the concrete security theorem.

4.1 Introducing a security parameter

Since all our definitions were done in locales, it is easy to introduce a security parameter after the fact. To that end, we define copies of all locales where their parameters now take the security parameter as an additional argument. We illustrate it for the locale *ind-cpa-pk*.

The **sublocale** command brings all the definitions and theorems of the original *ind-cpa-pk* into the copy and adds the security parameter where necessary.

```

locale ind-cpa-pk' =
  fixes key-gen :: security  $\Rightarrow$  ('pubkey  $\times$  'privkey, 'query, 'response) gpv
  and encrypt :: security  $\Rightarrow$  'pubkey  $\Rightarrow$  'plain  $\Rightarrow$  ('cipher, 'query, 'response) gpv
  and decrypt :: security  $\Rightarrow$  'privkey  $\Rightarrow$  'cipher  $\Rightarrow$  ('plain, 'query, 'response) gpv
  and valid-plains :: security  $\Rightarrow$  'plain  $\Rightarrow$  'plain  $\Rightarrow$  bool
begin
sublocale ind-cpa-pk key-gen  $\eta$  encrypt  $\eta$  decrypt  $\eta$  valid-plains  $\eta$  for  $\eta$  .
end

```

We do so similarly for *list-cdh*, *random-oracle*, and *hashed-ElGamal*.

```

locale hashed-ElGamal' =
  lcdh: list-cdh'  $\mathcal{G}$  +
  ro: random-oracle' len-plain
  for  $\mathcal{G}$  :: security  $\Rightarrow$  'grp cyclic-group
  and len-plain :: security  $\Rightarrow$  nat
begin
sublocale hashed-ElGamal  $\mathcal{G}$   $\eta$  len-plain  $\eta$  for  $\eta$  ..

```

4.2 Asymptotic security statements

For asymptotic security statements, CryptHOL defines the predicate *negligible*. It states that the given real-valued function approaches 0 faster than the inverse of any polynomial. A concrete security statement translates into an asymptotic one as follows:

- All advantages in the bound become negligibility assumptions.
- All side conditions of the concrete security theorems remain assumptions, but wrapped into an *eventually* statement. This expresses that the side condition holds eventually, i.e., there is a security parameter from which on it holds.
- The conclusion is that the bounded advantage is *negligible*.

```

theorem asymptotic-security-ElGamal:
  assumes negligible ( $\lambda \eta$ . lcdh.advantage  $\eta$  (ElGamal-reduction  $\eta$  ( $\mathcal{A}$   $\eta$ )))
  and eventually ( $\lambda \eta$ . ind-cpa.lossless ( $\mathcal{A}$   $\eta$ )) at-top
  shows negligible ( $\lambda \eta$ . ind-cpa.advantage  $\eta$  (ro.oracle  $\eta$ , ro.initial) ( $\mathcal{A}$   $\eta$ ))

```

The proof is canonical, too: Using the lemmas about *negligible* and Eberl's library for asymptotic reasoning [4], we transform the asymptotic statement into a concrete one and then simply use the concrete security statement.

```

apply(rule negligible-mono[OF assms(1)])
apply(rule landau-o.big-mono)
apply(rule eventually-rev-mp[OF assms(2)])

```

```

apply(intro eventuallyI impl)
apply(simp del: ind-cpa.advantage.simps add: ind-cpa.advantage-nonneg lcdh.advantage-nonneg)
by(rule concrete-security-ElGamal)

end

```

References

- [1] C. Ballarín. Locales: A module system for mathematical theories. *Journal of Automated Reasoning*, 52(2):123–153, Feb 2014.
- [2] D. A. Basin, A. Lochbihler, and S. R. Sefidgar. CryptHOL: Game-based proofs in higher-order logic. *IACR Cryptology ePrint Archive*, 2017:753, 2017.
- [3] W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [4] M. Eberl. Landau symbols. *Archive of Formal Proofs*, Jul 2015. http://isa-afp.org/entries/Landau_Symbols.html, Formal proof development.
- [5] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
- [6] S. Goldwasser and S. Micali. Probabilistic encryption & how to play mental poker keeping secret all partial information. In *ACM Symposium on Theory of Computing (STOC 1982), Proceedings*, pages 365–377, 1982.
- [7] B. Huffman and O. Kunčar. Lifting and Transfer: A modular design for quotients in Isabelle/HOL. In *Certified Programs and Proofs (CPP 2013), Proceedings*, pages 131–146, 2013.
- [8] A. Krauss. *Automating Recursive Definitions and Termination Proofs in Higher-Order Logic*. Dissertation, Technische Universität München, 2009.
- [9] A. Lochbihler. Probabilistic functions and cryptographic oracles in higher order logic. In *European Symposium on Programming (ESOP 2016), Proceedings*, pages 503–531, 2016.
- [10] A. Lochbihler. CryptHOL. *Archive of Formal Proofs*, May 2017. <http://isa-afp.org/entries/CryptHOL.html>, Formal proof development.
- [11] A. Lochbihler, S. R. Sefidgar, and B. Bhatt. Game-based cryptography in hol. *Archive of Formal Proofs*, 2017. http://isa-afp.org/entries/Game-Based_Crypto.shtml, Formal proof development.
- [12] J. C. Mitchell. Representation independence and data abstraction. In *ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL 1986), Proceedings*, pages 263–276, 1986.
- [13] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer, 2002.

- [14] J. Schneider, M. Eberl, and A. Lochbihler. Monad normalisation. *Archive of Formal Proofs*, May 2017. http://isa-afp.org/entries/Monad_Normalisation.html, Formal proof development.
- [15] G. Shafi and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
- [16] V. Shoup. Sequences of games: A tool for taming complexity in security proofs. *IACR Cryptology ePrint Archive*, 2004:332, 2004.
- [17] P. Wadler and S. Blott. How to make ad-hoc polymorphism less ad hoc. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 1989), Proceedings*, pages 60–76, 1989.
- [18] M. Wenzel. Isar — a generic interpretative approach to readable formal proof documents. In *International Conference on Theorem Proving in Higher Order Logics (TPHOL 1999), Proceedings*, pages 167–183, 1999.