

Saarland University
Faculty of Natural Sciences and Technology I
Department of Computer Science

Master's Thesis

Verifiable Security of Prefix-free Merkle-Damgård

submitted by

Malte Horst Arthur Skoruppa

on August 3, 2012

Supervisor

Prof. Dr. Michael Backes

Advisor

M.Sc. Matthias Berg

Reviewers

Prof. Dr. Michael Backes

Dr. Matteo Maffei

Eidesstattliche Erklärung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel verwendet habe.

Statement in Lieu of an Oath

I hereby confirm that I have written this thesis on my own and that I have not used any other media or materials than the ones referred to in this thesis.

Einverständniserklärung

Ich bin damit einverstanden, dass meine (bestandene) Arbeit in beiden Versionen in die Bibliothek der Informatik aufgenommen und damit veröffentlicht wird.

Declaration of Consent

I agree to make both versions of my thesis (with a passing grade) accessible to the public by having them added to the library of the Computer Science Department.

Saarbrücken, August 3, 2012

Malte Skoruppa

Acknowledgments

This thesis forms part of a project to carry out a formal verification of the security of some modern cryptographic hash designs in view of the ongoing SHA-3 competition. Several people were involved in this project, and an associated paper was published at the 25th IEEE Computer Security Foundations Symposium, *CSF 2012* [8].

I would like to thank Prof. Dr. Michael Backes for integrating me into such a fascinating and cutting-edge project on formal verification, as well as for allowing me to travel to Madrid (to work on the project) and to Boston (where the conference took place). I also want to express my gratitude to my advisor Matthias Berg, whose counsel and advice has always been of great value, and who proofread several drafts of this thesis. Further I would like to sincerely thank all my co-authors of aforementioned paper. Beside the two people just mentioned, these are Prof. Dr. Gilles Barthe, who initiated the project with Michael and welcomed us at Madrid; Dr. Benjamin Grégoire, one of the main developers of **EasyCrypt** who also performed a big part of the indistinguishability proof described in our paper; Dr. César Kunz, who implemented support for loops in **EasyCrypt** without which our work would not have been possible; and of course Dr. Santiago Zanella Béguelin, another main developer of **EasyCrypt**, whose intelligent advice and thorough work on the paper has been invaluable to all of us. I also thank Dr. Matteo Maffei, who kindly agreed to be the second reviewer of my thesis.

Finally, I want to thank my fellow students and good friends who made my studies so very enjoyable. Last, but not least, I deeply thank my family, who has faithfully supported me throughout all these years.

Abstract

In recent years, an increasingly popular approach to the game-playing technique in cryptographic proofs is to rigorously specify games as pieces of well-defined program code; this irons out potential ambiguities in their specification, and enables a formal analysis of those games and proofs with the help of automated tools, as envisioned by Halevi [49]. Barthe et al. recently developed **EasyCrypt** [13], a tool which comes with a fully-specified programming language suitable for the formalization of cryptographic games, as well as an associated probabilistic relational Hoare logic and built-in program verification techniques. **EasyCrypt** can automatically generate proof obligations arising within a game-playing proof, and solve these mechanically using state-of-the-art automated tools.

In this thesis, we use **EasyCrypt** to formally verify the indistinguishability of the prefix-free Merkle-Damgård construction, following a seminal proof by Coron et al. [39]. Merkle-Damgård is a cryptographic construction ubiquitously used to implement hash functions: These have received considerable attention from the cryptographic community in the last few years, motivated by the ongoing SHA-3 competition. Indistinguishability is a powerful and non-trivial security notion that yields many implications, and certainly constitutes a desirable security property to achieve when designing a modern cryptographic hash function.

More concretely, we specify a sensible sequence of games in **EasyCrypt**'s language, and discuss the arguments that were needed for machine-checking the validity of the transitions relating those games. We focus in particular on the underlying axiomatization and derived lemmas used to justify the validity of side-conditions that arose when proving invariants of those games. Our results provide a first, but significant step towards a machine-checked verification of the indistinguishability of the finalists of the SHA-3 competition.

Contents

Notations and Conventions	1
1 Introduction	3
2 Related Work	9
3 Background	13
3.1 Hash Functions	13
3.2 Indifferentiability	15
3.3 The SHA-3 Competition	17
4 EasyCrypt	19
4.1 Organization of a Proof	19
4.2 The Language	21
4.3 The Logic	22
4.4 Reasoning about Probabilities	23
4.5 Selected Techniques	24
4.5.1 Random Oracles	24
4.5.2 Lazy and Eager Sampling	25
5 High-level proof	29
5.1 From G_{real} to $G_{\text{real}'}$	34
5.2 From $G_{\text{real}'}$ to G_{realRO}	35
5.3 From G_{realRO} to $G_{\text{idealEager}}$	42
5.4 From $G_{\text{idealEager}}$ to $G_{\text{idealLazy}}$	45
5.5 From $G_{\text{idealLazy}}$ to G_{ideal}	48

6	The manual part	51
6.1	Predicates, Definitions, and Axioms	52
6.1.1	Maps	52
6.1.2	Operators <code>pad</code> and <code>unpad</code>	53
6.1.3	Operator <code>mapfst</code>	54
6.1.4	Operator <code>findseq</code>	54
6.1.5	Predicate <code>set_bad3</code>	56
6.1.6	Predicate <code>Claim5</code>	57
6.2	Padding Lemmas	57
6.3	Lemmas on <code>mapfst</code>	58
6.4	Lemmas on <code>ischain</code>	59
6.5	Lemmas on <code>findseq</code>	69
6.6	Lemmas on <code>valid_chain</code>	77
6.7	Lemmas on <code>set_bad3</code>	80
6.8	Lemmas on <code>Claim5</code>	82
7	Conclusions	89
7.1	Applicability to SHA-3 Finalists	89
7.2	Future Work	91
	References	91

Notations and Conventions

NAMING CONVENTIONS.

x – Usually a block, i.e. a bitstring of length k (see Section 3.1)

y – Usually a chaining value, i.e. a bitstring of length n (see Section 3.1)

m – Usually a message, i.e. a bitstring of arbitrary length (see Section 3.1)

p – Usually a padding, i.e. a list of blocks (see Section 3.1)

IV – A fixed chaining value used as initialization vector in the Merkle-Damgård iteration (see Section 3.1)

PRHL JUDGMENTS.

$\vdash c_1 \sim c_2 : \Psi \implies \Phi$ – **pRHL** judgment correlating commands c_1 and c_2 under the pre-condition Ψ and the post-condition Φ (see Section 4.3)

$e\langle 1 \rangle$ – Within a pre-condition or a post-condition of a **pRHL** judgment, the interpretation of e in the left command of that judgment (see Section 4.3)

$e\langle 2 \rangle$ – Within a pre-condition or a post-condition of a **pRHL** judgment, the interpretation of e in the right command of that judgment (see Section 4.3)

$=\{x\}$ – Short-hand notation for $x\langle 1 \rangle = x\langle 2 \rangle$ for some variable x (see Section 4.3)

PROBABILITIES.

$\Pr [c, m : A]$ – Probability of event A after the execution of command c under initial memory m (see Section 4.4)

$\Pr[c : A]$ – Probability of event A after the execution of command c under an arbitrary initial memory (see Section 4.4)

MAPS.

$\text{upd}(T, xy, z)$ – The map resulting from the update of map T with the pair (xy, z) (see Section 6.1.1)

LISTS.

nil – The empty list

$x::xs$ – Construction of a list

$xs_\ell \parallel xs_r$ – Concatenation of lists

OPTIONS.

$\text{Some}(x)$ – The option associated to object x

None – Empty option

π_o – The projection of option o , i.e. if $o = \text{Some}(x)$ for some object x , then $\pi_o = x$

TYPES.

$\tau \text{ list}$ – The type of a list over elements of type τ

$\tau \text{ option}$ – The option type over elements of type τ

THE GAME-PLAYING TECHNIQUE is a salient technique to prove security properties of cryptographic constructions in modern cryptography, advocated by Bellare and Rogaway [23] and Shoup [68]. In a game-based proof, one typically formulates a desired security property of a cryptographic construction as a probabilistic experiment, called the *initial game*. In this game, a *challenger* uses or initializes the construction in question and communicates with an *adversary*, described by a black-box function. The goal of the adversary is to break some security property of the construction, denoted in the initial game by some event S . Then, the probability of event S is the probability of the adversary to break the security property. This probability is usually expressed in terms of a *security parameter*, and the goal is to show that the probability of S is *negligible* in terms of this security parameter. For instance, the probability for some adversary to determine which of two ciphertexts is the encryption of a chosen plaintext may be negligible in the length of the key.

To derive such a bound in the game-based setting, one defines a sequence of intermediate games and proves that the probability of S in each two consecutive games is equal, or differs only by a negligible amount. The last game usually has a form where it is easy to see the probability of S , or where such a probability is assumed or already known. Finally, by summing up over all the differences of the probabilities of S of any two consecutive games, one obtains an upper bound on the difference of the probabilities of event S in the initial and final games, yielding a bound for any adversary to break the desired security property. That is, since the adversary cannot distinguish with more than negligible probability between the initial and final games, and furthermore it is assumed or clear that the adversary cannot perform a successful attack on the final game with more than negligible probability, then it cannot do so for the initial game either.

Whenever two consecutive games do not have the same probability of some event S (typically the result of the game), but whose respective probabilities of event S differ only by a negligible amount, one introduces *failure events* capturing low-probability events that may lead to the games differing in event S . Then, one shows that the games have the same behavior with respect to event S unless one of these failure events occurs. Finally,

the Fundamental Lemma of Game-Playing [23, 68] is used to conclude that the difference of the probabilities of event S in the two games is upper-bounded by the probability of the failure events, and the probability of the failure events is concretely calculated.

A CRISIS OF RIGOR. However, in practice the formal arguments used to justify the validity of such transformations can, more often than not, leave many things to be desired. Typically, games may be outlined in natural language at worst, or specified in some kind of ad-hoc pseudocode at best, and their precise mathematical definitions and formal inner workings left unclear. The proofs are then performed in a rather semi-formal way, along with a lot of “handwaving”. Yet, even though some transformation may be easily understood at an intuitive level, it may in fact require rigorous mathematical arguments to formally justify its validity. In [23], Bellare and Rogaway speak of a “*crisis of rigor*” concerning cryptographic proofs in general, and endorse the usage of the game-playing technique along with a pseudocode-based descriptive language to formalize the games, but they still assume the operational semantics of the pseudocode to be implicitly clear. In [49], Halevi argues that the problem arises, for the most part, from the technical difficulties of fully formal proofs, especially since they consist in great parts of “mundane” arguments (e.g., checking that a transformation is truly valid), as opposed to the creative part of the proof, such as describing a simulator or a reduction. Indeed, cryptographers can hardly be expected to manually verify each and every one of these mundane parts down to their last detail; on the other hand, this would be necessary to obtain a fully verified proof that can be believed without reservation. In Halevi’s words,

“The problem is that as a community, we generate more proofs than we carefully verify (and as a consequence some of our published proofs are incorrect). I became acutely aware of this when I wrote my EME paper [48]. After spending a considerable effort trying to simplify the proof, I ended up with a 23-page proof of security for a — err — marginally useful mode-of- operation for block ciphers. Needless to say, I do not expect anyone in his right mind to even read the proof, let alone carefully verify it. (On the other hand I was compelled to write the proof, since it was the only way that I could convince myself that the mode was indeed secure.)”*

— Shai Halevi, [49]

Therefore, Halevi takes the idea of code-based game-playing proofs even further, and advocates the definition of a fully-specified programming language used to describe games; further, he advocates the creation of a tool tuned to this language that helps to automate the verification of the mundane parts of the proof, using formal verification methods.

A POSSIBLE SOLUTION: EASYCRYPT. In the years that followed, several tools addressing this situation saw the light of day, outlined in greater detail in Section 2. Recently, a new tool has been developed, called **EasyCrypt** [13], which defines an imperative programming language with random assignments, suitable for formalizing games in cryptographic

proofs, and implements an associated probabilistic relational Hoare logic and program verification techniques. In this way, **EasyCrypt** enables an automated mechanical verification of game transformations with the help of state-of-the-art SMT solvers and automated theorem provers interacting with **EasyCrypt**. The cryptographer needs only concentrate on the essence of a proof by specifying a sequence of games, written as small programs, along with Hoare judgments specifying how the games are related to each other, and logical statements expressing properties of the operators used in the games. Then, the reduction can be mechanically verified, significantly increasing the trust that can be placed in the proof. **EasyCrypt** also delivers tools to support the reasoning with failure events, in particular a mechanical implementation of the Fundamental Lemma of Game-Playing as well as a lemma to bound the probability of failure events (see Section 4.4). Therefore, **EasyCrypt** seems to be an excellent candidate for achieving the goals described by Halevi’s ambitious program.

HASH FUNCTIONS. Motivated by the ongoing *SHA-3 competition*, discussed in more detail in Section 3.3, hash functions and their security properties have received considerable attention from the cryptographic community in recent years. Most hash functions, including all SHA-3 finalists, implement the Merkle-Damgård construction (or a variant thereof), put forward by the pioneering works of Merkle [60] and Damgård [42] in 1990. The plain Merkle-Damgård construction hashes a message as follows. First, a message is split into blocks of the same length using an appropriate *padding function*. Then, a *compression function* is iterated over the blocks, generating an intermediate *chaining value* in each step. The compression function takes as arguments a block and a chaining value. The chaining value used for the compression of the first block is a fixed value called the *initialization vector*; the last chaining value output by the compression function is the final hash. Variants of Merkle-Damgård include the chopping of trailing bits of the last chaining value before outputting the hash, or imposing certain restrictions on the padding function. Several classic notions for the security of hash functions exist, such as *collision resistance*, *preimage resistance* and *second preimage resistance* (see discussion in Section 3.1). Merkle and Damgård showed independently that when used in conjunction with a *suffix-free* padding function, then the Merkle-Damgård iteration yields a collision-resistant hash function [42, 60] – a suffix-free padding function is a padding function with the property that no padding of a message is a suffix of the padding of another message. Both authors considered a concrete suffix-free padding function (in which a message is appended with as many 0’s as necessary to obtain a multiple of the block size, and an additional block encodes the length of the original message), although their proofs are applicable more generally for any suffix-free padding function, a fact which at least Merkle also observed [60].

A new security notion that arose relatively recently is the notion of *indifferentiability*, introduced by Maurer et al. in 2004 [59]. Shortly thereafter, Coron et al. applied this notion to hash functions [39]. Essentially, a hash function that is indifferentiable from a random oracle behaves like a variable input length random oracle [22], i.e. an oracle which takes as input a bitstring of variable length and that returns a value chosen uniformly at random from its range for each fresh query that it receives; multiple identical queries are answered

with the same response. The proof is performed in an idealized model: For instance, it may be assumed that the underlying compression function is a fixed input length random oracle, or even that an underlying building block of the compression function, such as a block cipher or a permutation, is an ideal primitive. Therefore, an indistinguishability proof gives confidence in the higher level design of the hash function: If an underlying building block is assumed to behave like an ideal primitive, then the hash function as a whole also behaves like an ideal primitive. Coron et al. furthermore showed that the plain Merkle-Damgård construction is indistinguishable from a random oracle when the employed padding function is *prefix-free*, and the compression function is assumed to be a fixed input length random oracle (a prefix-free padding function is defined analogously to a suffix-free one). Indistinguishability of a hash function is a powerful notion, and yields many other security properties including collision resistance, as well as preimage and second preimage resistance of the hash function [4], however the proof is not in the standard model, and may yield worse bounds than a direct proof of a specific security property. Still, indistinguishability is an interesting property to consider and is certainly a desirable security property when designing a state-of-the-art cryptographic hash function.

CONTRIBUTIONS. We set out to verify in **EasyCrypt** the security property of indistinguishability for the finalists of the SHA-3 competition. Our first step was to provide a mechanically verified proof of the indistinguishability of the prefix-free Merkle-Damgård construction, closely following the proof by Coron et al. [39]. The resulting paper was published at the 25th IEEE Computer Security Foundations Symposium (CSF 2012) [8]. My main involvement in this project laid in the abstraction, generalization and formal verification of the side conditions arising within the **EasyCrypt** proof, in the **Coq** proof assistant [70]. The resulting **EasyCrypt** proof consisted of over 3700 lines of code, and the accompanying **Coq** file had over 4200 lines of code. These are available at:

<http://easycrypt.gforge.inria.fr/csf12/>

The main contributions of this thesis are a deeper look into the indistinguishability proof performed in **EasyCrypt** than it was done in [8], detailed in Chapter 5, and the development of the aforementioned **Coq** file, whose proofs are described in Chapter 6.

OUTLINE. The layout of this thesis is as follows. We begin with a look at related work in Chapter 2. Then, we review the background of the thesis by discussing hash functions and their security properties, formally defining the Merkle-Damgård construction, as well as discussing the SHA-3 competition in Chapter 3. Next, we investigate the **EasyCrypt** tool in some more detail in Chapter 4. Our new results are outlined in Chapters 5 and 6. The high-level **EasyCrypt** proof is described in Chapter 5; it goes into deeper detail than the corresponding section in [8], and can be considered as an extended version thereof. Chapter 6 details the operators, predicates and axioms that the **EasyCrypt** proof is based on. Then, it discusses the side conditions that needed to be proven manually to enable **EasyCrypt** to verify the indistinguishability proof, and further provides formal proofs for all of those side conditions, based solely on the axiomatization presented in the same chapter.

Finally, Chapter 7 concludes with a discussion of the applicability of our results to the finalists of the SHA-3 competition, and future work.

Related Work

PURSUING a code-based approach to the game-playing technique to perform cryptographic proofs dates back to early work by Kilian and Rogaway in 1996 [53], and was applied in many subsequent papers by Rogaway, e.g. [21, 29, 30, 50, 51, 66]. Bellare and Rogaway argue in [23] that regarding games as code benefits the technique of game-based proofs since it is less error-prone and imposes some discipline on the process, and describe a sample programming language suitable for formalizing games. In [49], Halevi advocates a rigorous specification of such a programming language and furthermore the creation of an automated tool to support the verification of code-based game-playing proofs. Thus, cryptographers could concentrate on the creative parts of a proof, leaving machines to verify the more mundane parts (checking that some transformations are actually sound).

CryptoVerif was the first tool aiming at the automatic verification of game-based proofs in cryptographic schemes. In [33], Blanchet and Pointcheval used it to prove the unforgeability of the Full-Domain Hash signature scheme under the trapdoor-one-wayness of some permutations. Additionally, it is also suitable to prove security of cryptographic protocols: It has been applied to prove authentication and key secrecy properties of Kerberos [32]. Games in **CryptoVerif** are specified in a probabilistic process calculus, inspired by the pi-calculus [31], and transformations are captured by probabilistic bisimulation relations. Its focus is on automation: According to Blanchet, security assumptions on cryptographic primitives are applied in a fairly direct way and the tool has limitations as it comes to deeper mathematical reasoning. Therefore it is mainly suited for the analysis of security protocols using high-level primitives (e.g. encryptions or signatures) [31].

Several authors set out to develop frameworks to formalize game-based cryptographic proofs in proof assistants:

- In [62, 63], Nowak reports on a framework to develop game-based cryptographic proofs in **Coq** using a shallow embedding, in which games are implemented directly as **Coq** functions. He uses it to give security proofs for the semantic security of ElGamal [62], and for the security of two other cryptographic primitives based on number theory [63]. However, his framework lacks a formal theory for syntactic program transformations, which considerably hampers the possibility for proof automation in

game transitions, and it cannot deal with complexity notions such as polynomial-time programs; Nowak and Zhang address these problems in [64] by using a typed probabilistic lambda-calculus, but do not implement the latter in a tool.

- In [1], Affeldt et al. provide a formal proof of the PRP/PRF switching lemma in the proof assistant `Coq`. In contrast to Nowak’s framework, their framework uses a deep embedding; they define a simple programming language following the game-playing framework by Bellare and Rogaway [23], and a specialized version of the Fundamental Lemma. However, their framework is tuned to their particular example, and lacks a general applicability.
- In [9], Backes et al. developed a framework for game-based proofs in the proof assistant `Isabelle/HOL` [61]. It formalizes a simply-typed probabilistic higher-order functional language with references and support for events fit to express all constructs used in cryptographic games, including random oracles, polynomial-time programs, and sampling from possibly continuous probability measures. It comes along with a theory for observational equivalence and computational indistinguishability, and a definition of the Fundamental Lemma. Since the language is higher-order, oracles can be dealt with in a very natural fashion, as they can simply be passed as arguments to procedures. However, its high level of complexity is costly, and no working examples have been completed. In [44, 69], Driedger and Skoruppa contemplate in their respective Bachelor’s theses pen-and-paper proofs of syntactic program transformations in that framework, resulting in example proofs concerning one-way permutations and the semantic security of ElGamal, yet these are not implemented in `Isabelle/HOL`.
- In [15], Barthe et al. presented `CertiCrypt`, a fully verified framework built on top of the `Coq` proof assistant that enables specifying and machine-checking game-based cryptographic proofs. They formalize an imperative programming language with random assignments, structured datatypes and procedure calls, and provide a set of certified tools to reason about equivalence of probabilistic programs, including a probabilistic relational Hoare logic, a theory of observational equivalence, verified program transformations and cryptographic techniques such as reasoning about failure events and lazy/eager sampling. A fair amount of proofs has been carried out in `CertiCrypt`, such as: the PRP/PRF switching lemma [15, 16]; existential unforgeability under adaptive chosen-message attacks of the FDH signature scheme [15, 76]; semantic security of ElGamal [12, 15] and of OAEP [15]; adaptive IND-CCA security of OAEP [14]. The authors also formalized a class of zero-knowledge proofs, known as Σ -protocols [41], and formally proved the security of many instances thereof from the literature [17]. Furthermore, they verified the semantic security of the Boneh-Franklin Identity-based encryption scheme [34] in [19]. More recently, they machine-checked a proof of the indistinguishability of a construction for hashing into elliptic curves by Brier et al. [35] from a random oracle in [11], and developed an extension of `CertiCrypt`, called `CertiPriv`, to reason about differential privacy [18]. The only noteworthy disadvantage of `CertiCrypt` is that the fully formal carrying out of proofs requires a high level of expertise and can be fairly time consuming, which is why the authors came up with `EasyCrypt` [13] to help in the automation of this task.

Courant et al. investigate a different approach for automatically proving several security notions such as IND-CPA and adaptive IND-CCA security in [40]: They define a Hoare logic for a simple programming language, and use it to implement an automated verification procedure for asymmetric encryption schemes in the random oracle model. Their focus is on full automation; while their logic is sound, it is not complete (i.e. it may fail to prove a secure encryption scheme), however they do achieve security proofs for several encryption schemes in the literature. Later, Barthe et al. design a general logic which captures typical reasoning patterns common in provable security [10], such as reductions, simulations, random oracles, adaptive adversaries, failure events, and lazy/eager sampling. The logic is called CIL (Computational Indistinguishability Logic) and subsumes the former logic. The authors applied it to prove the security of probabilistic signature schemes, the FDH signature scheme, and adaptive IND-CCA security of OAEP [10]. The rules of the logic have been formalized in Coq by Corbineau et al. [38]. Shortly ago, Daubignard et al. extended this logic such as to capture domain extenders, and presented a generic reduction theorem to prove the indistinguishability of several hash function constructions [43], applicable to four of the five SHA-3 finalists – a work closely related to the one presented in this thesis. However, their method has not yet been implemented.

In stark contrast to all the above work, Kobitz casts a radically different view at the concept of automated theorem-proving in reductionist security proofs [54, 55] and speaks of a “metaphorical race between human and machine”. He criticizes inherent weaknesses of this concept, e.g. that mechanically verified proofs are not, or only in a limited way, understandable and replicable by humans – citing Manin’s “a good proof is one which makes us wiser” [57] – and submits that computer-assisted theorem-proving has not yet reached a wide audience amongst mathematicians. However, he duly admits to an inchoate knowledge in the field of automated theorem-proving [55], having never constructed a game-based proof himself, and many of his claims are easily refuted. While he does raise a few valid concerns, he also seems to miss the point that machine-checked proofs intend to complete, rather than to replace, human-made proofs; and to increase, rather than to give in the first place, confidence in the correctness of these proofs.

Concerning indistinguishability, the notion was first introduced by Maurer et al. [59] and applied to hash functions by Coron et al. [39], who in the same paper performed a proof of the indistinguishability of the prefix-free Merkle-Damgård construction (as well as other variants of Merkle-Damgård), which this thesis heavily relies on. The ongoing SHA-3 competition has motivated further research on the security of hash functions, including indistinguishability for all the SHA-3 finalists [2, 3, 20, 24, 26, 37]. We discuss these in more detail in later sections.

The notion of indistinguishability is not uncontroversial [65], nor is the random oracle methodology [36]. Nonetheless, indistinguishability from a random oracle represents an increasingly accepted notion for hash functions in the cryptographic community, and remains a good heuristic to assess their security, since an indistinguishability proof demonstrates that a hash function’s high-level design preserves an ideal functionality, and implies bounds on all the classical security properties of hash functions [4].

THE purpose of this chapter is to recapitulate the basic notions that this work is based on. First, we give some background on hash functions: How they are usually built, and classical security properties of cryptographic hash functions. Then, we introduce the more recent security notion of indistinguishability (and some related results), and finally we discuss the ongoing SHA-3 competition.

3.1 Hash Functions

Generally speaking, hash functions are functions that take inputs of variable length (e.g., large messages) and map them onto small bitstrings of a fixed length. We call such an output a *hash value*, or simply a *hash* for short.

Cryptographic hash functions are hash functions that aim to achieve certain security properties. In practice, they are designed such that their output is seemingly random and unique for each input message, even though in reality it clearly is not, since the domain of a hash function is necessarily much larger than its range. For instance, it should be hard for an adversary to generate a message that has a given hash, or generate a message that has the same hash as another message, or find two messages that have the same hash. Such functions are frequently used in many cryptographic constructions, in particular for (but not limited to) providing authenticity, such as in digital signature schemes or message authentication codes. The output of a cryptographic hash function is also called the *digest* of a message; depending on its usage, one may also speak e.g. of a *fingerprint* or a *checksum*.

At the heart of most cryptographic hash functions realized today lies the so-called Merkle-Damgård construction (or a variant thereof), which was proposed independently at the same time by Ralph Merkle [60] and Ivan Damgård [42]. It works by first splitting a message into *blocks* (bitstrings of a fixed length), using an appropriate *padding function*. Then, a *compression function* is iterated on those blocks. The compression function takes two arguments: a block and a *chaining value* (also a bitstring of a fixed length), and outputs a new chaining value to be used in the next iteration. The final chaining value is

the digest of the message. The chaining value to be used in the first iteration is a fixed value called the *initialization vector*. Throughout this thesis, we use the following conventions:

1. A *block* is a bitstring of length k , and is usually denoted by x ;
2. a *chaining value* is a bitstring of length n , and is usually denoted by y ;
3. a *message* is a bitstring of arbitrary length, and is usually denoted by m ;
4. a *padding* is a list of blocks, and is usually denoted by p ; and
5. the *initialization vector* is an n -bit constant \mathbf{IV} .

Furthermore, a padding function \mathbf{pad} is a function of type $\{0, 1\}^* \rightarrow \{0, 1\}^k$ list, that maps a message to a padding. A compression function f is a function of type $\{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, that compresses a pair of a block and a chaining value to a new chaining value. Now, we formally define the Merkle-Damgård construction.

Definition 3.1 (Merkle-Damgård). Let \mathbf{pad} be a padding function, f be a compression function, and \mathbf{IV} be a (public) initialization vector. The Merkle-Damgård hash function $\mathbf{MD}_{\mathbf{pad},f}$ is defined as follows:

$$\begin{aligned} \mathbf{MD}_{\mathbf{pad},f} & : \{0, 1\}^* \rightarrow \{0, 1\}^n \\ \mathbf{MD}_{\mathbf{pad},f}(m) & := f^*(\mathbf{pad}(m), \mathbf{IV}) \end{aligned}$$

where $f^* : \{0, 1\}^k$ list $\times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is recursively defined by the equations

$$\begin{aligned} f^*(\mathbf{nil}, y) & := y \\ f^*(x::xs, y) & := f^*(xs, f(x, y)) \end{aligned}$$

Later on, we will often use the notion of an *MD-chain* to ease explanations. An MD-chain is a list $[(x_1, \mathbf{IV}), (x_2, y_2), \dots, (x_j, y_j)]$ of pairs of blocks and chaining values, where the blocks are associated to the padding of a given message, and the chaining values are the intermediate values generated by a given compression function as used in the Merkle-Damgård iteration. We give the formal definition below.

Definition 3.2 (MD-chain). Let \mathbf{pad} be a padding function, f be a compression function, and \mathbf{IV} be a (public) initialization vector. Then an MD-chain is a list $[(x_1, y_1), \dots, (x_j, y_j)]$ such that

1. $y_1 = \mathbf{IV}$,
2. $\forall i \in \{1, \dots, j-1\}. f(x_i, y_i) = y_{i+1}$, and
3. $[x_1, \dots, x_j]$ is in the range of \mathbf{pad} .

Remark. In the above definition we obtain $f(x_j, y_j) = \mathbf{MD}_{\mathbf{pad},f}(\mathbf{pad}^{-1}([x_1, \dots, x_j]))$.

Many variants of the Merkle-Damgård construction exist. For instance, the HAIFA construction [27] uses a compression function of the form $f(x, y, b, s)$, where the additional parameters b and s denote a counter and a salt, respectively. The counter ensures that each block is processed with a unique variant of its compression function, while the usage of a salt value means that the HAIFA construction defines a *family* of hash functions, rather than a single hash function. A popular variant of Merkle-Damgård is the so-called

chop-MD construction, which chops some trailing bits off the final hash; very similarly, the wide-pipe hash construction [56] prescribes to use larger internal chaining values than the final output, using a different compression function for the last iteration. Using a different compression function in the last iteration is also a widely used paradigm in its own right, and usually referred to as a *final transformation* (note that the chopping of trailing bits is typically not considered part of a final transformation).

All of these variants can be understood as heuristics aimed at improving the security properties of the plain Merkle-Damgård construction. Prominent classical security properties that cryptographic hash functions aim to achieve include the following.

Collision resistance. It is hard to find two distinct messages m_1, m_2 such that $H(m_1) = H(m_2)$.

Preimage resistance. For a given digest h , it is hard to find a message m such that $H(m) = h$.

Second preimage resistance. For a given message m_1 , it is hard to find a message $m_2 \neq m_1$ such that $H(m_1) = H(m_2)$.

Resistance to length extension attacks. Given a digest $H(m_1)$, it is hard to compute $H(m_1 \parallel m_2)$ for any non-empty message m_2 .

An extensive survey of these (and related) security properties and their relationships can be found in [67]; an overview of several iteration modes and their property-preserving traits appears in [6].

We point out here the following basic result, a special case of which was already described by both Merkle and Damgård in [42, 60]: if a padding function pad is suffix-free (see below), then the plain Merkle-Damgård-construction $\text{MD}_{\text{pad},f}$ is at least as collision-resistant as the compression function f . More precisely, they show that if a collision can be found for $\text{MD}_{\text{pad},f}$, then a collision can also be found for f .

Lastly, we formally define the notion of a suffix-free (resp. prefix-free) padding function.

Definition 3.3 (Prefix- and suffix-free padding). A padding function pad is *prefix-free* (resp. *suffix-free*) iff for any distinct messages m_1, m_2 and any padding p , it holds that $\text{pad}(m_1) \neq \text{pad}(m_2) \parallel p$ (resp. $\text{pad}(m_1) \neq p \parallel \text{pad}(m_2)$).

3.2 Indifferentiability

The notion of indifferentiability of two systems was introduced by Maurer et al. in [59], and constitutes a generalization of the well-established notion of indistinguishability. Similarly as for indistinguishability, it yields an elegant composition theorem; the indifferentiability of two components \mathcal{H} and \mathcal{F} implies that the security of any cryptosystem $\mathcal{C}(\mathcal{F})$ based on the component \mathcal{F} is preserved when replacing the component \mathcal{F} by the component \mathcal{H} in the cryptosystem $\mathcal{C}(\cdot)$, i.e. $\mathcal{C}(\mathcal{H})$ is as secure as $\mathcal{C}(\mathcal{F})$. The difference is that indifferentiability

may be applied to more general settings, namely where a distinguisher has access to some internal functions or variables of the considered components.

In [39], Coron et al. applied this notion to hash functions, considering in particular the Merkle-Damgård construction. Broadly speaking, if $\mathcal{H}^{\mathcal{G}}$ is a hash function based on an internal resource \mathcal{G} , and \mathcal{F} is a variable input length random oracle, then we say that $\mathcal{H}^{\mathcal{G}}$ is indistinguishable from \mathcal{F} (i.e., an ideal hash function) if there exists a simulator \mathcal{S} such that no adversary can distinguish the scenario where it is given access to the procedures $\mathcal{H}^{\mathcal{G}}$ and \mathcal{G} , and the scenario where it is given access to the procedures \mathcal{F} and $\mathcal{S}^{\mathcal{F}}$, with more than negligible probability (see Figure 3.1). The crucial difference to indistinguishability is that the adversary is given access to the internal resource \mathcal{G} that the hash function is based on; this makes it possible to perform a proof of the statement in the ideal model, where \mathcal{G} is assumed to be an ideal primitive, by constructing an appropriate simulator.

Definition 3.4 (Indistinguishability). A procedure \mathcal{H} with oracle access to an ideal primitive \mathcal{G} is $(t_{\mathcal{S}}, q, \epsilon)$ -indistinguishable from \mathcal{F} if there exists a simulator \mathcal{S} with oracle access to \mathcal{F} and executing within time $t_{\mathcal{S}}$, such that for any distinguisher \mathcal{D} that makes at most q oracle queries, the following inequality holds:

$$|\Pr [b \leftarrow \mathcal{D}^{\mathcal{H}, \mathcal{G}}() : b] - \Pr [b \leftarrow \mathcal{D}^{\mathcal{F}, \mathcal{S}}() : b]| \leq \epsilon$$




Figure 3.1: Indistinguishability of \mathcal{H} from an ideal functionality \mathcal{F}

In the remainder of this thesis, the procedure \mathcal{H} that we consider is the Merkle-Damgård construction, and the component \mathcal{G} is the underlying idealized compression function, implemented as a fixed input length random oracle. The procedure \mathcal{F} represents an ideal hash function, implemented as a variable input length random oracle. Therefore, the role of the simulator \mathcal{S} is to simulate the compression function, i.e. it should behave towards the ideal hash function \mathcal{F} similarly as the ideal compression function \mathcal{G} behaves towards the Merkle-Damgård construction \mathcal{H} .

Coron et al. showed in [39] that when used with a prefix-free padding function, the plain Merkle-Damgård construction is indistinguishable from a random oracle. In this thesis, we focus on the formal verification of this statement using `EasyCrypt` and `Coq`. Coron et al. also showed the indistinguishability of similar constructions, such as chop-MD (see previous section), for a non-trivial number of output bits chopped off, which however are of no concern in the remainder.

It has been shown in [65] that the indistinguishability composition theorem outlined above does not apply unconditionally, since it does not capture cryptosystems that involve multiple, disjoint adversarial stages. Nevertheless, indistinguishability from a random oracle remains a very strong property for a hash function. Indeed, it shows that an ideal functionality is preserved, rather than a specific property such as collision-resistance. In fact, it implies bounds on the properties of collision resistance, preimage resistance, second preimage resistance and resistance to length extension attacks (see [4]). However, the proof is performed not in the standard model, but in an ideal model (since an underlying primitive is assumed to be ideal), and may deliver looser bounds than a direct proof. The proof presented in this thesis assumes that the compression function used in the Merkle-Damgård construction is ideal; generally, indistinguishability proofs for similar hash function constructions may go deeper than that by assuming that some underlying building block of the compression function is ideal, but not necessarily the compression function itself. Note that the assumption that some primitive is ideal is usually not realistic; however, an indistinguishability proof ensures that the higher level design of the hash function has no structural weaknesses, and therefore gives confidence in this design. In other words, if a vulnerability is found, then the problem should be looked for in the underlying primitive that was considered ideal in the indistinguishability proof. The notion of indistinguishability continues to grow in its acceptance by the cryptographic community, and may very well be considered a must-have for any cryptographic hash function designed in the future.

3.3 The SHA-3 Competition

The U.S. National Institute for Standards and Technology (NIST) specifies seven – supposedly – secure hash algorithms in the Federal Information Processing Standard (FIPS) 180-4, one of them called SHA-1, and six algorithms that belong to the SHA-2 family of hash functions, which differ significantly from SHA-1 (SHA stands for Secure Hash Algorithm). In the light of theoretical attacks found against SHA-1 in 2005 [71, 72], NIST announced that they planned to phase out SHA-1 in favor of the SHA-2 family of hash functions, and later on, in November 2007, announced a public competition to develop a new secure hash algorithm, the winner of which shall be adopted in the FIPS standard and named SHA-3. This competition is commonly known as the *SHA-3 competition*. As of today, there are still no practical collisions found for the full SHA-1, but in 2011 theoretical attacks with an estimated complexity of 2^{51} hash function calls have been published [58].

The cryptographic community answered enthusiastically to NIST's call for a new SHA algorithm and NIST received 64 submissions by the end of October 2008. 51 of those algorithms were accepted for the first round of the competition. In July 2009, 14 were selected for the second round. After further analysis and discussion supported by the cryptographic community, NIST narrowed down the list of candidates to 5 finalists in December 2010.

The five SHA-3 finalists are called BLAKE [7], Grøstl [47], JH [74], Keccak [25], and Skein [45]. All finalists can be considered as variants of the Merkle-Damgård construction:

BLAKE is a HAIFA [27] construction (see Section 3.1). JH, Keccak, and Skein can be seen as a chop-MD construction (or a slight variant thereof in the case of Skein), and Grøstl can be seen as a chop-MD construction with a final transformation before chopping. The compression functions of all finalists are built either from a block cipher (BLAKE, Skein) or from a permutation (Grøstl, JH, Keccak).

The original NIST requirements for the candidates of the SHA-3 competition included the classical security properties of collision resistance, preimage resistance and second preimage resistance. All of the finalists satisfy these properties, and in most cases, they even achieve optimal bounds for them when the underlying block ciphers or permutations are assumed to be ideal [5]. The design and security of all candidates has been surveyed and discussed in depth in [4, 5] and we refer the reader there for more details on this subject.

Although the original requirements did not include the property of indistinguishability from a random oracle, this notion has also been considered in the literature and is achieved by all five finalists [2, 3, 20, 24, 26, 37]. Indeed, as discussed in the previous section, it seems an important property to consider. The proofs of indistinguishability of all five finalists are quite involved and complex. Our original motivation and main aim is the formal verification of those proofs. For several reasons, our proof of the indistinguishability of the prefix-free Merkle-Damgård construction is not immediately applicable to any of the SHA-3 finalists. In particular, only two of them use a prefix-free padding function, and the compression function cannot be assumed ideal for any of the finalists. We discuss this, and how our proof would have to be adapted to fit the SHA-3 finalists in more detail in the conclusion. Nevertheless, our mechanically verified proof of prefix-free Merkle-Damgård provides an excellent starting point to tackle the proofs of those finalists.

EASYCRYPT is a tool that aims to support cryptographers in performing fully formalized and mechanically verified game-based proofs. It was quite recently presented at *CRYPTO 2011*, where it won the Best Paper Award [13]. Using **EasyCrypt**, it is possible to specify games and the relations used to justify transitions between them in a strictly formal way; more importantly, **EasyCrypt** helps conducting game-based proofs by automatically generating the necessary verification conditions to derive the validity of desired claims – and solving these with the help of automated tools. It has been developed by the same people who previously published the tool **CertiCrypt** [15], and builds upon the same programming language **pWHILE** and its related probabilistic relational Hoare logic, called **pRHL**. The most notable difference to **CertiCrypt** is that **EasyCrypt** provides automation to solve proof obligations, i.e. it actively helps the user to verify the “mundane” parts of a proof, so that the user can focus on the more interesting creative parts. Hence, **EasyCrypt** makes a tremendous step towards the realization of Halevi’s program (see Introduction). In this chapter, we condense the most important information from [8] and [13] concerning **EasyCrypt** relevant for a better understanding of the thesis, including a description of the language **pWHILE**, the logic **pRHL**, and how to use **pRHL** judgments to derive probability claims. We also go into two specific techniques that will be important later on, namely how to model random oracles and how to deal with lazy/eager sampling. More details can be found in [8, 13]; many lemmas and techniques that apply to the **EasyCrypt** framework are also discussed in previous publications on **CertiCrypt**, such as [12, 15, 16, 75].

4.1 Organization of a Proof

The basic idea of an **EasyCrypt** proof is that the user needs only to provide a *proof sketch* representing the essence of a security proof. In a typical **EasyCrypt** proof, a user introduces the types and operators manipulated in a construction, states associated axioms and lemmas, provides a sequence of games, and writes **pRHL** judgments and probability bounds that relate these games. When parsing such a proof sketch, **EasyCrypt** automati-

cally generates verification conditions to derive the validity of the **pRHL** judgments, and solves these with the help of automated theorem provers and SMT solvers. In a final step, these **pRHL** judgments can be used to derive claims about the probability of certain events in the games.

More in detail, a proof in **EasyCrypt** is composed of 5 parts:

- 1. Declaration of types, constants, operators, and predicates.** Typically, one first declares the basic types and constants used in a scheme. Additionally to the built-in types, a user may declare arbitrary other types (e.g. group elements), or type aliases (e.g. variable-length bitstrings as lists of Booleans). Operators may also be declared or defined and can be used within the definitions of the games. Furthermore, predicates may be declared, providing an elegant way to describe invariants in a compact fashion.
- 2. Axioms and lemmas.** The second step is to specify axioms and lemmas that describe properties of objects manipulated by the scheme. These are used by the SMT solvers and automated theorem provers to check the validity of arising proof obligations. They are written as first-order formulae involving expressions over variables, constants, operators and predicates (declared previously). For instance, one may specify as an axiom that some constant is positive, or declare that an operator which decides a predicate is sound and/or complete. The difference between an axiom and a lemma in **EasyCrypt** is that the validity of a lemma is derived automatically, using either built-in axioms (e.g., properties of lists), or user-defined axioms declared previously. By contrast, the validity of an axiom is simply assumed. Thus, in order to believe the validity of an **EasyCrypt** proof, one only has to believe the axioms stated in this proof (and the soundness of the tool itself, obviously). In practice, when developing complex proofs, some claims intended as lemmas may be non-trivial, and lie out of the scope of the supported automated tools. Different automated tools may be tried, as **EasyCrypt** formats proof obligations in the intermediate format of the **Why** tool [46]; when all of these automated tools fail, it is still possible to export the corresponding statements to the **Coq** proof assistant and prove them manually. In Chapter 6, we detail the axioms and lemmas used in our proof of indifferenciability, and give human-readable proofs for all of the statements that needed to be derived from our underlying axiomatization.
- 3. Definition of a game sequence.** A game is specified as a collection of procedures, global variables manipulated by these procedures, and adversary specifications. Procedures are defined in the language **pWHILE**, defined in Section 4.2. Adversaries are specified as abstract procedures with oracle access; these oracles must be instantiated by other procedures defined within a game. An abstract procedure representing an adversary in a game may only call those procedures which it is given access to. Additionally, a game typically (though not necessarily) implements a main procedure describing the overall game.

- 4. pRHL judgments.** The next step is to provide logical judgments in the probabilistic relational Hoare logic implemented by **EasyCrypt**. These judgments establish logical equivalences between the games; for instance, the return value of a procedure P_1 in a game G_1 may always be equal to the return value of a procedure P_2 in a game G_2 . The form of a pRHL judgment in **EasyCrypt** is $\vdash G_1.P_1 \sim G_2.P_2 : \Psi \implies \Phi$, where G_1, G_2 are games, P_1, P_2 are procedures, and Ψ, Φ are relational first-order formulae representing pre-conditions (resp. post-conditions) that must hold before (resp. after) the execution of both P_1 and P_2 . This logic is briefly outlined in Section 4.3.
- 5. Probability claims.** Finally, one writes probability claims relating the probabilities of events in one or more games using arithmetic operators and mathematical relations. For instance, one may want to derive that $\Pr[G_1 : b] = \Pr[G_2 : b]$, i.e. the probability that the games G_1 and G_2 output the same bit b is equal. The validity of such claims can be automatically derived by **EasyCrypt**, using either previously verified pRHL judgments, or by applying some hard-coded rules to compute bounds on the probability of elementary events. The final claim is typically a bound on the probability of an adversary performing a successful attack on the cryptosystem in consideration. We investigate reasoning about probabilities in more detail in Section 4.4.

4.2 The Language

Games in **EasyCrypt** are defined in the language **pWHILE**, a strongly-typed, probabilistic, procedural, imperative programming language. The language **pWHILE** includes types ranging over Booleans, integers, fixed-length bitstrings, lists, finite maps, as well as product, option and sum types. The grammar of the language is defined inductively by the following commands:

$\mathcal{C} ::=$	skip	nop
	$\mathcal{V} \leftarrow \mathcal{E}$	deterministic assignment
	$\mathcal{V} \xleftarrow{\$} \mathcal{DE}$	probabilistic assignment
	if \mathcal{E} then \mathcal{C} else \mathcal{C}	conditional
	while \mathcal{E} do \mathcal{C}	loop
	$\mathcal{V} \leftarrow \mathcal{P}(\mathcal{E}, \dots, \mathcal{E})$	procedure call
	$\mathcal{C}; \mathcal{C}$	sequence

where \mathcal{V} is a set of variables, \mathcal{P} is a set of procedures, \mathcal{E} is a set of expressions, and \mathcal{DE} is a set of distribution expressions. Expressions are built from operators and predicates as usual. Distribution expressions represent distributions from which values can be randomly sampled. In this thesis, we consider only the distribution expression $\{0, 1\}^n$, which denotes the uniform distribution on bitstrings of length n . That is, the command $y \xleftarrow{\$} \{0, 1\}^n$ samples uniformly at random a bitstring of length n and assigns it to the variable y .

A *program memory* maps global and local variables to values. We denote by \mathcal{M} the set of program memories; then, the semantics of a command c is defined as a function $\llbracket c \rrbracket : \mathcal{M} \rightarrow \mathcal{D}(\mathcal{M})$ from program memories to sub-distributions on memories (if a command

does not always terminate, it generates a sub-distribution with total probability less than 1). The precise definition of the semantics is given in [15], but is not relevant in the remainder of the thesis.

4.3 The Logic

At the heart of *EasyCrypt* lies a probabilistic relational Hoare logic, called **pRHL** for short. Using this logic, one can express that a relation Φ over program memories must hold after the execution of two commands c_1 and c_2 , under the assumption that a relation Ψ over program memories held before their execution; the general form a **pRHL** judgment is a quadruple of the form

$$\vdash c_1 \sim c_2 : \Psi \Longrightarrow \Phi,$$

where c_1, c_2 are commands, and Ψ and Φ are first-order formulae built with the usual connectives:

$$\Psi, \Phi ::= e \mid \neg\Phi \mid \Psi \wedge \Phi \mid \Psi \vee \Phi \mid \Psi \Rightarrow \Phi \mid \forall x. \Phi \mid \exists x. \Phi$$

Here, e is a Boolean expression built from constants and variables that occur in the commands c_1 or c_2 (possibly involving previously defined predicates). These expressions are tagged with $\langle 1 \rangle$ or $\langle 2 \rangle$, denoting their interpretation in the left or right command, respectively. Additionally, an expression e may contain the keyword **res**, which denotes the return value of a procedure. Furthermore the short-hand notation $=\{x\}$ may be used to denote $x\langle 1 \rangle = x\langle 2 \rangle$.

The first-order formulae Ψ and Φ are interpreted as relations over memories: For example, the formula $x\langle 1 \rangle \leq y\langle 2 \rangle$ is interpreted as the relation $R = \{(m_1, m_2) \mid m_1(x) \leq m_2(y)\}$. Then, the validity of a **pRHL** judgment is defined as follows. For any memories m_1, m_2 that satisfy the pre-condition Ψ , the sub-distributions $\llbracket c_1 \rrbracket(m_1)$ and $\llbracket c_2 \rrbracket(m_2)$ must satisfy the lifting of post-condition Φ , i.e. it must hold $\llbracket c_1 \rrbracket(m_1) \mathcal{L}(\Phi) \llbracket c_2 \rrbracket(m_2)$ (recall that since **pWHILE** is probabilistic, its semantics maps program memories to *sub-distributions* over program memories). The lifting $\mu_1 \mathcal{L}(R) \mu_2$ of a relation $R \subseteq A \times B$ (for discrete A and B) to μ_1 and μ_2 is defined by

$$\exists \mu : \mathcal{D}(A \times B). \pi_1(\mu) = \mu_1 \wedge \pi_2(\mu) = \mu_2 \wedge \forall (a, b) : A \times B. \mu(a, b) > 0 \Rightarrow a R b,$$

where $\pi_1(\mu)$ and $\pi_2(\mu)$ denote respectively the projection of μ on its first and second components. We refer the reader to [13] for more details.

Finally, the above definition can be used to derive rules for valid **pRHL** judgments for commands, in the style of

$$\frac{\vdash c_1 \sim c_2 : \Phi \Longrightarrow \Phi' \quad \vdash c'_1 \sim c'_2 : \Phi' \Longrightarrow \Phi''}{\vdash c_1; c'_1 \sim c_2; c'_2 : \Phi \Longrightarrow \Phi''} [\text{Seq}]$$

We refer the reader to [8] or [15] for an ample selection of **pRHL** rules, and comments about them. These rules are hard-coded into *EasyCrypt*, but have been verified in the *CertiCrypt*

framework (improving the base of trust by making it possible to export an **EasyCrypt** proof into **CertiCrypt**).

To support the verification of **pRHL** judgments, these rules are implemented as tactics in **EasyCrypt** that can be applied in a goal-oriented fashion. These tactics may generate additional proof obligations that are passed to SMT solvers and automated theorem provers. Additionally, **EasyCrypt** implements a weakest pre-condition calculus that computes for a (deterministic loop-free) program and a set of post-conditions a set of sufficient verification conditions (see [13]), which must be implied by the pre-conditions. We conclude by observing that for the verification of a **pRHL** judgment, some preparation by the user is still required (specifying invariants, axioms and lemmas, sensible application of tactics, ...), but the involved effort is minimal when compared to a tool such as **CertiCrypt**.

4.4 Reasoning about Probabilities

As mentioned earlier, the last step in an **EasyCrypt** proof is to derive probability claims about the cryptosystem in question. These probabilities can be derived from previous **pRHL** judgments. First, we denote by $\Pr [c, m : A]$ the probability of event A after the execution of command c with initial memory m . We simply write $\Pr [c : A]$ when the initial memory is not relevant (i.e., the claim holds for any memory). The two basic rules applicable in **EasyCrypt** to derive probability claims from **pRHL** judgments are:

$$\frac{m_1 \Psi m_2 \quad \vdash c_1 \sim c_2 : \Psi \implies \Phi \quad \Phi \Rightarrow (A\langle 1 \rangle \Leftrightarrow B\langle 2 \rangle)}{\Pr [c_1, m_1 : A] = \Pr [c_2, m_2 : B]} [\text{PrEq}]$$

$$\frac{m_1 \Psi m_2 \quad \vdash c_1 \sim c_2 : \Psi \implies \Phi \quad \Phi \Rightarrow (A\langle 1 \rangle \Rightarrow B\langle 2 \rangle)}{\Pr [c_1, m_1 : A] \leq \Pr [c_2, m_2 : B]} [\text{PrLe}]$$

In particular, for $\Psi = \text{true}$ (interpreted as the total relation on memories), we obtain $\Pr [c_1 : A] = \Pr [c_2 : B]$ and $\Pr [c_1 : A] \leq \Pr [c_2 : B]$, respectively.

However, these rules alone would not be sufficient in many instances. Indeed, often cryptographic proofs make use of *failure events*; typically, one argues that unless some specific failure event F is triggered in a program c_1 and in a program c_2 , an event A occurs in c_1 if and only if an event B occurs in c_2 . Then, the difference in probability of event A in c_1 and event B in c_2 is bounded by the probability of this specific event F . The lemma which formalizes this argument is known as the *Fundamental Lemma of Game-Playing* [23, 68]. **EasyCrypt** implements this lemma using its probabilistic relational Hoare logic as follows [13].

Lemma 4.1 (Fundamental Lemma). *Let c_1 and c_2 be two terminating commands, m_1 and m_2 be two initial memories, and A, B, F events such that*

$$\vdash c_1 \sim c_2 : \Psi \implies (F\langle 1 \rangle \Leftrightarrow F\langle 2 \rangle) \wedge (\neg F\langle 1 \rangle \Rightarrow (A\langle 1 \rangle \Leftrightarrow B\langle 2 \rangle))$$

Then, if $m_1 \Psi m_2$,

1. $\Pr [c_1, m_1 : A \wedge \neg F] = \Pr [c_2, m_2 : B \wedge \neg F]$, and
2. $|\Pr [c_1, m_1 : A] - \Pr [c_2, m_2 : B]| \leq \Pr [c_1, m_1 : F] = \Pr [c_2, m_2 : F]$.

It still remains to compute a concrete bound on the probability of the failure event F . Typically, failure events are only triggered by oracle calls made by an adversary; when the adversary can only make a known bounded number of oracle queries, the following lemma can be used in EasyCrypt to compute such a bound [8].

Lemma 4.2 (Failure event lemma). *Consider a program $c_1; c_2$, an integer expression i , an event F , and $u \in \mathbb{R}$. Assume the following:*

1. *Free variables in F and i are only modified by c_1 or oracles in some set O ;*
2. *After executing c_1 , F does not hold and $0 \leq i$;*
3. *Oracles $\mathcal{O} \in O$ do not decrease i and strictly increase i when F is triggered;*
4. *For every oracle \mathcal{O} in O , $\neg F \Rightarrow \Pr [\mathcal{O} : F] \leq u$.*

Then, $\Pr [c_1; c_2 : F \wedge i \leq q] \leq q \cdot u$.

Lastly, we note that EasyCrypt can also derive probability claims using previously proven probability claims using standard arithmetics, and furthermore provides mechanisms to compute bounds on elementary events. For instance, it is possible to derive mechanically that the probability that a value v sampled uniformly at random from a set X belongs to a list of n values which are elements of X and independent of v is at most $n/|X|$.

4.5 Selected Techniques

Two techniques will be of particular interest later on. Since we will use random oracles, we will first discuss how these can be formalized in EasyCrypt. Furthermore, for a transition described in Section 5.4, we explain how to justify equivalences between games that sample random values across different procedures.

4.5.1 Random Oracles

It is easy to formalize random oracles in EasyCrypt. Indeed, a random oracle can simply be implemented as a stateful procedure that maintains a map storing answers for previous queries. Upon a fresh query, a random oracle samples uniformly at random a value from its output domain, stores this value in its map, and returns it; for a repeated query, the random oracle simply retrieves the old answer from its map, and returns this value.

Figure 4.1 shows an implementation of a random oracle $\mathcal{O} : X \rightarrow Y$ in EasyCrypt, where the variable \mathbf{T} is a globally defined and initially empty finite map, usually modified exclusively by the procedure \mathcal{O} .

In Chapter 5, we will consider *fixed input length* and *variable input length* random oracles. The only difference is that the input domain of \mathcal{O} is either the set of bitstrings of a fixed length, or the set of bitstrings of variable length. Concretely, we will use a fixed input length random oracle of type $\{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, which can equivalently be seen as a fixed input length random oracle that takes bitstrings of length $k + n$.

Oracle $\mathcal{O}(x)$:
 if $x \notin \text{dom}(\mathbf{T})$ then
 $\mathbf{T}[x] \stackrel{\$}{\leftarrow} Y$
 return $\mathbf{T}[x]$

Figure 4.1: A random oracle in EasyCrypt

4.5.2 Lazy and Eager Sampling

When considering in a transformation two games that both sample random values, the actual sampling of these values in the games may happen across different procedures, meaning that they may be sampled at different points in time. In particular, we distinguish scenarios where values are sampled *lazily*, and scenarios where values are sampled *eagerly*. Lazy sampling means that values are sampled only at that point in time where they are first needed, e.g. when an adversary queries a random oracle on a particular input for the first time. By contrast, eager sampling means that such values may be sampled much earlier, for instance at the beginning of the main procedure of a game, where the variables of the game are initialized. This is best explained by an example: Figure 4.2 illustrates the difference between both scenarios (a similar example is discussed in [16, 75]).

Game G_{lazy} :
 $\mathbf{T} \leftarrow \emptyset$;
 $b \leftarrow \mathcal{D}^{\mathcal{O}_{\text{lazy}}}()$;
 return b

Oracle $\mathcal{O}_{\text{lazy}}(x)$:
 if $x \notin \text{dom}(\mathbf{T})$ then
 $\mathbf{T}[x] \stackrel{\$}{\leftarrow} Y$
 return $\mathbf{T}[x]$

Game G_{eager} :
 $\mathbf{T} \leftarrow \emptyset$; $\hat{\mathbf{y}} \stackrel{\$}{\leftarrow} Y$;
 $b \leftarrow \mathcal{D}^{\mathcal{O}_{\text{eager}}}()$;
 return b

Oracle $\mathcal{O}_{\text{eager}}(x)$:
 if $x \notin \text{dom}(\mathbf{T})$ then
 if $x = \hat{x}$ then $\mathbf{T}[x] \leftarrow \hat{\mathbf{y}}$ else $\mathbf{T}[x] \stackrel{\$}{\leftarrow} Y$
 return $\mathbf{T}[x]$

Figure 4.2: An example illustrating lazy and eager sampling

In both games, a distinguisher is given access to some random oracle, and returns a bit which is the output of the game. The only difference between the two games is the following. In the eager game, for some constant \hat{x} , the return value $\hat{\mathbf{y}}$ of oracle $\mathcal{O}_{\text{eager}}$ is sampled eagerly, in this case in the initialization phase of the game. By contrast, in the lazy game, oracle $\mathcal{O}_{\text{lazy}}$ samples this value lazily, i.e. only upon the call $\mathcal{O}_{\text{lazy}}(\hat{x})$. We expect the two games to behave identically, since the value $\hat{\mathbf{y}}$ is *uniformly distributed and independent of the adversary's view* – that is, it remains unknown to the adversary until it makes the corresponding call, at which point it is sampled uniformly at random in the lazy scenario.

One of the main motivations for switching from an eager sampling scenario to a lazy sampling scenario is precisely to capture this kind of informal argument. Since the sampling of the value $\hat{\mathbf{y}}$ takes place in the main procedure of game G_{eager} , but in the random oracle

$\mathcal{O}_{\text{lazy}}$ in the game G_{lazy} , the problem of switching between these two scenarios can be seen as a problem of *inter-procedural* code motion; see [16] for further details.

To formally deal with this kind of a situation, EasyCrypt reduces this problem to a problem of swapping statements. Generally speaking (but continuing with the same names as in the above example for clarity), let $\mathcal{O}_{\text{eager}}$ be an oracle that uses a value \hat{y} that has been sampled eagerly, and $\mathcal{O}_{\text{lazy}}$ the same as $\mathcal{O}_{\text{eager}}$, except that it samples \hat{y} when first needed. To formalize in EasyCrypt the argument that \hat{y} is uniformly distributed and independent of the adversary's view, we show that if the adversary has obtained no information about \hat{y} , then re-sampling the value \hat{y} preserves the semantics of the program. More precisely, we identify some condition used that holds whenever the adversary has obtained some information about \hat{y} (e.g. in the above example, a sensible condition would be the condition $\hat{x} \in \text{dom}(\mathbf{T})$). Then, we show that the piece of code $S_{\hat{y}} := \text{if } \neg \text{used} \text{ then } \hat{y} \stackrel{\$}{\leftarrow} Y$, which re-samples the value \hat{y} if the adversary has obtained no information about it, swaps with calls to oracles $\mathcal{O}_{\text{lazy}}$ and $\mathcal{O}_{\text{eager}}$, namely

$$\vdash y \leftarrow \mathcal{O}_{\text{lazy}}(\vec{e}); S_{\hat{y}} \sim S_{\hat{y}}; y \leftarrow \mathcal{O}_{\text{eager}}(\vec{e}) : = \{\vec{e}\} \wedge \Phi \implies = \{y\} \wedge \Phi,$$

where Φ implies equality over all global variables.

Continuing with our above example, to show that the games G_{lazy} and G_{eager} behave identically, we would introduce the intermediate games shown in Figure 4.3.

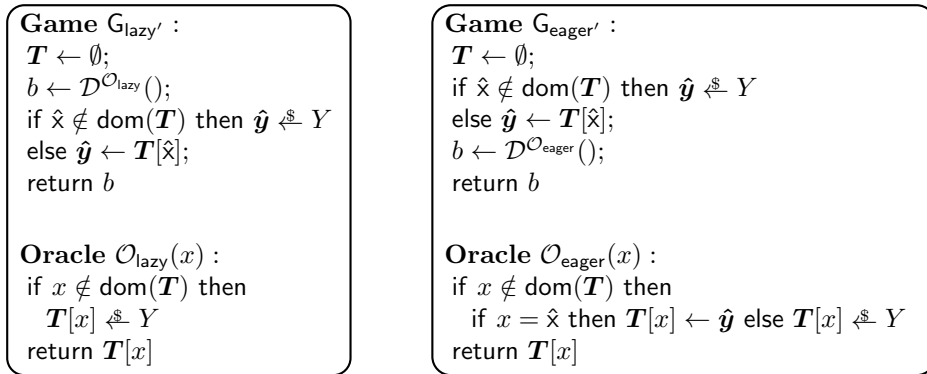


Figure 4.3: An example illustrating lazy and eager sampling, continued

Showing in EasyCrypt that $\Pr[G_{\text{lazy}} : b] = \Pr[G_{\text{lazy}'} : b]$ and $\Pr[G_{\text{eager}} : b] = \Pr[G_{\text{eager}'} : b]$ is easy and is not further detailed here. The step from a lazy sampling scenario to an eager sampling scenario happens in the transition between the games $G_{\text{lazy}'}$ and $G_{\text{eager}'}$. As outlined above, we show that

$$\vdash \begin{array}{l} y \leftarrow \mathcal{O}_{\text{lazy}}(x); \\ \text{if } \hat{x} \notin \text{dom}(\mathbf{T}) \text{ then } \hat{y} \stackrel{\$}{\leftarrow} Y \\ \text{else } \hat{y} \leftarrow \mathbf{T}[\hat{x}] \end{array} \sim \begin{array}{l} \text{if } \hat{x} \notin \text{dom}(\mathbf{T}) \text{ then } \hat{y} \stackrel{\$}{\leftarrow} Y \\ \text{else } \hat{y} \leftarrow \mathbf{T}[\hat{x}]; \\ y \leftarrow \mathcal{O}_{\text{eager}}(x) \end{array} : = \{x, \mathbf{T}, \hat{y}\} \implies = \{y, \mathbf{T}, \hat{y}\}.$$

By using a logic for swapping statements that is built into pRHL [16], this judgment then extends to the following judgment, which EasyCrypt can conclude automatically:

$$\vdash \begin{array}{l} b \leftarrow \mathcal{D}^{\mathcal{O}_{\text{lazy}}}(); \\ \text{if } \hat{x} \notin \text{dom}(\mathbf{T}) \text{ then } \hat{y} \leftarrow Y \\ \text{else } \hat{y} \leftarrow \mathbf{T}[\hat{x}] \end{array} \sim \begin{array}{l} \text{if } \hat{x} \notin \text{dom}(\mathbf{T}) \text{ then } \hat{y} \leftarrow Y \\ \text{else } \hat{y} \leftarrow \mathbf{T}[\hat{x}]; \\ b \leftarrow \mathcal{D}^{\mathcal{O}_{\text{eager}}}() \end{array} : =\{\mathbf{T}, \hat{y}\} \Longrightarrow =\{b, \mathbf{T}, \hat{y}\}$$

The latter then trivially implies

$$\vdash \mathbf{G}_{\text{lazy}'} \sim \mathbf{G}_{\text{eager}'} : \text{true} \Longrightarrow =\{b\}.$$

We conclude $\Pr[\mathbf{G}_{\text{lazy}'} : b] = \Pr[\mathbf{G}_{\text{eager}'} : b]$ (see rule [PrEq] in Section 4.4) and hence $\Pr[\mathbf{G}_{\text{lazy}} : b] = \Pr[\mathbf{G}_{\text{eager}} : b]$. We refer the reader to [16] for a formal discussion of this topic.

High-level proof

INDIFFERENTIABILITY from a random oracle for the prefix-free Merkle-Damgård construction was first shown by Coron et al. [39] (see Section 3.2), as a pen-and-paper proof. We closely followed that proof in our `EasyCrypt` formalization, which we will describe in detail in this chapter. Our proof slightly differs from that of Coron in that the latter is based on a Davies-Meyer compression function [73] and is performed in the ideal cipher model (where the underlying block cipher is assumed to be ideal), whereas our proof sees the entire compression function as an ideal primitive and is therefore in the random oracle model — that is, we assume that the compression function is a fixed-input-length random oracle. More precisely, we claim the following.

Theorem 5.1 (Indifferentiability of pfMD). *The Merkle-Damgård construction MD with an ideal compression function f , prefix-free padding function pad , and initialization vector IV is $(t_{\mathcal{S}}, q_{\mathcal{D}}, \epsilon)$ -indifferentiable from a variable-input-length random oracle $F_q : \{0, 1\}^* \rightarrow \{0, 1\}^n$, where*

$$\epsilon = \frac{3\ell^2 \cdot q_{\mathcal{D}}^2}{2^n} \quad t_{\mathcal{S}} = \mathcal{O}(\ell \cdot q_{\mathcal{D}}^2)$$

and ℓ is an upper bound on the block length of $\text{pad}(m)$ for any message m appearing in a query of the distinguisher.

The precise axiomatization of the function pad (and its inverse unpad) is given in Section 6.1.2.

To perform the proof in `EasyCrypt`, we first define two games \mathbf{G}_{real} and $\mathbf{G}_{\text{ideal}}$. The first game represents the scenario where the distinguisher \mathcal{D} is given access to the Merkle-Damgård construction, implemented by the procedure $F_q : \{0, 1\}^* \rightarrow \{0, 1\}^n$, and the compression function, implemented as a fixed input length random oracle by the procedure $f_q : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. We call this scenario the *real world* since it models the actual MD-construction (based on an idealized compression function). The second game represents the scenario where the hash function behaves like a random oracle itself. Therefore, here the procedure F_q implements a variable-input-length random oracle. Since this hash function is idealized, we call this scenario the *ideal world*. The games \mathbf{G}_{real} and

Game G_{real} : $\mathbf{q}_f \leftarrow 0;$ $\mathbf{T} \leftarrow \emptyset;$ $b \leftarrow \mathcal{D}^{F_q, f_q}();$ return b	Oracle $F_q(m)$: $xs \leftarrow \text{pad}(m); y \leftarrow \text{IV};$ if $\mathbf{q}_f + xs \leq q$ then $\mathbf{q}_f \leftarrow \mathbf{q}_f + xs ;$ while $xs \neq \text{nil}$ do $y \leftarrow f(\text{hd}(xs), y);$ $xs \leftarrow \text{tl}(xs)$ return y	Oracle $f(x, y)$: if $(x, y) \notin \text{dom}(\mathbf{T})$ then $z \xleftarrow{\$} \{0, 1\}^n;$ $\mathbf{T}[x, y] \leftarrow z$ return $\mathbf{T}[x, y]$	Oracle $f_q(x, y)$: if $\mathbf{q}_f + 1 \leq q$ then $\mathbf{q}_f \leftarrow \mathbf{q}_f + 1;$ $z \leftarrow f(x, y)$ else $z \leftarrow \text{IV};$ return z
Game G_{ideal} : $\mathbf{q}_f \leftarrow 0;$ $\mathbf{R}, \mathbf{T}' \leftarrow \emptyset;$ $b \leftarrow \mathcal{D}^{F_q, f_q}();$ return b	Oracle $F_q(m)$: $xs \leftarrow \text{pad}(m); y \leftarrow \text{IV};$ if $\mathbf{q}_f + xs \leq q$ then $\mathbf{q}_f \leftarrow \mathbf{q}_f + xs ;$ $z \leftarrow F(m)$ else $z \leftarrow \text{IV}$ return z	Oracle $F(m)$: if $m \notin \text{dom}(\mathbf{R})$ then $z \xleftarrow{\$} \{0, 1\}^n;$ $\mathbf{R}[m] \leftarrow z$ return $\mathbf{R}[m]$	Oracle $f_q(x, y)$: if $\mathbf{q}_f + 1 \leq q$ then if $(x, y) \notin \text{dom}(\mathbf{T}')$ then $xs \leftarrow \text{findseq}((x, y), \mathbf{T}');$ if $xs \neq \text{None}$ then $m \leftarrow \pi_{\text{unpad}}(\text{mapfst}(\pi_{xs})\ [xs]);$ $\mathbf{T}'[x, y] \leftarrow F(m)$ else $\mathbf{T}'[x, y] \xleftarrow{\$} \{0, 1\}^n;$ $z \leftarrow \mathbf{T}'[x, y]; \mathbf{q}_f \leftarrow \mathbf{q}_f + 1$ else $z \leftarrow \text{IV}$ return z

Figure 5.1: The games G_{real} and G_{ideal}

G_{ideal} are defined in Figure 5.1. Our goal is to show that the probabilities of the outcomes of the games G_{real} and G_{ideal} differ at most by a negligible amount, i.e. the advantage of the distinguisher in differentiating the two games is negligible, which proves the claim.

Note that in order to limit the number of oracle queries allowed to the simulator, we increase a global counter \mathbf{q}_f in both the procedures f_q and F_q . A call to the procedure f_q increments the value \mathbf{q}_f . A call to the procedure F_q increases the value of \mathbf{q}_f by the number of blocks associated with the padding of the parameter m – the idea is that one call to the Merkle-Damgård iteration results in that many calls to the compression function. Thus, the counter \mathbf{q}_f counts the number of evaluations of the FIL-RO f in game G_{real} . This is more permissive than the proof of Coron et al. [39], which uses separate counters for the two oracles accessible to the distinguisher. Indeed, our formalization gives the distinguisher the freedom to decide how to distribute his queries among F_q and f_q . The bound which we enforce on this counter is $q := \ell \cdot q_{\mathcal{D}}$; if the value of \mathbf{q}_f exceeds q , then all oracles accessible to the distinguisher in all games always simply return the value IV . This effectively means that if the distinguisher makes n_f queries to f_q and n_F queries to F_q , we require

$$\mathbf{q}_f \leq n_f + \ell \cdot n_F \leq \ell \cdot (n_f + n_F) = \ell \cdot q_{\mathcal{D}} = q.$$

For instance, if the distinguisher makes only queries to f_q but none to F_q , we obtain $\mathbf{q}_f = n_f = q_{\mathcal{D}} = q$. On the other hand, if it makes only queries to F_q (with a maximum query length of ℓ) but none to f_q , we get $\mathbf{q}_f \leq \ell \cdot n_F = \ell \cdot q_{\mathcal{D}} = q$. Therefore, the distinguisher

is always allowed at least $q_{\mathcal{D}}$ queries in the two games, as required in Theorem 5.1.

To enable us to perform the proof, the simulator f_q implemented in the game $\mathbf{G}_{\text{ideal}}$ must sensibly simulate the role of the idealized compression function f_q in the game \mathbf{G}_{real} . That is, the simulator $\mathbf{G}_{\text{ideal}} \cdot f_q$ has to behave towards the oracle $\mathbf{G}_{\text{ideal}} \cdot F_q$ similarly as the idealized compression function $\mathbf{G}_{\text{real}} \cdot f_q$ behaves towards the MD-construction $\mathbf{G}_{\text{real}} \cdot F_q$.

The connection between procedures F_q and f_q in the real world is the following. Let m be a message with $\text{pad}(m) = [x_1, \dots, x_j]$. Then, if the distinguisher makes the query $F_q(m)$ in the game \mathbf{G}_{real} , the result will be the return value of $f(x_j, y_j)$, where y_j is some chaining value sampled during the Merkle-Damgård iteration. Thus, the idealized compression function f_q will return the same value upon the query $f_q(x_j, y_j)$, since it relays to the FIL-RO f . The simulator needs to achieve the same thing in the ideal world. Whenever the distinguisher queries $F_q(m)$ in the game $\mathbf{G}_{\text{ideal}}$ and then reconstructs the MD-chain associated with m using queries to simulator f_q , the simulator should return the same answer as $F_q(m)$ upon the query $f_q(x_j, y_j)$. Conversely, if the distinguisher first builds an MD-chain using queries to f_q and later queries F_q for the associated message of this chain, the answer of F_q should be consistent with the last query to f_q . To achieve this in the ideal world, the simulator must detect when a fresh query it receives may be the *last* element of an MD-chain, and relay to the VIL-RO F in this case. We will see later that up to some negligible probability, the distinguisher can only construct an MD-chain using calls to f_q by querying all its elements in the correct order, because it cannot guess the intermediate chaining values. Therefore, if the distinguisher constructs an MD-chain using calls to f_q , when it queries for the last element, it must already have queried f_q for all preceding elements of this chain. Hence, the simulator is able to detect that a chain has been completed (as it maintains a table of its previous queries), and it can reconstruct the entire message m and forward it to oracle F to maintain consistency with F_q .

To better express this idea, we introduce the notion of a *complete chain*. Intuitively, a complete chain is an MD-chain whose elements are all in the domain of the map maintained by the simulator, except for the ultimate element of this chain. Formally, it is defined as follows.

Definition 5.2 (Complete chain). A complete chain in a map $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a list $[(x_1, y_1), \dots, (x_j, y_j)]$ such that

1. $y_1 = \mathbf{IV}$,
2. $\forall i \in \{1, \dots, j-1\}. (x_i, y_i) \in \text{dom}(T) \wedge T[x_i, y_i] = y_{i+1}$, and
3. $[x_1, \dots, x_j]$ is in the domain of unpad .

To detect complete chains, the simulator uses the function findseq (see Section 6.1.4 for its axiomatization). The function $\text{findseq}((x, y), \mathbf{T}')$ searches in the map \mathbf{T}' for a complete chain of the form $[(x_1, y_1), \dots, (x_{j-1}, y_{j-1}), (x, y)]$ and returns $[(x_1, y_1), \dots, (x_{j-1}, y_{j-1})]$, or None if no such chain is found. Thus, upon a query $f_q(x, y)$ in game $\mathbf{G}_{\text{ideal}}$, the simulator relays to F whenever there is a complete chain of the form $[(x_1, \mathbf{IV}), \dots, (x_{j-1}, y_{j-1}), (x, y)]$ in the simulator's map \mathbf{T}' . Otherwise, the simulator answers with a uniformly distributed random value.

An example illustrating how the simulator works is depicted in Figure 5.2. Here, the simulator receives a sequence of queries

$$y_2 \leftarrow f_q(x_1, IV); y_3 \leftarrow f_q(x_2, y_2); y_4 \leftarrow f_q(x_3, y_3)$$

where $[x_1, x_2, x_3] = \text{pad}(m)$ for some message m . The first two queries do not complete a chain – note that because of the prefix-freeness of pad there is no message m' such that $\text{pad}(m') = [x_1]$ or $\text{pad}(m') = [x_1, x_2]$. Therefore the simulator samples random values for these queries; indeed, in both worlds the corresponding chaining values are unknown to the distinguisher until it has queried f_q for them. The third query completes a chain, and so the simulator answers it by querying F for the hash of $\text{unpad}([x_1, x_2, x_3])$. Hence, the simulator's answer to query $f_q(x_3, y_3)$ is consistent with the random oracle's answer to query $F_q(m)$, exactly as in the real world.




Figure 5.2: An example illustrating how the simulator works

To show that the probabilities of the outcomes of the two games \mathbf{G}_{real} and $\mathbf{G}_{\text{ideal}}$ differ only by a negligible amount, we introduce a sequence of hybrid games. We then upper-bound the probability by which each two consecutive games may differ; by summing up over these probabilities, we obtain a concrete bound for the advantage of the distinguisher in telling apart the initial and final games. Specifically, we prove:

$$|\Pr[\mathbf{G}_{\text{real}} : b] - \Pr[\mathbf{G}_{\text{ideal}} : b]| \leq \frac{3q^2}{2^n} \quad (5.1)$$

Recall that $q = \ell \cdot q_D$, which justifies the bound stated in Theorem 5.1. The overall running time of the simulator is bounded by $\mathcal{O}(\ell \cdot q_D^2)$, as the function findseq runs in time $\mathcal{O}(\ell \cdot q_D)$ and the simulator may be called up to q_D times. The running time of the function findseq is justified since the length of any possible chain in \mathbf{T} is upper-bounded by q ; see Section 6.1.4 for a pseudocode implementation of findseq . The functions unpad and mapfst (see Sections 6.1.2 and 6.1.3 respectively for their definitions) are easily seen to run in time $\mathcal{O}(\ell \cdot q_D)$ as well.

To bound the amount by which the probabilities of the outcomes of two consecutive games may differ, we may introduce failure events in the hybrid games. We then apply the Fundamental Lemma of Game-Playing (see Theorem 4.1) to conclude that the probability of some event in the two games is bounded by the probability of this failure event. The event which we are interested in is usually the outcome of the game, but we may also be interested e.g. in relating the probability of some failure events in two games with each other.

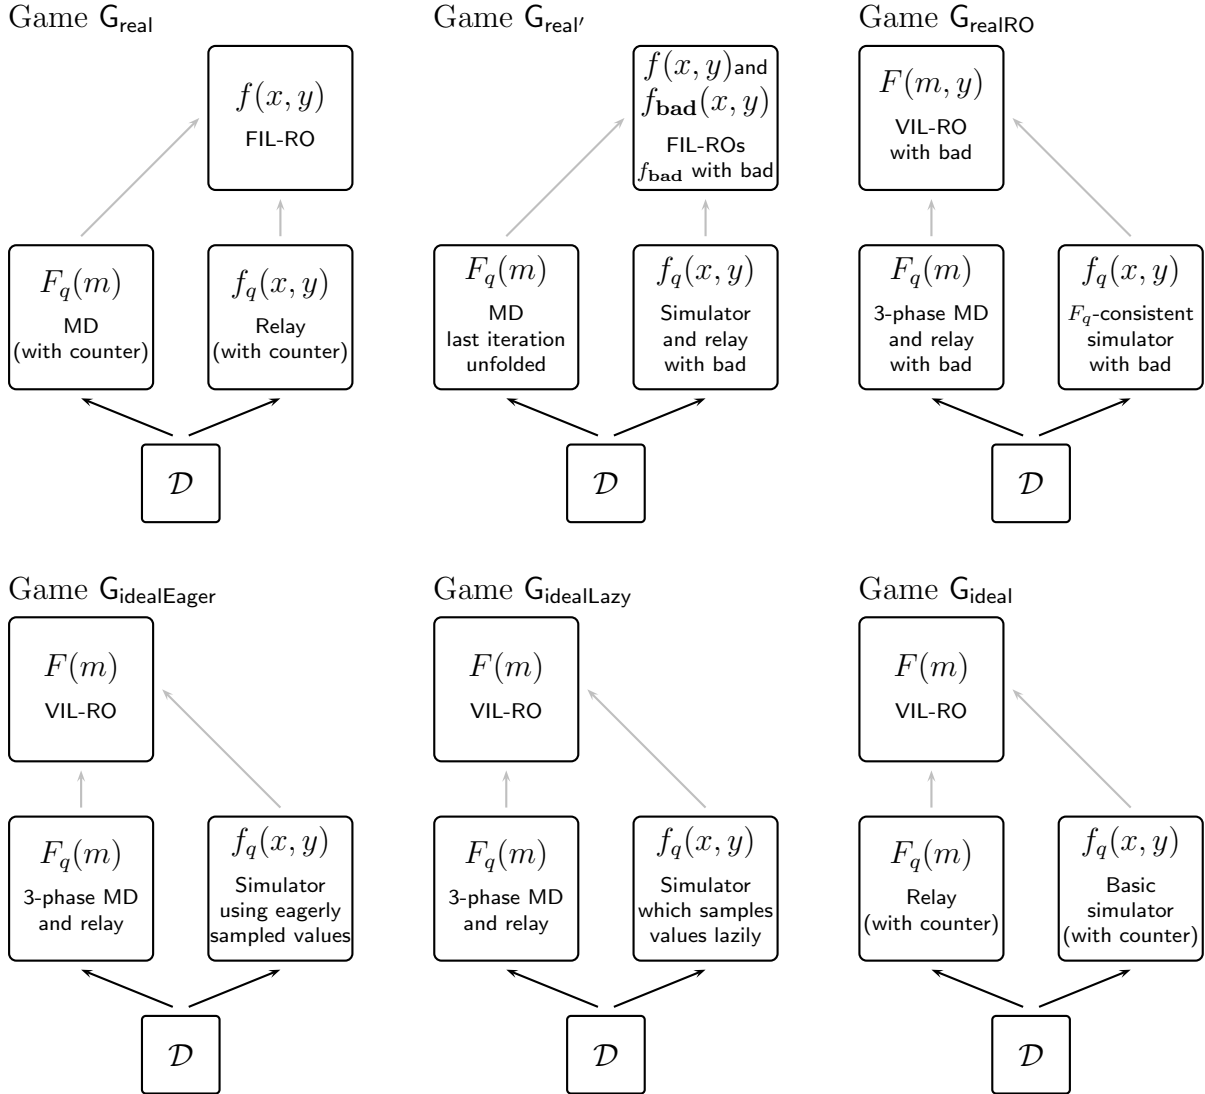


Figure 5.3: High-level overview of the games

Figure 5.3 shows a high-level overview over the sequence of games which we use. The arrows show in which way the procedures may call each other; the distinguisher is always given access to the two procedures $F_q : \{0, 1\}^* \rightarrow \{0, 1\}^n$ and $f_q : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, as depicted by the black arrows in the figure. In all cases, these two procedures may call further internal procedures, as depicted by the gray arrows. Note that the distinguisher has no knowledge of the existence of these internal procedures, or of the calls made by F_q and f_q .

The two most difficult steps are those from $G_{\text{real}'}$ to G_{realRO} , and from $G_{\text{idealEager}}$ to $G_{\text{idealLazy}}$. All other steps are easy and essentially concern the introduction or removing of some failure events, or some code cleanup. In the step from $G_{\text{real}'}$ to G_{realRO} , we perform the step from the actual Merkle-Damgård iteration using an idealized compression function to a variable-

Game $G_{\text{real}'}$: $\mathbf{q}_f \leftarrow 0;$ $\mathbf{T}, \mathbf{T}' \leftarrow \emptyset;$ $\mathbf{Y} \leftarrow \text{nil};$ $\mathbf{Z} \leftarrow \text{IV}::\text{nil};$ $\mathbf{bad}_1 \leftarrow \text{false};$ $\mathbf{bad}_2 \leftarrow \text{false};$ $\mathbf{bad}_3 \leftarrow \text{false};$ $b \leftarrow \mathcal{D}^{F_q, f_q}();$ return b	Oracle $F_q(m)$: $xs \leftarrow \text{pad}(m); y \leftarrow \text{IV};$ if $\mathbf{q}_f + xs \leq q$ then $\mathbf{q}_f \leftarrow \mathbf{q}_f + xs ;$ while $ xs > 1$ do $y \leftarrow f_{\text{bad}}(\text{hd}(xs), y);$ $xs \leftarrow \text{tl}(xs)$ $y \leftarrow f_{\text{bad}}(\text{hd}(xs), y)$ return y	Oracle $f(x, y)$: if $(x, y) \notin \text{dom}(\mathbf{T})$ then $z \stackrel{\$}{\leftarrow} \{0, 1\}^n;$ $\mathbf{Z} \leftarrow z::\mathbf{Z}; \mathbf{Y} \leftarrow y::\mathbf{Y};$ $\mathbf{T}[x, y] \leftarrow z$ return $\mathbf{T}[x, y]$ Oracle $f_{\text{bad}}(x, y)$: if $(x, y) \notin \text{dom}(\mathbf{T})$ then $z \stackrel{\$}{\leftarrow} \{0, 1\}^n;$ $\mathbf{bad}_1 \leftarrow \mathbf{bad}_1 \vee z \in \mathbf{Z};$ $\mathbf{Z} \leftarrow z::\mathbf{Z}; \mathbf{Y} \leftarrow y::\mathbf{Y};$ $\mathbf{bad}_2 \leftarrow \mathbf{bad}_2 \vee z \in \mathbf{Y};$ $\mathbf{T}[x, y] \leftarrow z$ return $\mathbf{T}[x, y]$	Oracle $f_q(x, y)$: if $\mathbf{q}_f + 1 \leq q$ then if $(x, y) \notin \text{dom}(\mathbf{T}')$ then $xs \leftarrow \text{findseq}((x, y), \mathbf{T}')$ if $xs \neq \text{None}$ then $\mathbf{T}'[x, y] \leftarrow f_{\text{bad}}(x, y)$ else if $\text{set_bad3}(y, \mathbf{T}', \mathbf{T})$ then $\mathbf{bad}_3 \leftarrow \text{true};$ $\mathbf{T}'[x, y] \leftarrow f(x, y)$ else $\mathbf{T}'[x, y] \leftarrow f_{\text{bad}}(x, y)$ $z \leftarrow \mathbf{T}'[x, y]; \mathbf{q}_f \leftarrow \mathbf{q}_f + 1$ else $z \leftarrow \text{IV}$ return z
--	--	---	---

Figure 5.4: The game $G_{\text{real}'}$

input-length random oracle and a consistent simulator. The step from $G_{\text{idealEager}}$ to $G_{\text{idealLazy}}$ concerns the step from eager sampling to lazy sampling, which we will address in more detail in the concerned section. For all the transformations except the one from $G_{\text{real}'}$ to G_{realRO} , we have that the probabilities of the outcomes are equal; the difference of the probabilities of the outcomes of $G_{\text{real}'}$ and G_{realRO} is bounded by $3q^2/2^n$, which leads to the final bound stated in Equation 5.1. In the following sections, we will consider each of these transformations in detail.

5.1 From G_{real} to $G_{\text{real}'}$

The first hybrid game that we consider is the game $G_{\text{real}'}$, defined in Figure 5.4. In this game, three easy but essential modifications to the game G_{real} are being made to prepare for the next transformation:

1. The events \mathbf{bad}_1 , \mathbf{bad}_2 and \mathbf{bad}_3 and some related changes are introduced;
2. the relay procedure f_q is transformed into a simulator that still behaves as a relay;
3. the last loop iteration in F_q is unrolled.

The first change means that we instrument the code such that under certain conditions, a failure event is raised. We will explain the precise meaning of those events in the next transformation, where they are needed. We also introduce a copy of oracle f , called f_{bad} , which behaves like f except that it may set the flags \mathbf{bad}_1 and \mathbf{bad}_2 . The lists \mathbf{Y} and \mathbf{Z} are introduced to allow us to appropriately detect these events. We note that modifying the global variables \mathbf{Y} , \mathbf{Z} , \mathbf{bad}_1 , \mathbf{bad}_2 or \mathbf{bad}_3 , all unknown to the distinguisher, does not change the observable behavior of the procedures F_q and f_q .

The second change, concerning procedure f_q , introduces a new map \mathbf{T}' of queries known to the distinguisher. Observe that $\mathbf{T}' \subseteq \mathbf{T}$, because queries to F_q result in entries being

added only to \mathbf{T} , whereas queries to f_q result in the same entries being added to both \mathbf{T} and \mathbf{T}' . Additionally, the procedure f_q now behaves in two ways depending on whether $\text{findseq}((x, y), \mathbf{T}') \neq \text{None}$; this distinction is a central element of the simulator. Indeed, if this condition holds, then there is a complete chain in $\text{map } \mathbf{T}'$ ending in (x, y) . This means that the simulator should call the VIL-RO F to maintain consistency with the oracle F_q in the ideal world. If it does not hold, then the simulator should simply sample a fresh random value, i.e. behave like a FIL-RO. In this game, oracle f_q still returns the same answer in both cases, and only sets $\mathbf{bad}_{\{1,2,3\}}$ accordingly.

The unrolling of the last loop iteration will also be important in the next transformation, where the last call to the procedure $f_{\mathbf{bad}}$ will be replaced by a call to the VIL-RO introduced there. Here, to be able to show that the unrolling of the last iteration is sound, **EasyCrypt** needs to know that after the assignment $xs \leftarrow \text{pad}(m)$, the variable xs contains at least one element, which follows immediately from Lemma 6.17.

Since none of the transformations change the behavior of the procedures F_q or f_q accessible to the distinguisher, it is easy to conclude in **EasyCrypt** that

$$\vdash G_{\text{real}} \sim G_{\text{real}'} : \text{true} \implies = \{b\},$$

which implies

$$\Pr [G_{\text{real}} : b] = \Pr [G_{\text{real}'} : b]. \quad (5.2)$$

5.2 From $G_{\text{real}'}$ to G_{realRO}

This transformation lies at the heart of our proof. Intuitively, we replace the FIL-ROs f and $f_{\mathbf{bad}}$ by a VIL-RO F , and modify the oracle F_q such that it behaves as a relay for F , while the simulator now maintains consistency with F_q by calling F whenever a new query completes a chain in \mathbf{T}' . To prove the validity of this transformation, we have to show that the simulator indeed behaves consistently with F_q .

Broadly speaking, we introduce the following changes:

1. The VIL-RO F is substituted for the FIL-ROs f and $f_{\mathbf{bad}}$;
2. the Merkle-Damgård iteration in F_q is split into three phases;
3. the last call to the procedure $f_{\mathbf{bad}}$ in oracle F_q of the game $G_{\text{real}'}$ is replaced by a call to the procedure F in the game G_{realRO} ;
4. the simulator also calls oracle F in case it gets a new query which completes a chain, so as to be consistent with oracle F_q ;
5. the new maps \mathbf{I} and \mathbf{T}'_i are introduced, along with a modification to $\text{map } \mathbf{T}$ and further related changes, to allow for a more powerful book-keeping.

Game G_{realRO} : $q_f \leftarrow 0;$ $q'_f \leftarrow 1;$ $T, T', T'_i, R, I \leftarrow \emptyset;$ $I[0] \leftarrow (IV, \text{false});$ $Y \leftarrow \text{nil};$ $Z \leftarrow IV::\text{nil};$ $\text{bad}_1 \leftarrow \text{false};$ $\text{bad}_2 \leftarrow \text{false};$ $\text{bad}_3 \leftarrow \text{false};$ $b \leftarrow \mathcal{D}^{F_q, f_q}();$ return b	Oracle $F_q(m)$: $xs \leftarrow \text{pad}(m); y \leftarrow IV;$ $i \leftarrow 0;$ if $q_f + xs \leq q$ then $q_f \leftarrow q_f + xs ;$ while $ xs > 1 \wedge$ $(\text{hd}(xs), y) \in \text{dom}(T')$ do $i \leftarrow T'_i[\text{hd}(xs), y];$ $y \leftarrow T'[\text{hd}(xs), y];$ $xs \leftarrow \text{tl}(xs);$ while $ xs > 1 \wedge$ $(\text{hd}(xs), i) \in \text{dom}(T)$ do $i \leftarrow T[\text{hd}(xs), i];$ $y \leftarrow \text{fst}(I[i]);$ $xs \leftarrow \text{tl}(xs);$ while $ xs > 1$ do $z \xleftarrow{\$} \{0, 1\}^n;$ $\text{bad}_1 \leftarrow \text{bad}_1 \vee z \in Z;$ $Z \leftarrow z::Z; Y \leftarrow y::Y;$ $\text{bad}_2 \leftarrow \text{bad}_2 \vee z \in Y;$ $T[\text{hd}(xs), i] \leftarrow q'_f;$ $I[q'_f] \leftarrow (z, \text{true});$ $i \leftarrow q'_f;$ $y \leftarrow z;$ $q'_f \leftarrow q'_f + 1;$ $xs \leftarrow \text{tl}(xs)$ $y \leftarrow \text{fst}(F(m, y))$ return y	Oracle $F(m, y)$: if $m \notin \text{dom}(R)$ then $z \xleftarrow{\$} \{0, 1\}^n;$ $\text{bad}_1 \leftarrow \text{bad}_1 \vee z \in Z;$ $Z \leftarrow z::Z; Y \leftarrow y::Y;$ $\text{bad}_2 \leftarrow \text{bad}_2 \vee z \in Y;$ $R[m] \leftarrow (z, q'_f)$ $I[q'_f] \leftarrow (z, \text{false})$ $q'_f \leftarrow q'_f + 1$ return $R[m]$	Oracle $f_q(x, y)$: if $q_f + 1 \leq q$ then if $(x, y) \notin \text{dom}(T')$ then $xs \leftarrow \text{findseq}((x, y), T')$ if $xs \neq \text{None}$ then $m \leftarrow \pi_{\text{unpad}}(\text{mapfst}(\pi_{xs})\ [x]);$ $(z, j) \leftarrow F(m, y);$ $T'[x, y] \leftarrow z; T'_i[x, y] \leftarrow j;$ else $\text{found}, \text{found_bad3} \leftarrow \text{false};$ $j, k' \leftarrow 0;$ while $k' < q'_f$ do if $\text{snd}(I[k'])$ then $\text{found_bad3} \leftarrow (\text{fst}(I[k']) = y);$ else if $\neg \text{found} \wedge \text{fst}(I[k']) = y \wedge$ $(x, k') \in \text{dom}(T) \wedge$ $\text{snd}(I[T[x, k']])$ then $\text{found} \leftarrow \text{true}; j \leftarrow T[x, k'];$ $k' \leftarrow k' + 1;$ if found then $z \leftarrow \text{fst}(I[j]); I[j] \leftarrow (z, \text{false});$ $T'[x, y] \leftarrow z; T'_i[x, y] \leftarrow j;$ else if found_bad3 then $\text{bad}_3 \leftarrow \text{true};$ $z \xleftarrow{\$} \{0, 1\}^n;$ $I[q'_f] \leftarrow (z, \text{false});$ $T'[x, y] \leftarrow z;$ $T'_i[x, y] \leftarrow q'_f;$ $q'_f \leftarrow q'_f + 1;$ else $z \xleftarrow{\$} \{0, 1\}^n;$ $\text{bad}_1 \leftarrow \text{bad}_1 \vee z \in Z;$ $Z \leftarrow z::Z; Y \leftarrow y::Y;$ $\text{bad}_2 \leftarrow \text{bad}_2 \vee z \in Y;$ $I[q'_f] \leftarrow (z, \text{false});$ $T'[x, y] \leftarrow z;$ $T'_i[x, y] \leftarrow q'_f;$ $q'_f \leftarrow q'_f + 1$ $z \leftarrow T'[x, y]; q_f \leftarrow q_f + 1$ else $z \leftarrow IV$ return z
--	--	--	--

 Figure 5.5: The game G_{realRO}

Replacing the FIL-ROs by a VIL-RO is a significant modification. Indeed, now the oracle F_q does not use a FIL-RO to perform a normal Merkle-Damgård iteration any longer. Instead, the last call to f_{bad} is replaced by a call to the new VIL-RO F . However, F_q still performs an iteration similar to Merkle-Damgård in order to sample values eagerly, as was also the case in game G_{real} – we will further discuss these loops later. As for the simulator f_q , it now behaves differently depending on the condition $\text{findseq}((x, y), T') \neq \text{None}$; specifically, whenever a query (x, y) to f_q completes a chain in T' , the simulator now calls F , to maintain consistency with F_q from the distinguisher's perspective.

We introduce several new maps in this game. Firstly, we introduce a map $I : \mathbb{N} \rightarrow \{0, 1\}^n \times \mathbb{B}$ which enumerates all values randomly sampled so far in game G_{realRO} through

means of the counter \mathbf{q}'_f . Each time when a new value is sampled at random, this newly sampled value is written to $\mathbf{I}[\mathbf{q}'_f]$, with an additional *tainted* flag that keeps track of values known to the distinguisher (here **false** means that a value is known to the distinguisher, and **true** means it is not). Then, the counter \mathbf{q}'_f is incremented. As before, the simulator maintains a map $\mathbf{T}' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ of previously answered queries; we additionally introduce the map $\mathbf{T}'_i : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \mathbb{N}$, also maintained by the simulator, to be able to relate the maps \mathbf{T}' and \mathbf{I} . The VIL-RO F maintains its previous queries in a table $\mathbf{R} : \{0, 1\}^* \rightarrow \{0, 1\}^n \times \mathbb{N}$, where the integer indicates the value of counter \mathbf{q}'_f at the point at which the corresponding value was sampled. Finally, the map $\mathbf{T} : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ from game G_{real} is transformed into a map $\mathbf{T} : \{0, 1\}^k \times \mathbb{N} \rightarrow \mathbb{N}$ in game G_{realRO} . Now, instead of immediately storing the query-answer pairs mapping pairs of blocks and chaining values to chaining values, it simply stores the indices of those chaining values enumerated in table \mathbf{I} . All of these modifications allow us to precisely keep track of the order in which queries were made, and of which values are known to the distinguisher. It is necessary to formulate explicitly several invariants about the relation of these maps first so as to allow **EasyCrypt** to derive more powerful invariants. We list them here for reference, and the purpose of completeness. We start off with very basic invariants:

1. $\mathbf{T}'\langle 1 \rangle \subseteq \mathbf{T}\langle 1 \rangle$
2. $1 \leq \mathbf{q}'_f\langle 2 \rangle$
3. $\mathbf{q}'_f\langle 2 \rangle \notin \text{dom}(\mathbf{I}\langle 2 \rangle)$
4. $\forall i \in \text{dom}(\mathbf{I}\langle 2 \rangle). 0 \leq i < \mathbf{q}'_f\langle 2 \rangle$
5. $\forall i. 0 \leq i < \mathbf{q}'_f\langle 2 \rangle \Rightarrow i \in \text{dom}(\mathbf{I}\langle 2 \rangle)$
6. $\text{fst}(\mathbf{I}\langle 2 \rangle[0]) = \text{IV}$

Further, relating the new maps $\mathbf{I}\langle 2 \rangle$, $\mathbf{T}'_i\langle 2 \rangle$, $\mathbf{R}\langle 2 \rangle$ and $\mathbf{T}\langle 2 \rangle$ between each other and with the old map $\mathbf{T}'\langle 1 \rangle$, we derive

7. $\forall (x, y) \in \text{dom}(\mathbf{T}'\langle 1 \rangle). (x, y) \in \text{dom}(\mathbf{T}'_i\langle 2 \rangle)$
8. $\forall (x, y) \in \text{dom}(\mathbf{T}'_i\langle 2 \rangle). (x, y) \in \text{dom}(\mathbf{T}'\langle 1 \rangle) \wedge \mathbf{T}'_i\langle 2 \rangle[(x, y)] \in \text{dom}(\mathbf{I}\langle 2 \rangle) \wedge \text{fst}(\mathbf{I}\langle 2 \rangle[\mathbf{T}'_i\langle 2 \rangle[(x, y)]]) = \mathbf{T}'\langle 1 \rangle[(x, y)]$
9. $\forall m \in \text{dom}(\mathbf{R}\langle 2 \rangle). \text{snd}(\mathbf{R}\langle 2 \rangle[m]) \in \text{dom}(\mathbf{I}\langle 2 \rangle) \wedge \text{fst}(\mathbf{I}\langle 2 \rangle[\text{snd}(\mathbf{R}\langle 2 \rangle[m])]) = \text{fst}(\mathbf{R}\langle 2 \rangle[m])$
10. $\forall (x, i) \in \text{dom}(\mathbf{T}\langle 2 \rangle). i \in \text{dom}(\mathbf{I}\langle 2 \rangle) \wedge (x, \text{fst}(\mathbf{I}\langle 2 \rangle[i])) \in \text{dom}(\mathbf{T}\langle 1 \rangle) \wedge \mathbf{T}\langle 2 \rangle[(x, i)] \in \text{dom}(\mathbf{I}\langle 2 \rangle) \wedge \text{fst}(\mathbf{I}\langle 2 \rangle[\mathbf{T}\langle 2 \rangle[(x, i)]]) = \mathbf{T}\langle 1 \rangle[(x, \text{fst}(\mathbf{I}\langle 2 \rangle[i]))]$

Certain (highly improbable) events may occur that may lead to the games \mathbf{G}_{real} and $\mathbf{G}_{\text{realRO}}$ being distinguishable; we capture these using appropriate failure events, and prove that the difference of the probabilities of the outcomes of the games is bounded by the probability that these events occur. We use three different failure events \mathbf{bad}_1 , \mathbf{bad}_2 and \mathbf{bad}_3 .

1. \mathbf{bad}_1 is triggered whenever oracle $f_{\mathbf{bad}}$ samples a random value that is either \mathbf{IV} or has already been sampled for a distinct query before. The role of this event is twofold: on the one hand, if \mathbf{IV} is sampled as a random value, then there could exist a complete chain in \mathbf{T} that is a suffix of another complete chain in \mathbf{T} as illustrated in the first example of Figure 5.6 (here $\mathbf{T}[x_2, y_2] = \mathbf{IV}$). The problem is that oracle F_q in the game \mathbf{G}_{real} will generate the same values for the two messages corresponding to those two chains, while F_q in the game $\mathbf{G}_{\text{ideal}}$ most likely will not. On the other hand, if a sampled value has been sampled for another query before, then there could exist two complete chains in \mathbf{T} that collide at some point and are identical from that point on as illustrated in the second example of Figure 5.6. Again the two corresponding messages would yield the same answer in \mathbf{G}_{real} but most likely not in $\mathbf{G}_{\text{ideal}}$ on queries to F_q .
2. \mathbf{bad}_2 is triggered whenever oracle $f_{\mathbf{bad}}$ samples a random value that has already been used as a chaining value in a previous query. This means that this query may be part of an MD-chain of which the distinguisher has already queried later points in the chain, which should never happen. The event also captures that no fixed-points (i.e. entries of the form $\mathbf{T}[x, y] = y$) should be sampled.
3. \mathbf{bad}_3 is triggered whenever a chaining value y in a query has already been sampled as a random value and is in the range of \mathbf{T} for some previous query (x', y') , but (x', y') does not appear in the domain of \mathbf{T}' and (x', y') is not the last element of a complete chain in \mathbf{T} . Intuitively, this means that y was never returned by f_q or F_q and hence the distinguisher managed to guess a previously randomly sampled value which should be unknown to it.




Figure 5.6: Two examples illustrating the necessity of event \mathbf{bad}_1

Event \mathbf{bad}_1 gives us two important invariants, namely

1. $\text{Injective}(\mathbf{T}\langle 1 \rangle)$ (see Definition 6.1 for the straightforward definition of this predicate)
2. $\text{IV} \notin \mathbf{T}\langle 1 \rangle$

These two conditions are the *well-formedness* conditions of table $\mathbf{T}\langle 1 \rangle$, and they play an important role for many lemmas described in Chapter 6. Before we get to the more involved invariants, we need to derive a few more rather technical invariants which hold up to some failure event, or are related in some other way to the failure events.

3. $\text{Injective_fst}(\mathbf{I}\langle 2 \rangle)$
4. $\forall i. (\text{IV}, i) \notin \mathbf{R}\langle 2 \rangle$
5. $\forall (x, y) \in \text{dom}(\mathbf{T}\langle 1 \rangle). y \in \mathbf{Y}\langle 1 \rangle$
6. $\forall z \in \text{ran}(\mathbf{T}\langle 1 \rangle). z \in \mathbf{Z}\langle 1 \rangle$
7. $\forall (z, i) \in \text{ran}(\mathbf{R}\langle 2 \rangle). z \in \mathbf{Z}\langle 1 \rangle$
8. $\forall i \in \text{dom}(\mathbf{I}\langle 2 \rangle). \text{fst}(\mathbf{I}\langle 2 \rangle[i]) \in \mathbf{Z}\langle 1 \rangle$
9. $\forall z \in \mathbf{Z}\langle 1 \rangle. \exists i. i \in \text{dom}(\mathbf{I}\langle 2 \rangle) \wedge \text{fst}(\mathbf{I}\langle 2 \rangle[i]) = z$
10. $\forall (x, y) \in \text{dom}(\mathbf{T}\langle 1 \rangle). (x, y) \in \text{dom}(\mathbf{T}'\langle 1 \rangle) \vee y = \text{IV} \vee y \in \text{ran}(\mathbf{T}\langle 1 \rangle)$
11. $\forall x, y, i. (x, y) \in \text{dom}(\mathbf{T}\langle 1 \rangle) \wedge (x, y) \notin \text{dom}(\mathbf{T}'\langle 1 \rangle) \wedge i \in \text{dom}(\mathbf{I}\langle 2 \rangle) \wedge \text{fst}(\mathbf{I}\langle 2 \rangle[i]) = y \wedge (x, i) \notin \text{dom}(\mathbf{T}\langle 2 \rangle) \Rightarrow \text{findseq}(x, y, \mathbf{T}\langle 1 \rangle) \neq \text{None}$
12. $\forall i \in \text{dom}(\mathbf{I}\langle 2 \rangle). \text{set_bad3}(\text{fst}(\mathbf{I}\langle 2 \rangle[i]), \mathbf{T}'\langle 1 \rangle, \mathbf{T}\langle 1 \rangle) \Leftrightarrow \text{snd}(\mathbf{I}\langle 2 \rangle[i])$

The predicate Injective_fst is defined as

$$\text{Injective_fst}(\mathbf{I}) := \forall i, j. i \in \text{dom}(\mathbf{I}) \Rightarrow j \in \text{dom}(\mathbf{I}) \Rightarrow \text{fst}(\mathbf{I}[i]) = \text{fst}(\mathbf{I}[j]) \Rightarrow i = j.$$

Now, we show in `EasyCrypt` an invariant of elementary importance. Intuitively, we want to say that the distinguisher can only make queries to f_q “in order”. That is, it can only make queries associated with an MD-chain by first querying for the first chaining value, then for the second, and so on, since the intermediate chaining values are always unknown to it until it has made a query to f_q for the preceding element in the MD-chain. The simulator checks for this ordering in game G_{real} by using the predicate `set_bad3` (see Definition 6.16); in game G_{realRO} , it uses the map \mathbf{I} to iterate over all preceding chaining values. Furthermore the failure event `bad2` in both games ensures that a randomly sampled value was not accidentally already used as a chaining value in a previous query from the distinguisher to f_q , since this may otherwise destroy the correct ordering of a possible chain (which the simulator might not notice through the use of event `bad3`).

Knowing this, we now turn our attention to the three loops in oracle F_q . Indeed, although oracle F_q eventually behaves as a simple relay for procedure F , it still samples random

values eagerly as was the case in game $\mathbf{G}_{\text{real}'}$. The first loop in F_q reconstructs chaining values that may already have been sampled by previous queries of the distinguisher to f_q and are therefore in the domain of \mathbf{T}' . Since the distinguisher can only make queries that belong to an MD-chain in order, these values can only occur in the prefix of such a chain. The second and third loops then take care of the rest of the MD-chain, i.e. the suffix not yet queried by the distinguisher, except for the very last element (which is always taken care of by a call to F at the end). The purpose of the second loop is to re-use chaining values that might already have been sampled by a previous call to F_q , e.g. because the padding of the input from a previous call and the padding of the current input have the same prefix, or simply because the same query was made more than once. Finally, the third loop samples fresh chaining values if needed. Lastly, observe that in the case that *found* is set to true in the procedure f_q of game $\mathbf{G}_{\text{realRO}}$, the corresponding chaining value pre-sampled eagerly by the third loop in F_q is re-used by reading it from the map \mathbf{I} , and the *tainted* flag is updated in map \mathbf{I} to reflect that the current value is no longer unknown to the distinguisher. Thus f_q and F_q use consistent intermediate chaining values, as is the case in game $\mathbf{G}_{\text{real}'}$, where both procedures call oracle f_{bad} (resp. f).

The invariant we derive to express that the simulator can only make queries in order is

$$\text{Claim5}(\mathbf{T}'\langle 1 \rangle, \mathbf{T}\langle 1 \rangle).$$

The predicate **Claim5**, named thus in reminiscence of the *Claim 5* from the proof of ind-differentiability of the prefix-free Merkle-Damgård construction appearing in the extended version of Coron's paper [39], is defined as follows.

Definition 5.3 (Predicate **Claim5**). Let \mathbf{T} and \mathbf{T}' be two maps of type $\{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. Then

$$\begin{aligned} \text{Claim5}(\mathbf{T}', \mathbf{T}) := & \forall (x, y), (x', y'). (x', y') \in \text{dom}(\mathbf{T}) \wedge (x, y) \in \text{dom}(\mathbf{T}') \wedge \\ & \mathbf{T}[(x', y')] = y \wedge \text{findseq}(x', y', \mathbf{T}) = \text{None} \Rightarrow (x', y') \in \text{dom}(\mathbf{T}'). \end{aligned}$$

Intuitively, this means that whenever we have some element (x, y) in the domain of map \mathbf{T}' that succeeds an element (x', y') of an MD-chain in \mathbf{T} , this preceding element (x', y') must also be in the domain of \mathbf{T}' , except if (x', y') is already the last element of a complete chain in \mathbf{T} (in that case, y is known to the distinguisher). The latter restriction is expressed by the premise $\text{findseq}(x, y, \mathbf{T}) = \text{None}$. We refer to Section 6.1.6 for a more detailed discussion of this predicate.

The restriction that the distinguisher may only query chaining values in order has one essential implication. Namely, it means that any chain in the table \mathbf{T}' maintained by the simulator will always be completed by the last element of that chain (both in the games $\mathbf{G}_{\text{real}'}$ and $\mathbf{G}_{\text{realRO}}$). Therefore, in the game $\mathbf{G}_{\text{realRO}}$, upon a query $f_q(x, y)$ where (x, y) is the last element of an MD-chain, the simulator will always detect that this query corresponds to the last element of this MD-chain (by the completeness of the function findseq , see Axiom 6.14). Hence, for any query $F_q(m)$ that the distinguisher could make, the simulator

will always be consistent with F_q 's answer whenever the distinguisher tries to construct the MD-chain associated to m by repeatedly calling f_q . We formalize this in a final invariant:

$$\text{RinT}(\mathbf{R}\langle 2 \rangle, \mathbf{T}\langle 1 \rangle),$$

and we define the predicate RinT as follows.

Definition 5.4 (Predicate RinT). Let $\mathbf{T} : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $\mathbf{R} : \{0, 1\}^* \rightarrow \{0, 1\}^n \times \mathbb{N}$ be two maps. Then

$$\begin{aligned} \text{RinT}(\mathbf{R}, \mathbf{T}) := & (\forall m \in \text{dom}(\mathbf{R}). \exists (x, y). \text{findseq}(x, y, \mathbf{T}) \neq \text{None} \wedge \\ & m = \pi_{\text{unpad}}(\text{mapfst}(\pi_{\text{findseq}(x, y, \mathbf{T})} \parallel [x])) \wedge \\ & (x, y) \in \text{dom}(\mathbf{T})) \wedge \\ & (\forall (x, y). \text{findseq}(x, y, \mathbf{T}) \neq \text{None} \Rightarrow \\ & \pi_{\text{unpad}}(\text{mapfst}(\pi_{\text{findseq}(x, y, \mathbf{T})} \parallel [x])) \in \text{dom}(\mathbf{R}) \Rightarrow \\ & (x, y) \in \text{dom}(\mathbf{T})) \wedge \\ & (\forall (x, y). \text{findseq}(x, y, \mathbf{T}) \neq \text{None} \Rightarrow \\ & (x, y) \in \text{dom}(\mathbf{T}) \Rightarrow \\ & \pi_{\text{unpad}}(\text{mapfst}(\pi_{\text{findseq}(x, y, \mathbf{T})} \parallel [x])) \in \text{dom}(\mathbf{R}) \wedge \\ & \text{fst}(\mathbf{R}[\pi_{\text{unpad}}(\text{mapfst}(\pi_{\text{findseq}(x, y, \mathbf{T})} \parallel [x])]) = \mathbf{T}[(x, y)]). \end{aligned}$$

Intuitively, the invariant says that the map $\mathbf{R}\langle 2 \rangle$ contains every message associated to every complete chain in map $\mathbf{T}\langle 1 \rangle$, and nothing else. More precisely, for every message m in the domain of $\mathbf{R}\langle 2 \rangle$, there is a corresponding complete chain of the form $c \parallel [(x, y)]$ in the map $\mathbf{T}\langle 1 \rangle$ with $(x, y) \in \text{dom}(\mathbf{T}\langle 1 \rangle)$. Conversely, for every complete chain of the form $c \parallel [(x, y)]$ in the map $\mathbf{T}\langle 1 \rangle$ with $(x, y) \in \text{dom}(\mathbf{T}\langle 1 \rangle)$, its associated message m is in the domain of $\mathbf{R}\langle 2 \rangle$. Thus, in oracle f_q , in the case where $\text{findseq}((x, y), \mathbf{T}') \neq \text{None}$, the transformation can be proven valid because EasyCrypt can infer from the above invariant that $(x, y) \in \text{dom}(\mathbf{T}\langle 1 \rangle)$ in game $G_{\text{real}'}$ if and only if $m \in \text{dom}(\mathbf{R}\langle 2 \rangle)$ in game G_{realRO} .

Finally, we can show in EasyCrypt the following pRHL judgment:

$$\vdash G_{\text{real}'} \sim G_{\text{realRO}} : \text{true} \Longrightarrow =\{\mathbf{bad}_1, \mathbf{bad}_2, \mathbf{bad}_3\} \wedge \neg(\mathbf{bad}_1 \vee \mathbf{bad}_2 \vee \mathbf{bad}_3)\langle 1 \rangle \Rightarrow =\{b\}$$

We now apply Theorem 4.1 to conclude that the advantage of the distinguisher in differentiating between the games $G_{\text{real}'}$ and G_{realRO} is upper bounded by the probability of any of the failure events occurring in game G_{realRO} , i.e.

$$|\Pr[G_{\text{real}'} : b] - \Pr[G_{\text{realRO}} : b]| \leq \Pr[G_{\text{realRO}} : \mathbf{bad}_1 \vee \mathbf{bad}_2 \vee \mathbf{bad}_3]. \quad (5.3)$$

We are left to bound the probability of the failure events \mathbf{bad}_1 , \mathbf{bad}_2 and \mathbf{bad}_3 . Here, the bounds for \mathbf{bad}_1 and \mathbf{bad}_2 can be computed in the same fashion. Both events are triggered when a freshly sampled random value is already in a list of size at most q (list \mathbf{Z} for event \mathbf{bad}_1 , and list \mathbf{Y} for event \mathbf{bad}_2). Therefore, the probability of either event is

upper-bounded by $q/2^n$, for each of the possible q queries. To derive this in **EasyCrypt**, we first observe that the code which samples random values and raises the events **bad**₁ and **bad**₂ is the same in the oracles f_q , F_q and F of game $\mathbf{G}_{\text{realRO}}$, namely

$$\begin{aligned} z &\leftarrow^{\$} \{0, 1\}^n; \\ \mathbf{bad}_1 &\leftarrow \mathbf{bad}_1 \vee z \in \mathbf{Z}; \\ \mathbf{Z} &\leftarrow z::\mathbf{Z}; \mathbf{Y} \leftarrow y::\mathbf{Y}; \\ \mathbf{bad}_2 &\leftarrow \mathbf{bad}_2 \vee z \in \mathbf{Y}; \end{aligned}$$

In a meta-game $\hat{\mathbf{G}}_{\text{realRO}}$, we outline this code fragment in all three procedures into a new procedure *sample_z*. We can then apply Lemma 4.2 in this game by taking $u = q/2^n$ and $i = |\mathbf{Z}| - 1$ (resp. $i = |\mathbf{Y}|$) to bound the probability of event **bad**₁ (resp. **bad**₂). We obtain

$$\Pr[\mathbf{G}_{\text{realRO}} : \mathbf{bad}_1] = \Pr[\hat{\mathbf{G}}_{\text{realRO}} : \mathbf{bad}_1] \leq \frac{q^2}{2^n} \quad (5.4)$$

$$\Pr[\mathbf{G}_{\text{realRO}} : \mathbf{bad}_2] = \Pr[\hat{\mathbf{G}}_{\text{realRO}} : \mathbf{bad}_2] \leq \frac{q^2}{2^n} \quad (5.5)$$

The bounding of event **bad**₃ is more complex, since it relates to values that have already been sampled eagerly, and therefore the distribution of those values is not locally known. Hence, we leave the bounding of this event to the next sections, until after the step to a scenario where values are sampled lazily. Nevertheless, we already reveal here that the upper bound that we will ultimately obtain for event **bad**₃ in game $\mathbf{G}_{\text{realRO}}$ is equal to those for **bad**₁ and **bad**₂, namely $q^2/2^n$. Summing up over Equations 5.3, 5.4, 5.5 and 5.13, we finally get:

$$|\Pr[\mathbf{G}_{\text{real}'} : b] - \Pr[\mathbf{G}_{\text{realRO}} : b]| \leq \frac{3q^2}{2^n} \quad (5.6)$$

5.3 From $\mathbf{G}_{\text{realRO}}$ to $\mathbf{G}_{\text{idealEager}}$

The transformation considered in this section is an easy one, since it simply serves as a preparation for the next transformation, which will be concerned with the step from an eager sampling scenario to a lazy sampling scenario. In the game $\mathbf{G}_{\text{idealEager}}$, defined in Figure 5.7, the following changes are introduced:

1. The main body of the game $\mathbf{G}_{\text{realRO}}$ is instrumented with a loop which re-samples chaining values that are unknown to the adversary;
2. the failure events **bad**₁, **bad**₂ and **bad**₃ are removed, and the resulting code is simplified;
3. the oracle $F(m, y)$ is transformed into an oracle $F(m)$;
4. a new event **bad**₄ is introduced that will be used to bound the probability of event **bad**₃ in game $\mathbf{G}_{\text{realRO}}$.

<p>Game $G_{\text{idealEager}}$:</p> <p>Game $G_{\text{idealLazy}}$:</p> <p>$\mathbf{q}_f \leftarrow 0;$ $\mathbf{q}'_f \leftarrow 1;$ $\mathbf{T}, \mathbf{T}', \mathbf{T}'_i, \mathbf{R}, \mathbf{I} \leftarrow \emptyset;$ $\mathbf{I}[0] \leftarrow (\text{IV}, \text{false});$ $\mathbf{Y}' \leftarrow \text{nil};$ $\text{bad}_4 \leftarrow \text{false};$ $l \leftarrow 0;$</p> <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p>$l \leftarrow 0;$ while $l < \mathbf{q}'_f$ do if $\text{snd}(\mathbf{I}[l])$ then $z \stackrel{\\$}{\leftarrow} \{0, 1\}^n;$ $\mathbf{I}[l] \leftarrow (z, \text{true});$ $l \leftarrow l + 1;$</p> </div> <p>$b \leftarrow \mathcal{D}^{F_q, f_q}();$</p> <div style="border: 1px dashed black; padding: 5px; margin: 5px 0;"> <p>$l \leftarrow 0;$ while $l < \mathbf{q}'_f$ do if $\text{snd}(\mathbf{I}[l])$ then $z \stackrel{\\$}{\leftarrow} \{0, 1\}^n;$ $\mathbf{I}[l] \leftarrow (z, \text{true});$ $l \leftarrow l + 1;$</p> </div> <p>return b</p>	<p>Oracle $F_q(m)$:</p> <p>$xs \leftarrow \text{pad}(m); y \leftarrow \text{IV};$ $i \leftarrow 0;$ if $(0 < \mathbf{q}'_f \wedge$ $\mathbf{q}_f + xs \leq q)$ then $\mathbf{q}_f \leftarrow \mathbf{q}_f + xs ;$ while $xs > 1 \wedge$ $(\text{hd}(xs), y) \in \text{dom}(\mathbf{T}')$ do $i \leftarrow \mathbf{T}'_i[\text{hd}(xs), y];$ $y \leftarrow \mathbf{T}'[\text{hd}(xs), y];$ $xs \leftarrow \text{tl}(xs);$ while $xs > 1 \wedge$ $(\text{hd}(xs), i) \in \text{dom}(\mathbf{T})$ do $i \leftarrow \mathbf{T}[\text{hd}(xs), i];$ $xs \leftarrow \text{tl}(xs);$ while $xs > 1$ do $z \stackrel{\\$}{\leftarrow} \{0, 1\}^n;$ $\mathbf{T}[\text{hd}(xs), i] \leftarrow \mathbf{q}'_f;$ $\mathbf{I}[\mathbf{q}'_f] \leftarrow (z, \text{true});$ $i \leftarrow \mathbf{q}'_f;$ $\mathbf{q}'_f \leftarrow \mathbf{q}'_f + 1;$ $xs \leftarrow \text{tl}(xs);$ $y \leftarrow \text{fst}(F(m));$ return y</p>	<p>Oracle $F(m)$:</p> <p>if $m \notin \text{dom}(\mathbf{R})$ then $z \stackrel{\\$}{\leftarrow} \{0, 1\}^n;$ $\mathbf{R}[m] \leftarrow (z, \mathbf{q}'_f)$ $\mathbf{I}[\mathbf{q}'_f] \leftarrow (z, \text{false})$ $\mathbf{q}'_f \leftarrow \mathbf{q}'_f + 1;$ return $\mathbf{R}[m]$</p>	<p>Oracle $f_q(x, y)$:</p> <p>if $\mathbf{q}_f + 1 \leq q$ then if $(0 < \mathbf{q}'_f \wedge$ $(x, y) \notin \text{dom}(\mathbf{T}'))$ then $xs \leftarrow \text{findseq}((x, y), \mathbf{T}')$ if $xs \neq \text{None}$ then $m \leftarrow \pi_{\text{unpad}}(\text{mapfst}(\pi_{xs})\ xs);$ $(z, j) \leftarrow F(m);$ $\mathbf{T}'[x, y] \leftarrow z; \mathbf{T}'_i[x, y] \leftarrow j;$ else $\text{found} \leftarrow \text{false}; j, k' \leftarrow 0;$ while $(k' < \mathbf{q}'_f \wedge \neg \text{found})$ do if $(\mathbf{I}[k'] = (y, \text{false}) \wedge$ $(x, k') \in \text{dom}(\mathbf{T}) \wedge$ $\text{snd}(\mathbf{I}[\mathbf{T}[x, k']]) \wedge$ $k' < \mathbf{T}[x, k'] \wedge$ $\mathbf{T}[x, k'] < \mathbf{q}'_f)$ then $\text{found} \leftarrow \text{true}; j \leftarrow \mathbf{T}[x, k'];$ else $k' \leftarrow k' + 1;$ if found then <div style="border: 1px solid black; padding: 2px; display: inline-block;"> $z \leftarrow \text{fst}(\mathbf{I}[j]);$ </div> <div style="border: 1px dashed black; padding: 2px; display: inline-block; margin-left: 10px;"> $z \stackrel{\\$}{\leftarrow} \{0, 1\}^n;$ </div> $\text{bad}_4 \leftarrow \text{bad}_4 \vee z \in \mathbf{Y}';$ $\mathbf{I}[j] \leftarrow (z, \text{false});$ $\mathbf{T}'[x, y] \leftarrow z; \mathbf{T}'_i[x, y] \leftarrow j;$ else $z \stackrel{\\$}{\leftarrow} \{0, 1\}^n;$ $\mathbf{I}[\mathbf{q}'_f] \leftarrow (z, \text{false});$ $\mathbf{T}'[x, y] \leftarrow z;$ $\mathbf{T}'_i[x, y] \leftarrow \mathbf{q}'_f;$ $\mathbf{q}'_f \leftarrow \mathbf{q}'_f + 1;$ $\mathbf{Y}' \leftarrow y::\mathbf{Y}';$ $z \leftarrow \mathbf{T}'[x, y]; \mathbf{q}_f \leftarrow \mathbf{q}_f + 1;$ else $z \leftarrow \text{IV};$ return z</p>
--	--	---	--

Figure 5.7: The games $G_{\text{idealEager}}$ and $G_{\text{idealLazy}}$

The loop which re-samples chaining values in the main body of game $G_{\text{idealEager}}$ will be discussed in more depth in the next section, since it is important for the transition from eager to lazy sampling. Here, simply observe that the loop has no effect in this game, except that it increments the value of l (unknown to the distinguisher) once: at the point where the loop is executed, we have $\mathbf{q}'_f = 1$ and $\text{snd}(\mathbf{I}[0]) = \text{false}$. Therefore there is only one iteration and the guard of the if-statement is not fulfilled.

Next, we remove events $\text{bad}_{\{1,2,3\}}$ and their related lists \mathbf{Y} and \mathbf{Z} . Because of this, we can also merge the two branches under the “if $\text{found_bad3} \dots$ ”-statement of the simulator f_q in game G_{realRO} . Additionally, we transform the oracle $F(m, y)$ into an oracle $F(m)$. Indeed, the only reason why we needed to pass y as an argument to F in the previous game was so that we could appropriately set event bad_2 – by removing this event we can now get rid of the superfluous parameter y . Lastly, because F does not need this parameter any longer, we do not need to keep track of the intermediate chaining values y in the second

and third loops of oracle F_q ; it becomes even more clear now that F_q behaves as a relay and that the three loops have no effect on its return value.

Similarly to the transformation described in Section 5.1, we observe that the removal or introduction of failure events (unknown to the distinguisher) does not otherwise influence the behavior of either the procedures f_q or F_q .

Finally, because we still have to bound the probability of event \mathbf{bad}_3 in game $\mathbf{G}_{\text{realRO}}$, we introduce a new event \mathbf{bad}_4 along with an associated list \mathbf{Y}' (named so to distinguish it from the list \mathbf{Y}). This event plays a similar role to event \mathbf{bad}_3 used in the previous game, albeit not the same. It ensures that a random value pre-sampled by procedure F_q and re-used by the distinguisher has not already been used as a chaining value in a previous query to f_q ; event \mathbf{bad}_3 did the same in game $\mathbf{G}_{\text{realRO}}$, but for random values sampled by the simulator itself, not for those pre-sampled by F_q . Therefore, event \mathbf{bad}_4 on its own is not enough to bound the probability of event \mathbf{bad}_3 . We need it because we want to bound the probability of \mathbf{bad}_3 by the probability of a more general event in game $\mathbf{G}_{\text{idealEager}}$: We show that if \mathbf{bad}_3 occurs in game $\mathbf{G}_{\text{realRO}}$, then either \mathbf{bad}_4 is triggered in game $\mathbf{G}_{\text{idealEager}}$, or there exists a positive value $i \leq \mathbf{q}'_f$ such that $\mathbf{I}\langle 2\rangle[i] = (y, \text{true})$ and $y \in \mathbf{Y}\langle 2\rangle$. To achieve this, we first define the following predicate.

Definition 5.5 (Predicate \mathbf{I}_\exists). Let $\mathbf{q}'_f \in \mathbb{N}$ be an integer, $\mathbf{I} : \mathbb{N} \rightarrow \{0, 1\}^n \times \mathbb{B}$ a map and $\mathbf{Y}' : \{0, 1\}^n$ list a list of chaining values. Then

$$\mathbf{I}_\exists(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}') := \exists i. 0 \leq i \leq \mathbf{q}'_f \wedge \text{snd}(\mathbf{I}[i]) \wedge \text{fst}(\mathbf{I}[i]) \in \mathbf{Y}'.$$

Intuitively, this means that there is a chaining value appearing in map \mathbf{I} which should be unknown to the distinguisher, but the distinguisher used it as part of a query to f_q – this should be impossible. We now show in **EasyCrypt**:

$$\vdash \mathbf{G}_{\text{realRO}} \sim \mathbf{G}_{\text{idealEager}} : \text{true} \implies =\{b\} \wedge (\mathbf{bad}_3\langle 1\rangle \Rightarrow (\mathbf{bad}_4\langle 2\rangle \vee \mathbf{I}_\exists(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')\langle 2\rangle))$$

This yields the following (in-)equations:

$$\Pr[\mathbf{G}_{\text{realRO}} : b] = \Pr[\mathbf{G}_{\text{idealEager}} : b] \tag{5.7}$$

$$\Pr[\mathbf{G}_{\text{realRO}} : \mathbf{bad}_3] \leq \Pr[\mathbf{G}_{\text{idealEager}} : \mathbf{bad}_4 \vee \mathbf{I}_\exists(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')] \tag{5.8}$$

To derive the above **pRHL** judgment, for technical reasons we need to show a few generic invariants that hold in game $\mathbf{G}_{\text{realRO}}$. We list them here for the sake of completeness.

1. $0 < \mathbf{q}'_f\langle 1\rangle$
2. $\forall m \in \text{dom}(\mathbf{R}\langle 1\rangle). \text{snd}(\mathbf{R}\langle 1\rangle[m]) < \mathbf{q}'_f$
3. $\forall (x, y) \in \text{dom}(\mathbf{T}'\langle 1\rangle). (x, y) \in \text{dom}(\mathbf{T}'_i\langle 1\rangle)$
4. $\forall (x, y) \in \text{dom}(\mathbf{T}'_i\langle 1\rangle). \mathbf{T}'_i\langle 1\rangle[x, y] < \mathbf{q}'_f$
5. $\forall (x, i) \in \text{dom}(\mathbf{T}\langle 1\rangle). i < \mathbf{T}\langle 1\rangle[x, i] < \mathbf{q}'_f$

We postpone the computation of the probability of $\mathbf{bad}_4 \vee \mathbf{I}_\exists(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')$ in game $\mathbf{G}_{\text{idealEager}}$ to the next section, since again this will be easier when the values are sampled lazily and their distribution is locally known.

5.4 From $G_{\text{idealEager}}$ to $G_{\text{idealLazy}}$

In this step, we will make the transition from a simulator which re-uses eagerly sampled random values to a simulator which lazily samples its own random values upon fresh queries. In the new game $G_{\text{idealLazy}}$, defined in Figure 5.7, we introduce the following two changes:

1. The simulator f_q now samples a new random value in the case where *found* is true, instead of re-using the corresponding value in map \mathbf{I} ;
2. the loop in the main body introduced in the last section is swapped with the call to the distinguisher \mathcal{D} .

We have investigated the difference between eager and lazy sampling in Section 4.5.2. At this point, we highlight that the initial game G_{real} samples values eagerly, while the final game G_{ideal} samples values lazily. To clarify this, let m be a message with $\text{pad}(m) = [x_1, \dots, x_n]$, and assume for simplicity that the distinguisher has made no calls to oracles f_q or F_q yet in the two games. Now, observe the following fundamental difference: If the distinguisher makes the query $F_q(m)$ in game G_{real} , the FIL-RO f is called repeatedly before an answer is returned, and therefore n random values are sampled. If the distinguisher makes the query $F_q(m)$ in game G_{ideal} , then F_q calls the VIL-RO F once, and therefore only a single random value is sampled. Now, if the distinguisher queries $f_q(x_1, \mathbf{IV})$ in game G_{real} , oracle f_q will query f and therefore the corresponding random value which was pre-sampled by the call $F_q(m)$ will be used. In game G_{ideal} , the simulator f_q will sample a new random value instead. That is, in game G_{real} , some random values may be sampled *earlier* than in game G_{ideal} . Additionally, some values which were randomly sampled in game G_{real} may not be sampled at all in game G_{ideal} – the distinguisher needs not reconstruct the entire MD-chain associated to m by queries to f_q . Notwithstanding, the delaying of this sampling of random values does not make any difference from the distinguisher’s perspective; the cryptographical argument is that these values are *uniformly distributed and independent of the adversary’s view*. Indeed, the intermediate chaining values generated in a Merkle-Damgård iteration remain unknown to the distinguisher until it queries the oracle f_q for them, at which point they are sampled uniformly at random if needed. Therefore we see that a sequence of games from G_{real} to G_{ideal} must involve a transition of lazy/eager sampling, as described in Section 4.5.2. The purpose of the present transformation is to tackle this particular issue.

The main difference between the games $G_{\text{idealEager}}$ and $G_{\text{idealLazy}}$ considered here is that the oracle f_q in game $G_{\text{idealEager}}$ re-uses values that were eagerly pre-sampled by the oracle F_q , while the oracle f_q in game $G_{\text{idealLazy}}$ re-samples those random values. Intuitively, this is sound because these eagerly sampled values have not been revealed to the distinguisher; we know this because at that point it holds that $\text{snd}(\mathbf{I}[j]) = \text{true}$ (the second component of the pairs in map \mathbf{I} keeps track of whether a value is unknown to the distinguisher). Now, we would like to show that the re-sampling of these eagerly sampled values, unknown to the distinguisher, preserves the semantics of the program.

However, the values pre-sampled by F_q are not locally known in the context of the procedure f_q , such that it is difficult if not impossible to argue about them in the local context of procedure f_q . This is where the loop of the main body (depicted in the inner boxes of the two games in Figure 5.7) comes into play. This loop eagerly samples all random values for \mathbf{I} in the game $\mathbf{G}_{\text{idealEager}}$ (i.e., it is executed *before* the call to the distinguisher). In game $\mathbf{G}_{\text{idealLazy}}$, this loop is swapped with the call to the distinguisher (i.e., it is executed *after* the call to the distinguisher). In game $\mathbf{G}_{\text{idealLazy}}$, the purpose of the loop is to sample all random values that have not been lazily sampled during the execution of the distinguisher. For readability, let `loop` denote the code of the loop. Focusing now on the pieces of code `loop; GidealEager.Fq(m)` and `GidealLazy.Fq(m); loop` (respectively `loop; GidealEager.fq(x, y)` and `GidealLazy.fq(x, y); loop`), we can address all randomly sampled variables locally, and by showing that this loop swaps with the calls to the eager and lazy oracles in a semantic-preserving way, we prove that the re-sampled values are independent of the view of an adversary with access to these oracles. This is the main trick in a transition of lazy/eager sampling; referring to the discussion in Section 4.5.2, the condition used here is the condition $\text{snd}(\mathbf{I}[j]) = \text{false}$ for any value j in the domain of \mathbf{I} (meaning that the value $\text{fst}(\mathbf{I}[j])$ has been revealed to the distinguisher), and the statement $S_{\hat{g}}$ which re-samples values unknown to the distinguisher is `loop`.

Concretely, we show

$$\vdash \begin{array}{l} l \leftarrow 0; \\ \text{while } l < \mathbf{q}'_f \text{ do} \\ \quad \text{if } \text{snd}(\mathbf{I}[l]) \text{ then} \\ \quad \quad z \stackrel{\$}{\leftarrow} \{0, 1\}^n; \\ \quad \quad \mathbf{I}[l] \leftarrow (z, \text{true}); \\ \quad \quad l \leftarrow l + 1; \\ r \leftarrow \mathbf{G}_{\text{idealEager}}.F_q(m); \end{array} \sim \begin{array}{l} r \leftarrow \mathbf{G}_{\text{idealLazy}}.F_q(m); \\ l \leftarrow 0; \\ \text{while } l < \mathbf{q}'_f \text{ do} \\ \quad \text{if } \text{snd}(\mathbf{I}[l]) \text{ then} \\ \quad \quad z \stackrel{\$}{\leftarrow} \{0, 1\}^n; \\ \quad \quad \mathbf{I}[l] \leftarrow (z, \text{true}); \\ \quad \quad l \leftarrow l + 1; \end{array} : = \{m, l\} \wedge =\Phi \implies = \{r, l\} \wedge =\Phi$$

and

$$\vdash \begin{array}{l} l \leftarrow 0; \\ \text{while } l < \mathbf{q}'_f \text{ do} \\ \quad \text{if } \text{snd}(\mathbf{I}[l]) \text{ then} \\ \quad \quad z \stackrel{\$}{\leftarrow} \{0, 1\}^n; \\ \quad \quad \mathbf{I}[l] \leftarrow (z, \text{true}); \\ \quad \quad l \leftarrow l + 1; \\ r \leftarrow \mathbf{G}_{\text{idealEager}}.f_q(x, y); \end{array} \sim \begin{array}{l} r \leftarrow \mathbf{G}_{\text{idealLazy}}.f_q(x, y); \\ l \leftarrow 0; \\ \text{while } l < \mathbf{q}'_f \text{ do} \\ \quad \text{if } \text{snd}(\mathbf{I}[l]) \text{ then} \\ \quad \quad z \stackrel{\$}{\leftarrow} \{0, 1\}^n; \\ \quad \quad \mathbf{I}[l] \leftarrow (z, \text{true}); \\ \quad \quad l \leftarrow l + 1; \end{array} : = \{x, y, l\} \wedge =\Phi \implies = \{r, l\} \wedge =\Phi,$$

where $\Phi := \{\mathbf{q}_f, \mathbf{q}'_f, \mathbf{T}, \mathbf{T}', \mathbf{T}'_i, \mathbf{R}, \mathbf{I}, \mathbf{Y}', \text{bad}_4\}$ is the set of all global variables. From this, we immediately conclude

$$\vdash \begin{array}{l} l \leftarrow 0; \\ \text{while } l < \mathbf{q}'_f \text{ do} \\ \quad \text{if } \text{snd}(\mathbf{I}[l]) \text{ then} \\ \quad \quad z \stackrel{\$}{\leftarrow} \{0, 1\}^n; \\ \quad \quad \mathbf{I}[l] \leftarrow (z, \text{true}); \\ \quad \quad l \leftarrow l + 1; \\ b \leftarrow D^{\mathbf{G}_{\text{idealEager}}.f_q, \mathbf{G}_{\text{idealEager}}.F_q}(); \end{array} \sim \begin{array}{l} b \leftarrow D^{\mathbf{G}_{\text{idealLazy}}.f_q, \mathbf{G}_{\text{idealLazy}}.F_q}(); \\ l \leftarrow 0; \\ \text{while } l < \mathbf{q}'_f \text{ do} \\ \quad \text{if } \text{snd}(\mathbf{I}[l]) \text{ then} \\ \quad \quad z \stackrel{\$}{\leftarrow} \{0, 1\}^n; \\ \quad \quad \mathbf{I}[l] \leftarrow (z, \text{true}); \\ \quad \quad l \leftarrow l + 1; \end{array} : = \{l\} \wedge =\Phi \implies = \{b, l\} \wedge =\Phi.$$

An additional instruction $l \leftarrow 0$ situated before the loop was introduced along with said loop in $G_{\text{idealEager}}$ (see Figure 5.7). This is needed to satisfy the pre-condition $=\{l\}$ in the latter pRHL judgment, since we want to use this judgment to conclude equivalence of the two entire games. Indeed, using this it is now a trivial task for **EasyCrypt** to derive

$$\vdash G_{\text{idealEager}} \sim G_{\text{idealLazy}} : \text{true} \implies =\{b, \mathbf{bad}_4, \mathbf{I}, \mathbf{q}'_f, \mathbf{Y}'\}.$$

This implies that the games $G_{\text{idealEager}}$ and $G_{\text{idealLazy}}$ have the same probabilities of outcome, and the same probability of event $\mathbf{bad}_4 \vee I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')$:

$$\Pr [G_{\text{idealEager}} : b] = \Pr [G_{\text{idealLazy}} : b] \quad (5.9)$$

$$\Pr [G_{\text{idealEager}} : \mathbf{bad}_4 \vee I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')] = \Pr [G_{\text{idealLazy}} : \mathbf{bad}_4 \vee I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')] \quad (5.10)$$

The hardest part here is the proof of the above judgment for oracle f_q (i.e., the second judgment) in the case where $\text{found} = \text{true}$, since here $G_{\text{idealEager}}.f_q$ uses the chaining value $\mathbf{I}[j]$, whereas $G_{\text{idealLazy}}.f_q$ re-samples a fresh random value. In this case, observe the following. In the lazy program (the one on the right-hand side), since we have $\text{found} = \text{true}$, we know (from the while loop which was responsible for setting $\text{found} = \text{true}$) that $j \in \text{dom}(\mathbf{I})$ and $\text{snd}(\mathbf{I}[j]) = \text{true}$. Let us fix the value j at this point in the lazy program as \hat{j} . Now, since \mathbf{I} was not modified in this program before, we must already have had $\hat{j} \in \text{dom}(\mathbf{I})$ and $\text{snd}(\mathbf{I}[\hat{j}]) = \text{true}$ at the start of the lazy program. In the pRHL judgment, we have as a pre-condition that all global variables are equal at the start of the two programs, in particular \mathbf{I} . So, we must also have $\hat{j} \in \text{dom}(\mathbf{I})$ and $\text{snd}(\mathbf{I}[\hat{j}]) = \text{true}$ at the start of the eager program (the one on the left-hand side). Furthermore, we know – using an analogous argument – that $\hat{j} < \mathbf{q}'_f$ at the start of the eager program. Thus, the value $\mathbf{I}[\hat{j}]$ will be sampled at the start of the eager program by the loop which samples random values. We can “cancel out” this random sampling against the fresh random sampling performed in procedure f_q of the lazy program (the other random samplings in the loop at the start of the eager program can be canceled out against the random samplings in the loop at the end of the lazy program). Technically, we *derandomize* the two random samplings of $\mathbf{I}[\hat{j}]$ on both sides, that is, we assume that \hat{z} is a uniformly distributed random value and we replace these two samplings on both sides by deterministic assignments of the value \hat{z} . After the execution of both programs, we need to show that $\text{upd}(\text{upd}(\mathbf{I}\langle 1 \rangle, \hat{j}, (\hat{z}, \text{true})), \hat{j}, (\hat{z}, \text{false})) = \text{upd}(\mathbf{I}\langle 2 \rangle, \hat{j}, (\hat{z}, \text{false}))$, as the value $\mathbf{I}[\hat{j}]$ has been updated twice in the eager game, but only once in the lazy game. Note that we have $\mathbf{I}\langle 1 \rangle = \mathbf{I}\langle 2 \rangle$, so the only thing we need to do to finish this argument is to formalize an extensionality axiom:

Axiom 5.6 (Extensionality of \mathbf{I}). *Let $\mathbf{I}, \mathbf{I}' : \mathbb{N} \rightarrow \{0, 1\}^n \times \mathbb{B}$ be two maps. Then*

$$\begin{aligned} \mathbf{I} = \mathbf{I}' \Leftrightarrow & ((\forall i. i \in \text{dom}(\mathbf{I}) \Leftrightarrow i \in \text{dom}(\mathbf{I}')) \wedge \\ & (\forall i \in \text{dom}(\mathbf{I}). \mathbf{I}[i] = \mathbf{I}'[i])). \end{aligned}$$

Lastly, we want to bound the probability of event $\mathbf{bad}_4 \vee I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')$ in game $G_{\text{idealLazy}}$. We define a meta-game $G'_{\text{idealLazy}}$, whose only modification is that the loop which re-samples

values in the main body of the game now also sets an event \mathbf{bad}_4 . More precisely, we replace the instruction $z \xleftarrow{\$} \{0, 1\}^n$ in the loop of the main body with

$$\begin{aligned} z &\xleftarrow{\$} \{0, 1\}^n; \\ \mathbf{bad}_4 &\leftarrow \mathbf{bad}_4 \vee z \in \mathbf{Y}' \end{aligned}$$

Because of this modification, we can say the following: If $I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')$ holds at the end of game $G_{\text{idealLazy}}$, then clearly $I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')$ also holds at the end of $G'_{\text{idealLazy}}$. If $I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')$ holds at the end of game $G'_{\text{idealLazy}}$, then necessarily \mathbf{bad}_4 must have been set to true in this game, either by the simulator or later by the loop in the main body. Therefore, the event $I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')$ in game $G_{\text{idealLazy}}$ implies $\mathbf{bad}_4 = \text{true}$ in game $G'_{\text{idealLazy}}$. Further, clearly the event $\mathbf{bad}_4 = \text{true}$ in game $G_{\text{idealLazy}}$ also implies $\mathbf{bad}_4 = \text{true}$ in game $G'_{\text{idealLazy}}$. Thus we can easily derive in EasyCrypt that

$$\Pr [G_{\text{idealLazy}} : \mathbf{bad}_4 \vee I_{\exists}(\mathbf{q}'_f, \mathbf{I}, \mathbf{Y}')] \leq \Pr [G'_{\text{idealLazy}} : \mathbf{bad}_4]. \quad (5.11)$$

Finally, we use the same technique to bound the probability of event \mathbf{bad}_4 in game $G'_{\text{idealLazy}}$ that we also used to bound the probability of events \mathbf{bad}_1 and \mathbf{bad}_2 in game G_{realRO} . We outline the code fragment $z \xleftarrow{\$} \{0, 1\}^n; \mathbf{bad}_4 \leftarrow \mathbf{bad}_4 \vee z \in \mathbf{Y}'$ in both the procedures f_q and the main body of $G'_{\text{idealLazy}}$ into an oracle sample_z to obtain a new meta-game $\hat{G}'_{\text{idealLazy}}$. We apply Lemma 4.2 in this game by taking $u = q/2^n$ and $i = |\mathbf{Y}'|$, and obtain:

$$\Pr [G'_{\text{idealLazy}} : \mathbf{bad}_4] = \Pr [\hat{G}'_{\text{idealLazy}} : \mathbf{bad}_4] \leq \frac{q^2}{2^n} \quad (5.12)$$

Pulling Equations 5.8, 5.10, 5.11 and 5.12 together, we finally obtain the bound for the probability of event \mathbf{bad}_3 in the game G_{realRO} :

$$\Pr [G_{\text{realRO}} : \mathbf{bad}_3] \leq \Pr [\hat{G}'_{\text{idealLazy}} : \mathbf{bad}_4] \leq \frac{q^2}{2^n} \quad (5.13)$$

5.5 From $G_{\text{idealLazy}}$ to G_{ideal}

Finally, we are only left to show that the probabilities of the outcomes of the game $G_{\text{idealLazy}}$ and of the final game G_{ideal} (see Figure 5.1) are equal. This only amounts to some rather simple code cleanup:

1. The loop which re-samples values in the main body of $G_{\text{idealLazy}}$ is removed;
2. the three loops in procedure F_q are also removed;
3. in the simulator, the loop which looks for values pre-sampled in map \mathbf{I} and may set found to true is removed;
4. the global variables \mathbf{bad}_4 , \mathbf{Y}' , \mathbf{q}'_f , \mathbf{I} , \mathbf{T}'_i , and \mathbf{T} and all associated instructions are removed;

5. the two branches under the conditional statement “if *found* . . .” in the simulator are merged.

The modification which we introduced in the last section means that the simulator is no longer dependent of the pre-sampled values in map \mathbf{I} . Therefore, the simulator does not need to look for these pre-sampled values any longer, and we can remove the loop with the assignment $\textit{found} \leftarrow \text{true}$. Further, it is now easy to remove all loops in the main body and in procedure F_q . Indeed, the variables modified in these loops are disregarded by the procedures F_q and f_q ; F_q always behaves as a relay for F , while the simulator f_q always forwards calls to F in the case where $\text{findseq}((x, y), \mathbf{T}') \neq \text{None}$ and samples a uniformly distributed random value otherwise.

Next, we remove all global variables which are of no further use. Specifically, we remove the global variables \mathbf{bad}_4 , \mathbf{Y}' , \mathbf{q}'_f , \mathbf{I} , \mathbf{T}'_i , and \mathbf{T} . Furthermore we remove all associated assignments from or to these variables. Since we removed \mathbf{q}'_f , we modify the types of procedure $F : \{0, 1\}^* \rightarrow \{0, 1\}^n \times \mathbb{N}$ and of map $\mathbf{R} : \{0, 1\}^* \rightarrow \{0, 1\}^n \times \mathbb{N}$ to simply $F : \{0, 1\}^* \rightarrow \{0, 1\}^n$ and $\mathbf{R} : \{0, 1\}^* \rightarrow \{0, 1\}^n$, and adapt the associated code accordingly. Finally, the code which remains under the two branches of the statement “if *found* . . .” in procedure f_q is only “ $z \xleftarrow{\$} \{0, 1\}^n$; $\mathbf{T}'[x, y] \leftarrow z$ ”, and we merge these branches.

The resulting game is the game G_{ideal} . Because none of the above modifications influence the behavior of the oracles accessible to the distinguisher, we can show in **EasyCrypt** that

$$\vdash G_{\text{idealLazy}} \sim G_{\text{ideal}} : \text{true} \implies = \{b\}$$

and hence

$$\Pr [G_{\text{idealLazy}} : b] = \Pr [G_{\text{ideal}} : b]. \quad (5.14)$$

Here we need the same generic invariants as those used to prove the equivalence in Section 5.3. Finally, using Equations 5.2, 5.6, 5.7, 5.9 and 5.14 we obtain Equation 5.1, which finishes the proof of Theorem 5.1.

IN this chapter, we will consider the axiomatization of the functions used in the games described in Chapter 5, and build increasingly complex lemmas about their properties. Indeed, to enable the SMT solvers and automated theorem provers used by **EasyCrypt** to successfully verify proof obligations arising within the proof, **EasyCrypt** needs to feed these automated tools with appropriate characteristics of those functions. As an obvious example, in the step from $G_{\text{real}'}$ to G_{realRO} , where the case distinction on $\text{findseq}(x, y, T') \neq \text{None}$ plays a central role, it is often necessary to provide the automated tools with the knowledge that the function `findseq` is both sound and complete (we will formalize this in two axioms); in practice, the lemmas needed to prove the validity of logical side conditions can become much more complex than that.

As discussed in Section 4.1, those axioms and lemmas are stated in **EasyCrypt** before the proof of some **pRHL** judgment is performed; during the proof of this judgment, they can then be passed to the automated tools as needed. We normally start by postulating some basic axioms on which a proof relies; from there, we state more involved lemmas. **EasyCrypt** is often able to derive the validity of some easier lemmas automatically from those axioms and/or from prior lemmas. However, as the lemmas get more involved, this does not typically work; even the validity of some easier statements cannot normally be derived if their proof involves so much as a simple induction (inductive proofs seem to be hard for automated tools). In these cases, the concerned lemmas can be exported to the **Coq** proof assistant, and proven there. Then, they can simply be stated as axioms in **EasyCrypt**, as they are proven elsewhere.

In this chapter, we will describe the axioms that we needed to state for our proof of indifferentiability, and the lemmas that we derived from them. In many cases, these lemmas had to be proven manually in the **Coq** proof assistant: The corresponding **Coq** file contained over 4200 lines of code, and is available upon request, or from:

<http://easycrypt.gforge.inria.fr/csf12/>

Here, we additionally describe all of those proofs in a human-readable manner.

6.1 Predicates, Definitions, and Axioms

In this section, we outline the basic definitions and axioms that we will build on later in the chapter.

6.1.1 Maps

Since many lemmas in this chapter concern the maps maintained in the games described in Chapter 5, we introduce here a few notions to deal with them more easily. The type of maps that we are interested in these lemmas always have the type $\{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, i.e. maps that map pairs of a block and a chaining value to a chaining value; these maps are used to keep track of the previous answers of oracle f_q (the simulator) in the games.

In the games G_{real} and G_{realRO} , whose equivalence is discussed in Section 5.2, we use several failure events that are raised whenever such a map may be updated with a value that exhibits certain dependencies to values already contained in this map. Therefore, the maps in these games possess certain properties (up to the failure events) which are essential for many lemmas.

Two of these properties are particularly important, and are both guaranteed by the failure event \mathbf{bad}_1 . The first is that these maps are injective. We define a straight-forward predicate in `EasyCrypt` to express this as follows.

Definition 6.1 (Predicate Injective). Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map. The predicate `Injective` is defined as

$$\text{Injective}(T) := \forall xy, xy'. xy \in \text{dom}(T) \Rightarrow xy' \in \text{dom}(T) \Rightarrow T[xy] = T[xy'] \Rightarrow xy = xy'.$$

The second property is that the constant \mathbf{IV} does not appear in the range of a map. Together, we call these properties the *well-formedness* conditions of a map T .

Definition 6.2 (Well-formedness). Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map. T is called *well-formed* iff it holds that `Injective`(T) and $\mathbf{IV} \notin \text{ran}(T)$.

Recall from Section 5.2 that in game G_{real} , two maps T' and T are maintained such that it holds that $T' \subseteq T$, i.e. all elements in the domain of T' are also in the domain of T , and they map to the same values. We define the following predicate to express this property.

Definition 6.3 (Predicate Inclusion). Let T and T' be two maps of type $\{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. Then the predicate `Inclusion` is defined as

$$\text{Inclusion}(T', T) := \forall xy. xy \in \text{dom}(T') \Rightarrow xy \in \text{dom}(T) \wedge T[xy] = T'[xy].$$

We use the statements $T' \subseteq T$ and `Inclusion`(T', T) equivalently. `EasyCrypt` also provides a polymorphic update function `upd` such that

$$\text{upd}(T, xy, z)[xy'] = \text{if } xy = xy' \text{ then return } z \text{ else return } T[xy'],$$

which is internally defined by the elementary properties that one expects of such an operator. In the above statement we say that T gets *updated* with the pair (xy, z) . It is easy to automatically derive additional properties related to our predicates, such as:

1. If $\text{Inclusion}(T', T)$ and $\text{Injective}(T)$, then $\text{Injective}(T')$;
2. If $\text{Inclusion}(T', T)$ and $xy \notin \text{dom}(T)$, then $xy \notin \text{dom}(T')$;
3. If $\text{Inclusion}(T', T)$ and $z \notin \text{ran}(T)$, then $z \notin \text{ran}(T')$.

Later, we will often need such properties and we note them here once and for all; they will be implicit in the proofs performed in this chapter.

Lastly, we introduce the notion of a *benign* update. Intuitively, a benign update of a well-formed map T is an update which preserves the well-formedness of T , and does not modify values that were already present in T . Formally, we define it as follows.

Definition 6.4 (Benign update). Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, xy be a pair of a block and a chaining value, and z be a chaining value. Assume that T is well-formed. The update of the map T with the pair (xy, z) is *benign* iff it holds that $T \subseteq \text{upd}(T, xy, z)$, and $\text{upd}(T, xy, z)$ is still well-formed.

Remark. The definition does not exclude the case $xy \in \text{dom}(T)$. Indeed, in this case the update of T with the pair $(xy, T[xy])$ is benign.

In the games described in Section 5.2, all updates of T are benign up to the failure events described there. Similarly as above, we note that we can automatically derive certain related properties of the upd function.

1. If $xy \notin \text{dom}(T)$, then $\forall z. \text{Inclusion}(T, \text{upd}(T, xy, z))$;
2. If $\text{Injective}(T)$ and $z \notin \text{ran}(T)$, then $\forall xy. \text{Injective}(\text{upd}(T, xy, z))$;
3. If $z' \notin \text{ran}(T)$ and $z \neq z'$, then $\forall xy. z' \notin \text{ran}(\text{upd}(T, xy, z))$.

Thus it is easy to see that if a map T is well-formed and gets updated with a pair (xy, z) such that $xy \notin \text{dom}(T)$, $z \notin \text{ran}(T)$ and $z \neq \text{IV}$, then this update is benign. Similarly, if we consider two well-formed maps $T' \subseteq T$ and some element $xy \in \text{dom}(T)$, then the update of T' with the pair $(xy, T[xy])$ is benign.

6.1.2 Operators pad and unpad

The first operators that we define are the functions pad and unpad , which are used to decompose a message into a padding, respectively to reassemble a message from a padding.

Definition 6.5 (Operators pad and unpad). The operator pad is a function which maps messages to paddings, i.e. a function of type $\{0, 1\}^* \rightarrow \{0, 1\}^k \text{ list}$. Conversely, the operator unpad is a function of type $\{0, 1\}^k \text{ list} \rightarrow \{0, 1\}^* \text{ option}$ which maps paddings in the range of pad back to messages.

We characterize the desired behavior of the functions pad and unpad and their mutual relation using the following axioms.

Axiom 6.6. *It holds that $\text{unpad}(\text{nil}) = \text{None}$.*

Axiom 6.7 (Soundness of unpad). *Let p be a padding such that $\text{unpad}(p) \neq \text{None}$. Then there exists a message m such that $\text{pad}(m) = p$.*

Axiom 6.8 (Completeness and correctness of `unpad`). *Let m be a message. Then it holds true that $\text{unpad}(\text{pad}(m)) = \text{Some}(m)$.*

Further, we need the obvious property that the function `unpad` is injective (in other words, no message has two distinct paddings).

Axiom 6.9 (Injectivity of `unpad`). *Let p and p' be paddings. Assume that $\text{unpad}(p) \neq \text{None}$ and that $\text{unpad}(p) = \text{unpad}(p')$. Then it holds that $p = p'$.*

Lastly, we formalize our assumption that the padding function is prefix-free as follows.

Axiom 6.10 (Prefix-freeness). *Let m and m' be two messages with $m \neq m'$. Then for all paddings p , it holds that $\text{pad}(m) \neq \text{pad}(m') \parallel p$.*

6.1.3 Operator `mapfst`

The operator `mapfst` is a simple operator which maps the operator `fst` to a list. It is used to recover a padding from a list returned by the function `findseq`, since the function `findseq` returns a chain of the form $[(x_1, y_1), \dots, (x_j, y_j)]$, consisting of pairs of blocks and chaining values.

Definition 6.11 (Operator `mapfst`). The operator `mapfst` is a function of type $(\{0, 1\}^k \times \{0, 1\}^n) \text{ list} \rightarrow \{0, 1\}^k \text{ list}$ defined by the following equations:

$$\begin{aligned} \text{mapfst}(\text{nil}) &:= \text{nil} \\ \text{mapfst}((x, y)::c) &:= x::\text{mapfst}(c), \end{aligned}$$

where x is a block, y a chaining value and $c : (\{0, 1\}^k \times \{0, 1\}^n) \text{ list}$ is a list of pairs of blocks and chaining values.

6.1.4 Operator `findseq`

As discussed in Chapter 5, the operator `findseq` is used to find a complete chain (see Definition 5.2) in a map T . That is, the function call `findseq(x, y, T)` searches in map T for a complete chain of the form $c \parallel [(x, y)]$. It returns `Some(c)`, or `None` if there is no such chain. All elements of c are required to be in the domain of T , however the last element of the complete chain, namely (x, y) , is not.

To define this operator, we first define an auxiliary predicate `ischain`.

Definition 6.12 (Predicate `ischain`). Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, y and z be chaining values, and xy and xy' be pairs of blocks and chaining values. Furthermore let $c : (\{0, 1\}^k \times \{0, 1\}^n) \text{ list}$ be a list. Then the predicate `ischain` is recursively defined by

the following equations:

$$\begin{aligned}
 \text{ischain}(T, y, \text{nil}, z) &:= (y = z) \\
 \text{ischain}(T, y, xy::\text{nil}, z) &:= (xy \in \text{dom}(T) \wedge y = \text{snd}(xy) \wedge \\
 &\quad T[xy] = z) \\
 \text{ischain}(T, y, xy::xy'::c, z) &:= (xy \in \text{dom}(T) \wedge y = \text{snd}(xy) \wedge \\
 &\quad T[xy] = \text{snd}(xy') \wedge \\
 &\quad \text{ischain}(T, \text{snd}(xy'), xy'::c, z))
 \end{aligned}$$

Now, instead of giving a concrete implementation for the function `findseq`, we simply define it by the properties we expect it to have. On the one hand, this yields a higher level of abstraction as we do not commit to a specific implementation; on the other, this also means that we have to believe that there indeed exists a function which exhibits these properties. Specifically, we characterize the function `findseq` by two axioms stating that it is sound and complete, respectively.

Axiom 6.13 (Soundness of `findseq`). *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x a block and y a chaining value. Assume that $\text{findseq}(x, y, T) \neq \text{None}$. Then, it holds that $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x, y, T)}, y)$ and $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, T)}) \parallel [x]) \neq \text{None}$.*

Axiom 6.14 (Completeness of `findseq`). *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x a block, y a chaining value, and $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list a list. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $\text{ischain}(T, \text{IV}, c, y)$
4. $\text{unpad}(\text{mapfst}(c) \parallel [x]) \neq \text{None}$

Then, it holds that $\text{findseq}(x, y, T) \neq \text{None}$.

Note that for the completeness, we require that T is injective and that $\text{IV} \notin \text{ran}(T)$, i.e. T is well-formed. Requiring the well-formedness of T for the completeness of `findseq` is not strictly necessary to perform our indifferenciability proof, but it makes the function `findseq` much easier to implement in practice, since `findseq` needs only be successful in well-formed maps. We will also see later that we can use these well-formedness conditions to prove that given the premises of the completeness axiom, we can not only say that $\text{findseq}(x, y, T) \neq \text{None}$, but indeed that $\text{findseq}(x, y, T) = \text{Some}(c)$ (see Lemma 6.34). A possible pseudocode implementation of `findseq` that fulfills the above axioms is given below.

```

function findseq(x, y, T) =
  c ← nil;
  len ← 0;
  while len < q ∧ y ∈ ran(T) ∧ y ≠ IV do
    (x', y') ← T-1[y];
    c ← (x', y')::c; len ← len + 1; y ← y';
  if y = IV ∧ unpad(mapfst(c) ∥ [x]) ≠ None then
    return c
  else return None
    
```

As in Chapter 5, we use $q := \ell \cdot q_D$ to bound the maximal length of a chain (recall that the number of oracle queries that the distinguisher can perform is upper-bounded by q). This ensures termination even in the highly unlikely event where there are loops in a chain, and yields a running time of $\mathcal{O}(\ell \cdot q_D)$ for the function `findseq`.

Lastly, we also define a related predicate `valid_chain`. Intuitively, `valid_chain(T, c)` holds whenever c is a complete chain in T whose last element is also in the domain of T .

Definition 6.15 (Predicate `valid_chain`). Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map and $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list a list. Then the predicate `valid_chain` is defined by

$$\text{valid_chain}(T, c) := \exists z. \text{ischain}(T, IV, c, z) \wedge \text{unpad}(\text{mapfst}(c)) \neq \text{None}$$

6.1.5 Predicate `set_bad3`

The predicate `set_bad3` is used in game $\mathbf{G}_{\text{real}'}$ to detect the failure event `bad3`, as discussed in Section 5.2. In this chapter, we will see several lemmas involving this predicate which are necessary to prove the validity of the transformation described in that section.

Definition 6.16 (Predicate `set_bad3`). Let y be a chaining value and let T and T' be two maps of type $\{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. Then

$$\text{set_bad3}(y, T', T) := y \in \text{ran}(T) \wedge T^{-1}[y] \notin \text{dom}(T') \wedge \forall c. \neg \text{valid_chain}(c \parallel [T^{-1}[y]], T).$$

The idea is the following. Whenever the distinguisher makes a query $f_q(x, y)$ in game $\mathbf{G}_{\text{real}'}$ and there is some preceding element $(x', y') := T^{-1}[y]$ in the domain of T that maps to the given chaining value y , then the current query is a continuation of a chain in T . But if this preceding element (x', y') does not appear in the domain of the map T' , then this means that the distinguisher has never made the query $f_q(x', y')$; now if additionally, (x', y') is not the last element of a complete chain in T , then y should be unknown to the distinguisher. Indeed, the distinguisher cannot guess the intermediate chaining values generated in a Merkle-Damgård iteration until it queries for them, which it can only do in the correct order; the distinguisher may only learn the last value generated in a Merkle-Damgård iteration beforehand, by querying the oracle implementing the hash function, since the last chaining value is the final hash. Thus if `set_bad3(y, T', T)` evaluates to `true` in the procedure $\mathbf{G}_{\text{real}'}.f_q$, then the distinguisher has used in the current query a chaining value which should be unknown to it, and `bad3` is set to `true`.

Note that we use a map T^{-1} in this predicate. This is possible because in the game $\mathbf{G}_{\text{real}'}$ where this predicate is used, we have that the map T is injective. In our actual `EasyCrypt` implementation, we update a map T_{inv} along with the map T , prove as an invariant that T_{inv} is the inverse of T throughout the game $\mathbf{G}_{\text{real}'}$, and pass T_{inv} as an additional parameter to `set_bad3` – we omitted these technical details in the description of the games for the sake of the reader. In all the lemmas described in the remainder of this chapter that involve an inversion on T , we will have as a premise that T is injective, such that it makes sense to speak about the map T^{-1} .

6.1.6 Predicate Claim5

Lastly, we recall the definition of the predicate **Claim5** as presented in Definition 5.3 of Section 5.2. We will see several lemmas concerning this predicate at the end of this chapter.

$$\begin{aligned} \text{Claim5}(T', T) = \forall(x, y), (x', y'). (x', y') \in \text{dom}(T) \wedge (x, y) \in \text{dom}(T') \wedge \\ T[(x', y')] = y \wedge \text{findseq}(x', y', T) = \text{None} \Rightarrow (x', y') \in \text{dom}(T'). \end{aligned}$$

Recall that in game $\mathbf{G}_{\text{real}'}$, the simulator f_q maintains a map T' of previously answered queries, while the FIL-RO f (resp. f_{bad}) maintains a map T , and we have that $T' \subseteq T$: Indeed, a fresh call of the distinguisher to procedure F_q generates an MD-chain whose elements are all contained in map T , while a fresh call to f_q generates a single element in the maps T' and T . Then, the statement **Claim5**(T', T) means that the distinguisher has only made queries to f_q associated to an MD-chain in the correct order; essentially, for all elements (x, y) that belong to a chain in map T , it holds that the preceding element $(x', y') \in \text{dom}(T)$, for which $T[(x', y')] = y$, is also in the domain of T' whenever (x, y) is in the domain of T' . That is, for any element (x, y) that belongs to a chain in T' , all preceding elements of the chain are also in the domain of T' , meaning that the distinguisher has made the corresponding queries.

In the corner case where (x, y) is the *first* element of a chain in a well-formed map T , the statement is trivially fulfilled where we consider it. Indeed, in this case it holds that $y = \mathbf{IV}$, and therefore the premises $(x', y') \in \text{dom}(T)$ and $T[(x', y')] = y$ imply that $\mathbf{IV} \in \text{ran}(T)$, which is false if T is well-formed.

Similarly, the case where (x', y') is the *last* element of a chain in T also represents a corner case, since in this case the element $T[(x', y')] = y$ corresponds to the final hash of a message (instead of an intermediate chaining value) and is therefore known to the distinguisher. So, the distinguisher could *extend* this chain by making additional queries (consider e.g. a length extension attack). This is not a problem in our setup as we assume the padding function to be prefix-free; however, for the statement **Claim5**(T', T) to hold as an invariant, we need to explicitly exclude this case, since here the fact that an element (x, y) is in the domain of T' does not necessarily mean that all preceding elements of the original chain are also in the domain of T' . This is why the predicate uses the premise $\text{findseq}(x', y', T) = \text{None}$, meaning that (x', y') is not the end of a complete chain in T .

In the game $\mathbf{G}_{\text{real}'}$ where this predicate is used, we always have the invariants that $T' \subseteq T$, and that T is well-formed. Therefore we are able to eventually obtain that **Claim5**(T', T) is also an invariant of that game.

6.2 Padding Lemmas

In this section, we will prove two simple lemmas concerning the padding function, based on the axioms stated in Section 6.1.2. The first lemma states that `nil` is not a valid padding. This is needed to be able to conclude that any padding contains at least one element, which is relevant e.g. in the transformation described in Section 5.1.

Lemma 6.17. *There is no message m such that $\text{pad}(m) = \text{nil}$.*

Proof. We prove the statement by contradiction. Let m be a message, and assume for contradiction that $\text{pad}(m) = \text{nil}$. By Axiom 6.8, we know that $\text{unpad}(\text{pad}(m)) = \text{Some}(m)$ and hence we obtain $\text{unpad}(\text{nil}) = \text{Some}(m)$, contradicting Axiom 6.6. \square

The second padding lemma that we want to show is an alternative description of the prefix-freeness of the padding function, based on unpad instead of pad . This will later be useful as a helper lemma.

Lemma 6.18. *Let p_1 and p_2 be two paddings. Assume that $\text{unpad}(p_1 \parallel p_2) \neq \text{None}$ and that $p_2 \neq \text{nil}$. Then it holds that $\text{unpad}(p_1) = \text{None}$.*

Proof. We know by Axiom 6.7 that there exists a message m such that $\text{pad}(m) = p_1 \parallel p_2$. Now, assume for contradiction that $\text{unpad}(p_1) \neq \text{None}$. Then we also know that there exists a message m' such that $\text{pad}(m') = p_1$. Therefore we get $\text{pad}(m) = \text{pad}(m') \parallel p_2$. We now perform a case distinction:

Case $m = m'$. If $m = m'$, then $\text{pad}(m) = \text{pad}(m')$, and thus $p_2 = \text{nil}$, contradicting our assumptions.

Case $m \neq m'$. In this case, we obtain a contradiction to prefix-freeness: Using Axiom 6.10, we conclude that $\text{pad}(m) \neq \text{pad}(m') \parallel p_2$, which immediately yields the contradiction. \square

6.3 Lemmas on mapfst

We note here that the operator mapfst is distributive.

Lemma 6.19 (Distributivity of mapfst). *Let c_1 and c_2 be two lists of type $(\{0,1\}^k \times \{0,1\}^n)$ list. Then it holds that $\text{mapfst}(c_1 \parallel c_2) = \text{mapfst}(c_1) \parallel \text{mapfst}(c_2)$.*

Proof. We prove the claim by structural induction over the list c_1 .

Base case: $c_1 = \text{nil}$. Trivial.

Induction step: $c_1 = a::al$.

Induction hypothesis:

$$\text{mapfst}(al \parallel c_2) = \text{mapfst}(al) \parallel \text{mapfst}(c_2)$$

We need to show that $\text{mapfst}(a::al \parallel c_2) = \text{mapfst}(a::al) \parallel \text{mapfst}(c_2)$. This is easy:

$$\begin{aligned} \text{mapfst}(a::al \parallel c_2) &= \text{mapfst}(a::(al \parallel c_2)) \\ &= \text{fst}(a)::\text{mapfst}(al \parallel c_2) && \text{(by Definition 6.11)} \\ &= \text{fst}(a)::(\text{mapfst}(al) \parallel \text{mapfst}(c_2)) && \text{(by IH)} \\ &= \text{fst}(a)::\text{mapfst}(al) \parallel \text{mapfst}(c_2) \\ &= \text{mapfst}(a::al) \parallel \text{mapfst}(c_2) && \text{(by Definition 6.11)} \end{aligned}$$

□

In the remainder of this chapter, we will assume an intuitive understanding of this operator. For instance, we assume that it is clear that a list $\text{mapfst}(c \parallel [(x, y)])$ can be rewritten as $\text{mapfst}(c) \parallel [x]$. We strive to write such lists in their simplest possible form, except where it makes sense to do otherwise for the purpose of clarity.

6.4 Lemmas on ischain

In this section, we will discuss several lemmas concerning the predicate `ischain`. Many of these lemmas are not immediately needed for the proof of indifferentiability performed in `EasyCrypt`; rather, they serve as helper lemmas on which to build more complex lemmas later on.

We start with a very basic lemma that makes explicit a property of the statement $\text{ischain}(T, y, c, z)$ where c is not empty: the value y is the first chaining value used in c .

Lemma 6.20. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list a list, x be a block, and y, y' and z be chaining values. Assume that $\text{ischain}(T, y, (x, y')::c, z)$ holds. Then we have $y = y'$.*

Proof. We only need to distinguish the cases where $c = \text{nil}$ and where c has the form $a::al$. In each case the statement follows immediately from Definition 6.12. □

The next lemma can be thought of as an alternative definition of the predicate `ischain`, in the case where the chain in question is not `nil`. It is an extremely helpful tool in many of the lemmas that follow.

Lemma 6.21. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c_1, c_2 : (\{0, 1\}^k \times \{0, 1\}^n)$ list be lists, x' be a block, and y, y' and z be chaining values. Then it holds:*

$$\begin{aligned} \text{ischain}(T, y, c_1 \parallel (x', y')::c_2, z) &\Leftrightarrow \text{ischain}(T, y, c_1, y') \wedge \\ &\quad (x', y') \in \text{dom}(T) \wedge \\ &\quad \text{ischain}(T, T[(x', y')], c_2, z) \end{aligned}$$

Proof. To prove the claim, we prove both directions of the claim separately. In both cases we use a structural induction over the list c_1 . In the respective base cases, we distinguish the cases where $c_2 = \text{nil}$ and where c_2 has the form $b::bl$ and apply Definition 6.12. For the direction “ \Leftarrow ”, we additionally need to apply Lemma 6.20 to obtain $T[(x', y')] = \text{snd}(b)$ in the case where $c_2 = b::bl$, since the definition of `ischain` requires that the second argument of the predicate `ischain` is equal to the second projection of the first element of the chain. In the respective induction steps, we distinguish the cases where $c_1 = \text{nil}$ and where c_1 has the form $a::al$ and use Definition 6.12 on the induction premises to fulfill the premises of the respective induction hypotheses, and easily obtain the conclusion. □

The previous lemma gives us two corollaries. The first corollary permits the automated tools used in `EasyCrypt` to perform a case distinction on the statement $\text{ischain}(T, \text{IV}, c, z)$: If $c = \text{nil}$, then it holds that $z = \text{IV}$; if c is not empty, then there is an element in the domain of T (namely, the last element of c) that maps to the chaining value z .

Corollary 6.22. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, and z be a chaining value. Assume that $\text{ischain}(T, \text{IV}, c, z)$ holds. Then it holds true that $z = \text{IV} \vee \exists xy. xy \in \text{dom}(T) \wedge T[xy] = z$.*

Proof. If $c = \text{nil}$, then we immediately get $z = \text{IV}$ from Definition 6.12. If c has the form $al \parallel [a]$, then Lemma 6.21 applies and we conclude $a \in \text{dom}(T)$ and $\text{ischain}(T, T[a], \text{nil}, z)$ which implies $T[a] = z$. \square

The second corollary, useful later, states that any element that belongs to a chain in a map T is in the domain of T .

Corollary 6.23. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, xy be a pair of a block and a chaining value and z be a chaining value. Assume the following:*

1. $\text{ischain}(T, \text{IV}, c, z)$
2. $xy \in c$

Then it holds that $xy \in \text{dom}(T)$.

Proof. Since $xy \in c$ then c has the form $c_\ell \parallel xy::c_r$, for appropriate c_ℓ and c_r . Hence Lemma 6.21 applies and we find $xy \in \text{dom}(T)$. \square

In the transformations presented in Chapter 5, to show that some invariant of a game relating to maps maintained in that game is valid, it is often necessary to state under which conditions a map may be updated without afflicting the validity of the statement expressed by this invariant. In turn, the validity of these conditions is usually guaranteed by the failure events. The invariants described in Section 5.2 involve predicates such as `set_bad3` or `Claim5`, which in turn involve the predicates `valid_chain` or the operator `findseq`. This are built upon the predicate `ischain`. Hence, when breaking down the necessary lemmas that ensure that some assertion in a game holds after a map update when it held before the map update, it often comes to show that under sensible conditions, a chain in a map T remains a chain when T gets updated. This is what the next two lemmas are used for. The first states that when a map T contains a chain and gets updated with a value not part of this chain, then this chain is still contained in the updated map.

Lemma 6.24. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, xy be a pair of a block and a chaining value and z and z' be chaining values. Assume the following:*

1. $\text{ischain}(T, \text{IV}, c, z)$
2. $xy \notin c$

Then it holds that $\text{ischain}(\text{upd}(T, xy, z'), \text{IV}, c, z)$.

Proof. We first *generalize* the constant IV as an arbitrary chaining value y ; that is, we show the claim for arbitrary values y where the claim focuses only on the instance where $y = \text{IV}$. This is necessary so as to obtain a strong enough induction hypothesis. We prove the resulting generalized statement by a structural induction over the list c .

Base case: $c = \text{nil}$. We know $\text{ischain}(T, y, \text{nil}, z)$ which implies $y = z$. This implies $\text{ischain}(\text{upd}(T, xy, z'), y, \text{nil}, z)$ (see Definition 6.12).

Induction step: $c = a::al$.

Induction hypothesis:

$$\forall y. \text{ischain}(T, y, al, z) \Rightarrow xy \notin al \Rightarrow \text{ischain}(\text{upd}(T, xy, z'), y, al, z)$$

We must show that $\text{ischain}(\text{upd}(T, xy, z'), y, a::al, z)$. By Lemma 6.21, this is equivalent to the following three conditions:

- (i) $\text{ischain}(\text{upd}(T, xy, z'), y, \text{nil}, \text{snd}(a))$: By Definition 6.12 we have to show that $y = \text{snd}(a)$. This is immediately implied by the premise $\text{ischain}(T, y, a::al, z)$ (Lemma 6.20).
- (ii) $a \in \text{dom}(\text{upd}(T, xy, z'))$: By Lemma 6.21, $\text{ischain}(T, y, a::al, z)$ implies that $a \in \text{dom}(T)$, and therefore it also holds that $a \in \text{dom}(\text{upd}(T, xy, z'))$.
- (iii) $\text{ischain}(\text{upd}(T, xy, z'), \text{upd}(T, xy, z')[a], al, z)$: Here we need the induction hypothesis. First note that $xy \neq a$ since we know that $xy \notin a::al$. Therefore it holds true that $\text{upd}(T, xy, z')[a] = T[a]$. Hence, it suffices to show $\text{ischain}(\text{upd}(T, xy, z'), T[a], al, z)$. This is implied by taking $y := T[a]$ in the induction hypothesis, and we are only left to show the two premises of the induction hypothesis: First, we know $\text{ischain}(T, y, a::al, z)$, and thus it holds that $\text{ischain}(T, T[a], al, z)$ by Lemma 6.21. Second, we know $xy \notin a::al$, which implies $xy \notin al$.

□

Next, we prove that a chain in a map T' remains a chain in an updated map T as long as it holds that $T \supseteq T'$. That is, a map T' containing a chain can be updated with arbitrarily many values outside of its domain and still contain the same chain. The proof is mostly analogous to the proof of the previous lemma, except in the case where the induction hypothesis needs to be applied.

Lemma 6.25. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be two maps, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, and y and z be chaining values. Assume the following:*

1. $\text{ischain}(T', y, c, z)$
2. $\text{Inclusion}(T', T)$

Then it holds that $\text{ischain}(T, y, c, z)$.

Proof. As in Lemma 6.24, we perform a structural induction over the list c .

Base case: $c = \text{nil}$. As in the proof of Lemma 6.24.

Induction step: $c = a::al$.

Induction hypothesis:

$$\forall y. \text{ischain}(T', y, al, z) \Rightarrow \text{ischain}(T, y, al, z)$$

As in the previous lemma, we show the three necessary and sufficient conditions according to Lemma 6.21:

- (i) $\text{ischain}(T, y, \text{nil}, \text{snd}(a))$: We know $\text{ischain}(T', y, a::al, z)$, and therefore it holds that $y = \text{snd}(a)$, which yields the conclusion by Definition 6.12.
- (ii) $a \in \text{dom}(T)$: We know that $\text{ischain}(T', y, a::al, z)$, which implies $a \in \text{dom}(T')$ by Lemma 6.21. Since $\text{Inclusion}(T', T)$ this further implies $a \in \text{dom}(T)$.
- (iii) $\text{ischain}(T, T[a], al, z)$: This is the case where we apply the induction hypothesis. We instantiate the induction hypothesis taking $y := T'[a]$. From the premise $\text{ischain}(T', y, a::al, z)$, we obtain using Lemma 6.21 that $a \in \text{dom}(T')$ and that $\text{ischain}(T', T'[a], al, z)$; the latter fulfills the premise of the induction hypothesis. Therefore we obtain that $\text{ischain}(T, T'[a], al, z)$, which is almost our goal. It only remains to show that $T'[a] = T[a]$. But this is clear since we have $a \in \text{dom}(T')$ and $\text{Inclusion}(T', T)$.

□

The previous lemma also yields an immediate corollary that allows a single update to a map T with a value outside of the domain of T .

Corollary 6.26. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, xy be a pair of a block and a chaining value, and y, z, z' be chaining values. Assume the following:*

1. $\text{ischain}(T, y, c, z)$
2. $xy \notin \text{dom}(T)$

Then it holds that $\text{ischain}(\text{upd}(T, xy, z'), y, c, z)$.

Proof. This follows immediately from Lemma 6.25 since $xy \notin \text{dom}(T)$ clearly implies that $\text{Inclusion}(T, \text{upd}(T, xy, z'))$. □

A dual of Lemma 6.24 is the following statement. If a list c is *not* a chain in some map T , but becomes a chain after a single update of this map, then the element which was used to update the map must be part of the chain.

Lemma 6.27. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, xy be a pair of a block and a chaining value, and z and z' be chaining values. Assume:*

1. $\neg \text{ischain}(T, \text{IV}, c, z)$.
2. $\text{ischain}(\text{upd}(T, xy, z'), \text{IV}, c, z)$.

Then it holds that $xy \in c$.

Proof. We generalize the constant IV as an arbitrary variable y , and perform a structural induction over the list c .

Base case: $c = \text{nil}$. In this case we get a contradiction since the first premise implies $y \neq z$, but the second premise implies $y = z$ (see Definition 6.12).

Induction step: $c = a::al$.

Induction hypothesis:

$$\forall y. \neg \text{ischain}(T, y, al, z) \Rightarrow \text{ischain}(\text{upd}(T, xy, z'), y, al, z) \Rightarrow xy \in al$$

We must show that $xy \in a::al$. That is, we have to show that $xy = a \vee xy \in al$. To accomplish this, assume that $xy \neq a$; under this hypothesis we will prove that $xy \in al$. We perform a case distinction on the structure of al .

Case $al = \text{nil}$. In this case, we obtain a contradiction. On the one hand, we know $\neg \text{ischain}(T, y, [a], z)$. On the other hand, we know $\text{ischain}(\text{upd}(T, xy, z'), y, [a], z)$ and $xy \neq a$. We show below that this implies $\text{ischain}(T, y, [a], z)$, yielding the contradiction. In fact, the latter is easily seen, since it suffices to unfold the definition of ischain and show the following:

- (i) $a \in \text{dom}(T)$: This follows from $a \in \text{dom}(\text{upd}(T, xy, z'))$ and $xy \neq a$.
- (ii) $y = \text{snd}(a)$: This immediately follows from $\text{ischain}(\text{upd}(T, xy, z'), y, [a], z)$.
- (iii) $T[a] = z$: This follows from $\text{upd}(T, xy, z')[a] = z$ and $xy \neq a$.

Case $al = a'::al'$. We have to show $xy \in a'::al'$. We instantiate the induction hypothesis taking $y := \text{upd}(T, xy, z')[a]$, which yields the goal. It remains to show the two premises of the induction hypothesis. The second premise of the induction hypothesis, namely $\text{ischain}(\text{upd}(T, xy, z'), \text{upd}(T, xy, z')[a], a'::al', z)$, follows quickly since it holds that $\text{ischain}(\text{upd}(T, xy, z'), y, a::a'::al', z)$, and we apply Lemma 6.21. We are only left to show the first premise of the induction hypothesis, namely $\neg \text{ischain}(T, \text{upd}(T, xy, z')[a], a'::al', z)$. Recall that $\neg \text{ischain}(T, y, a::a'::al', z)$, that is

$$\neg(a \in \text{dom}(T) \wedge y = \text{snd}(a) \wedge T[a] = \text{snd}(a') \wedge \text{ischain}(T, \text{snd}(a'), a'::al', z)).$$

Next, recall that $\text{ischain}(\text{upd}(T, xy, z'), y, a::a'::al', z)$. Hence, by Definition 6.12 we know (i) $a \in \text{dom}(\text{upd}(T, xy, z'))$ and hence $a \in \text{dom}(T)$ since $xy \neq a$; (ii) $y = \text{snd}(a)$, and (iii) $\text{upd}(T, xy, z')[a] = \text{snd}(a')$, implying $T[a] = \text{snd}(a')$ since $xy \neq a$. Hence, it must hold that $\neg \text{ischain}(T, \text{snd}(a'), a'::al', z)$. Lastly, assume for contradiction that $\text{ischain}(T, \text{upd}(T, xy, z')[a], a'::al', z)$ were true. By Lemma 6.20 this would imply $\text{ischain}(T, \text{snd}(a'), a'::al', z)$, yielding a contradiction. Therefore, we obtain $\neg \text{ischain}(T, \text{upd}(T, xy, z')[a], a'::al', z)$, which finishes the proof.

□

Next, we show that any chain in a map T and the value that it hashes to are uniquely defined by the padding associated to it.

Lemma 6.28. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c_1, c_2 : (\{0, 1\}^k \times \{0, 1\}^n)$ list be lists, and let z_1 and z_2 be chaining values. Assume the following:*

1. $\text{ischain}(T, y, c_1, z_1)$
2. $\text{ischain}(T, y, c_2, z_2)$
3. $\text{mapfst}(c_1) = \text{mapfst}(c_2)$

Then it holds that $c_1 = c_2$ and $z_1 = z_2$.

Proof. We perform a structural induction over c_1 .

Base case: $c_1 = \text{nil}$. If $c_2 \neq \text{nil}$, we obtain a contradiction to $\text{mapfst}(c_1) = \text{mapfst}(c_2)$. If $c_2 = \text{nil}$, it only remains to show that $z_1 = z_2$. From the first and second premises, we obtain $y = z_1$ and $y = z_2$ (unfolding Definition 6.12), which finishes this case.

Induction step: $c_1 = a::al$.

Induction hypothesis:

$$\forall c_2. \text{mapfst}(al) = \text{mapfst}(c_2) \Rightarrow \forall y. \text{ischain}(T, y, al, z_1) \Rightarrow \text{ischain}(T, y, c_2, z_2) \Rightarrow al = c_2 \wedge z_1 = z_2$$

If $c_2 = \text{nil}$, we obtain a contradiction to $\text{mapfst}(c_1) = \text{mapfst}(c_2)$. Therefore assume that c_2 has the form $b::bl$. We instantiate the induction hypothesis with the list bl . From $\text{mapfst}(a::al) = \text{mapfst}(b::bl)$, we obtain that $\text{fst}(a) = \text{fst}(b)$ and that $\text{mapfst}(al) = \text{mapfst}(bl)$. The former hypothesis will yet be important later; the latter can be used to fulfill the first premise of the induction hypothesis. We obtain a new induction hypothesis:

$$\forall y. \text{ischain}(T, y, al, z_1) \Rightarrow \text{ischain}(T, y, bl, z_2) \Rightarrow al = bl \wedge z_1 = z_2$$

Next, we show that $\text{snd}(a) = \text{snd}(b)$. Indeed, from the first and second premises we obtain (using Lemma 6.21) that $\text{ischain}(T, y, \text{nil}, \text{snd}(a))$ and $\text{ischain}(T, y, \text{nil}, \text{snd}(b))$. By unfolding Definition 6.12, we obtain $y = \text{snd}(a)$ and $y = \text{snd}(b)$. Since we now know $\text{snd}(a) = \text{snd}(b)$ and we already knew $\text{fst}(a) = \text{fst}(b)$, we conclude $a = b$. Further, we conclude from the first and second premises (again using Lemma 6.21) that $\text{ischain}(T, T[a], al, z_1)$ and $\text{ischain}(T, T[b], bl, z_2)$. As we have $a = b$, we rewrite the latter hypothesis as $\text{ischain}(T, T[a], bl, z_2)$. We can now instantiate the induction hypothesis taking $y := T[a]$, and we have already proven the two resulting premises of the induction hypothesis. Thus we conclude $al = bl$ and $z_1 = z_2$. Since $a = b$, we immediately get $c_1 = a::al = b::bl = c_2$.

□

Now, we want to show a central property of chains arising from the well-formedness conditions of maps (see Section 6.1.1). Namely, we will show that any chain in a well-formed map T is uniquely defined by the value that it hashes to. Formally, two chains in a well-formed map T that hash to the same value are, in fact, equal. This lemma also has many important implications that we will see later. However, the induction needed for the proof of this statement is somewhat tricky. To make it go through easier, we first show two closely related claims, considering the cases where the two chains have the same length or a different length, respectively.

Lemma 6.29. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c_1, c_2 : (\{0, 1\}^k \times \{0, 1\}^n)$ list be lists, and let y_1, y_2 and z be chaining values. Assume the following:*

1. $|c_1| = |c_2|$
2. $\text{Injective}(T)$
3. $\text{ischain}(T, y_1, c_1, z)$
4. $\text{ischain}(T, y_2, c_2, z)$

Then it holds that $c_1 = c_2$.

Proof. Fix c_1 . Given the premise $\text{Injective}(T)$, we show that

$$\forall c_2, y_1, y_2. |c_1| = |c_2| \Rightarrow \text{ischain}(T, y_1, c_1, z) \Rightarrow \text{ischain}(T, y_2, c_2, z) \Rightarrow c_1 = c_2,$$

using a structural induction over c_1 .

Base case: $c_1 = \text{nil}$.

If $c_2 = \text{nil}$, the claim is trivial. If c_2 has the form $b::bl$, we get a contradiction since we assume $|c_1| = |c_2|$, i.e. $|\text{nil}| = |b::bl|$.

Induction step: $c_1 = a::al$.

Induction hypothesis:

$$\forall c_2, y_1, y_2. |al| = |c_2| \Rightarrow \text{ischain}(T, y_1, al, z) \Rightarrow \text{ischain}(T, y_2, c_2, z) \Rightarrow al = c_2$$

We must show $a::al = c_2$. If $c_2 = \text{nil}$, we get a contradiction to the premise $|c_1| = |c_2|$. Therefore assume that c_2 has the form $b::bl$. We can now sensibly instantiate the induction hypothesis as follows:

$$|al| = |bl| \Rightarrow \text{ischain}(T, T[a], al, z) \Rightarrow \text{ischain}(T, T[b], bl, z) \Rightarrow al = bl.$$

The first premise of this hypothesis follows easily since we know that $|a::al| = |b::bl|$. Further, we know that $\text{ischain}(T, y_1, a::al, z)$ and $\text{ischain}(T, y_2, b::bl, z)$. Hence, the second and third premises follow by using Lemma 6.21. We conclude that $al = bl$. Thus it only remains to show that $a = b$. Note that we now have $\text{ischain}(T, y_1, a::al, z)$ and $\text{ischain}(T, y_2, b::al, z)$. Using Lemma 6.21 we obtain that $a \in \text{dom}(T)$ and $b \in \text{dom}(T)$. Further, we also know that $T[a] = T[b]$: Indeed, in the case where $al \neq \text{nil}$, we have $T[a] = \text{hd}(al)$ and $T[b] = \text{hd}(al)$ (from Lemma 6.20), and in the case where $al = \text{nil}$, we have $T[a] = z$ and $T[b] = z$ (from Definition 6.12). Thus, we can use the injectivity of T to conclude that $a = b$, which finishes the proof.

□

Lemma 6.30. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c_1, c_2 : (\{0, 1\}^k \times \{0, 1\}^n)$ list be lists, and let y_1, y_2 and z be chaining values. Assume the following:*

1. $|c_1| > |c_2|$
2. $\text{Injective}(T)$
3. $\text{ischain}(T, y_1, c_1, z)$
4. $\text{ischain}(T, y_2, c_2, z)$

Then there exists a list pl and a chaining value p such that $c_1 = pl \parallel p::c_2$.

Proof. Fix c_2 . Given the premise $\text{Injective}(T)$, we show that

$$\forall c_1, y_1, y_2. |c_1| > |c_2| \Rightarrow \text{ischain}(T, y_1, c_1, z) \Rightarrow \text{ischain}(T, y_2, c_2, z) \Rightarrow \exists p', pl'. c_1 = pl' \parallel p'::c_2$$

using a structural induction over c_2 .

Base case: $c_2 = \text{nil}$.

If $c_1 = \text{nil}$, we get a contradiction to $|c_1| > |c_2|$. Therefore assume that c_1 has the form $a::al$. We must show that there exist pl, p such that $a::al = pl \parallel [p]$, which is easily seen.

Induction step: $c_2 = b::bl$.

Induction hypothesis:

$$\forall c_1, y_1, y_2. |c_1| > |bl| \Rightarrow \text{ischain}(T, y_1, c_1, z) \Rightarrow \text{ischain}(T, y_2, bl, z) \Rightarrow \exists p', pl'. c_1 = pl' \parallel p'::bl$$

We must show that there exist p, pl such that $c_1 = pl \parallel p::b::bl$. If $c_1 = \text{nil}$, we get a contradiction to the premise $|c_1| > |c_2|$. Therefore assume that c_1 has the form $a::al$. We can now sensibly instantiate the induction hypothesis as follows:

$$|al| > |bl| \Rightarrow \text{ischain}(T, T[a], al, z) \Rightarrow \text{ischain}(T, T[b], bl, z) \Rightarrow \exists p', pl'. al = pl' \parallel p'::bl$$

The first premise of this hypothesis follows easily since we know that $|a::al| > |b::bl|$. Further, we know that $\text{ischain}(T, y_1, a::al, z)$ and $\text{ischain}(T, y_2, b::bl, z)$. Hence, the second and third premises follow by using Lemma 6.21. We conclude that there exist p', pl' such that $al = pl' \parallel p'::bl$, and we fix p' and pl' . We still have to show that there exist p, pl such that $a::al = pl \parallel p::b::bl$. Since we have $al = pl' \parallel p'::bl$, this means that we have to show that there exist p, pl such that $a::pl' \parallel p'::bl = pl \parallel p::b::bl$. It is easily seen that there exist p, pl such that $a::pl' = pl \parallel [p]$. Thus it only remains to show that $p' = b$. Note that we have $\text{ischain}(T, y_1, a::al, z)$, which can be rewritten as $\text{ischain}(T, y_1, a::pl' \parallel p'::bl, z)$. Using Lemma 6.21, this implies $\text{ischain}(T, T[p'], bl, z)$ and $p' \in \text{dom}(T)$. On the other hand, we also know that $\text{ischain}(T, y_2, b::bl, z)$, which similarly implies $\text{ischain}(T, T[b], bl, z)$ and $b \in \text{dom}(T)$. Hence we obtain $T[p'] = T[b]$ (using Definition 6.12 if $bl = \text{nil}$, and Lemma 6.20 if $bl \neq \text{nil}$). Finally, using the injectivity of T , we conclude that $p' = b$, which finishes the proof.

□

The previous two lemmas only included one of two well-formedness conditions, namely that T is injective. Basically, we saw that if an injective map T contains two chains c_1 and c_2 that hash to the same value, then by the injectivity of T this means that the shorter chain is a suffix of the longer chain. At this point we can exploit the other well-formedness condition, namely that $\text{IV} \notin \text{ran}(T)$. Indeed, if two chains using IV as a first chaining value hash to the same value, and one of the two chains is longer than the other, then the longer chain includes an element that maps to the chaining value IV at the boundary of its suffix that is equal to the shorter chain. However, as IV is not in the range of a well-formed map, there cannot be two chains (starting with the chaining value IV) of different length in a well-formed map. We formally derive the contradiction in the following corollary.

Corollary 6.31. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c_1, c_2 : (\{0, 1\}^k \times \{0, 1\}^n)$ list be lists, and let z be a chaining value. Assume the following:*

1. $|c_1| > |c_2|$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $\text{ischain}(T, \text{IV}, c_1, z)$
5. $\text{ischain}(T, \text{IV}, c_2, z)$

Then we obtain a contradiction, i.e. the above premises cannot hold all at once.

Proof. Using all but the third premise, we obtain using Lemma 6.30 a chaining value p and a list pl such that $c_1 = pl \parallel p::c_2$. Thus, the fourth premise can be rewritten as $\text{ischain}(T, \text{IV}, pl \parallel p::c_2, z)$. We now use a case distinction on the structure of c_2 .

Case $c_2 = \text{nil}$. By applying Lemma 6.21 we obtain $p \in \text{dom}(T)$ and $\text{ischain}(T, T[p], \text{nil}, z)$. By Definition 6.12, the latter implies $T[p] = z$. Therefore $z \in \text{ran}(T)$. By the fifth premise we know $\text{ischain}(T, \text{IV}, \text{nil}, z)$ which implies $z = \text{IV}$, and thus $\text{IV} \in \text{ran}(T)$. Contradiction to the third premise.

Case $c_2 = b::bl$. This case is analogous, but the value $T[p]$ is equal to $\text{snd}(b)$ instead of z , and we use that $\text{snd}(b)$ is equal to IV , yielding a contradiction.

□

Finally, using both Lemma 6.29 and Corollary 6.31, it is straightforward to conclude the uniqueness lemma mentioned above: We show that any chain in a well-formed map is uniquely defined by its final hash.

Lemma 6.32 (Uniqueness of chains). *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c_1, c_2 : (\{0, 1\}^k \times \{0, 1\}^n)$ list be lists, and let z be a chaining value. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $\text{ischain}(T, \text{IV}, c_1, z)$
4. $\text{ischain}(T, \text{IV}, c_2, z)$





Figure 6.1: Colliding chains, an impossibility because of *Claim 4* Figure 6.2: Length extension, an impossibility because of prefix-freeness

Then it holds that $c_1 = c_2$.

Proof. It cannot hold that either $|c_1| > |c_2|$ or $|c_2| > |c_1|$, since in both cases we obtain a contradiction using Corollary 6.31. Therefore it must hold that $|c_1| = |c_2|$. Thus Lemma 6.29 applies and we obtain $c_1 = c_2$. \square

The last lemma of this section is another important one, with many consequences seen in later sections. Essentially, it states that there can be no colliding chains in a well-formed map T , as illustrated in Figure 6.1. It is named *Claim 4* in reminiscence of the same claim appearing in the proof by Coron [39], which our proof of indifferentiability is based on. It is a consequence of the uniqueness of chains described above, and the assumption that the padding function is prefix-free. Figures 6.1 and 6.2 illustrate the idea. They depict respectively two complete chains c_1 and c_2 , ending respectively in (x_1, y_2) and (x_2, y_2) , in a well-formed map T . By the uniqueness of chains, the chain c_1 is a prefix of chain c_2 , but this contradicts our assumption of prefix-freeness.

Lemma 6.33 (Claim 4). *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c_1, c_2 : (\{0, 1\}^k \times \{0, 1\}^n)$ list be lists, let x_1, x_2 be blocks, and let y_1, y_2 be chaining values. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $\text{ischain}(T, \text{IV}, c_1, y_1)$
4. $\text{unpad}(\text{mapfst}(c_1) \parallel [x_1]) \neq \text{None}$
5. $\text{ischain}(T, \text{IV}, c_2, y_2)$
6. $\text{unpad}(\text{mapfst}(c_2) \parallel [x_2]) \neq \text{None}$

Then it holds that $(x_1, y_1) \notin c_2$.

Proof. We prove the claim by contradiction. Assume $(x_1, y_1) \in c_2$. Then, clearly c_2 has the form $c_2^\ell \parallel (x_1, y_1) :: c_2^r$, for appropriate c_2^ℓ and c_2^r . Thus, we can rewrite the fifth premise as $\text{ischain}(T, \text{IV}, c_2^\ell \parallel (x_1, y_1) :: c_2^r, y_2)$. By Lemma 6.21 this implies $\text{ischain}(T, \text{IV}, c_2^\ell, y_1)$. Now, since we also know $\text{ischain}(T, \text{IV}, c_1, y_1)$ and additionally the well-formedness conditions on map T hold, Lemma 6.32 applies and we conclude that $c_2^\ell = c_1$. Hence, we can rewrite

the sixth premise as $\text{unpad}(\text{mapfst}(c_1 \parallel (x_1, y_1)::c_2^r) \parallel [x_2]) \neq \text{None}$. From this we will obtain a contradiction to the prefix-freeness of pad . By the distributivity of mapfst , the latter statement can be rewritten as $\text{unpad}(\text{mapfst}(c_1) \parallel [x_1] \parallel \text{mapfst}(c_2^r) \parallel [x_2]) \neq \text{None}$. Now, taking $p_1 := \text{mapfst}(c_1) \parallel [x_1]$ and $p_2 := \text{mapfst}(c_2^r) \parallel [x_2]$ in Lemma 6.18, we get $\text{unpad}(\text{mapfst}(c_1) \parallel [x_1]) = \text{None}$, yielding a contradiction to the fourth premise. \square

6.5 Lemmas on findseq

We will now discuss several lemmas concerning the function findseq . Many of these are immediately relevant to enable the automated tools used in the **EasyCrypt** proof to solve arising side conditions, since the function findseq is directly used in the games presented in Chapter 5 – by contrast, the lemmas discussed in the previous section are more often relevant as a basis to prove the lemmas in this section, or in later sections.

The first statement presented here is a more powerful variant of the completeness of the function findseq (see Axiom 6.14). Indeed, the definition of findseq 's completeness postulates only that whenever there is a chain in a well-formed map T ending in (x, y) , then $\text{findseq}(x, y, T)$ will find *something*. Now, the uniqueness of chains presented in the previous section gives us that there can be but a single chain ending in (x, y) in a well-formed map T . Therefore, it can be derived from both the soundness and completeness of findseq that the chain returned by $\text{findseq}(x, y, T)$ is indeed the very chain ending in (x, y) .

Lemma 6.34 (Completeness and uniqueness of findseq). *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x a block, y a chaining value, and $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list a list. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $\text{ischain}(T, \text{IV}, c, y)$
4. $\text{unpad}(\text{mapfst}(c) \parallel [x]) \neq \text{None}$

Then it holds that $\text{findseq}(x, y, T) = \text{Some}(c)$.

Proof. By the completeness of findseq , we obtain that $\text{findseq}(x, y, T) \neq \text{None}$. By the soundness of findseq , this implies $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x, y, T)}, y)$. Now, Lemma 6.32 applies and we obtain $\pi_{\text{findseq}(x, y, T)} = c$, which implies $\text{findseq}(x, y, T) = \text{Some}(c)$. \square

Dually, we observe that if there is a chain ending in (x, y) in a well-formed map T which does *not* correspond to a valid padding, then $\text{findseq}(x, y, T)$ will return nothing at all.

Lemma 6.35. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x a block, y a chaining value, and $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list a list. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $\text{ischain}(T, \text{IV}, c, y)$
4. $\text{unpad}(\text{mapfst}(c) \parallel [x]) = \text{None}$

Then it holds that $\text{findseq}(x, y, T) = \text{None}$.

Proof. Assume for contradiction that $\text{findseq}(x, y, T) \neq \text{None}$. Then, by the soundness of findseq , we get $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x,y,T)}, y)$ and $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x,y,T)}) \parallel [x]) \neq \text{None}$. Furthermore, Lemma 6.32 implies that $\pi_{\text{findseq}(x,y,T)} = c$. Therefore, we obtain that $\text{unpad}(\text{mapfst}(c) \parallel [x]) \neq \text{None}$, which yields the contradiction. \square

We also show that the padding associated to a chain returned by findseq uniquely determines this chain (similarly as in Lemma 6.28), and further uniquely determines the arguments of findseq that must have been used to return this chain. Therefore, any complete chain can only be retrieved using uniquely defined arguments to the function findseq .

Lemma 6.36. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x_1, x_2 be blocks, and y_1, y_2 be chaining values. Assume the following:*

1. $\text{findseq}(x_1, y_1, T) \neq \text{None}$
2. $\text{findseq}(x_2, y_2, T) \neq \text{None}$
3. $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x_1,y_1,T)}) \parallel [x_1]) = \text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x_2,y_2,T)}) \parallel [x_2])$

Then it holds that $(x_1, y_1) = (x_2, y_2)$.

Proof. First, since we have that the function unpad is injective (see Axiom 6.9) and that $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x_1,y_1,T)}) \parallel [x_1]) \neq \text{None}$ (by the soundness of findseq), we obtain that $\text{mapfst}(\pi_{\text{findseq}(x_1,y_1,T)}) \parallel [x_1] = \text{mapfst}(\pi_{\text{findseq}(x_2,y_2,T)}) \parallel [x_2]$. Therefore, it holds that $\text{mapfst}(\pi_{\text{findseq}(x_1,y_1,T)}) = \text{mapfst}(\pi_{\text{findseq}(x_2,y_2,T)})$ and $x_1 = x_2$. Additionally, the soundness of findseq implies that $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x_1,y_1,T)}, y_1)$ and $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x_2,y_2,T)}, y_2)$. Hence, Lemma 6.28 applies and we obtain $\pi_{\text{findseq}(x_1,y_1,T)} = \pi_{\text{findseq}(x_2,y_2,T)}$ and $y_1 = y_2$, which finishes the proof. \square

The following lemma makes explicit a relation between the function findseq and the predicate ischain . Namely, as already discussed in Section 6.1.4, the function findseq searches in a map T for a specific complete chain, but a complete chain does not require its last element to be in T (see Definition 5.2); by contrast, the predicate ischain requires all the elements of its list argument to be in the domain of T . Thus, for the predicate $\text{ischain}(T, \text{IV}, c, z)$ to hold for a complete chain c in T , we need that the last element of c is also in the domain of T . We state this in a slightly generalized fashion.

Lemma 6.37. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y be a chaining value. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{findseq}(x, y, T') \neq \text{None}$
3. $(x, y) \in \text{dom}(T)$

Then it holds that $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x,y,T')} \parallel [(x, y)], T[(x, y)])$.

Proof. To show the conclusion of the statement, we show the three necessary and sufficient conditions required by Lemma 6.21:

- (i) $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x,y,T')}, y)$: The soundness of the function findseq immediately implies that $\text{ischain}(T', \text{IV}, \pi_{\text{findseq}(x,y,T')}, y)$. Therefore, we obtain this subgoal by applying Lemma 6.25.

- (ii) $(x, y) \in \text{dom}(T)$: We require this as a premise in the main statement.
- (iii) $\text{ischain}(T, T[(x, y)], \text{nil}, T[(x, y)])$: This is obvious since this statement is equivalent to $T[(x, y)] = T[(x, y)]$ by Definition 6.12.

□

Similarly, if $\text{findseq}(x, y, T')$ is successful in a well-formed map $T' \subseteq T$, and there is an element in T which maps to y , then this element must be the penultimate element of the chain found by findseq , and therefore it is also in the domain of T' .

Lemma 6.38. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x, x' be blocks, and y, y' be chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $(x', y') \in \text{dom}(T)$
5. $T[(x', y')] = y$
6. $\text{findseq}(x, y, T') \neq \text{None}$

Then it holds that $(x', y') \in \text{dom}(T')$.

Proof. First, by the soundness of findseq we know that $\text{ischain}(T', \text{IV}, \pi_{\text{findseq}(x, y, T')}, y)$. Next, we will show that $\pi_{\text{findseq}(x, y, T')}$ has the form $c \parallel [(x', y')]$, for an appropriate list c . To do this, we perform a case distinction on the structure of the list $\pi_{\text{findseq}(x, y, T')}$.

Case $\pi_{\text{findseq}(x, y, T')} = \text{nil}$. Here we obtain a contradiction. Indeed, in this case it holds $\text{ischain}(T', \text{IV}, \text{nil}, y)$ and thus $y = \text{IV}$. Recall that we have $(x', y') \in \text{dom}(T)$ and $T[(x', y')] = y = \text{IV}$, contradicting the assumption $\text{IV} \notin \text{ran}(T)$.

Case $\pi_{\text{findseq}(x, y, T')} = a \parallel [a]$. It suffices to show that $a = (x', y')$. From the assumption $\text{ischain}(T', \text{IV}, a \parallel [a], y)$ we conclude (using Lemma 6.21) that $a \in \text{dom}(T')$ and $\text{ischain}(T', T'[a], \text{nil}, y)$, implying $T'[a] = y$ by Definition 6.12. Therefore, it also holds $a \in \text{dom}(T)$ and $T[a] = y$. So by the injectivity of T , we conclude $(x', y') = a$.

Now we know that $\pi_{\text{findseq}(x, y, T')}$ has the form $c \parallel [(x', y')]$. We fix c , and finally obtain $\text{ischain}(T', \text{IV}, c \parallel [(x', y')], y)$. By Lemma 6.21, we therefore obtain $(x', y') \in \text{dom}(T')$. □

In the next few lemmas, we will compare the behavior of the calls $\text{findseq}(x, y, T')$ and $\text{findseq}(x, y, T)$ where $T' \subseteq T$. More precisely, we are interested to see under which conditions we have $\text{findseq}(x, y, T') = \text{findseq}(x, y, T)$, or if not, what conclusions we can draw from this. We first introduce the notion of a *completing* update.

Definition 6.39 (Completing update). Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x, x' be blocks, and y, y' and z' be chaining values. Assume $\text{findseq}(x, y, T) = \text{None}$ and $\text{findseq}(x, y, \text{upd}(T, (x', y'), z')) \neq \text{None}$. Then the update of T with the pair $((x', y'), z')$ is called *completing* for the chain ending in (x, y) .

We show below that when an update is completing for some chain in a well-formed map, then the element which the map was updated with is part of this chain.

Lemma 6.40. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x, x' be blocks, and y, y' and z' be chaining values. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $\text{findseq}(x, y, T) = \text{None}$
4. $\text{findseq}(x, y, \text{upd}(T, (x', y'), z')) \neq \text{None}$

Then it holds that $(x', y') \in \pi_{\text{findseq}(x, y, \text{upd}(T, (x', y'), z'))}$.

Proof. First, note that the fourth premise implies by the soundness of the function findseq that $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, \text{upd}(T, (x', y'), z'))} \parallel [x])) \neq \text{None}$. Therefore, it must hold that $\neg \text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x, y, \text{upd}(T, (x', y'), z'))}, y)$: Assume for contradiction that it held true that $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x, y, \text{upd}(T, (x', y'), z'))}, y)$, then the completeness of findseq would imply that $\text{findseq}(x, y, T) \neq \text{None}$, contradicting the third premise. Thus, it must indeed hold that $\neg \text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x, y, \text{upd}(T, (x', y'), z'))}, y)$. On the other hand, the fourth premise implies that $\text{ischain}(\text{upd}(T, (x', y'), z'), \text{IV}, \pi_{\text{findseq}(x, y, \text{upd}(T, (x', y'), z'))}, y)$ by the soundness of findseq . Hence we can apply Lemma 6.27 to conclude $(x', y') \in \pi_{\text{findseq}(x, y, \text{upd}(T, (x', y'), z'))}$. \square

As mentioned above, we will now see under which conditions we can indeed say that $\text{findseq}(x, y, T') = \text{findseq}(x, y, T)$. Indeed, it is often relevant in the proof outlined in Section 5.2 to know that whenever findseq finds a complete chain for some arguments x, y in a well-formed map T' , then it will still find the same chain after benign updates to T' . Similarly, if it does not find such a complete chain in T' , then it should still not find any chain after benign updates that are not completing. We will now see several lemmas that describe different conditions under which we expect the behavior of findseq to be the same before and after updates to a map.

We begin with the following helper lemma: Whenever the function call $\text{findseq}(x, y, T')$ finds a chain in a well-formed map T' for a block x and chaining value y , then it will still find the same chain in a well-formed map $T \supseteq T'$. This lemma is also interesting on its own. For instance, in the transformation described in Section 5.2, this lemma intuitively implies that any complete chain in the map T' maintained by the simulator is also a complete chain in the bigger map T maintained by the FIL-RO f (resp. f_{bad}) in the game \mathbf{G}_{real} .

Lemma 6.41. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y be a chaining value. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $\text{findseq}(x, y, T') \neq \text{None}$

Then it holds that $\text{findseq}(x, y, T) = \text{findseq}(x, y, T')$.

Proof. By the soundness of function findseq , we find $\text{ischain}(T', \text{IV}, \pi_{\text{findseq}(x, y, T')}, y)$ and $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, T')} \parallel [x])) \neq \text{None}$. By Lemma 6.25, the former moreover implies

that $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x,y,T')}, y)$. Hence, by the completeness and uniqueness of findseq (see Lemma 6.34), we conclude $\text{findseq}(x, y, T) = \text{Some}(\pi_{\text{findseq}(x,y,T')})$, which implies the main conclusion. \square

Using the above lemma, we can now easily derive a corollary which explicitly imposes a single benign update, but allows arbitrarily many of them (as we speak about a well-formed map $T \supseteq T'$, which may be obtained from T' by a number of appropriate updates).

Corollary 6.42. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x, x' be blocks, and y, y' and z' be a chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $(x', y') \notin \text{dom}(T)$
5. $z' \notin \text{ran}(T)$
6. $z' \neq \text{IV}$
7. $\text{findseq}(x, y, T') \neq \text{None}$

Then it holds that $\text{findseq}(x, y, \text{upd}(T, (x', y'), z')) = \text{findseq}(x, y, T')$.

Proof. The corollary quickly follows from Lemma 6.41. To be able to apply the latter lemma, we need to fulfill its premises by showing that (i) $\text{Inclusion}(T', \text{upd}(T, (x', y'), z'))$, (ii) $\text{Injective}(\text{upd}(T, (x', y'), z'))$ and (iii) $\text{IV} \notin \text{ran}(\text{upd}(T, (x', y'), z'))$. The first condition is easily fulfilled by the first and fourth premises. The second condition is fulfilled by the second and fifth premises. Finally, the last condition is fulfilled by the third and sixth premises. \square

In the above corollary, we spoke about the *same* parameters x and y for the call of findseq in the premises and the conclusion. We will now state two lemmas stating sensible conditions under which findseq will yield the same return values for *arbitrary* parameters x', y' in the conclusion. The first considers the case where a map T' is updated with a value that corresponds to the last element of an already complete chain in T' . It makes use of Lemma 6.33: this lemma essentially stated that the end of any complete chain in a map T' may not be part of any other complete chain in T' . Therefore, if we know that $\text{findseq}(x, y, T')$ is successful, and therefore (x, y) is the end of a complete chain in T' , then extending the domain of T' with the element (x, y) will not influence the return value of findseq for arbitrary parameters x', y' : Neither in the case $(x, y) = (x', y')$, since the end of a complete chain is not required to be in T' for findseq to be successful; nor in the case $(x, y) \neq (x', y')$, since (x, y) cannot be part of any other chain in T' .

Lemma 6.43. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x, x' be blocks, and y, y' be chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $\text{findseq}(x, y, T') \neq \text{None}$

5. $(x, y) \in \text{dom}(T)$

Then it holds that $\text{findseq}(x', y', \text{upd}(T', (x, y), T[(x, y)])) = \text{findseq}(x', y', T')$.

Proof. First, we perform a case distinction on the return value of $\text{findseq}(x', y', T')$.

Case $\text{findseq}(x', y', T') \neq \text{None}$. In this case, the conclusion follows from Lemma 6.41, since it is easily seen that the update with the pair $((x, y), T[(x, y)])$ is benign, i.e. it holds $\text{Inclusion}(T', \text{upd}(T', (x, y), T[(x, y)]))$, and the map $\text{upd}(T', (x, y), T[(x, y)])$ is well-formed.

Case $\text{findseq}(x', y', T') = \text{None}$. In this case, we have to show that the update of the map T' with the pair $((x, y), T[(x, y)])$ is not completing for any chain in T' , i.e. it also holds that $\text{findseq}(x', y', \text{upd}(T', (x, y), T[(x, y)])) = \text{None}$. For the sake of readability in the remainder of the proof, let us define $T'' := \text{upd}(T', (x, y), T[(x, y)])$. Note that $T' \subseteq T'' \subseteq T$, and that all these maps are well-formed. Now, assume for contradiction that $\text{findseq}(x', y', T'') \neq \text{None}$. Then, on the one hand, the update of T' was completing for the chain ending in (x', y') : We can use Lemma 6.40 to conclude that $(x, y) \in \pi_{\text{findseq}(x', y', T'')}$. On the other hand, we also know (fourth premise) that $\text{findseq}(x, y, T') \neq \text{None}$. Therefore, by Lemma 6.41 it also holds that $\text{findseq}(x, y, T'') \neq \text{None}$. Hence, by the soundness of findseq it holds that $\text{ischain}(T'', \text{IV}, \pi_{\text{findseq}(x, y, T'')}, y)$ and that $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, T'')}) \parallel [x]) \neq \text{None}$. Furthermore, since we assume that $\text{findseq}(x', y', T'') \neq \text{None}$, it similarly holds that $\text{ischain}(T'', \text{IV}, \pi_{\text{findseq}(x', y', T'')}, y')$ and that $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x', y', T'')}) \parallel [x']) \neq \text{None}$. Hence, Lemma 6.33 applies and we obtain $(x, y) \notin \pi_{\text{findseq}(x', y', T'')}$, which yields the contradiction. Thus, we must indeed have $\text{findseq}(x', y', T'') = \text{None}$, which finishes the proof. □

In a corollary, we state the restriction that the update has to be benign in a more intuitive fashion, similarly as in Corollary 6.42.

Corollary 6.44. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x, x' be blocks, and y, y' and z be chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $(x, y) \notin \text{dom}(T)$
5. $z \notin \text{ran}(T)$
6. $z \neq \text{IV}$
7. $\text{findseq}(x, y, T') \neq \text{None}$

Then it holds that $\text{findseq}(x', y', \text{upd}(T', (x, y), z)) = \text{findseq}(x', y', T')$.

Proof. By instantiating Lemma 6.43 with the two maps T' and $\text{upd}(T', (x, y), z)$, we obtain the conclusion $\text{findseq}(x', y', \text{upd}(T', (x, y), \text{upd}(T', (x, y), z)[(x, y)])) = \text{findseq}(x', y', T')$, which is equivalent to the conclusion of this corollary, as clearly $\text{upd}(T', (x, y), z)[(x, y)] = z$. We are left to fulfill the premises of Lemma 6.43. This is not difficult:

- (i) $\text{Inclusion}(T', \text{upd}(T', (x, y), z))$: By the premises 1 and 4.
- (ii) $\text{Injective}(\text{upd}(T', (x, y), z))$: By the premises 1, 2 and 5.
- (iii) $\text{IV} \notin \text{ran}(\text{upd}(T', (x, y), z))$: By the premises 1, 3 and 6.
- (iv) $\text{findseq}(x, y, T') \neq \text{None}$: By premise 7.
- (v) $(x, y) \in \text{dom}(\text{upd}(T', (x, y), z))$: Trivial.

□

Lemma 6.43 considered the situation where an update to a map T with a pair $((x, y), z)$ did not complete a chain in T by asserting that (x, y) is already the end of some complete chain in T . The next lemma also ensures that the update is not completing, however it does not require (x, y) to be the end of some complete chain; instead, it requires that (x, y) does not fill a gap in an incomplete chain ending in (x', y') . More precisely, it requires the update to be benign along with the usual well-formedness conditions for T . Additionally, it also requires that $((x, y), z)$ is not a fixed-point (in the transformation described in Section 5.2, this is guaranteed by the failure event \mathbf{bad}_2), and that (x, y) is not the predecessor of another element of a chain in T (this is guaranteed by the failure events \mathbf{bad}_2 and \mathbf{bad}_3). Now, observe that if (x, y) is the penultimate element of an incomplete chain ending in (x', y') , and (x', y') is not in the domain of $\text{upd}(T, (x, y), z)$, then the update with the pair (x, y) may indeed be completing for the chain ending in (x', y') without violating any of the other conditions (and therefore the results of the calls to findseq in the conclusion of the following lemma would differ); the last assumption of the lemma allows us to exclude this case.

Lemma 6.45. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x, x' be blocks, and y, y' and z be chaining values. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $(x, y) \notin \text{dom}(T)$
4. $z \notin \text{ran}(T)$
5. $z \neq \text{IV}$
6. $z \neq y$
7. $\nexists x_0. (x_0, z) \in \text{dom}(T)$
8. $(x', y') \in \text{dom}(\text{upd}(T, (x, y), z))$

Then it holds that $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) = \text{findseq}(x', y', T)$.

Proof. In the case where $\text{findseq}(x', y', T) \neq \text{None}$, the statement immediately follows from Corollary 6.42. Therefore, let us concentrate on the case where $\text{findseq}(x', y', T) = \text{None}$. Then, we need to show that also $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) = \text{None}$; that is, the update of T is not completing for any chain. We prove the claim by contradiction: Assume that $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) \neq \text{None}$. We expect (and show below) that this yields a contradiction to the premises of the statement. Since we assume $\text{findseq}(x', y', T) = \text{None}$ and $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) \neq \text{None}$, Lemma 6.40 applies and we conclude $(x, y) \in \pi_{\text{findseq}(x', y', \text{upd}(T, (x, y), z))}$. Furthermore, by the soundness of findseq we obtain that $\text{ischain}(\text{upd}(T, (x, y), z), \text{IV}, \pi_{\text{findseq}(x', y', \text{upd}(T, (x, y), z))}, y')$. For the remainder of the proof, let

us define $c := \pi_{\text{findseq}(x', y', \text{upd}(T, (x, y), z))}$. Now, using only the premises $(x, y) \in c$ and $\text{ischain}(\text{upd}(T, (x, y), z), \text{IV}, c, y')$ and the premises δ through δ of the main statement, we derive a contradiction using an induction over c , generalizing the constant IV as an arbitrary value \hat{y} . That is, under the premises δ through δ we prove the following statement:

$$\forall \hat{y}. \text{ischain}(\text{upd}(T, (x, y), z), \hat{y}, c, y') \Rightarrow (x, y) \in c \Rightarrow \perp$$

Base case: $c = \text{nil}$. This yields a contradiction immediately since we assume $(x, y) \in \text{nil}$.

Induction step: $c = a::al$.

Induction hypothesis:

$$\forall \hat{y}. \text{ischain}(\text{upd}(T, (x, y), z), \hat{y}, al, y') \Rightarrow (x, y) \in al \Rightarrow \perp$$

We treat the cases $a \neq (x, y)$ and $a = (x, y)$ separately.

Case $a \neq (x, y)$. This is the easier case, since in this case we can use the induction hypothesis to obtain the contradiction. Indeed, the two premises of the induction hypothesis are easily fulfilled: we have $\text{ischain}(\text{upd}(T, (x, y), z), \hat{y}, a::al, y')$, which immediately implies $\text{ischain}(\text{upd}(T, (x, y), z), \text{upd}(T, (x, y), z)[a], al, y')$ by Lemma 6.21. Furthermore, since we have $a \neq (x, y)$ and $(x, y) \in a::al$, we conclude $(x, y) \in al$.

Case $a = (x, y)$. This case is slightly more difficult, since we cannot use the induction hypothesis. However, since $a = (x, y)$, we now have the assumption $\text{ischain}(\text{upd}(T, (x, y), z), \hat{y}, (x, y)::al, y')$. We will now show that this yields a contradiction to the premises δ through δ . To achieve this, we need another case distinction to distinguish whether there is another element in this chain after the element (x, y) .

Case $al = \text{nil}$. By Definition 6.12, we obtain that $\text{upd}(T, (x, y), z)[(x, y)] = y'$, which immediately implies $z = y'$. Now, we quickly obtain a contradiction using a last case distinction.

Case $(x, y) = (x', y')$. Then $y' = y$, and thus $z = y$, contradiction to the sixth premise.

Case $(x, y) \neq (x', y')$. We have that $(x', y') \in \text{dom}(\text{upd}(T, (x, y), z))$ (eighth premise), and since $(x, y) \neq (x', y')$ we know $(x', y') \in \text{dom}(T)$. Since $z = y'$, we obtain $(x', z) \in \text{dom}(T)$, contradiction to the seventh premise.

Case $al = a'::al'$. By Definition 6.12, we get $\text{upd}(T, (x, y), z)[(x, y)] = \text{snd}(a')$, implying $z = \text{snd}(a')$, and $\text{ischain}(\text{upd}(T, (x, y), z), \text{snd}(a'), a'::al', y')$. Again, we perform a last case distinction to derive the contradiction.

Case $(x, y) = a'$. Then $\text{snd}(a') = y$, and thus $z = y$, contradiction to the sixth premise.

Case $(x, y) \neq a'$. We have $\text{ischain}(\text{upd}(T, (x, y), z), \text{snd}(a'), a'::a', y')$. By Lemma 6.21 this implies $a' \in \text{dom}(\text{upd}(T, (x, y), z))$. Since $(x, y) \neq a'$, we conclude $a' \in \text{dom}(T)$. Because $z = \text{snd}(a')$, we obtain $(\text{fst}(a'), z) \in \text{dom}(T)$, contradiction to the seventh premise.

Finally, our original assumption that $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) \neq \text{None}$ led to a contradiction. Thus, it must indeed hold that $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) = \text{None}$ and therefore $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) = \text{findseq}(x', y', T)$, which finishes the proof. \square

6.6 Lemmas on valid_chain

The predicate `valid_chain` (see Definition 6.15) is closely related to the function `findseq`. Namely, $\text{valid_chain}(T, c \parallel [(x, y)])$ means that $c \parallel [(x, y)]$ is a complete chain in T and additionally $(x, y) \in \text{dom}(T)$. By contrast, $\text{findseq}(x, y, T) = \text{Some}(c)$ only means that $c \parallel [(x, y)]$ is a complete chain in T , but the element (x, y) is not required to be in the domain of T . Therefore the former statement (involving `valid_chain`) is a slightly stronger statement than the latter statement (involving `findseq`); we use this predicate to speak about complete chains that are entirely contained in a map, i.e. in particular their last element is also in the domain of T .

We begin by observing two basic properties concerning the predicate `valid_chain`. First, we formally state the relation between `findseq` and `valid_chain` described above in a generalized fashion as follows.

Lemma 6.46. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y be a chaining value. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{findseq}(x, y, T') \neq \text{None}$
3. $(x, y) \in \text{dom}(T)$

Then it holds that $\text{valid_chain}(T, \pi_{\text{findseq}(x, y, T')} \parallel [(x, y)])$.

Proof. By Lemma 6.37, we obtain $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x, y, T')} \parallel [(x, y)], T[(x, y)])$. Furthermore, by the soundness of `findseq` we know $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, T')} \parallel [x])) \neq \text{None}$. Thus it holds $\text{valid_chain}(T, \pi_{\text{findseq}(x, y, T')} \parallel [(x, y)])$ by the definition of `valid_chain`. \square

We also observe that the predicate cannot hold for a non-empty list in the empty map (clearly, it never holds for the empty list because of Axiom 6.6). This is needed in the game transition outlined in Section 5.2 to conclude that there are no valid chains in the initial games (i.e., before the simulator has made any calls).

Lemma 6.47. *Let $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list and xy be a pair of a block and a chaining value. Then it holds that $\neg \text{valid_chain}(\emptyset, c \parallel [xy])$.*

Proof. Assume $\text{valid_chain}(\emptyset, c \parallel [xy])$, then we know $\text{ischain}(\emptyset, \text{IV}, c \parallel [xy], z)$ for some z . By Lemma 6.21 we obtain $xy \in \text{dom}(\emptyset)$, yielding a contradiction. \square

We also state an alternative variant of the completeness and uniqueness of the function `findseq` (see Lemma 6.34), using the predicate `valid_chain`.

Lemma 6.48 (Completeness and uniqueness of `findseq` under `valid_chain`). *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, x a block, y a chaining value, and $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list a list. Assume the following:*

1. `Injective(T)`
2. $\text{IV} \notin \text{ran}(T)$
3. `valid_chain(T, c || [(x, y)])`

Then it holds that $\text{findseq}(x, y, T) = \text{Some}(c)$.

Proof. We obtain from the premise `valid_chain(T, c || [(x, y)])` that for some z it holds that `ischain(T, IV, c || [(x, y)], z)`, and that `unpad(mapfst(c) || [x]) ≠ None`. The former implies `ischain(T, IV, c, y)` by Lemma 6.21. Hence, the conclusion follows by the completeness and uniqueness of `findseq` (Lemma 6.34). \square

Finally, we consider the behavior of predicate `valid_chain` when the map that it relates to gets updated. The first lemma simply states that we can extend a map T with an element not present in its domain without afflicting the validity of the predicate.

Lemma 6.49. *Let $T : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a map, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, xy be a pair of a block and a chaining value, and z be a chaining value. Assume:*

1. `valid_chain(T, c)`
2. $xy \notin \text{dom}(T)$

Then it holds that `valid_chain(upd(T, xy, z), c)`.

Proof. We have to show $\exists z'. \text{ischain}(\text{upd}(T, xy, z), \text{IV}, c, z') \wedge \text{unpad}(\text{mapfst}(c)) \neq \text{None}$. From `valid_chain(T, c)` we obtain a z' such that `ischain(T, IV, c, z')` and `unpad(mapfst(c)) ≠ None`. It remains only to show `ischain(upd(T, xy, z), IV, c, z')`. This follows quickly from Lemma 6.24; to apply it, we only need to show $xy \notin c$. But this is clear: If we had $xy \in c$, then by Corollary 6.23 it would hold $xy \in \text{dom}(T)$, yielding a contradiction. \square

The last two lemmas of this section are variants of Lemma 6.45. Intuitively, we would like to say that when a well-formed map T gets updated with a pair $((x, y), z)$, and this update is benign and not completing, then for any list c , we have `valid_chain(T, c)` if and only if `valid_chain(upd(T, (x, y), z), c)`. However, we need an additional condition for this to hold: Indeed, it could happen that (x, y) is the *last* element of c . In this case, the latter statement may hold, but the former may not, even though the update was not completing. This is the main difference to Lemma 6.45; there, an additional assumption was needed to ensure that the element (x, y) was not the *penultimate* element of the chain returned by `findseq`. Both of the following two lemmas are needed at different points in the `EasyCrypt` proof outlined in Section 5.2. The first uses the same premises as Lemma 6.45, and additionally assumes `findseq(x, y, T) = None`, thereby restricting the statement to the case where (x, y) is not the last element of a complete chain in T .

Lemma 6.50. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x, x' be blocks, and y, y' and z be chaining values. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $(x, y) \notin \text{dom}(T)$
4. $z \notin \text{ran}(T)$
5. $z \neq \text{IV}$
6. $z \neq y$
7. $\nexists x_0. (x_0, z) \in \text{dom}(T)$
8. $(x', y') \in \text{dom}(\text{upd}(T, (x, y), z))$
9. $\text{findseq}(x, y, T) = \text{None}$

Then it holds that $\text{valid_chain}(\text{upd}(T, (x, y), z), c \parallel [(x', y')]) \Leftrightarrow \text{valid_chain}(T, c \parallel [(x', y')])$.

Proof. The direction “ \Leftarrow ” of the conclusion follows immediately from Lemma 6.49. The interesting case is the direction “ \Rightarrow ”: Assume that $\text{valid_chain}(\text{upd}(T, (x, y), z), c \parallel [(x', y')])$. Then, we show that $\text{valid_chain}(T, c \parallel [(x', y')])$. First, we apply the completeness and uniqueness of findseq under valid_chain to conclude that $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) \neq \text{None}$ and $\pi_{\text{findseq}(x', y', \text{upd}(T, (x, y), z))} = c$ (note that $\text{upd}(T, (x, y), z)$ is still well-formed since the update is benign). Now, since Lemma 6.45 implies $\text{findseq}(x', y', \text{upd}(T, (x, y), z)) = \text{findseq}(x', y', T)$, we also know that $\text{findseq}(x', y', T) \neq \text{None}$, and what we need to show is $\text{valid_chain}(T, \pi_{\text{findseq}(x', y', T)} \parallel [(x', y')])$. We perform a case distinction.

Case $(x, y) = (x', y')$. This is the case where the new element used for the update is the last element of a complete chain. It yields a contradiction to the ninth premise, since we have $\text{findseq}(x', y', T) \neq \text{None}$.

Case $(x, y) \neq (x', y')$. In this case, the premise $(x', y') \in \text{dom}(\text{upd}(T, (x, y), z))$ implies $(x', y') \in \text{dom}(T)$. Then, the claim follows by Lemma 6.46.

□

The next lemma uses the same premises as Lemma 6.45, except that the premise $(x', y') \in \text{dom}(\text{upd}(T, (x, y), z))$ is strengthened as $(x', y') \in \text{dom}(T)$: Thereby, as in the previous lemma we exclude the case where the new element is the last element of a complete chain.

Lemma 6.51. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x, x' be blocks, and y, y' and z be chaining values. Assume the following:*

1. $\text{Injective}(T)$
2. $\text{IV} \notin \text{ran}(T)$
3. $(x, y) \notin \text{dom}(T)$
4. $z \notin \text{ran}(T)$
5. $z \neq \text{IV}$
6. $z \neq y$
7. $\nexists x_0. (x_0, z) \in \text{dom}(T)$
8. $(x', y') \in \text{dom}(T)$

Then it holds that $\text{valid_chain}(\text{upd}(T, (x, y), z), c \parallel [(x', y')]) \Leftrightarrow \text{valid_chain}(T, c \parallel [(x', y')])$.

Proof. Note that $(x', y') \in \text{dom}(T)$ implies $(x', y') \in \text{dom}(\text{upd}(T, (x, y), z))$. Now, the proof is analogous to the proof of Lemma 6.50, except in the case where $(x, y) = (x', y')$: In this case, here we obtain a contradiction between the premises $(x, y) \notin \text{dom}(T)$ and $(x', y') \in \text{dom}(T)$. \square

6.7 Lemmas on set_bad3

In this section, we discuss lemmas concerning the predicate `set_bad3`, described in Section 6.1.5. More precisely, we will see under which conditions the maps T and T' in game $\mathbf{G}_{\text{real}'}$ may be updated with the same pair $((x, y), z)$ (this may happen upon calls to the simulator f_q) such that $\text{set_bad3}(z_0, T', T) = \text{set_bad3}(z_0, \text{upd}(T', (x, y), z), \text{upd}(T, (x, y), z))$. Essentially, as for the Lemmas 6.43 and 6.45 discussed at the end of Section 6.5, we need to ensure that the update with $((x, y), z)$ is not completing for any chain in T . Firstly, we need that T is well-formed and that the update is benign. Then, we again consider two cases: either the pair (x, y) is the last element of a complete chain in T' , but then (x, y) cannot be part of any other chain in T' (this is the first lemma); or the pair (x, y) does not correspond to the last element of any complete chain, but in this case we require that it does not fill a gap by requiring that the pair $((x, y), z)$ is not a fixed-point and that z was not already used as a chaining value in a query to the simulator (this is the second lemma); in both cases, the update cannot complete a chain.

In the first lemma, to ensure that the newly added value is the end of a complete chain, we require that $\text{findseq}(x, y, T') \neq \text{None}$, similarly as was the case in Corollary 6.44.

Lemma 6.52. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y, z and z_0 be chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $(x, y) \notin \text{dom}(T)$
5. $z \notin \text{ran}(T)$
6. $z \neq \text{IV}$
7. $\text{findseq}(x, y, T') \neq \text{None}$

Then it holds that $\text{set_bad3}(z_0, \text{upd}(T', (x, y), z), \text{upd}(T, (x, y), z)) \Leftrightarrow \text{set_bad3}(z_0, T', T)$.

Proof. For the sake of readability of the proof, let us define $T'_{xyz} := \text{upd}(T', (x, y), z)$ and $T_{xyz} := \text{upd}(T, (x, y), z)$.

We begin with the direction “ \Rightarrow ”, which is the easier one. By unfolding Definition 6.16, we obtain the assumptions $z_0 \in \text{ran}(T_{xyz})$, $T_{xyz}^{-1}[z_0] \notin \text{dom}(T'_{xyz})$ and $\forall c. \neg \text{valid_chain}(c \parallel [T_{xyz}^{-1}[z_0]], T_{xyz})$. Furthermore, observe that $z_0 \neq z$: otherwise, we would have $T_{xyz}^{-1}[z_0] = (x, y)$, and the assumption $(x, y) \notin \text{dom}(T'_{xyz})$ in itself is contradictory. Therefore, we respectively obtain $z_0 \in \text{ran}(T)$, $T^{-1}[z_0] \notin \text{dom}(T')$, and $\forall c. \neg \text{valid_chain}(c \parallel [T^{-1}[z_0]], T_{xyz})$.

From the latter, we obtain $\forall c. \neg \text{valid_chain}(c \parallel [T^{-1}[z_0]], T)$ using Lemma 6.49. Hence, we conclude $\text{set_bad3}(z_0, T', T)$.

For the direction “ \Leftarrow ”, we have the assumptions $z_0 \in \text{ran}(T)$, $T^{-1}[z_0] \notin \text{dom}(T')$ and $\forall c. \neg \text{valid_chain}(c \parallel [T^{-1}[z_0]], T)$. Again, observe that $z_0 \neq z$, since otherwise we obtain a contradiction to the premise $z \notin \text{ran}(T)$. Similarly, note that $T^{-1}[z_0] \neq (x, y)$, since we have $z_0 \in \text{ran}(T)$, but $(x, y) \notin \text{dom}(T)$. We have to show the following:

- (i) $z_0 \in \text{ran}(T_{xyz})$: This is clear, since we have $z_0 \in \text{ran}(T)$ and $z_0 \neq z$.
- (ii) $T_{xyz}^{-1}[z_0] \notin \text{dom}(T'_{xyz})$: This is also clear, since we have $T^{-1}[z_0] \notin \text{dom}(T')$, $z_0 \neq z$ and $T^{-1}[z_0] \neq (x, y)$.
- (iii) $\forall c. \neg \text{valid_chain}(c \parallel [T_{xyz}^{-1}[z_0]], T_{xyz})$: Note that $T_{xyz}^{-1}[z_0] = T^{-1}[z_0]$ since $z_0 \neq z$. To show the goal, assume for contradiction that $\text{valid_chain}(c \parallel [T^{-1}[z_0]], T_{xyz})$ for some list c . This implies $\text{ischain}(T_{xyz}, \text{IV}, c \parallel [T^{-1}[z_0]], z_0)$ and $\text{unpad}(\text{mapfst}(c \parallel [T^{-1}[z_0]])) \neq \text{None}$. Now, on the one hand, we have $\neg \text{valid_chain}(c \parallel [T^{-1}[z_0]], T)$, and therefore it must hold that $\neg \text{ischain}(T, \text{IV}, c \parallel [T^{-1}[z_0]], z_0)$. Thus we can apply Lemma 6.27 to conclude $(x, y) \in c \parallel [T^{-1}[z_0]]$, and since $T^{-1}[z_0] \neq (x, y)$, we get $(x, y) \in c$. On the other hand, we have $\text{findseq}(x, y, T') \neq \text{None}$, which implies $\text{findseq}(x, y, T'_{xyz}) \neq \text{None}$ by Corollary 6.42. Hence, we know from the soundness of findseq that $\text{ischain}(T_{xyz}, \text{IV}, \pi_{\text{findseq}(x, y, T_{xyz})}, y)$ and $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, T_{xyz})})) \parallel [x] \neq \text{None}$. From $\text{ischain}(T_{xyz}, \text{IV}, c \parallel [T^{-1}[z_0]], z_0)$, we obtain by Lemma 6.21 that $\text{ischain}(T_{xyz}, \text{IV}, c, \text{snd}(T^{-1}[z_0]))$. Recall that $\text{unpad}(\text{mapfst}(c) \parallel \text{fst}([T^{-1}[z_0]])) \neq \text{None}$. Finally, Lemma 6.33 applies and we obtain $(x, y) \notin c$, which yields the contradiction. □

In the second lemma, as mentioned above we ensure that the update of T with the pair $((x, y), z)$ is not completing by requiring that the pair $((x, y), z)$ is not a fixed-point and that z was not previously used as a chaining value, similarly as for Lemma 6.45.

Lemma 6.53. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y, z and z_0 be chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $(x, y) \notin \text{dom}(T)$
5. $z \notin \text{ran}(T)$
6. $z \neq \text{IV}$
7. $z \neq y$
8. $\nexists x_0. (x_0, z) \in \text{dom}(T)$

Then it holds that $\text{set_bad3}(z_0, \text{upd}(T', (x, y), z), \text{upd}(T, (x, y), z)) \Leftrightarrow \text{set_bad3}(z_0, T', T)$.

Proof. As in the previous lemma, let us define $T'_{xyz} := \text{upd}(T', (x, y), z)$ and $T_{xyz} := \text{upd}(T, (x, y), z)$ for the sake of readability. The proof for the direction “ \Rightarrow ” is analogous to the previous proof, and the proof for direction “ \Leftarrow ” is analogous for the cases **(i)** $z_0 \in \text{ran}(T_{xyz})$ and **(ii)** $T_{xyz}^{-1}[z_0] \notin \text{dom}(T'_{xyz})$. For the case **(iii)** $\forall c. \neg \text{valid_chain}(c \parallel [T_{xyz}^{-1}[z_0]], T_{xyz})$, we note that $T_{xyz}^{-1}[z_0] = T^{-1}[z_0]$ since $z_0 \neq z$. Again, assume for contradiction $\text{valid_chain}(c \parallel [T^{-1}[z_0]], T_{xyz})$ for some list c . From this, as in the previous proof, we eventually obtain $(x, y) \in c$. Therefore, c has the form $c_\ell \parallel (x, y)::c_r$ for appropriate c_ℓ and c_r . Hence we have $\text{valid_chain}(c_\ell \parallel (x, y)::c_r \parallel [T^{-1}[z_0]], T_{xyz})$, which implies $\text{ischain}(T_{xyz}, \text{IV}, c_\ell \parallel (x, y)::c_r \parallel [T^{-1}[z_0]], z_0)$. We will now show that the latter statement leads to a contradiction to the premises 7 and 8. We begin with a case distinction on the structure of the list c_r .

Case $c_r = \text{nil}$. Using Lemma 6.21, we obtain $\text{ischain}(T_{xyz}, T_{xyz}[(x, y)], [T^{-1}[z_0]], z_0)$, and therefore $\text{ischain}(T_{xyz}, z, [T^{-1}[z_0]], z_0)$. By Definition 6.12, this implies $T^{-1}[z_0] \in \text{dom}(T_{xyz})$ and $z = \text{snd}(T^{-1}[z_0])$. We use another case distinction.

Case $T^{-1}[z_0] = (x, y)$. In this case we obtain $\text{snd}(T^{-1}[z_0]) = y$ and therefore $z = y$, contradiction to the seventh premise.

Case $T^{-1}[z_0] \neq (x, y)$. In this case we obtain $T^{-1}[z_0] \in \text{dom}(T)$, and therefore $(\text{fst}(T^{-1}[z_0]), z) \in \text{dom}(T)$, contradiction to the eighth premise.

Case $c_r = a::al$. By Lemma 6.21, it holds $\text{ischain}(T_{xyz}, T_{xyz}[(x, y)], a::al \parallel [T^{-1}[z_0]], z_0)$, and therefore $\text{ischain}(T_{xyz}, z, a::al \parallel [T^{-1}[z_0]], z_0)$. We perform another case distinction on the structure of al .

Case $al = \text{nil}$. Then, we have $\text{ischain}(T_{xyz}, z, a::T^{-1}[z_0]::\text{nil}, z_0)$. By Definition 6.12, this implies $a \in \text{dom}(T_{xyz})$ and $z = \text{snd}(a)$. We use another case distinction.

Case $a = (x, y)$. Then it holds that $\text{snd}(a) = y$, and therefore $z = y$, contradiction to the seventh premise.

Case $a \neq (x, y)$. Then it holds that $a \in \text{dom}(T)$, and therefore $(\text{fst}(a), z) \in \text{dom}(T)$, contradiction to the eighth premise.

Case $al = a'::al'$. In this case we have $\text{ischain}(T_{xyz}, z, a::a'::al' \parallel T^{-1}[z_0]::\text{nil}, z_0)$, and the rest is analogous to the case where $al = \text{nil}$.

Hence, the assumption $\text{valid_chain}(c \parallel [T^{-1}[z_0]], T_{xyz})$ leads to a contradiction and thus, since $T_{xyz}^{-1}[z_0] = T^{-1}[z_0]$ we must indeed have $\forall c. \neg \text{valid_chain}(c \parallel [T_{xyz}^{-1}[z_0]], T_{xyz})$, which finishes the last case, and the proof. \square

6.8 Lemmas on Claim5

In the last section of this chapter, we will discuss a few lemmas concerning the predicate Claim5, described in Section 6.1.6. Recall that Claim5(T', T), for well-formed maps $T' \subseteq T$, means that for any element (x, y) that belongs to a chain in T and that is also in the domain

of T' , all elements of this chain preceding the element (x, y) are also in the domain of T' , provided that there is no complete chain in T whose last element maps to the value y .

In a first lemma, we make this relation explicit: Assume that there is a complete chain of the form $c \parallel [(x', y'), (x, y)]$ in T , and (x', y') is also in the domain of T' . Then $c \parallel [(x', y')]$ is a chain in T' , since we can reason inductively that all elements of c are in the domain of T' .

Lemma 6.54. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, $c : (\{0, 1\}^k \times \{0, 1\}^n)$ list be a list, x, x' be blocks, and y, y' be chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $\text{ischain}(T, \text{IV}, c \parallel [(x', y')], y)$
5. $\text{unpad}(\text{mapfst}(c \parallel [(x', y')]) \parallel [x]) \neq \text{None}$
6. $(x', y') \in \text{dom}(T')$
7. $\text{Claim5}(T', T)$

Then it holds that $\text{ischain}(T', \text{IV}, c \parallel [(x', y')], y)$.

Proof. We will prove the claim by a structural induction over the list c . However, the premise $\text{unpad}(\text{mapfst}(c \parallel [(x', y')]) \parallel [x]) \neq \text{None}$ is not suitable for such an induction, since in general the fact that some chain is in the domain of unpad does not imply that the same holds for a suffix of this chain. Therefore, we replace this premise with a premise more suitable for induction; namely,

$$\forall x_0, y_0. (x_0, y_0) \in c \Rightarrow \text{findseq}(x_0, y_0, T) = \text{None}.$$

We eventually obtain the validity of this premise using Lemma 6.33: Fix an arbitrary element $(x_0, y_0) \in c$, and assume for contradiction that $\text{findseq}(x_0, y_0, T) \neq \text{None}$. By the soundness of function findseq , this implies that $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x_0, y_0, T)}, y_0)$ and that $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x_0, y_0, T)}) \parallel [x_0]) \neq \text{None}$. Hence, Lemma 6.33 applies and we obtain $(x_0, y_0) \notin c \parallel [(x', y')]$, which yields the contradiction since we have $(x_0, y_0) \in c$.

We will now show the claim using an induction over c , generalizing the constant IV as an arbitrary value \hat{y} . That is, under the premises 1 through 3, 6 and 7 we show:

$$\forall \hat{y}. \text{ischain}(T, \hat{y}, c \parallel [(x', y')], y) \Rightarrow (\forall x_0, y_0. (x_0, y_0) \in c \Rightarrow \text{findseq}(x_0, y_0, T) = \text{None}) \Rightarrow \text{ischain}(T', \hat{y}, c \parallel [(x', y')], y)$$

Base case: $c = \text{nil}$. Since we have $\text{ischain}(T, \hat{y}, [(x', y')], y)$ we know by Definition 6.12 that $\hat{y} = y'$ and $T[(x', y')] = y$, and we have to show $(x', y') \in \text{dom}(T')$, $\hat{y} = y'$ and $T'[(x', y')] = y$. This is trivial using premises 1 and 6.

Induction step: $c = a::al$.

Induction hypothesis:

$$\begin{aligned} \forall \hat{y}. \text{ischain}(T, \hat{y}, al \parallel [(x', y')], y) \Rightarrow \\ (\forall x_0, y_0. (x_0, y_0) \in al \Rightarrow \text{findseq}(x_0, y_0, T) = \text{None}) \Rightarrow \\ \text{ischain}(T', \hat{y}, al \parallel [(x', y')], y) \end{aligned}$$

We use a case distinction on the structure of the list al , and concentrate first on the case where $al = a'::al'$.

Case $al = a'::al'$. First, from the assumption $\text{ischain}(T, \hat{y}, a'::al' \parallel [(x', y')], y)$ we immediately obtain by Definition 6.12 that $a \in \text{dom}(T)$, $\hat{y} = \text{snd}(a)$, $T[a] = \text{snd}(a')$ and $\text{ischain}(T, \text{snd}(a'), a'::al' \parallel [(x', y')], y)$. The latter statement fulfills the first premise of the induction hypothesis (instantiated with $\hat{y} := \text{snd}(a')$). The second premise of the induction hypothesis is also easily fulfilled: since we know that $\forall x_0, y_0. (x_0, y_0) \in a::al \Rightarrow \text{findseq}(x_0, y_0, T) = \text{None}$, clearly it holds that $\forall x_0, y_0. (x_0, y_0) \in al \Rightarrow \text{findseq}(x_0, y_0, T) = \text{None}$. Hence, we obtain from the induction hypothesis that $\text{ischain}(T', \text{snd}(a'), a'::al' \parallel [(x', y')], y)$. Since we also have $\hat{y} = \text{snd}(a)$, to prove the goal $\text{ischain}(T', \hat{y}, a'::al' \parallel [(x', y')], y)$ it only remains to show that $a \in \text{dom}(T')$ and $T'[a] = \text{snd}(a')$ (see Definition 6.12). Note that we have $\text{Inclusion}(T', T)$, $a \in \text{dom}(T)$ and $T[a] = \text{snd}(a')$, so if $a \in \text{dom}(T')$ then $T'[a] = \text{snd}(a')$ immediately follows. Therefore we are only left to prove $a \in \text{dom}(T')$. This is where we need the premise $\text{Claim5}(T', T)$, which yields the conclusion since it implies that

$$\begin{aligned} a \in \text{dom}(T) \Rightarrow \\ T[a] = \text{snd}(a') \Rightarrow \\ a' \in \text{dom}(T') \Rightarrow \\ \text{findseq}(\text{fst}(a), \text{snd}(a), T) = \text{None} \Rightarrow a \in \text{dom}(T'). \end{aligned}$$

We have already fulfilled the first and second premises. Moreover, since we have $\text{ischain}(T', \text{snd}(a'), a'::al' \parallel [(x', y')], y)$ then $a' \in \text{dom}(T')$ follows immediately by Lemma 6.21. Finally, the last premise also holds since we have that $\forall x_0, y_0. (x_0, y_0) \in a::al \Rightarrow \text{findseq}(x_0, y_0, T) = \text{None}$, and we take $x_0 = \text{fst}(a)$ and $y_0 = \text{snd}(a)$. Thus, we obtain $a \in \text{dom}(T')$, which finishes this case.

Case $al = \text{nil}$. This case is analogous to the case $al = a'::al'$: Simply replace in the above proof the list $a'::al'$ with nil , respectively the variable a' (when on its own) with the pair (x', y') .

□

In a sense, both the predicates `set_bad3` and `Claim5` ensure that queries of the distinguisher can only be made in order. The difference is that `set_bad3` does this locally, i.e. for

a given chaining value y used in a query from the distinguisher, while **Claim5** guarantees a correct ordering of all queries globally, i.e. for all elements in map T' . Therefore we expect, and state formally below, the following relation between these two predicates. If **Claim5**(T', T) holds for given well-formed maps $T' \subseteq T$, then both maps can be updated with a pair $((x, y), z)$ such that **Claim5**($\text{upd}(T', (x, y), z), \text{upd}(T, (x, y), z)$) still holds, provided that the update is benign, and event **bad₃** is not triggered for the chaining value y (i.e., it holds that $\neg \text{set_bad3}(y, T', T)$). This lemma constitutes an important element to derive that as long as none of the failure events is triggered, the statement **Claim5**(T', T) is an invariant of game \mathbf{G}_{real} , meaning that the distinguisher can only make queries in order up to the failure events – this is an essential property needed in the transformation described in Section 5.2.

Lemma 6.55. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y and z be chaining values. Assume the following:*

1. **Inclusion**(T', T)
2. **Injective**(T)
3. $\text{IV} \notin \text{ran}(T)$
4. $(x, y) \notin \text{dom}(T)$
5. $z \notin \text{ran}(T)$
6. $z \neq \text{IV}$
7. $\neg \text{set_bad3}(y, T', T)$
8. **Claim5**(T', T)

*Then it holds that **Claim5**($\text{upd}(T', (x, y), z), \text{upd}(T, (x, y), z)$).*

Proof. We want to show the following:

$$\begin{aligned} \forall (x_0, y_0), (x'_0, y'_0). & (x'_0, y'_0) \in \text{dom}(\text{upd}(T, (x, y), z)) \wedge \\ & (x_0, y_0) \in \text{dom}(\text{upd}(T', (x, y), z)) \wedge \\ & \text{upd}(T, (x, y), z)[(x'_0, y'_0)] = y_0 \wedge \\ & \text{findseq}(x'_0, y'_0, \text{upd}(T, (x, y), z)) = \text{None} \Rightarrow \\ & (x'_0, y'_0) \in \text{dom}(\text{upd}(T', (x, y), z)). \end{aligned}$$

Fix x_0, y_0, x'_0 and y'_0 . Assume that the four premises of the above statement hold. Then, we show that $(x'_0, y'_0) \in \text{dom}(\text{upd}(T', (x, y), z))$. If $(x'_0, y'_0) = (x, y)$, this is trivial. Therefore assume that $(x'_0, y'_0) \neq (x, y)$. Now, since we have $(x_0, y_0) \in \text{dom}(\text{upd}(T', (x, y), z))$, it either holds that $(x_0, y_0) = (x, y)$, or $(x_0, y_0) \in \text{dom}(T')$. The basic idea of the proof is that in the former case, we obtain the conclusion because we know $\neg \text{set_bad3}(y, T', T)$, and in the latter case, we obtain the conclusion because we know **Claim5**(T', T). We perform a case distinction.

Case $(x_0, y_0) = (x, y)$. Then, we obtain $\text{upd}(T, (x, y), z)[(x'_0, y'_0)] = y$. Since $(x'_0, y'_0) \neq (x, y)$, we furthermore conclude $(x'_0, y'_0) \in \text{dom}(T)$ and $T[(x'_0, y'_0)] = y$. Moreover, we also know $\text{findseq}(x'_0, y'_0, T) = \text{None}$: If we had $\text{findseq}(x'_0, y'_0, T) \neq \text{None}$, we

would obtain by Corollary 6.42 that $\text{findseq}(x'_0, y'_0, \text{upd}(T, (x, y), z)) \neq \text{None}$, yielding a contradiction. We will now use the premise $\neg \text{set_bad3}(y, T', T)$ to conclude $(x'_0, y'_0) \in \text{dom}(T')$. Unfolding the definition of set_bad3 , this premise states:

$$\neg (y \in \text{ran}(T) \wedge T^{-1}[y] \notin \text{dom}(T') \wedge \forall c. \neg \text{valid_chain}(c \parallel [T^{-1}[y]], T))$$

First, note that since we know $(x'_0, y'_0) \in \text{dom}(T)$ and $T[(x'_0, y'_0)] = y$, we obtain $T^{-1}[y] = (x'_0, y'_0)$. Therefore we can rewrite the above statement as

$$y \notin \text{ran}(T) \vee (x'_0, y'_0) \in \text{dom}(T') \vee \exists c. \text{valid_chain}(c \parallel [(x'_0, y'_0)], T)$$

However, we know that $y \in \text{ran}(T)$ since $(x'_0, y'_0) \in \text{dom}(T)$ and $T[(x'_0, y'_0)] = y$. Furthermore, if there existed a c such that $\text{valid_chain}(c \parallel [(x'_0, y'_0)], T)$ were true, then by the definition of valid_chain (Definition 6.15) and the completeness of findseq (Axiom 6.14) we would obtain $\text{findseq}(x'_0, y'_0, T) \neq \text{None}$, contradicting the fact that $\text{findseq}(x'_0, y'_0, T) = \text{None}$. Hence, it must hold $(x'_0, y'_0) \in \text{dom}(T')$. Thus it also holds $(x'_0, y'_0) \in \text{dom}(\text{upd}(T', (x, y), z))$, which finishes this case.

Case $(x_0, y_0) \in \text{dom}(T')$. Since $(x'_0, y'_0) \neq (x, y)$, we obtain that $(x'_0, y'_0) \in \text{dom}(T)$ and $T[(x'_0, y'_0)] = y_0$. Furthermore, we also know that $\text{findseq}(x'_0, y'_0, T) = \text{None}$ by an analogous argument as in the previous case. Therefore, all premises of $\text{Claim5}(T', T)$ are fulfilled, implying that $(x'_0, y'_0) \in \text{dom}(T')$. Thus, it clearly also holds that $(x'_0, y'_0) \in \text{dom}(\text{upd}(T', (x, y), z))$, which finishes the proof. \square

The final two statements described in this section concern the behavior of the function findseq in the maps T' and T , where $T' \subseteq T$. Firstly, assume that $\text{findseq}(x, y, T)$ is successful in the map T . Then, if $\text{Claim5}(T', T)$ holds and event \mathbf{bad}_3 is not triggered by the simulator for a query using y as a chaining value, all elements of the chain returned by findseq must already be in the domain of T' , and therefore $\text{findseq}(x, y, T')$ will also be successful.

Lemma 6.56. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y be a chaining value. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $\neg \text{set_bad3}(y, T', T)$
5. $\text{Claim5}(T', T)$
6. $\text{findseq}(x, y, T) \neq \text{None}$

Then it holds that $\text{findseq}(x, y, T') \neq \text{None}$.

Proof. First, since $\text{findseq}(x, y, T) \neq \text{None}$, by the soundness of function findseq , it holds that $\text{ischain}(T, \text{IV}, \pi_{\text{findseq}(x, y, T)}, y)$ and $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, T)}) \parallel [x]) \neq \text{None}$. First we will consider the case where $\pi_{\text{findseq}(x, y, T)} = \text{nil}$, which is easy. Next, we will prove the claim in the case where $\pi_{\text{findseq}(x, y, T)}$ has the form $c \parallel [(x', y')]$.

Case $\pi_{\text{findseq}(x,y,T)} = \text{nil}$. From $\text{ischain}(T, \text{IV}, \text{nil}, y)$, we get that $\text{IV} = y$ by Definition 6.12, and therefore also $\text{ischain}(T', \text{IV}, \text{nil}, y)$. Then, since furthermore $\text{unpad}(\text{mapfst}([x])) \neq \text{None}$, and $T' \subseteq T$ is well-formed, the completeness of function findseq implies that $\text{findseq}(x, y, T') \neq \text{None}$.

Case $\pi_{\text{findseq}(x,y,T)} = c \parallel [(x', y')]$. In this case, we perform a proof by contradiction: Assume that $\text{findseq}(x, y, T') = \text{None}$. Then, we show that event bad_3 would have been triggered for the value y , i.e. we show that $\text{set_bad3}(y, T', T)$ holds true, contradicting the fourth assumption. Unfolding the definition of set_bad3 , this means that we have to prove the following:

$$y \in \text{ran}(T) \wedge T^{-1}[y] \notin \text{dom}(T') \wedge \forall c. \neg \text{valid_chain}(c \parallel [T^{-1}[y]], T)$$

We prove each of the operands in the above conjunction below; the assumption $\text{Claim5}(T', T)$ will only be needed for the proof of the second operand. In each of the three instances, we will need the following. Since $\pi_{\text{findseq}(x,y,T)}$ has the form $c \parallel [(x', y')]$, we have $\text{ischain}(T, \text{IV}, c \parallel [(x', y')], y)$ and $\text{unpad}(\text{mapfst}(c) \parallel [x', x]) \neq \text{None}$. By Lemma 6.21, the former implies that $(x', y') \in \text{dom}(T)$ and $T[(x', y')] = y$. Thus, we get $T^{-1}[y] = (x', y')$. We will henceforth speak only about (x', y') .

- (i) $y \in \text{ran}(T)$: Since $(x', y') \in \text{dom}(T)$ and $T[(x', y')] = y$, we immediately conclude $y \in \text{ran}(T)$.
- (ii) $(x', y') \notin \text{dom}(T')$: Assume for contradiction that $(x', y') \in \text{dom}(T')$. Then, we can make use of the assumption that $\text{Claim5}(T', T)$: by Lemma 6.54, we obtain $\text{ischain}(T', \text{IV}, c \parallel [(x', y')], y)$. Since furthermore $\text{unpad}(\text{mapfst}(c) \parallel [x', x]) \neq \text{None}$ and the map $T' \subseteq T$ is well-formed, by the completeness of findseq we obtain that $\text{findseq}(x, y, T') \neq \text{None}$, contradicting our assumption. Therefore, $(x', y') \notin \text{dom}(T')$.
- (iii) $\forall c. \neg \text{valid_chain}(c \parallel [(x', y')], T)$: Assume for contradiction that there exists a chain c' such that $\text{valid_chain}(c' \parallel [(x', y')], T)$ holds. That is, for some z' we have $\text{ischain}(T', \text{IV}, c' \parallel [(x', y')], z')$ and $\text{unpad}(\text{mapfst}(c') \parallel [x']) \neq \text{None}$ (see Definition 6.15). Then, by Lemma 6.21 it also holds that $\text{ischain}(T', \text{IV}, c', y')$. Now, on the other hand, we also have that $\text{ischain}(T, \text{IV}, c \parallel [(x', y')], y)$ and $\text{unpad}(\text{mapfst}(c \parallel [(x', y')]) \parallel [x]) \neq \text{None}$. This yields a contradiction since now Lemma 6.33 implies that $(x', y') \notin c \parallel [(x', y')]$, which is absurd. Hence, we must indeed have $\forall c. \neg \text{valid_chain}(c \parallel [(x', y')], T)$, and this finishes the proof.

□

A corollary of the above lemma is that, if $\text{findseq}(x, y, \text{upd}(T, (x, y), z))$ is successful, and the update with $((x, y), z)$ is benign, then $\text{findseq}(x, y, T)$ must also have been successful (since the last element of a complete chain is not required to be in the domain of the chain's map), and thus, as above, $\text{findseq}(x, y, T')$ must also be successful.

Corollary 6.57. *Let $T, T' : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be maps, x be a block, and y and z be chaining values. Assume the following:*

1. $\text{Inclusion}(T', T)$
2. $\text{Injective}(T)$
3. $\text{IV} \notin \text{ran}(T)$
4. $(x, y) \notin \text{dom}(T)$
5. $z \notin \text{ran}(T)$
6. $z \neq \text{IV}$
7. $\neg \text{set_bad3}(y, T', T)$
8. $\text{Claim5}(T', T)$
9. $\text{findseq}(x, y, \text{upd}(T, (x, y), z)) \neq \text{None}$

Then it holds that $\text{findseq}(x, y, T') \neq \text{None}$.

Proof. We will first show that it holds that $\text{findseq}(x, y, T) \neq \text{None}$. For this, assume for contradiction that $\text{findseq}(x, y, T) = \text{None}$. But then, by Lemma 6.40 we obtain that $(x, y) \in \pi_{\text{findseq}(x, y, \text{upd}(T, (x, y), z))}$, i.e. the last element of the complete chain found by findseq also appears at another point in this chain. This cannot hold true. To see this, note that from the soundness of findseq we get $\text{ischain}(\text{upd}(T, (x, y), z), \text{IV}, \pi_{\text{findseq}(x, y, \text{upd}(T, (x, y), z))}, y)$ and $\text{unpad}(\text{mapfst}(\pi_{\text{findseq}(x, y, \text{upd}(T, (x, y), z))}) \parallel [x]) \neq \text{None}$. By instantiating Lemma 6.33 twice with the same chain, we now find out that $(x, y) \notin \pi_{\text{findseq}(x, y, \text{upd}(T, (x, y), z))}$, which yields the contradiction. Therefore, it must indeed hold $\text{findseq}(x, y, T) \neq \text{None}$. Now, we can apply Lemma 6.56 to immediately conclude that $\text{findseq}(x, y, T') \neq \text{None}$, which finishes the proof. \square

Conclusions

WE formally showed that the plain Merkle-Damgård construction, when used in conjunction with a prefix-free padding function, is indistinguishable from a variable input length random oracle, under the assumption that the underlying compression function behaves like a fixed input length random oracle. We learned about **EasyCrypt** and how it comfortably enables the formalization and mechanical verification of game-based proofs. We followed the proof of indistinguishability previously published by Coron et al. [39], and implemented a similar sequence of game transformations in **EasyCrypt**. However, our game sequence turned out to be slightly more complex, to account for the technical difficulties encountered when formally verifying equivalence statements between games. For instance, the difference between eager and lazy sampling exhibited by the initial and final games, described in Section 5.4, was hardly tackled at all in [39]. Moreover, the latter paper also described the games only in natural language, and accordingly that proof was conducted in a rather semi-formal way. In our proof, all transformations were validated using the built-in probabilistic relational Hoare logic of **EasyCrypt**, and the logical side conditions that were needed to enable the supported automated tools to solve the necessary proof obligations were verified in the **Coq** proof assistant. This result strengthens our belief that **EasyCrypt** is a formidable tool to tackle the formal verification of real-world cryptographic constructions, and that **EasyCrypt** makes a tremendous step towards the realization of Halevi’s vision [49]. Our chief aim is to formally verify the indistinguishability of the finalists of the SHA-3 competition; let us discuss how our present work applies to them.

7.1 Applicability to SHA-3 Finalists

First, we consider how closely our formalization of the Merkle-Damgård iteration matches the actual implementation of the five SHA-3 finalists. Second, we discuss if and how well the assumption that the compression function behaves like a fixed input length random oracle is applicable to the real-world compression functions of those algorithms.

As mentioned in Section 3.3, all finalists implement variants of the Merkle-Damgård

construction (see Definition 3.1). All finalists, except for **BLAKE**, additionally chop a number of bits off the last chaining value, with the purpose to discard some information; this makes the compression function's computations even more difficult to invert. Furthermore, **Grøstl** also performs a final transformation before chopping. For a formal verification, those operations would have to be explicitly implemented in the initial and final games. In fact, such operations may actually be needed to perform an indifferenciability proof; for instance, Coron et al. also showed in [39] that instead of using a prefix-free padding function in the plain Merkle-Damgård construction, one may alternatively chop a non-trivial number of bits off the last chaining value, and this also yields an indifferenciability hash function.

In our formalization of Merkle-Damgård, the compression function takes only two arguments (a block and a chaining value). This is indeed the case for **Grøstl**, **JH**, and **Keccak**. However, **BLAKE**'s compression function additionally takes (as specified by the HAIFA framework [27]) a counter and a salt. **Skein**'s compression function is built on a *tweakable* block cipher, and takes as additional input for its compression function a tweak (which essentially encodes the number of bytes hashed so far, as well as some additional information [45]). The counters in **BLAKE** and the tweaks in **Skein** play a similar role, namely, each block is processed with a unique variant of its compression function. This is supposed to increase their security, since it makes it much harder to exploit e.g. fixed-points. We observe that it is possible to formalize the salt and counters (respectively the tweaks) as an integral part of the padding rule of these hash functions, i.e. the padding function can precompute the appropriate values and append them to the message blocks. Then, these compression functions can be seen as taking only the usual two arguments, since the additional arguments are simply encoded within the blocks. This makes it possible to also map the compression functions of **BLAKE** and **Skein** onto our formalization of the Merkle-Damgård iteration.

Lastly, while all of the finalists use suffix-free padding rules, only the padding rules of **BLAKE** and **Skein** are additionally prefix-free [5]. Factoring in this fact, we observe that the closest candidate where our proof may apply is **BLAKE**; as for **Skein**, as mentioned above the final chopping would have to be formalized. Unfortunately, the assumption that the compression function is ideal is too strong in any case; we discuss this next.

As mentioned in Section 3.3, all of the SHA-3 finalists have been proven (on paper) to be indifferenciability from a random oracle [2, 3, 20, 24, 26, 37]. Yet, as opposed to the proof of indifferenciability presented here, all of those proofs are based on the assumption that the underlying building blocks of the compression functions are ideal; these underlying building blocks are block ciphers or permutations in all cases. As those underlying primitives constitute lower building blocks than the compression function itself, assuming their ideality is a weaker assumption than the assumption that the whole compression function is ideal. In fact, assuming ideality of the whole compression function is not appropriate for most of the finalists:

- The compression functions of **JH** and **Keccak** are trivially insecure, as collisions and preimages can be found in only one query to the underlying permutation [5, 28];
- It is trivial to find fixed-points for the compression function of **Grøstl** [47];

- The compression function of **BLAKE** has recently been shown to exhibit non-random behavior: It is differentiable from a (fixed input length) random oracle even under the assumption that its underlying block cipher is ideal [2, 37].

For **Skein**, there are no immediate results invalidating the assumption that its compression function is indifferentiable from a random oracle. It should be noted however that non-randomness has been proven for reduced round versions of Threefish [52], the tweakable block cipher which **Skein** is based on.

Since our proof of indifferentiability of the prefix-free Merkle-Damgård construction relies on the assumption that the underlying compression function behaves like an ideal primitive, it cannot be applied to **BLAKE**, as this assumption has been invalidated by aforementioned result. As for **Skein**, although no such result is known, the assumption that its compression function is ideal is also seriously weakened by the attacks on Threefish mentioned above.

7.2 Future Work

Our next aim is to formally verify the indifferentiability of the candidates of the SHA-3 competition. To achieve this, our formalization and proof has to be adapted so as to match the actual implementations of those algorithms. Since only **BLAKE** and **Skein** employ a prefix-free padding function, they constitute the most appealing candidates to concentrate on first. For this, first the formalization of the respective iteration modes has to be adapted as outlined in the previous section. Second, as explained above the proof has to be based on weaker assumptions, namely that the underlying building blocks of the respective compression functions are ideal. More precisely, both **BLAKE** and **Skein** are based on a block cipher [7, 45], and their respective compression functions have to be formalized up to the point where the respective block cipher is used. Then, the proof has to be performed in the ideal cipher model, i.e. only under the assumption that the underlying block cipher is ideal. Since a block cipher provides not only an encryption algorithm (used in the compression function), but also a decryption algorithm, the distinguisher must additionally be given access to a decryption oracle, which may constitute the main challenge.

Several (pen-and-paper) indifferentiability proofs of hash designs under such weaker assumptions are known. In [39], Coron et al. proved the indifferentiability of the prefix-free Merkle-Damgård construction based on the Davies-Meyer compression function [73], defined as $f(x, y) := E_x(y) \oplus y$, in the ideal cipher model for E , and for similar constructions as well. For all the finalists, as mentioned above, direct indifferentiability proofs are available which only assume the ideality of the underlying block ciphers or permutations.

To conclude, our proof constitutes a non-trivial result about the prefix-free Merkle-Damgård construction, and a first, but significant step to formalize the indifferentiability proofs of the SHA-3 finalists in **EasyCrypt** as well. Indeed, these proofs essentially use the same techniques, albeit they are slightly more involved since they are based on weaker assumptions. Our result provides a good starting point for formalizing these proofs, and we will further this line of research in the future.

References

- [1] Reynald Affeldt, Miki Tanaka, and Nicolas Marti. Formal proof of provable security by game-playing in a proof assistant. In *1st International Conference on Provable Security – ProvSec 2007*, volume 4784 of *Lecture Notes in Computer Science*, pages 151–168. Springer, 2007.
- [2] Elena Andreeva, Atul Luykx, and Bart Mennink. Provable security of BLAKE with non-ideal compression function. Cryptology ePrint Archive, Report 2011/620, November 2011. <http://eprint.iacr.org/>.
- [3] Elena Andreeva, Bart Mennink, and Bart Preneel. On the indifferentiability of the Grøstl hash function. In *7th International Conference on Security in Communication Networks – SCN 2010*, volume 6280 of *Lecture Notes in Computer Science*, pages 88–105. Springer, 2010.
- [4] Elena Andreeva, Bart Mennink, and Bart Preneel. Security reductions of the second round SHA-3 candidates. In *13th International Conference on Information Security – ISC 2010*, volume 6531 of *Lecture Notes in Computer Science*, pages 39–53. Springer, 2011.
- [5] Elena Andreeva, Bart Mennink, Bart Preneel, and Marjan Škrobot. Security analysis and comparison of the SHA-3 finalists BLAKE, Grøstl, JH, Keccak, and Skein. In *Progress in Cryptology – AFRICACRYPT 2012*, volume 7374 of *Lecture Notes in Computer Science*, pages 287–305. Springer, 2012.
- [6] Elena Andreeva, Gregory Neven, Bart Preneel, and Thomas Shrimpton. Seven-property-preserving iterated hashing: ROX. In *Advances in Cryptology – ASIACRYPT 2007*, volume 4833 of *Lecture Notes in Computer Science*, pages 130–146. Springer, 2007.
- [7] Jean-Philippe Aumasson, Luca Henzen, Willi Meier, and Raphael C.-W. Phan. SHA-3 proposal BLAKE. Online – <http://131002.net/blake/>, December 2010.

- [8] Michael Backes, Gilles Barthe, Matthias Berg, Benjamin Grégoire, César Kunz, Malte Skoruppa, and Santiago Zanella Béguelin. Verified security of Merkle-Damgård. In *Proceedings of the 25th IEEE Computer Security Foundations Symposium – CSF 2012*, pages 354–368. IEEE Computer Society, June 2012. Also in *Grande Region Security and Reliability Day – GRSRD 2012*.
- [9] Michael Backes, Matthias Berg, and Dominique Unruh. A formal language for cryptographic pseudocode. In *15th International Conference on Logic for Programming, Artificial Intelligence and Reasoning – LPAR 2008*, volume 5330 of *Lecture Notes in Computer Science*, pages 353–376. Springer, 2008.
- [10] Gilles Barthe, Marion Daubignard, Bruce Kapron, and Yassine Lakhnech. Computational indistinguishability logic. In *17th ACM Conference on Computer and Communications Security – CCS 2010*, pages 375–386. ACM, 2010.
- [11] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, Federico Olmedo, and Santiago Zanella Béguelin. Verified indifferentiable hashing into elliptic curves. In *1st Conference on Principles of Security and Trust – POST 2012*, volume 7215 of *Lecture Notes in Computer Science*, pages 209–228. Springer, 2012.
- [12] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. Formal certification of ElGamal encryption. A gentle introduction to CertiCrypt. In *5th International Workshop on Formal Aspects in Security and Trust – FAST 2008*, volume 5491 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2009.
- [13] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. Computer-aided security proofs for the working cryptographer. In *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 71–90. Springer, 2011.
- [14] Gilles Barthe, Benjamin Grégoire, Yassine Lakhnech, and Santiago Zanella Béguelin. Beyond provable security. Verifiable IND-CCA security of OAEP. In *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 180–196. Springer, 2011.
- [15] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. Formal certification of code-based cryptographic proofs. In *36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages – POPL 2009*, pages 90–101. ACM, 2009.
- [16] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. Programming language techniques for cryptographic proofs. In *1st International Conference on Interactive Theorem Proving – ITP 2010*, volume 6172 of *Lecture Notes in Computer Science*, pages 115–130. Springer, 2010.
- [17] Gilles Barthe, Daniel Hedin, Santiago Zanella Béguelin, Benjamin Grégoire, and Sylvain Heraud. A machine-checked formalization of Sigma-protocols. In *23rd IEEE*

-
- Computer Security Foundations Symposium – CSF 2010*, pages 246–260. IEEE Computer Society, 2010.
- [18] Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. Probabilistic reasoning for differential privacy. In *39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages – POPL 2012*, pages 97–110. ACM, 2012.
- [19] Gilles Barthe, Federico Olmedo, and Santiago Zanella Béguelin. Verifiable security of Boneh-Franklin identity-based encryption. In *5th International Conference on Provable Security – ProVSec 2011*, volume 6980 of *Lecture Notes in Computer Science*, pages 68–83. Springer, 2011.
- [20] Mihir Bellare, Tadayoshi Kohno, Stefan Lucks, Niels Ferguson, Bruce Schneier, Doug Whiting, Jon Callas, and Jesse Walker. Provable security support for the Skein hash family. Online – <http://skein-hash.info>, April 2009.
- [21] Mihir Bellare, Ted Krovetz, and Phillip Rogaway. Luby-Rackoff backwards: Increasing security by making block ciphers non-invertible. In *Advances in Cryptology – EUROCRYPT 1998*, volume 1403 of *Lecture Notes in Computer Science*, pages 266–280. Springer, 1998.
- [22] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *1st ACM Conference on Computer and Communications Security – CCS 1993*, pages 62–73. ACM, 1993.
- [23] Mihir Bellare and Phillip Rogaway. Code-based game-playing proofs and the security of triple encryption. Cryptology ePrint Archive, Report 2004/331, November 2004. <http://eprint.iacr.org/>.
- [24] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indistinguishability of the sponge construction. In *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 181–197. Springer, 2008.
- [25] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. The Keccak reference. Online – <http://keccak.noekeon.org>, January 2011.
- [26] Rishiraj Bhattacharyya, Avradip Mandal, and Mridul Nandi. Security analysis of the mode of JH hash function. In *17th International Workshop on Fast Software Encryption – FSE 2010*, volume 6147 of *Lecture Notes in Computer Science*, pages 168–191. Springer, 2010.
- [27] Eli Biham and Orr Dunkelman. A framework for iterative hash functions - HAIFA. Cryptology ePrint Archive, Report 2007/278, July 2007. <http://eprint.iacr.org/>.
- [28] John Black, Martin Cochran, and Thomas Shrimpton. On the impossibility of highly-efficient blockcipher-based hash functions. *Journal of Cryptology*, 22:311–329, 2009.

- [29] John Black and Phillip Rogaway. CBC MACs for arbitrary-length messages: The three-key constructions. In *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 197–215. Springer, 2000.
- [30] John Black and Phillip Rogaway. A block-cipher mode of operation for parallelizable message authentication. In *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 384–397. Springer, 2002.
- [31] Bruno Blanchet. Mechanizing game-based proofs of security protocols. In *Software Safety and Security - Tools for Analysis and Verification*, volume 33 of *NATO Science for Peace and Security Series – D: Information and Communication Security*, pages 1–25. IOS Press, 2012.
- [32] Bruno Blanchet, Aaron D. Jaggard, Andre Scedrov, and Joe-Kai Tsay. Computationally sound mechanized proofs for basic and public-key Kerberos. In *15th ACM Conference on Computer and Communications Security – ASIACCS 2008*, pages 87–99. ACM, 2008.
- [33] Bruno Blanchet and David Pointcheval. Automated security proofs with sequences of games. In *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 537–554. Springer, 2006.
- [34] Dan Boneh and Matthew Franklin. Identity-based encryption from the Weil pairing. *SIAM Journal on Computing*, 32(3):586–615, 2003.
- [35] Eric Brier, Jean-Sébastien Coron, Thomas Icart, David Madore, Hugues Randriam, and Mehdi Tibouchi. Efficient indifferentiable hashing into ordinary elliptic curves. In *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 237–254. Springer, 2010.
- [36] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *Journal of the ACM*, 51(4):557–594, 2004.
- [37] Donghoon Chang, Mridul Nandi, and Moti Yung. Indifferentiability of the hash algorithm BLAKE. Cryptology ePrint Archive, Report 2011/623, November 2011. <http://eprint.iacr.org/>.
- [38] Pierre Corbineau, Mathilde Duclos, and Yassine Lakhnech. Certified security proofs of cryptographic protocols in the computational model: An application to intrusion resilience. In *1st International Conference on Certified Programs and Proofs – CPP 2011*, volume 7086 of *Lecture Notes in Computer Science*, pages 378–393. Springer, 2011.
- [39] Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 430–448. Springer, 2005.

-
- [40] Judicaël Courant, Marion Daubignard, Cristian Ene, Pascal Lafourcade, and Yassine Lakhnech. Towards automated proofs for asymmetric encryption schemes in the random oracle model. In *15th ACM Conference on Computer and Communications Security – CCS 2008*, pages 371–380. ACM, 2008.
- [41] Ronald Cramer. *Modular Design of Secure yet Practical Cryptographic Protocols*. PhD thesis, CWI and University of Amsterdam, 1996.
- [42] Ivan Damgård. A design principle for hash functions. In *Advances in Cryptology – CRYPTO 1989*, volume 435 of *Lecture Notes in Computer Science*, pages 416–427. Springer, 1990.
- [43] Marion Daubignard, Pierre-Alain Fouque, and Yassine Lakhnech. Generic indistinguishability proofs of hash designs. In *25th IEEE Computer Security Foundations Symposium – CSF 2012*, pages 340–353. IEEE Computer Society, June 2012.
- [44] Jonathan Driedger. Formalization of game-transformations. Saarland University, January 2010. Bachelor’s Thesis.
- [45] Niels Ferguson, Stefan Lucks, Bruce Schneier, Doug Whiting, Mihir Bellare, Tadayoshi Kohno, Jon Callas, and Jesse Walker. The Skein hash function family. Online – <http://skein-hash.info>, November 2008.
- [46] Jean-Christophe Filliâtre. The WHY verification tool: Tutorial and Reference Manual Version 2.31. Online – <http://why.lri.fr>, July 2012.
- [47] Praveen Gauravaram, Lars R. Knudsen, Krystian Matusiewicz, Florian Mendel, Christian Rechberger, Martin Schläffer, and Søren S. Thomsen. Grøstl – a SHA-3 candidate. Online – <http://www.groestl.info>, March 2011.
- [48] Shai Halevi. EME*: Extending EME to handle arbitrary-length messages with associated data. In *5th International Conference on Cryptology in India – INDOCRYPT’04*, volume 3348 of *Lecture Notes in Computer Science*, pages 315–327. Springer, 2004.
- [49] Shai Halevi. A plausible approach to computer-aided cryptographic proofs. Cryptology ePrint Archive, Report 2005/181, June 2005. <http://eprint.iacr.org/>.
- [50] Shai Halevi and Phillip Rogaway. A tweakable enciphering mode. In *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 482–499. Springer, 2003.
- [51] Shai Halevi and Phillip Rogaway. A parallelizable enciphering mode. In *Topics in Cryptology – CT-RSA 2004*, volume 2964 of *Lecture Notes in Computer Science*, pages 292–304. Springer, 2004.

- [52] Dmitry Khovratovich, Ivica Nikolić, and Christian Rechberger. Rotational rebound attacks on reduced Skein. In *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2010.
- [53] Joe Kilian and Phillip Rogaway. How to protect DES against exhaustive key search. In *Advances in Cryptology – CRYPTO 1996*, volume 1109 of *Lecture Notes in Computer Science*, pages 252–267. Springer, 1996.
- [54] Neal Koblitz. Another look at automated theorem-proving. *Journal of Mathematical Cryptology*, 1(4):385–403, 2008.
- [55] Neal Koblitz. Another look at automated theorem-proving. II. *Journal of Mathematical Cryptology*, 5(3-4):205–224, 2012.
- [56] Stefan Lucks. A failure-friendly design principle for hash functions. In *Advances in Cryptology – ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, pages 474–494. Springer, 2005.
- [57] Yuri I. Manin. *A Course in Mathematical Logic*. Springer, 1977. Translated by Neal Koblitz.
- [58] Stéphane Manuel. Classification and generation of disturbance vectors for collision attacks against SHA-1. *Designs, Codes and Cryptography*, 59:247–263, 2011.
- [59] Ueli Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In *1st Theory of Cryptography Conference – TCC 2004*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2004.
- [60] Ralph Merkle. One way hash functions and DES. In *Advances in Cryptology – CRYPTO 1989*, volume 435 of *Lecture Notes in Computer Science*, pages 428–446. Springer, 1990.
- [61] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL – A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
- [62] David Nowak. A framework for game-based security proofs. In *9th International Conference on Information and Communications Security – ICICS 2007*, volume 4861 of *Lecture Notes in Computer Science*, pages 319–333. Springer, 2007.
- [63] David Nowak. On formal verification of arithmetic-based cryptographic primitives. In *11th International Conference on Information Security and Cryptology – ICISC 2008*, volume 5461 of *Lecture Notes in Computer Science*, pages 368–382. Springer, 2009.
- [64] David Nowak and Yu Zhang. A calculus for game-based security proofs. In *4th International Conference on Provable Security – ProvSec 2010*, volume 6402 of *Lecture Notes in Computer Science*, pages 35–52. Springer, 2010.

-
- [65] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferenciability framework. In *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 487–506. Springer, 2011.
- [66] Phillip Rogaway, Mihir Bellare, John Black, and Ted Krovetz. OCB: A block-cipher mode of operation for efficient authenticated encryption. In *8th ACM Conference on Computer and Communications Security – CCS 2001*, pages 196–205. ACM, 2001.
- [67] Phillip Rogaway and Thomas Shrimpton. Cryptographic hash-function basics: Definitions, implications, and separations for preimage resistance, second-preimage resistance, and collision resistance. In *11th International Workshop on Fast Software Encryption – FSE 2004*, volume 3017 of *Lecture Notes in Computer Science*, pages 371–388. Springer, 2004.
- [68] Victor Shoup. Sequences of games: A tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332, November 2004. <http://eprint.iacr.org/>.
- [69] Malte Skoruppa. Formal verification of ElGamal encryption using a probabilistic Lambda-calculus. Saarland University, January 2010. Bachelor’s Thesis.
- [70] The Coq development team. The Coq Proof Assistant Reference Manual Version 8.3. Online – <http://coq.inria.fr>, 2010.
- [71] Xiaoyun Wang, Yiqun Yin, and Hongbo Yu. Finding collisions in the full SHA-1. In *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 17–36. Springer, 2005.
- [72] Xiaoyun Wang and Hongbo Yu. How to break MD5 and other hash functions. In *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 561–561. Springer, 2005.
- [73] Robert S. Winternitz. A secure one-way hash function built from DES. In *5th IEEE Symposium on Security and Privacy – S&P 1984*, pages 88–90. IEEE Computer Society, 1984.
- [74] Hongjun Wu. The hash function JH. Online – <http://www3.ntu.edu.sg/home/wuhj/research/jh/>, January 2011.
- [75] Santiago Zanella Béguelin. *Formal Certification of Game-Based Cryptographic Proofs*. PhD thesis, Ecole Nationale Supérieure des Mines de Paris – Mines ParisTech, 2010.
- [76] Santiago Zanella Béguelin, Benjamin Grégoire, Gilles Barthe, and Federico Olmedo. Formally certifying the security of digital signature schemes. In *30th IEEE Symposium on Security and Privacy – S&P 2009*, pages 237–250. IEEE Computer Society, 2009.