

Seedless Fruit is the Sweetest: Random Number Generation, Revisited

Sandro Coretti*
New York University
corettis@nyu.edu

Yevgeniy Dodis†
New York University
dodis@cs.nyu.edu

Harish Karthikeyan‡
New York University
karthik@cs.nyu.edu

Stefano Tessaro§
University of Washington
tessaro@cs.washington.edu

April 11, 2019

Abstract

The need for high-quality randomness in cryptography makes random-number generation one of its most fundamental tasks.

A recent important line of work (initiated by Dodis et al., CCS '13) focuses on the notion of *robustness* for *pseudorandom number generators (PRNGs) with inputs*. These are primitives that use various sources to accumulate sufficient entropy into a state, from which pseudorandom bits are extracted. Robustness ensures that PRNGs remain secure even under state compromise and adversarial control of entropy sources. However, the achievability of robustness inherently depends on a *seed*, or, alternatively, on an ideal primitive (e.g., a random oracle), independent of the source of entropy. Both assumptions are problematic: seed generation requires randomness to start with, and it is arguable whether the seed or the ideal primitive can be kept independent of the source.

This paper resolves this dilemma by putting forward new notions of robustness which enable both (1) *seedless* PRNGs and (2) *primitive-dependent* adversarial sources of entropy. To bypass obvious impossibility results, we make a realistic compromise by requiring that the source produce sufficient entropy *even* given its evaluations of the underlying primitive. We also provide natural, practical, and provably secure constructions based on hash-function designs from compression functions, block ciphers, and permutations. Our constructions can be instantiated with minimal changes to industry-standard hash functions SHA-2 and SHA-3, or HMAC (as used for the key derivation function HKDF), and can be downgraded to (*online*) *seedless randomness extractors*, which are of independent interest.

On the way we consider both a *computational* variant of robustness, where attackers only make a bounded number of queries to the ideal primitive, as well as a new *information-theoretic* variant, which dispenses with this assumption to a certain extent, at the price of requiring a high rate of injected weak randomness (as it is, e.g., plausible on Intel's on-chip RNG). The latter notion enables applications such as everlasting security.

Finally, we show that the CBC extractor, used by Intel's on-chip RNG, is provably insecure in our model.

*Supported by NSF grants 1314568 and 1619158.

†Partially supported by gifts from VMware Labs, Facebook and Google, and NSF grants 1314568, 1619158, 1815546.

‡Supported by NSF grant 1619158.

§Partially supported by NSF grants CNS-1553758 (CAREER), CNS-1719146, CNS-1528178, and IIS-1528041, and by a Sloan Research Fellowship.

Contents

1	Introduction	3
1.1	Previous Theoretical Models for PRNGs: Seeds	3
1.2	<i>Seedless</i> PRNGs and Extractors from Cryptographic Hashing	4
1.3	Toy Case: Monolithic Seedless Extraction from Oracle-Dependent Sources	6
1.4	Our Results	8
1.5	Other Related Work	10
2	Preliminaries	12
2.1	Statistical Distance and Min-Entropy	12
2.2	Security Games	12
3	Seedless Extraction and Key Derivation	12
3.1	Definition	13
3.2	Seedless Extraction with a Monolithic Random Oracle	14
3.3	Online Extraction	14
3.4	CBC-Based Extractors Are Insecure	16
3.5	Seedless HKDF	17
4	Pseudorandom Number Generators with Input	19
4.1	Syntax	20
4.2	Security Game	20
5	Constructions of PRNGs	22
5.1	PRNGs from Merkle-Damgård	23
5.2	PRNGs from Merkle-Damgård with Davies-Meyer	25
5.3	PRNGs from Sponges	26
6	Security Proofs for Computational Constructions	28
6.1	The H-Coefficient Technique	28
6.2	Monolithic Extractor	29
6.3	Intermediate PRNG Security Notions	32
6.4	PRNGs from Merkle-Damgård	34
6.5	PRNGs from Merkle-Damgård with Davies-Meyer	40
6.6	PRNGs from Sponges	45
7	Security Proofs for IT Constructions	51
7.1	Information-Theoretic Preliminaries	51
7.2	Monolithic Extractor	52
7.3	Intermediate IT PRNG Security Notions	54
7.4	IT PRNGs from Merkle-Damgård	55
7.5	IT PRNGs from Merkle-Damgård with Davies-Meyer	62
7.6	IT PRNGs from Sponge	65
A	Comparison to Previous PRNG Security Notions	73

B	Constructions of Online Extractors	74
B.1	Extractors from Merkle-Damgård	74
B.2	Extractors from Merkle-Damgård with Davies-Meyer	75
B.3	Extractors from Sponges	76
B.4	Extractors from HMAC	77
C	Hybrid Proofs	81
C.1	Recovering and Preserving Imply Robustness	81
C.2	IT-Recovering Implies IT-Robustness	84

1 Introduction

Good random number generation is essential for cryptography and beyond. In practice, this difficult task is solved by a primitive called *pseudorandom number generator with input (PRNG)*, whose aim is to quickly accumulate entropy from various physical sources in the environment (such as keyboard presses, timing of interrupts, etc.) into the state of the PRNG and then convert this high-entropy state into (pseudo) random bits. In particular, entropy accumulation should never stop since one may need to recover from occasional compromises of the PRNG state. PRNGs are ubiquitous and have extensive applications. For example, virtually all operating systems come equipped with a PRNG; e.g., `/dev/random` [47] for Linux, Yarrow [33] for MacOS/iOS/FreeBSD, and Fortuna [23] for Windows [22], where the latter two make use of standard cryptographic primitives as part of their design. Still, as we will argue below in a much broader context, even these widely used PRNGs lack adequate theoretical understanding and analysis, which are critical if such PRNGs or their future tweaks continue to be used ubiquitously.

The situation is not better in terms of standardization efforts, where existing PRNG standards [31, 34, 21, 5] are less mature than those for most other cryptographic primitives. For starters, there has not been any rigorous competition soliciting PNRG designs, and big parts of the existing standards concentrate on the difficult (ad-hoc and non-cryptographic) problem of entropy estimation rather than entropy accumulation and extraction. More importantly, standardized cryptographic PRNG constructions are rather ad-hoc, have no clear security definition/model, often have confusing syntax, and sometimes have been broken by subsequent analyses of the cryptographic community. The most widely known example is the DualEC PRNG, which appeared in the first version of the NIST SP 800-90A standard [5] in 2005 and remained there for years—despite early warnings by [41, 43] and allowing potential exploitation [12]—until Snowden’s revelations finally led to its deprecation. Recent work [48] identified a lot of gaps and imprecision (sometimes leading to attacks or security concerns) in the existing analyses and deployment for the other 3 PRNGs from the NIST SP 800-90A standard. In a similar vein, [42] found several gaps and misconceptions in previous analyses and security justifications for the popular Intel Secure Key Hardware PRNG introduced in 2011.

One of the main goals of this work is to reverse this poor state of affairs and to design a rigorous, theoretically sound model of PRNGs. This model should be general enough to incorporate practical entropy sources available in the real world, as well as to formally prove security of “good,” widely used PRNGs against realistic attackers.

1.1 Previous Theoretical Models for PRNGs: Seeds

In view of their practical importance, we are certainly not the first to formally study PRNGs through a theoretical lens. Indeed, several theoretical models and analyses of PRNGs have appeared in the literature [1, 18, 42, 20, 25, 27]. While differing in various details, these important works share two key principles:

- (a) The PRNG should work even against *adversarial* entropy sources, as long as such sources eventually provide enough entropy (such sources are called “legitimate” [18]);
- (b) assuming more structure beyond entropy is undesirable and brittle,¹ as this requires a rather detailed understanding of one’s entropy sources.

However, while such extremely minimalist assumptions make these PRNG models applicable to a wide variety of entropy sources, they also come with a subtle, but very important caveat: the

¹We do, however, later discuss an interesting approach suggested by [3].

randomness extraction module *cannot* be deterministic, as deterministic extraction from *general* entropy sources is impossible [15]. As a result, the PRNGs studied by these works are *seeded* (with the seed somehow chosen at initialization), but the entropy sources are assumed to be *independent of the seed*. This modeling is inherited from the underlying problem of randomness extraction, where *seeded extractors* [39] indeed overcome the impossibility of deterministic (or “seedless”) extraction from general entropy sources.

While natural and sufficient for some applications of extractors, we argue that the need for a seed seems rather problematic in the deployment of PRNGs. First, if the reason for random number generation is the lack of access to high-quality random bits, then we may not have any way to generate the seed. More importantly, even if we *can* generate a uniformly random seed, it is crucial for the analysis that (potentially adversarial) entropy sources remain *independent* of the seed, for otherwise the extractor guarantees are lost. For example, if physical entropy sources inside the computer are used, these sources may be affected by the internal computations of the PRNG itself, and thus there may be correlations between the seed and the sources. Moreover, for many seeded PRNGs, the attacker could obtain information about the seed by either directly reading it from memory, or indirectly when the recently compromised or rebooted RNG is called on “low-entropy” inputs (so the output is no longer random and leaks information about the seed; this is called “premature next” attack by [20]).

This means that it is certainly an issue if the seed is just generated once and for all (perhaps using an expensive source of randomness) and hard-coded within implementations to be used for all future randomness extractions. Moreover, if multiple entropy sources are used, it is natural that some of these sources are adversarial and could depend on the seed (which is hard to protect against with a dedicated attacker). Somewhat paradoxically (considering the common belief that “more entropy cannot hurt”), the mixing of such seed-dependent sources once again invalidates all the provable guarantees of seeded PRNGs, even if all the entropy is obtained from other, seed-independent sources.

We thus face a dilemma:

We want to support general entropy sources, for which seedless extraction is impossible, and seeded extraction is only possible under very dangerous and hard-to-ensure independence assumptions, which we would rather avoid.

The goal of this work is to provide a meaningful solution to this dilemma, by keeping the PRNG design *seedless* while respecting properties (a) and (b) mentioned above.

1.2 *Seedless* PRNGs and Extractors from Cryptographic Hashing

We will achieve this goal by using popular *cryptographic hash functions* (CHF) as our technical tool, and by carefully defining the notion of entropy in the setting when certain components of these CHFs are assumed idealized.

Why cryptographic hashing? Before describing our solution in more detail, we explain why using CHF appears essential for the design of *seedless*² PRNGs. For starters, all general-purpose software PRNGs used today, as well as all recommendations in existing PRNG standards, are based on CHF. Hence, this setting must definitely be understood in order to provide results useful in the real world.

²Or, in the non-uniform setting, “seed-dependent”

However, there is a more glaring theoretical reason as well. The key component of any PRNG is the shrinking function `refresh` which takes the current PRNG state S as well as a new entropic input X and produces a new state $S' \leftarrow \text{refresh}(S, X)$. The goal of this function is to absorb the potential entropy of X into the PRNG state S , in which case the entropy of S' should be higher than the original entropy of S . In the extraction literature, this property is called *condensing*. If one uses a seed, building such condensers is easy to accomplish information-theoretically. For example, in the PRNG design of [18], the refresh function is linear: $S' = aS + X$, where a is a seed independent of X .

In the seedless/seed-dependent setting, it is not hard to see [19] that condensers must be built cryptographically, as they require at least some form of preimage- and collision-resistance.³ For example, when used in iteration, the simple $aS+X$ condenser function above—which yields (together with other building blocks) a provably secure seeded PRNG construction [18]—can be broken *in a catastrophic way* if the distribution of the input blocks X_1, X_2, \dots could depend on the constant a : it is not hard to see that an attacker knowing a can rather easily produce *high-entropy inputs* such that if the condenser is applied to it, the resulting would have no entropy at all. In practice, one cannot imagine a PRNG system which would risk such a catastrophic failure by critically depending on the fact that the constant a must remain hidden for the lifetime of the PRNG. Therefore, not surprisingly, all real-world PRNG designs—including those used by Windows, MacOS, and FreeBSD—critically rely on CHF’s, despite lacking adequate theoretical justification.

Cryptographically secure condensers, which at an absolute minimum seedless PRNGs have to be, can be built using a (very strong form of) collision-resistance [19]. However, the types of condensers needed for applications, called average-case seedless condensers, seem to require non-standard cryptographic assumptions. For example, a relatively weak form of such average-case condensers (called “condensers for leaky sources”) are already sufficient for instantiating the Fiat-Shamir heuristic for public-coin proof systems [19]—and it is a major open problem to provide such an instantiation under standard cryptographic assumptions.

To put it differently, even ignoring the fact that we want our PRNGs to be full-blown *seedless extractors*—a problem we will address next—just achieving provably secure entropy accumulation appears to require the use of CHF’s as well as either (1) non-standard cryptographic assumptions (making the results appear somewhat tautologous) or (2) some supporting justification argument in an idealized model of computation, which is the approach taken by this work.

Our approach: new min-entropy notion. To describe our approach, it is instructive to recall the basic impossibility of seedless extraction for general entropy sources. Given any candidate (seedless) extractor G , an adversary can perform a so-called *extractor-fixing attack* by sampling a random input X several times until the first bit of $G(X)$ is 0. The resulting distribution X has very high entropy, but $G(X)$ is clearly not uniformly random. Observe that with a strong enough CHF G , one might be able to formally argue that the extractor-fixing attack is the “most damaging” attack possible; for example by showing, that $G(X)$ has almost full entropy (i.e., is a good condenser) for any efficiently samplable source X , as was done by [19]. In other words, using CHF’s will protect against the completely devastating attacks possible with information-theoretic extractors.

However, our goal is to have a meaningful model where real randomness *extraction* is possible, so that we can later extend it all the way to the full PRNG system. Our solution will be to define a elegant and practically motivated refinement of general min-entropy *in settings where CHF’s exist*, so that:

³For example, the ability to compute a random preimage of a given element, which is known to imply one-way functions [30], allows the attacker to produce entropic inputs whose entropy is completely lost by the refresh procedure.

- (a) somewhat artificial sources resulting from intentionally performing extractor-fixing will not have much entropy according to our notion (meaning they are no longer “legitimate”); in fact, seedless extraction will become *possible* for our notion of min-entropy;
- (b) most natural entropy sources, including those used by major operating systems, will likely have good entropy according to our new measure.

While our final constructions and interpretation of our security analyses will apply to real-world CHFs, such as those derived from SHA-2, SHA-3, or HKDF, at present the only rigorous way we know how to achieve our ambitious goals (a) and (b) will be by going to the idealized models of computation, such as the random oracle, the ideal cipher or the random permutation model. This is quite standard for many areas of symmetric-key cryptography, and we already indicated that doing provably secure (non-tautologous) seedless PRNG constructions in the “standard model” appears beyond our current capabilities, even for much simpler building blocks, such as (average-case) seedless condensers.

1.3 Toy Case: Monolithic Seedless Extraction from Oracle-Dependent Sources

We start by presenting our new entropy notion for the simpler problem of “monolithic randomness extraction,” where the entropy source X is assumed to come in one piece (rather than slowly accumulated using a fixed-length PRNG state), and a monolithic CHF G —modeled as a monolithic random oracle—is used to output the value $R = G(X)$ (so that we temporarily ignore any fine-grained structure inside G , such as Merkle-Damgård or Sponge [8] iteration).

At first, it appears that we solved our problem in a totally trivial (and uninteresting) way, even without refining standard min-entropy. Namely, in the random oracle model, the following folklore proof (see [17]) appears to show that a (seedless) random oracle G is a good extractor: For any min-entropy γ^* source X , the probability the distinguisher D can distinguish $G(X)$ is upper bounded by the probability D queries G on X , which is at most $q \cdot 2^{-\gamma^*}$, where q is the number of random oracle queries allowed to \mathcal{A} .⁴

Implicit in this simple proof, however, is the key assumption that the distribution X is independent of the random oracle G , meaning that our (potentially adversarial) sampler producing X is not allowed to call the random oracle G . Thus, modeling G as a random oracle is but a fancy way of introducing an exponentially long seed that is independent of the source, making extraction trivial.⁵ Indeed, to capture PRNG sources X arising in the real world, we must allow the source X to depend on the ideal primitive G . For example, if the timing of computer interrupts is used as our entropy source X —which is the most common source of randomness in software PRNGs—it seems unreasonable to assume that none of these interrupts could be affected by frequent hash function computations done inside and outside the operating system.

Oracle-dependent sources. To fix this problem, in Section 3 we will explicitly model our source as part of the attacker \mathcal{A} , so that $\mathcal{A}^G = (\mathcal{A}_1^G, \mathcal{A}_2^G)$, where \mathcal{A}_1^G outputs the *oracle-dependent* source X and passes state Σ to the second state attacker $\mathcal{A}_2^G(\Sigma)$, whose goal is to distinguish $R = G(X)$ from uniform. Of course, for this definition to make sense, we must require that X is “legitimate,”

⁴In fact, if the length of $G(X)$ is slightly less than γ^* , we can even let \mathcal{A} query all of G and use leftover-hash lemma [29] to get information-theoretic security.

⁵Prior to our work, the above modeling of sources as being independent of the ideal primitive, was the only way to overcome extractor-fixing attacks. Examples of this approach include [17, 35, 48] and many others. While these results are non-trivial due to the “non-monolithic” structure of their extractors G , none of these works model the setting where *the source could depend on the ideal primitive*.

meaning it has entropy at least γ^* given the state information Σ (for some parameter γ^*). In the standard model, this could be formalized by requiring $H_\infty(X|\Sigma) \geq \gamma^*$ (see Section 2). But this is too weak, as this still allows for extractor-fixing attacks, by sampling a long random X and remembering a few bits of $G(X)$ in the leakage Σ . In fact, this extractor-fixing attack still works even if we condition on the entire random oracle G (i.e., require $H_\infty(X|\Sigma, G) \geq \gamma^*$). This leads to a central question of this work:

What is the “right” notion of entropy for oracle-dependent sources X ?

The key insight of our work comes from the fact that while it is reasonable to assume that the source X could *depend* on the random oracle G , the natural sources of entropy we want to extract from do not natively evaluate cryptographic hash functions, but somehow add extra entropy *in addition* to all hash function evaluations around them. For example, it is unreasonable to assume that the timing of interrupts could not depend, even slightly, on various hash function evaluations inside the computer. However, it seems that the real entropy of interrupt timings comes from the fact that the attacker cannot perfectly predict the exact lower order bits of the timing measurements, *even if the attacker knew all the hash function evaluations*. Indeed, instead of only requiring that $H_\infty(X|\Sigma, G) \geq \gamma^*$, our approach will make a stronger requirement that

$$H_\infty(X|(\Sigma, \mathcal{L})) \geq \gamma^* , \tag{1}$$

where \mathcal{L} is the input-output list of random oracle queries made by the sampler \mathcal{A}_1 to the random oracle. Another, equivalent way to interpret this legitimacy condition is to mandate that \mathcal{A}_1 cannot “forget” any of its random oracle queries when passing its state Σ to \mathcal{A}_2 , but must forget some other useful information about X , to ensure that X has entropy conditioned on Σ and \mathcal{L} .

Notice, our solution places a more stringent requirement than conditioning on the entire G , as \mathcal{A}_1 did not touch anything outside \mathcal{L} , so these un-queried values do not reduce entropy of X beyond what is done by \mathcal{L} . Also, when the number of queries q is not too large, the extractor-fixing is no longer a legitimate attack, since X will not have much entropy when conditioned on \mathcal{L} (which contains the pair $(X, G(X))$). In fact, we can easily show *full extraction* (see Theorems 1 and 2), along the lines of the folklore proof for oracle-independent sources mentioned above. The basic intuition comes from the fact that our conditioning on the list \mathcal{L} ensures that with overwhelming probability the sampler \mathcal{A}_1 did not himself evaluate $G(X)$, which is essential for the extractor-fixing attack to succeed.

Did we go too far? Of course, the main question is whether the legitimacy requirement $H_\infty(X|(\Sigma, \mathcal{L})) \geq \gamma^*$ does not overly limit the class of high-entropy sources from which we want to extract. We believe the answer is negative. First, in the restrictive “folklore case” when X is independent of G (meaning $\mathcal{L} = \emptyset$), we get the best-possible min-entropy condition $H_\infty(X|\Sigma) \geq \gamma^*$ we had in the standard (non-random-oracle) model. Namely, our notion of min-entropy relative to G includes all general min-entropy sources.⁶

Second, while we certainly allow the source X to substantially depend on G , we ensure that non-trivial bulk of entropy must come from *outside* of the actual oracle evaluation queries. In other words, while “nature,” who outputs X , could conceivably be influenced by a couple of hash function evaluations, it should generate some intrinsic entropy *in addition to* (but possibly dependent on!) these evaluations. We feel that all practically used physical sources (timing of interrupts, temperature, keystroke dynamics, etc.) have very little to do with hash functions, and should easily satisfy this requirement.

⁶Of course, when we instantiate G with a real-world hash function, this is no longer the case, as we discuss below.

Thus, we believe that our technical restriction on the legitimacy for extraction using CHF’s—by conditioning min-entropy on the list of hash function evaluations—strikes the right balance between allowing for seedless extraction, and yet keeping the family of high-entropy sources large and realistic for applications.

1.4 Our Results

While the above toy example (analyzed in Section 3.2) illustrated the key technical insight behind our approach, in practice it is uncommon to assume access to a monolithic random oracle G . Instead, practical hash functions are usually built from (public) compression functions, ciphers, or permutations. These underlying primitives P have limited input length and will therefore not be able to process inputs of arbitrary length m . Therefore, extractors and PRNGs should be designed in such a way that they can process short m -bit input blocks (e.g., $m = 256, 512, 1600$) and accumulate their entropy in the internal state.

Online extractors and insecurity of CBC. Thus, in Section 3.3 we formalize the more realistic notion of *online (seedless) extractors*, which slowly accumulate their long input into a fixed-length state (using access to a P), and then finalize their output once the whole input is processed. We also define both computational and information-theoretic (IT) notions of online extractor security, where in the latter notion the attacker is allowed to read the entire ideal primitive P after it finished generating the oracle-dependent source X .

Turning to natural and widely used examples of such online extractors, we show that the popular CBC mode of operation is insecure as a seedless extractor in our framework. The details of our attack are given in Section 3.4, but the result is a somewhat unexpected, since CBC is used as the extractor underlying the CTR_DRBG construction in the NIST PRNG standard NIST SP 800-90A Rev. 1 [4], and also as the extractor for Intel’s on-chip RNG [37]. Moreover, its security was formally shown by Dodis et al. [17], but in the setting where the entropy source X was *independent* of the random permutation π . In contrast, we show that once the latter assumption is relaxed to our oracle-dependent sources, the CBC extractor is no longer secure (unless one generalizes it to the Sponge construction in Section 5.3, where the input is only XORed to *part* of the state). Of course, our attack is somewhat theoretical, and does not directly translate to attacking the Intel on-chip RNG, for example. However, coupled with our positive results, we feel our attack suggests using a different online extractor, if possible.

On a positive side, in Appendix B we show several other (both computational and information-theoretic) online extractors based on popular modes of operations used inside hash functions SHA-2 and SHA-3, which are provably secure in our framework: from Merkle-Damgård with a random compression function, from Merkle-Damgård with the Davies-Meyer compression function, and from Sponges. Hence, for the first time practitioners can use seedless extractors which are both practical and have firm theoretical foundation. The security of these natural online extractors follows as special cases of more general PRNG security results, which we describe next. Due to its widespread use, we also analyze HMAC as a seedless extractor. The four resulting online extractor constructions are pictorially represented in Figure 1.

In Section 3.5, we also briefly discuss *seedless key-derivation functions (seedless KDFs)*, and extend the Extract-then-PRF paradigm of Krawczyk [35] to the seedless setting. Using our new seedless extraction results for HMAC, we get theoretical justification of the widely used HKDF construction of [35] in the seedless setting.

Full-scale seedless PRNGs. Finally, we take all our ideas together to solve our main problem: defining and building practical, yet provably secure seedless PRNGs. In Section 4 we introduce a novel security definition for PRNGs that differs from previous notions [18, 1, 25] in several crucial ways. The detailed comparison appears in Appendix A, but we present the highlights here.

First and foremost, our design is seedless. This is accomplished by carefully defining the legitimacy condition (relative to the fixed-length ideal primitive P), by conditioning our entropy notion on the list \mathcal{L} of the queries to P made by the attacker. Second, our seedless design allows us to merge the “distribution sampler” and the distinguisher used by [18, 25] into a single attacker \mathcal{A} ,⁷ making our notion much simpler to describe. Third, the works of [18, 25] used a much weaker notion of worst-case min-entropy; moreover, the final entropy of the the source X was defined as sum of individual worst-case min-entropies of the individual blocks of X conditioned on all the other blocks (before and after). In contrast, we use a much better notion of *average-case* min-entropy, and only look at the global average-case min-entropy of the entire (long) vector X . Thus, our notion of entropy is much less conservative: realistic entropy sources are likely to have *much higher* entropy according to our definition, even when conditioning on the list \mathcal{L} . Fourth, the notion of [18, 25] had explicit “entropy estimates” that the attacker had to provide. Our notion gets rid of these estimates. Finally, and somewhat surprisingly, we still managed to define our notion of legitimacy of the entropy source in a manner which is *construction-independent*. This means that one can potentially study the entropy properties of the source in a manner independent of the PRNG used on this source.

We also define both computational and information-theoretic (IT) notions of PRNG security. As with on-line extractors, for IT-PRNGs the attacker is allowed to read the entire ideal primitive P after it finished generating the last block of it’s oracle-dependent source X used for extraction. Such a notion is important for applications where privacy must hold well after the PRNG is finished its operations, or where information-theoretic security is important.

Our PRNG constructions. In Section 5 we then present three main PRNGs which are provably secure in our framework: based on Merkle-Damgård with a random compression function (see Figure 3), based on Merkle-Damgård with the Davies-Meyer compression function (see Figure 4), and based on Sponges (see Figure 5). All these constructions are extremely natural and practical, as Merkle-Damgård-based functions abstract SHA-2, while Sponges abstract SHA-3—two most widely used cryptographic hash functions. Moreover, for all three constructions we proved their security both in the computational and in the information-theoretic settings. See Figures 3, 4, 5 to see how the corresponding constructions change between the computational and the IT settings.

Thus, our work (including new notion of oracle-dependent entropy) could be used as theoretical justifications why these popular hash functions yield good *seedless* PRNGs (as well as online randomness extractors) even for a wide class of oracle-dependent entropy sources. Our three computational proofs heavily use the “coefficient-H” technique [40, 13], while our three information-theoretic proofs extend the framework of so-called “graph-counting” proofs [17, 7, 24] to bound the collision probability of iterated hash constructions. One novel challenge we had to solve here comes from the fact that theinput source could depend on the list \mathcal{L} of the ideal primitive queries, which breaks the “source-primitive” independence assumption crucially used in these already subtle proofs.

We also showed numeric examples of how we propose to use our constructions. Overall, we believe all of them are deployment ready, and we hope this work will start influencing future PRNG deployments, and will be incorporated into next RNG standards.

⁷Since we no longer need to hide the seed from the distribution sampler, forcing us to separate it from the attacker.

Implications to standard model. To overcome the impossibility of seedless extraction, our entropy notion is defined relative to the ideal primitive P . As we argued in detail in Section 1.2, working in the idealized model seems somewhat inherent to our approach, provided we wish to avoid highly non-standard, and likely tautological, cryptographic assumptions about the CHF we are using in the standard model. Still, it is good to ask what one might expect from our extractor and PRNG constructions with real-world CHFs, such as those based on SHA-2, SHA-3, HMAC, HKDF, etc.

As we already mentioned, we believe these constructions are secure for real-world entropy sources, because our idealized notion of entropy informally corresponds to sources which have fresh entropy, even given all the hash function evaluations happening around the source, which seems to be true for all distributions from the class $\mathcal{F}_{\text{practice}}$ of sources encountered in practice. To state the counter-positive, we believe that any real-world attack against our constructions with existing hash functions will either require a highly artificial entropy source (outside the class $\mathcal{F}_{\text{practice}}$), or will find a surprising weakness in the corresponding CHF. In other words, we heuristically conjecture:

All seedless extractor, PRNG and KDF constructions studied in this work are secure for all all high-entropy distributions in the class $\mathcal{F}_{\text{practice}}$.

1.5 Other Related Work

We mention some important categories of related works, in particular with respect to seedless extraction, PRNGs, and their security.

Seeded extractors and PRNGs. We already mentioned the extensive work on seeded extractors started by the seminal paper of Nisan and Zuckerman [38], and why they are problematic in our context. In the context of PRNGs, the first seeded PRNG notion was defined and constructed by Dodis et al. [18], who extended the prior “monolithic PRNG” definition of Barak and Halevi [1] (which did not explicitly talk about the seed, assuming the extraction module is “good enough” for the class of distributions produced by the entropy source). This line of work was extended in various ways by [20, 25, 28], where the latter two works were also analyzed in the random permutation model (in addition to the seed). However, none of these works considered a seedless setting for general entropy sources.

Extractors and PRNGs in ideal models. Extractors and PRNGs were also studied in the ideal models by several works [17, 9, 42, 48]. While not having explicit seeds, these works nevertheless modeled the *entropy source as being independent of the ideal primitive*. As we argued above, such oracle-independent modeling seems to be too restrictive for many realistic scenarios. Also, from a theory point of view, it effectively allows an exponentially long seed (the randomness used to sample the corresponding ideal primitive), making the positive results less interesting theoretically than the above-mentioned work on seeded extractors and PRNGs.

Indeed, the main motivation of all these papers was not to design theoretically optimal extractors and PRNGs, but to analyze the heuristic use of various cryptographic hash functions and popular modes of operations (such as CBC, HMAC, etc.) for randomness generation and extraction—a task these objects were not natively designed for. From this perspective, and given their widespread use, analyzing their extraction properties was an important first step in understanding their security, even under the restrictive oracle-independence assumption. Our work could be viewed as making a critical leap forward, by dropping—for the first time—the oracle-independence assumption, but

instead carefully modeling what constitutes entropy in the much more realistic, oracle-dependent setting.

Restricting the class of entropy sources. This line of work has primarily focused on the question of extraction, by assuming that the source X has more structure beyond entropy. Early examples [46, 16, 10, 36, 14] include various bit-fixing and limited dependence sources, culminating with the question of extracting from several independent sources [2, 11]. While mathematically very elegant, the types of sources studied by these works appear “too structured” to be realistic in the PRNGs scenario.

A different kind of restriction on the entropy source was studied by Barak et al. [3]. Rather than restrict sources by some property of their distribution, the work of [3] allows for arbitrary min-entropy sources, but assumes they come from an a-priori bounded number of distributions. While potentially promising for the setting of PRNGs, there are two disadvantages of the work of [3] as compared to this work. First, the work of [3] concentrated on the “monolithic” extraction setting, and did not address the question of entropy accumulation, where the entropy in X might come slowly from a large number of samples, and has to be accumulated into bounded state. In particular, it is unclear how to extend their constructions to address entropy accumulation with a fixed-length state. Second, the particular solutions offered by [3] used so called t -wise independent hash functions for a large values of t (at least as large as the overall source length). These functions are quite inefficient, and might not be fast enough for general purpose PRNGs.

We note that our work could also be viewed as overcoming impossibility of extraction by restricting the type of the source. However, we feel that our modeling is more natural for (and, thus, applicable to) the existing entropy sources, as used by the current PRNGs.

Low-Complexity Samplers. Introduced by Trevisan and Vadhan [45] and later extended by [32], here one assumes that the entropy source producing input X is unable to run the extractor/PRNG even once, thus making it impossible to do extractor-fixing. While this might be useful for situations where the entropy source is extremely simple, it is too restrictive for most applications, such as general purpose PRNG design studied in this work. In contrast, in this work the entropy source can easily run the extractor, but the legitimacy condition is defined in a way that doing the trivial extractor-fixing attack—by running the extractor—will result in a low-entropy, “illegitimate” source.

Randomness condensers. This approach, formalized by Dodis, Ristenpart and Vadhan [19], relaxes the security guarantees of the randomness extractor to only ensure that the output of the (seedless or “source-dependent-on-seed”) condenser is almost full entropy, despite not being perfectly uniform. Indeed, this weaker security turns out to be sufficient for several applications, such as key derivation schemes for signature schemes. Unfortunately, if we want an extractor rather than a condenser—which is essential for general purpose PRNGs—this approach is not sufficient.

UCEs and public-seed pseudorandomness. The notion of universal computational extractors (UCEs) [6], and its generalizations [44], study a complementary problem to the one studied here: how to extract from any entropy source which is only *computationally-hard-to-predict*, so it only has “computational entropy”. On a positive note, and similar to this work, when instantiated with constructions from an ideal primitive P , a UCE hash function yields a good extractor even if the inputs to it (the actual source) can be sampled depending on the ideal primitive. The issue, however, is that the current UCE notion inherently *requires a seed*, making it inapplicable for the PRNG scenario. An interesting direction for future research could be to extend our work to deal

with computational entropy, by defining and constructing *seedless* UCEs in idealized models, and possibly extending them to full-blown seedless PRNGs for computational entropy.

2 Preliminaries

2.1 Statistical Distance and Min-Entropy

The *statistical distance* of two random variables X and Y is $\text{SD}(X, Y) = \frac{1}{2} \sum_x |\mathbb{P}[X = x] - \mathbb{P}[Y = x]|$. The *prediction probability* of a random variable X is $\text{Pred}(X) := \max_x \mathbb{P}[X = x]$, and we also denote $\text{Pred}(X|y) := \max_x \mathbb{P}[X = x|Y = y]$. The conditional version of prediction probability is defined as

$$\text{Pred}(X|Y) := \mathbf{E}_{y \leftarrow Y} [\text{Pred}(X|y)] .$$

The (*average-case*) *conditional min-entropy* is $H_\infty(X|Y) = -\log(\text{Pred}(X|Y))$.

2.2 Security Games

All of the security properties considered in this paper are captured by considering a game between a challenger and an attacker \mathcal{A} , both of which may have access to an ideal primitive P . The goal of the attacker is to guess a random bit b chosen by the challenger, who offers a set of oracles to the attacker to aid with this task. The *advantage* of \mathcal{A} is defined as

$$2 \cdot | \mathbb{P}[\mathcal{A} \text{ wins}] - 1/2 | ,$$

where the probability is over the randomness of \mathcal{A} , of the challenger, and of the ideal primitive. The cases where $b = 0$ and $b = 1$ are referred to as the *real world* and the *ideal world*, respectively. One may equivalently consider \mathcal{A} 's advantage at telling these two worlds apart, i.e.,

$$| \mathbb{P}[\mathcal{A} = 1|b = 0] - \mathbb{P}[\mathcal{A} = 1|b = 1] | .$$

3 Seedless Extraction and Key Derivation

As a warm-up for full-fledged seedless PRNGs, this section considers the simpler property of *extraction*, i.e., producing uniformly random bits from weak high-entropy sources. Extraction can be seen as corresponding to the post-compromise security of PRNGs, and as such it will be implied by PRNG *robustness* (as defined in Section 4.2). The definition of extraction security in Section 3.1 considers the entropy of the attacker's input to the extractor conditioned on the attacker's *state* and the *queries* made to an ideal primitive P . A definition is provided for *computational* or *information-theoretic* security. IT extractors differ from computational ones in that the output of the extractor remains random even if the attacker, *after providing the input*, is given the entire function table of the underlying ideal primitive. That is, IT extractors achieve so-called *everlasting security* (cf. works in the hybrid bounded-storage model by Harnik and Naor [26]).

Section 3.2 considers extracting with a *monolithic random oracle*. The corresponding security proofs (for the computational and IT cases) are instructive for understanding the actual PRNG constructions provided in Section 5. Since considering a monolithic oracle is not motivated by any hash function used in practice, Section 3.3 introduces the concept of *online* extraction. An online extractor accumulates the entropy of its inputs in an internal state, from which uniform randomness can be produced. Finally, in order to illustrate the non-triviality of online extraction, Section 3.4 shows that extractors based on the popular CBC mode are not suitable for extraction.

Due to their importance, Section 3.5 briefly discusses *key-derivation functions (KDFs)* and analyzes the Extract-then-PRF paradigm (on which the widely used HKDF is based) put forth by Krawczyk [35] in the seedless setting.

3.1 Definition

In a model with idealized primitive P (chosen from some set \mathcal{P}), seedless extractors are algorithms $\text{ext}^P : \mathcal{X} \rightarrow \mathcal{Y}$ with oracle access to P . The security definition for such extractors considers a two-stage attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, where both parts have access to P . The first stage \mathcal{A}_1 outputs a value x and some state information σ for \mathcal{A}_2 . The second stage takes an input $y \in \mathcal{Y}$ and outputs a single bit (i.e., it acts as a distinguisher).

For an attacker \mathcal{A} , denote by \mathcal{L}_1 and \mathcal{L}_2 the (random variables corresponding to) the lists of the P -queries made by \mathcal{A}_1 and \mathcal{A}_2 , respectively.

Definition 1. *An attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is called a q -attacker if $|\mathcal{L}_1 \cup \mathcal{L}_2| \leq q$ always; it is called a q -IT-attacker if $|\mathcal{L}_1| \leq q$ always.*

That is, for IT-attackers the second stage \mathcal{A}_2 may make an arbitrary number of queries to P . Equivalently, \mathcal{A}_2 can be thought of as being given the entire function table of P .

The security game for seedless extractors in the P -model roughly requires that if the extractor is given a high-entropy input by \mathcal{A}_1 , then \mathcal{A}_2 cannot tell the extractor output apart from a random value in \mathcal{Y} , even given the state information σ and access to P . Formally, it proceeds as follows:

1. The challenger chooses $b \leftarrow \{0, 1\}$ and $P \leftarrow \mathcal{P}$ uniformly at random.
2. \mathcal{A}_1 gets access to P and produces $(\sigma, x) \leftarrow \mathcal{A}_1^P$.
3. The output of the extractor is computed as $y_0 \leftarrow \text{ext}^P(x)$. Moreover, the challenger picks a value $y_1 \leftarrow \mathcal{Y}$ uniformly at random.
4. The second-stage attacker \mathcal{A}_2 is given σ and y_b and outputs a decision bit $b' \leftarrow \mathcal{A}_2^P(\sigma, y_b)$. The attacker wins if and only if $b' = b$.

The advantage of \mathcal{A} in this extraction game is denoted by $\text{Adv}_{\text{ext}}^{\text{ext}, P}(\mathcal{A})$.

An attacker has to satisfy a legitimacy condition. Intuitively, this condition requires that the output X of \mathcal{A}_1 have high min-entropy even conditioned on the state information Σ and the list of queries \mathcal{L}_1 .⁸

Definition 2. *An attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is said to be γ^* -legitimate if, in the extraction game above,*

$$\mathbb{H}_\infty(X|\Sigma\mathcal{L}_1) \geq \gamma^* .$$

The above finally leads to the following definition of seedless extractor in the P -model:

Definition 3. *An algorithm $\text{ext}^P : \mathcal{X} \rightarrow \mathcal{Y}$ with oracle access to P is a seedless $(\gamma^*, q, \varepsilon)$ -(IT-)extractor in the P -model if for every γ^* -legitimate q -(IT-)attacker \mathcal{A} ,*

$$\text{Adv}_{\text{ext}}^{\text{ext}, P}(\mathcal{A}) \leq \varepsilon .$$

⁸Note, in the extraction game the definition of \mathcal{L}_1 is the same in the real and the ideal worlds. For our future definitions of PRNGs, however, it will be important that the notion of legitimacy is defined in the ideal world (i.e., conditioned on $b = 1$).

3.2 Seedless Extraction with a Monolithic Random Oracle

For instructive purposes it is useful to consider monolithic extraction, i.e., the case where the ideal primitive P itself is used as an extractor. To exemplify this, assume P is a random oracle, i.e., a function $G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ chosen uniformly at random. Then, the monolithic extractor is defined as follows:

Construction 1 (Monolithic extractor). The monolithic seedless extractor $\text{mono}^G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ using a random oracle $G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined by

$$\text{mono}^G(x) := G(x) .$$

Theorem 1 (Monolithic seedless extraction). *Construction mono is a $(\gamma^*, q, \varepsilon)$ -extractor in the G -model for*

$$\varepsilon \leq \frac{q}{2^{\gamma^*}} .$$

The proof of Theorem 1 is deferred and can be found in Section 6.2.

Theorem 2 (Monolithic seedless IT-extraction). *Construction mono is a $(\gamma^*, q, \varepsilon)$ -IT-extractor in the G -model for*

$$\varepsilon \leq \frac{1}{2} \sqrt{\frac{2^{-(\gamma^* - n)}}{1 - \rho}} + \rho ,$$

where $\rho = q/2^{\gamma^*}$.

The proof of Theorem 2 is deferred and can be found in Section 7.2.

Parameter choices. In terms of concrete parameters, observe the following for the constructions towards monolithic seedless extraction from above:

- **Computational:** If we let $n = 512$ and $q = 2^{80}$. We would need $\gamma^* \approx 160$ to get 80 bits of security.
- **Information Theoretic:** We let $n = 512$. We also approximate $1/(1-\rho) \leq 2$, very generously. Then, if we set for example $q = 2^{80}$. We would need the entropy loss, i.e., $\gamma^* = 160$ for 80 bits of security.

3.3 Online Extraction

In practice it is uncommon to assume access to a monolithic random oracle. Instead, practical hash functions are usually built from (public) compression functions, ciphers, or permutations. These underlying primitives P have limited input length and will therefore not be able to process inputs of arbitrary length m . Therefore, extractors (and PRNGs) should be designed in such a way that they can process short m -bit input blocks (e.g., $m = 256, 512, 1600$) and accumulate their entropy in the internal state.

An “accumulating” extractor ext satisfies the same security Definition 3, but its syntax can be thought of as two algorithms $\text{ext} = (\text{refresh}, \text{finalize})$, where refresh accumulates entropy in an internal state and finalize produces the extractor output from the current state.

Definition 4. *An online extractor construction consists of two algorithms $\text{ext} = (\text{refresh}, \text{finalize})$, where*

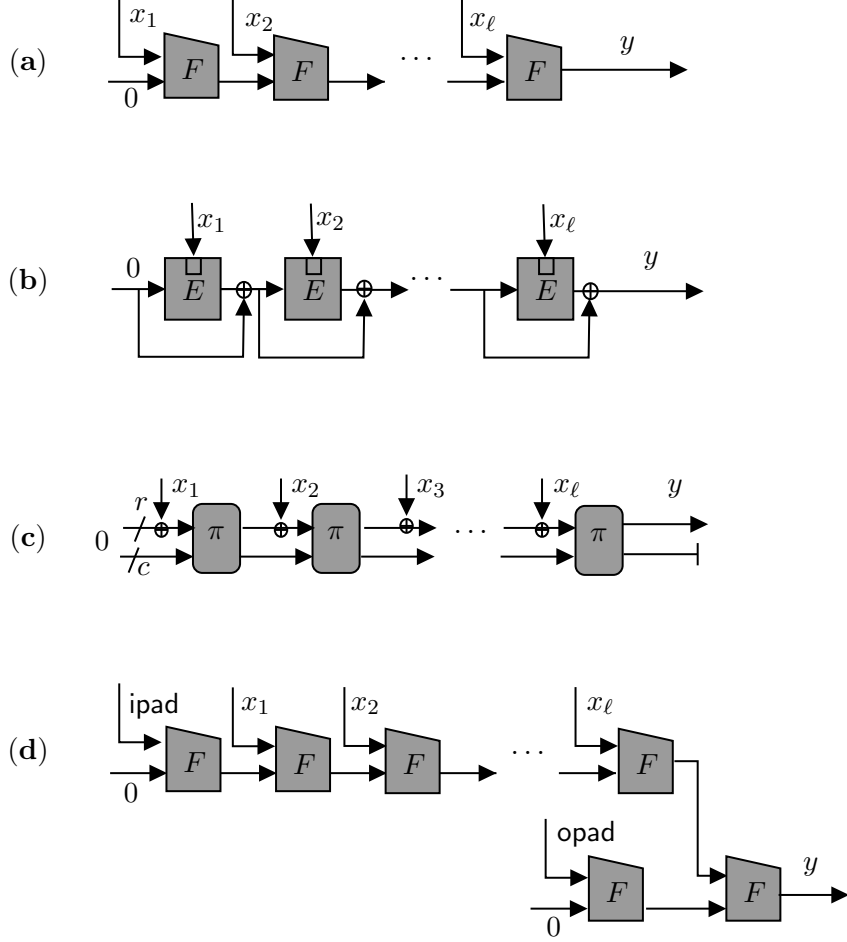


Figure 1: Four online computational extractors are represented in this diagram. These are based on: (a) Merkle-Damgård with a random compression function; (b) Merkle-Damgård with Davies-Meyer; (c) Sponge; and (d) HMAC.

Each extractor is shown to process inputs $x_1 \dots x_\ell$ (calls to refresh) to compute the output y (call to finalize). The IT variant for constructions (a), (b), (d) truncates the output y to the first r bits. For the Sponge construction (c), the IT variant is identical, as even the computational variant needs a truncated output.

- refresh takes a state s and an input $x \in \{0, 1\}^m$ and produces a new state $s' \leftarrow \text{refresh}^P(s, x)$, and
- finalize takes a state s and produces an output $y \in \{0, 1\}^r$, i.e., $y \leftarrow \text{finalize}^P(s)$.

An online extractor processing m -bit inputs and producing r -bit output is called a (m, r) -online extractor.

The security definition for online extractors additionally considers the number ℓ of times refresh is called by the attacker, i.e., it considers (q, ℓ) -attackers.

Definition 5. An algorithm $\text{ext}^P : \mathcal{X} \rightarrow \mathcal{Y}$ defined by two algorithms $\text{ext} = (\text{refresh}, \text{finalize})$ with oracle access to P is an $(\gamma^*, q, \ell, \varepsilon)$ -(IT)-online extractor in the P -model if for every γ^* -legitimate (q, ℓ) -(IT)-attacker \mathcal{A} ,

$$\text{Adv}_{\text{ext}}^{\text{ext}, P}(\mathcal{A}) \leq \varepsilon.$$

Online extractors can be built just like the PRNG constructions in Section 5, and, in fact, the corresponding security results follow as a special case of PRNG security. Correspondingly, their treatment

is deferred until Section 5, where such online extractors (and, in fact, full-fledged PRNGs) can be obtained from Merkle-Damgård with a random compression function, from Merkle-Damgård with the Davies-Meyer compression function, and from Sponges. For the reader’s convenience, Appendix B contains the online extractor constructions along with the security bounds—for applications where extraction is sufficient. In addition, said section in the appendix also considers the HMAC construction as a seedless extractor. HMAC is roughly based on Merkle-Damgård, but it has a few modifications/additions that are unnecessary for extraction. However, due to its wide-spread use, we point out how to modify the Merkle-Damgård proofs to obtain a security statement for HMAC. These constructions are pictorially represented in Figure 1.

In contrast to Merkle-Damgård and Sponges, as shown in the next section, using the CBC paradigm (which can be thought of as an “extreme sponge”) will not lead to a secure online extractor.

3.4 CBC-Based Extractors Are Insecure

A natural candidate for an online seedless extractor is using a permutation in CBC mode. A CBC-based extractor construction uses a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to absorb n -bit inputs. Its refresh function is defined as

$$\text{refresh}^\pi(s, x) = \pi(s \oplus x) .$$

However, it turns out that this approach does not lead to a secure extractor. This section presents a simple attack against CBC-based extractors. The attack works irrespective of how the finalization function is defined.

Theorem 3 (Attack against CBC Extractors). *Let refresh as defined above. There exists an ℓ -legitimate q -attacker \mathcal{A} with black-box access to a function finalize, such that for all CBC = (refresh, finalize)*

$$\text{Adv}_{\text{CBC}}^{\text{ext}, \pi}(\mathcal{A}) = 1 - 2^{-(r-1)} ,$$

where r is the output length of the extractor, $q = 2\ell + 2\alpha$, and α is the query complexity of finalize.

The idea of the attack is to have the attacker create the i^{th} input block as either $\pi^i(0^n) \oplus \pi^i(1^n)$ or 0, each with probability $1/2$.⁹ After ℓ such steps, the attacker will have provided ℓ bits of entropy (even conditioned on its π -queries), but only a single bit will have accumulated in the state, which will be $\pi^i(0^n)$ or $\pi^i(1^n)$, each with probability $1/2$.

Proof. Consider a two-stage attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. \mathcal{A}_1 works as follows:

1. Initially, set $a = 0^n$ and $b = 1^n$.
2. For blocks $i = 1$ to ℓ :
 - (a) Set $x_{i,0} = 0$ and $x_{i,1} = a \oplus b$. Choose random bit $\beta \leftarrow \{0, 1\}$ and set $x_i = x_{i,\beta}$.
 - (b) Set $a \leftarrow \pi(a)$ and $b \leftarrow \pi(b)$.
3. Set state information to $\sigma \leftarrow (a, b)$. In particular, *forget* all values $x_{i,b}$ and all values of β . Return σ and $x = (x_1, \dots, x_\ell)$.

\mathcal{A}_2 , given $\sigma = (a, b)$ and y as input, proceeds as follows:

1. Compute $\text{finalize}(a)$ and $\text{finalize}(b)$. If either of them equals y , return 0; else, return 1.

⁹Here, π^i denotes the i -fold application of π .

Assume the initial state of the extractor is 0^n . In order to understand the attack, consider the state of the extractor after XORing the first input x_1 : it is either 0^n or 1^n , each with probability $1/2$. After applying `refresh`, which consists of just applying π to the state, the new state is either $\pi(0^n)$ or $\pi(1^n)$, again with probability $1/2$ each.

The second input x_2 is either 0^n or $\pi(0^n) \oplus \pi(1^n)$; XORing it to the state results in either $\pi(0^n)$ or $\pi(1^n)$. After applying `refresh`, the new state is $\pi^2(0^n)$ or $\pi^2(1^n)$. Extending this argument to all ℓ inputs, the state of the extractor after absorbing them is either $a = \pi^\ell(0^n)$ or $a = \pi^\ell(1^n)$. After calling `finalize`, the state is either `finalize(a)` or `finalize(b)`. Therefore, in the real world ($b = 0$), \mathcal{A}_2 will always output 0, whereas in the ideal world ($b = 1$) this will only happen with probability at most $2^{-(r-1)}$.

In terms of efficiency, observe that \mathcal{A}_1 makes 2ℓ queries to π , and \mathcal{A}_2 makes twice the number α of π -queries required to evaluate a call to `finalize`.

Finally, in order to show that \mathcal{A} is legitimate, observe first that the state σ can be computed from the queries made by \mathcal{A}_1 . It easily seen that given only the queries, the vector $x = (x_1, \dots, x_\ell)$ has ℓ bits of entropy. \square

3.5 Seedless HKDF

This section argues about the security of the well-known HKDF [35] key-derivation function (KDF), where recall that KDFs are used in practice to derive random-looking key material from high-entropy sources. The work of [35] shows that such a KDF can be generically constructed by composing an (online) extractor with a (variable-output-length) pseudorandom function (PRF). HKDF is then a particular instantiation of this general paradigm, where both the extractor and the PRF are instantiated with HMAC. The original HKDF construction is seeded, since the extractor used is seeded. This section presents an *unseeded* version of HKDF, by using a *seedless* extractor instead.

Standard or Ideal Primitive Security? Before presenting our results, we must decide if we will consider “standard” or “ideal primitive” security of KDFs when analyzing HKDF. At first a glance, the latter option seems advantageous, since HMAC is well known to be a PRF in the ideal model with a random compression function (see [35]), and we also show in Section B.4 that HMAC is a seedless extractor with respect to the same idealized model (and our new modeling of entropy).

Technically, however, we cannot automatically argue the security of HKDF as a KDF in the ideal compression function model¹⁰ by composing these two individual proofs. This is because the HKDF construction uses the same compression function for the extractor and the PRF, and it is well known that standard composition does not generically hold when the same ideal primitive is reused. While we have no doubt that we could do one long and tedious proof showing that no such composition problem arises in the case of HKDF, we feel such direct proof is not particularly illuminating, and does not add much as compared to doing everything in the standard model, which we explain next.

Namely, the standard model definition of KDFs (and the analysis of HKDF) would be parameterized by the family of input distributions \mathcal{F} that we care about for our KDF setting. We will then show that the composition of \mathcal{F} -secure (seedless) extractor and (standard model) PRF, advocated by [35], still gives an \mathcal{F} -secure (seedless) KDF in the standard model. Then, we will use our ideal model results mentioned above to heuristically argue that HMAC is simultaneously a secure standard-model PRF (something which is widely believed well before this work), but also $\mathcal{F}_{\text{practice}}$ -secure seedless extractor for the class of sources “encountered in practice”. Combined, we get a convincing heuristic showing $\mathcal{F}_{\text{practice}}$ -security of HKDF, without formally defining $\mathcal{F}_{\text{practice}}$.

¹⁰Something we would have to first define with respect to our new modeling of entropy.

To compare the two approaches outlined above, we see that they would achieve the same end result: a conjectured security of HKDF against “practically relevant” entropy sources $\mathcal{F}_{\text{practice}}$. However, the ideal model approach would have a huge cost — a new definition of KDFs with respect to our new entropy notion, and a long/tedious direct proof for HKDF, largely repeating two existing ideal model proofs for HMAC. In contrast, the standard model approach described above achieves the same effect without needing any new proofs or definitions; by applying the “random oracle heuristic” one level earlier twice (at the level of extractor and PRF) rather than one level later once (at the level of KDF). In fact, one can even argue that the “standard model approach” achieves a stronger counter-positive which is not achieved by using the higher-level KDF abstraction. If the composed construction is not a secure standard-model KDF, then this would constitute a natural counterexample to the random-oracle methodology: by virtue of Theorem 4, either HMAC is not a good standard-model extractor (but secure in the idealized model), or HMAC is not a good standard-model PRF (but secure in the idealized model). Either one of such non-contrived counterexamples would be a remarkable breakthrough in the cryptanalysis of HMAC.

Hence, the rest of this section will deviate from the rest of the paper, and follow the “standard model” approach for our treatment of KDFs.

Syntax. A *seedless key-derivation function (SL-KDF)* KDF is an algorithm that takes as input

- source material $x \in \{0, 1\}^m$,
- context information c , as well as
- a desired output length o

and outputs an o -bit string

$$y \leftarrow \text{KDF}(x, c, o) .$$

Security. The security definition for seedless KDFs—in the standard model—considers a two-stage attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. The first stage \mathcal{A}_1 outputs a value x and some state information σ for \mathcal{A}_2 . The second stage attacker \mathcal{A}_2 , based on σ , may make particular queries (described below) and outputs a single bit at the end. Formally, the game proceeds as follows:

1. The challenger chooses $b \leftarrow \{0, 1\}$.
2. \mathcal{A}_1 produces $(\sigma, x) \leftarrow \mathcal{A}_1$.
3. The second-stage attacker \mathcal{A}_2 is given σ and may make construction queries (c, o) , which are answered by $\text{KDF}(x_1, \dots, x_\ell, c, o)$ if $b = 0$, and by a uniformly random value in $\{0, 1\}^o$ if $b = 1$. \mathcal{A}_2 is restricted to a single query for every value c .
4. At the end, \mathcal{A}_2 outputs a decision bit b' and wins the game if and only if $b' = b$.

The advantage of \mathcal{A} in this extraction game is denoted by $\text{Adv}_{\text{KDF}}^{\text{kdf}}(\mathcal{A})$.

An attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is called an (\mathcal{F}, T, q) -attacker if (1) the distribution of the values (σ, x) produced by \mathcal{A}_1 is from a class of distributions \mathcal{F} , (2) \mathcal{A}_2 runs in time at most T , and (3) makes at most q queries to the KDF. The above leads to the following definition of seedless key-derivation functions in the standard model:

Definition 6. An algorithm KDF is a $(\mathcal{F}, T, q, \varepsilon)$ -SL-KDF if for every (\mathcal{F}, T, q) -attacker \mathcal{A} ,

$$\text{Adv}_{\text{KDF}}^{\text{kdf}}(\mathcal{A}) \leq \varepsilon .$$

Pseudorandom functions. A *variable-length pseudorandom function (VL-PRF)* is a function $\text{PRF}(k, c, o)$ that, on a uniformly random key, is computationally indistinguishable from a truly random function that produces outputs of the desired length. Specifically, the security of a VL-PRF can be captured in the P -model by giving an attacker \mathcal{A} access to P and allowing it to make *construction* queries to either $\text{PRF}(k, \cdot, \cdot)$ on a random key or to a truly random function that outputs a random and independent value from $\{0, 1\}^o$ for every input (c, o) ; \mathcal{A} is restricted to make at most one construction query for every value c . Function PRF is called a (p, q, ε) -VL-PRF in the P -model if an attacker making at most p queries to P and at most q construction queries has advantage at most ε of in telling PRF apart from a random function in the above experiment. Similarly, one can define a (T, q, ε) -VL-PRF in the standard model, where T is the running time of \mathcal{A} .

Extract-then-expand KDFs and HKDF. Seedless KDFs can be generically built by composing a seedless extractor ext with a variable-length PRF PRF as follows:

$$\text{KDF}(x, c, o) := \text{PRF}(\text{ext}(x), c, o) .$$

Observe that the definition of seedless extractor (Definition 3) can be modified to a standard-model definition of $(\mathcal{F}, T, \varepsilon)$ -extractors (similarly to Definition 6):

Definition 7. An algorithm ext is a seedless $(\mathcal{F}, T, \varepsilon)$ -extractor if for every (\mathcal{F}, T) -attacker \mathcal{A} ,

$$\text{Adv}_{\text{ext}}^{\text{ext}}(\mathcal{A}) \leq \varepsilon .$$

The following theorem by [35] is easy to prove (for most reasonable function classes \mathcal{F}).

Theorem 4. Let ext be a seedless $(\mathcal{F}, T, \varepsilon)$ -extractor and let PRF be a (T', q', ε') -VL-PRF. Then, KDF as defined above is a seedless $(\mathcal{F}, \min(T, T'), q, \varepsilon + \varepsilon')$ -KDF.

Recall that HKDF instantiates both the extractor and the VL-PRF with HMAC (cf. Section B.4). The security of HKDF as a KDF can be argued heuristically in the standard model as follows: First, one infers the standard-model security of HMAC as an extractor—from Theorem 55. In particular, the heuristic assumption is that the extractor security of HMAC in the random-compression-function model implies security of HMAC in the standard model against all distributions from the class $\mathcal{F}_{\text{practice}}$ of sources encountered in practice. Second, one infers the standard-model security of HMAC as a VL-PRF—from the well-known fact that HMAC is a good VL-PRF in the random-compression-function model.

4 Pseudorandom Number Generators with Input

A *pseudorandom number generator with input (PRNG)* is a stateful cryptographic primitive. It gradually accumulates entropy in its state by absorbing inputs and can be used to output pseudorandom bits once the entropy of the state is sufficiently high. Moreover, it is both forward and backward secure, i.e., past outputs remain random upon future state compromise, and, by absorbing sufficient amounts of entropy, a PRNG can recover from state compromise.

This section introduces a novel security definition for PRNGs that differs from previous notions in several crucial ways. Specifically, a comparison to the original *robustness* notion by Dodis *et al.* [18], based on work by Barak and Halevi [1], as well as to an adaptation of it by Gazi and Tessaro [25] for idealized models is provided in Appendix A.

This paper considers two notions of PRNGs: computational PRNGs and information-theoretically secure (IT) PRNGs. IT PRNGs differ from computational PRNGs in that once the attacker stops

interacting with the PRNG, the output of the PRNG remains random even if the attacker is given the entire function table of the underlying ideal primitive. That is, IT PRNGs achieve so-called *everlasting security* (cf. works in the hybrid bounded-storage model by Harnik and Naor [26]). This distinction is analogous to that between seedless extractors and IT seedless extractors (cf. Section 3).

4.1 Syntax

A PRNG consists of two algorithms: one for absorbing new inputs and one for producing pseudo-random outputs. Formally, it is defined as follows:

Definition 8 (Syntax of PRNGs). *A pseudorandom number generator with input (PRNG) is a pair of algorithms $\text{PRNG} = (\text{refresh}, \text{next})$ having access to an ideal primitive P and sharing an n -bit state s , where*

- *refresh takes a state s and an input $x \in \{0, 1\}^m$ and produces a new state $s' \leftarrow \text{refresh}^P(s, x)$, and*
- *next takes a state s and produces a new state and an output $y \in \{0, 1\}^r$, i.e., $(s', y) \leftarrow \text{next}^P(s)$.*

A PRNG processing m -bit inputs and producing r -bit output is called a (m, r) -PRNG.

4.2 Security Game

Robustness game. PRNGs are expected to satisfy the so-called *robustness* property, which captures the properties discussed at the beginning of Section 4. The corresponding security game is depicted in Figure 2. The game initially chooses a random bit b and initializes the state of the PRNG to 0^n . Subsequently, it offers the following oracles to \mathcal{A} :

- **adv-refresh**(x) calls the **refresh** procedure to absorb $x \in \{0, 1\}^n$ into the internal state of the PRNG;
- **get-next** and **get-next*** allow the attacker to get pseudorandom outputs by calling the **next** procedure on the current state and returning the output y . The difference between the two oracles is that **get-next** is supposed to be called only when the state has high entropy, whereas **get-next*** can be called *prematurely*, i.e., before the state has absorbed enough randomness for the **next** function to output pseudorandom values (cf. definition of legitimate attackers below).
- **next-ror** works like the **get-next**-oracle, except that it creates a challenge, i.e., if $b = 1$, it outputs a uniform random value $y_1 \in \{0, 1\}^r$ instead of the PRNG output y_0 .
- **get-state** and **set-state** model state compromises by letting the attacker learn the current state or set it to an arbitrary value, respectively.

The advantage of \mathcal{A} in the robustness game is denoted by $\text{Adv}_{\text{PRNG}}^{\text{rob}, P}(\mathcal{A})$.

Canonical attackers. It will be useful to define the following notion of canonical attackers: Consider the interaction of an attacker \mathcal{A} with the robustness game. The following events are called *entropy drains*:

- the beginning of the game,
- calls to **get-state** or **set-state**, and

The PRNG Robustness Game

init $s \leftarrow 0^n$ $b \leftarrow \{0, 1\}$	adv-refresh (x) $s \leftarrow \text{refresh}^P(s, x)$	get-state $\text{return } s$
next-ror $(s, y_0) \leftarrow \text{next}^P(s)$ $y_1 \leftarrow \{0, 1\}^r$ $\text{return } y_b$	get-next/get-next* $(s, y) \leftarrow \text{next}^P(s)$ $\text{return } y$	set-state (s^*) $s \leftarrow s^*$

Figure 2: Oracles for the PRNG robustness game.

- calls to **get-next***.

In other words, entropy drains are the events that cause the PRNG state to lose its entropy, which includes premature calls to **next**. An attacker \mathcal{A} is said to be *canonical* if it does not make **get-next*** queries nor the following query pattern: an entropy drain followed by one or more **adv-refresh** queries, followed by a **get-state** query.

Considering canonical attackers only is without loss of generality. This is because the above sequence of queries can be simulated by the attacker by making a **get-state** query right away and computing the output of **get-state** or **get-next*** itself. In particular, for every attacker \mathcal{A} , there exists a canonical attacker \mathcal{A} with the same advantage. All attackers in the remainder of this work are therefore assumed to be canonical.

Legitimate attackers. In order to obtain a sensible definition devoid of trivial attacks, attackers must satisfy a “legitimacy” condition. The condition roughly requires that an attacker only ask for challenges when it has sufficient amount of uncertainty about the PRNG’s internal state.

Towards formalizing the legitimacy condition, consider the interaction of \mathcal{A} with a *variant* of the robustness game defined as follows: Whenever oracles **next-ror** or **get-next** are called, instead of evaluating **next**, the game simply uses two uniformly random and independent values (s, y) as the output of **next**.

Observe that this variant of the robustness game, called the *legitimacy game* corresponds to an interaction between \mathcal{A} and an *ideal* PRNG, which produces perfect randomness. Moreover, the legitimacy game is *construction-independent*.

In the legitimacy game, define now the following random variables immediately before \mathcal{A} makes the i^{th} call to oracle **get-next** or **next-ror**:

- \mathcal{L}_i : the list of P -queries by \mathcal{A} and the corresponding answers;
- Σ_i : the state of \mathcal{A} ;
- \overline{X}_i : vector of inputs provided by \mathcal{A} since the *the most recent entropy drain (MRED)*; and
- S_i : the state of the PRNG immediately after the MRED.

The legitimacy condition requires that \mathcal{A} provide inputs that have high min-entropy even conditioned on its current state, the queries so far, and the state of the PRNG after the MRED.

Definition 9 (Legitimate attackers). *An attacker \mathcal{A} is said to be γ^* -legitimate if for all i ,*

$$H_\infty(\overline{X}_i | \Sigma_i \mathcal{L}_i S_i) \geq \gamma^* ,$$

where MREDs are defined as above.

In order to capture IT-legitimate attackers (against IT PRNGs), the set of entropy drains is extended to include

- calls to **get-next** and **next-ror**.

With this definition of MRED and notation analogous to that in the previous definition, IT-legitimate attackers are defined as follows:

Definition 10 (Legitimate IT attackers). *An attacker \mathcal{A} is said to be γ^* -IT-legitimate if for all i ,*

$$H_\infty(\overline{X}_i | \Sigma_i \mathcal{L}_i S_i) \geq \gamma^* ,$$

w.r.t. the extended definition of MRED.

Robust PRNGs. We are now ready to quantify the efficiency of attacker \mathcal{A} , and to define our final notion of PRNG robustness.

Definition 11 (Attacker efficiency). *An attacker is called a (q, t, ℓ) -attacker if*

- q is the maximum number of P -queries it makes,
- ℓ is the maximum number of **adv-refresh** calls between any entropy drain and successive call to either **next-ror** or **get-next**, and
- t is the maximum total number of calls to any oracle in the robustness game other than **adv-refresh**.

An attacker is called a (q, t, ℓ) -IT-attacker if it satisfies the above conditions but makes an arbitrary number of queries to P after the interaction with the challenger ends.

Definition 12 (Robustness of PRNGs). *A PRNG construction $\text{PRNG} = (\text{refresh}, \text{next})$ with oracle access an ideal primitive P is $(\gamma^*, q, t, \ell, \varepsilon)$ -IT-robust in the P -model if for every γ^* -IT-legitimate (q, t, ℓ) -IT-attacker,*

$$\text{Adv}_{\text{PRNG}}^{\text{rob}, P}(\mathcal{A}) \leq \varepsilon .$$

Observe that online extractors (cf. Definition 4) are a special case of robust PRNGs. In terms of construction, the PRNG **next** algorithm can be replaced by **finalize**, which simply discards the state output by **next**. If then the PRNG robustness game is relaxed such that the only queries the attacker can make are (a) arbitrarily many queries to **adv-refresh** followed by (b) $t = 1$ query to **next-ror**, one obtains a notion equivalent to Definition 3.

5 Constructions of PRNGs

This section presents three simple, intuitive, and—most importantly—practical PRNG constructions:

- a construction based on the *Merkle-Damgård paradigm* using a public *fixed-length compression function*;
- a construction based on the *Merkle-Damgård paradigm* using the *Davies-Meyer compression function* (as in SHA-2), which is built from any public block cipher; and
- a construction based on the *Sponge paradigm* (as in SHA-3), which uses a public permutation.

For PRNGs based on the MD paradigm, there are in fact two constructions: one achieving normal, computational PRNG security and one achieving information-theoretic (IT) security. The security analyses of all three constructions are provided in Sections 6 (computational PRNGs) and 7 (IT security). There is also an IT candidate based on Sponges, but its security analysis is left for future work.

For the reader’s convenience, Appendix B states the corresponding online extractor constructions along with the security bounds—for applications where extraction is sufficient.

5.1 PRNGs from Merkle-Damgård

A PRNG can be obtained from a compression function F as follows (cf. Figure 3):¹¹

Construction 2 (PRNG from Merkle-Damgård). The (m, r) -PRNG construction MD = (refresh, next) based on Merkle-Damgård with a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined as follows:¹²

- $\text{refresh}^F(s, x) = F(s, x)$, and
- $\text{next}^F(s) = (F(s, 0), F(s, 1) \parallel \dots \parallel F(s, r/n))$.

The security of Construction 2 is proved in the F -model, where F is a uniformly random function. The proof of the following theorem is deferred to Section 6.4.

Theorem 5 (Robustness of Merkle-Damgård PRNGs). *Construction 2 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -robust PRNG in the F -model for*

$$\varepsilon_{\text{rob}} \leq 2t \cdot \left(\frac{\tilde{q}^2 + \tilde{q}\ell + \ell^2}{2^n} + \frac{\tilde{q}}{2^{\gamma^*}} \right),$$

where $\tilde{q} = q + r/n + 1$.

An IT-robust PRNG based on Merkle-Damgård can be obtained if the next function simply outputs the truncated state (and outputs 0^n as the new state):

Construction 3 (IT-PRNG from Merkle-Damgård). The (m, r) -PRNG construction MD _{r} = (refresh, next) based on Merkle-Damgård with a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined as follows:

- $\text{refresh}^F(s, x) = F(s, x)$, and
- $\text{next}^F(s) = (0^n, s[1..r])$.

¹¹To reduce notational clutter, the algorithms `refresh` and `next` of the PRNG constructions are not “branded” with the design name. There will be no ambiguity as to which construction is meant in any place in this paper.

¹²The integer arguments to the compression function are to be naturally mapped to $\{0, 1\}^n$.

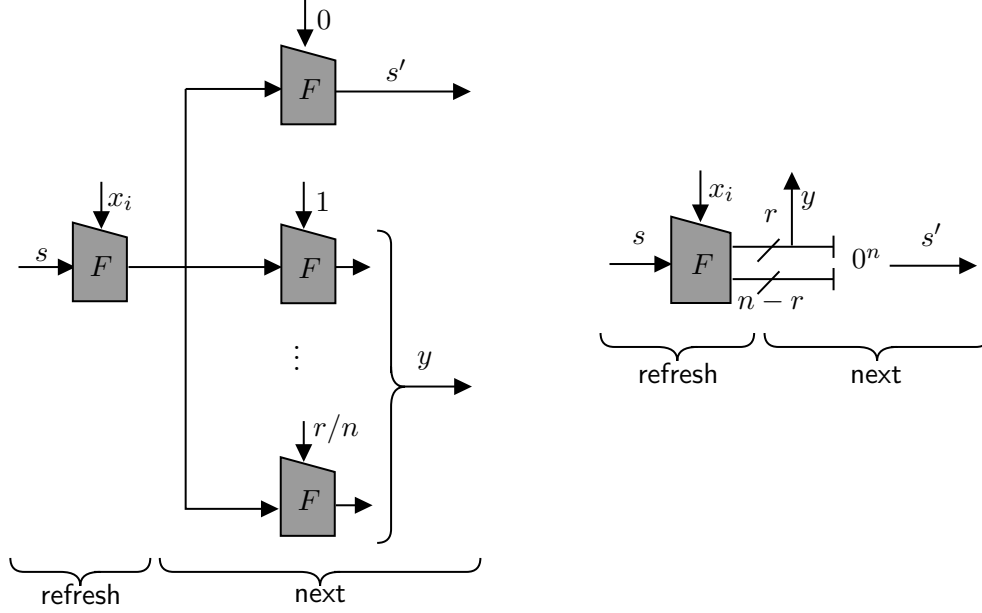


Figure 3: Procedures refresh (processing a single-block input x_i) and next of Merkle-Damgård PRNG constructions with compression function F . Left: Computationally secure Construction 2; right: IT secure Construction 3.

The security of Construction 3 is proved in the F -model, where F is a uniformly random function. To state the theorem for the IT construction, for an integer ℓ , let

$$d'(\ell) = \max_{\ell' \in \{1, \dots, \ell\}} |\{d \in \mathbb{N} : d|\ell'\}|.$$

Observe that, asymptotically, $d'(\ell)$ grows very slowly, i.e., as $\ell^{o(1)}$. Furthermore, let F be a random compression function. The proof of the following theorem is deferred to Section 7.4.

Theorem 6 (IT-Robustness of Merkle-Damgård PRNGs). *Construction 3 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust PRNG in the F -model, where*

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 \cdot \frac{\tilde{q}^2 2^r}{2^{2n}}} + t\rho,$$

for $\rho = \frac{\tilde{q}^2}{2^r}$ where $\tilde{q} = q + t\ell$.

Parameter choices. In terms of concrete parameters, observe the following for the Merkle-Damgård constructions above:

- **Computational PRNG:** If one were to use SHA-512 as compression function with $n = 512$, and, moreover, choose $r = n$. We let $t = 1, q = 2^{80}$ and let $\gamma^* = \ell$. This assumes that we get at least one bit of entropy from each block. We would need $\gamma^* \approx 160$ to get 80 bits of security.
- **IT PRNG:** For example, assume SHA-512's compression function is used, i.e., $n = 512$. If we let $r = 256$, then we get (we also approximate $1/(1-\rho) \leq 2$, very generously)

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{2^{257-\gamma^*} + \frac{\ell \cdot d'(\ell)}{2^{256}}} + t \frac{q^2}{2^{256}},$$

We let $\ell = \gamma^*$. Then, if we set for example $q = 2^{80}$. We would need the entropy loss, i.e, $\gamma^* - r = 162$ for 80 bits of security.

5.2 PRNGs from Merkle-Damgård with Davies-Meyer

The Davies-Meyer compression function maps two inputs $a \in \{0, 1\}^m$ and $b \in \{0, 1\}^n$ to an n -bit string

$$E(b, a) \oplus a ,$$

where E is an arbitrary block cipher (where b is the key and a the input).¹³ Correspondingly, a PRNG can be obtained from E as follows (cf. Figure 4):

Construction 4 (PRNG from MD-DM). The (n, r) -PRNG construction $\text{DM} = (\text{refresh}, \text{next})$ based on Merkle-Damgård with Davies-Meyer (MD-DM) uses a cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and is defined as follows:¹⁴

- $\text{refresh}^E(s, x) = E(x, s) \oplus s$, and
- $\text{next}^E(s) = (E(0, s) \oplus s, E(1, s) \oplus s \parallel \dots \parallel E(r/n, s) \oplus s)$.

The security of Construction 4 is proved in the E -model, where E is a cipher chosen uniformly at random from the set of all ciphers and can be queried in both the forward and backward direction. The proof of the following theorem is deferred to Section 6.5.

Theorem 7 (Robustness of MD-DM PRNGs). *Construction 4 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -robust PRNG in the E -model for*

$$\varepsilon_{\text{rob}} \leq 4t \cdot \left(\frac{\tilde{q}^2 + \tilde{q}\ell + \ell^2}{2^n} + \frac{\tilde{q}}{2^{\gamma^*}} \right),$$

where $\tilde{q} = q + r/n + 1$.

In the IT-secure variant of the MD-DM construction, refresh remains the same, but next will truncate the input state to r bits, which it outputs, and then zero out the state.

Construction 5 (IT-PRNG from MD-DM). The (n, r) -PRNG construction $\text{DM}_r = (\text{refresh}, \text{next})$ using Merkle-Damgård with Davies-Meyer (MD-DM) uses a block cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and is defined as follows:

- $\text{refresh}^E(s, x) = E(x, s) \oplus s$, and
- $\text{next}^E(s) = (0^n, s[1..r])$.

The security of Construction 5 is proved in the E -model, where E is a cipher chosen uniformly at random from the set of all ciphers and can be queried in both the forward and backward direction. Let $d'(\ell)$ be defined as in Section 5.1. The proof of the following theorem is deferred to Section 7.5.

Theorem 8 (IT-Robustness of MD-DM PRNGs). *Construction 5 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust PRNG in the E -model, where*

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + 64\ell^4 \cdot \frac{2^r}{2^{2n-2}} + 16\ell^2 \tilde{q}^2 \cdot \frac{2^r}{2^{2n-2}}} + t\rho ,$$

for $\rho = \frac{\tilde{q}^2}{2^r}$ where $\tilde{q} = q + t\ell$

¹³A (block) cipher is an efficiently computable and invertible permutation $E(k, \cdot) : \{0, 1\}^n \rightarrow \{0, 1\}^n$ for every key $k \in \{0, 1\}^k$.

¹⁴The integer arguments to the cipher are to be naturally mapped to $\{0, 1\}^n$.




Figure 4: Procedures *refresh* (processing a single-block input x_i) and *next* of Merkle-Damgård PRNG constructions with the Davies-Meyer compression function based on a block cipher E . Left: Computationally secure Construction 4; right: IT secure Construction 5.

Parameter choices. In terms of concrete parameters, observe the following for the PRNG constructions from Merkle-Damgård with Davies-Meyer above:

- **Computational PRNG:** SHA-512 is a 512-bit block cipher algorithm that encrypts 512 bit hash value using the input as key. Therefore, we let $n = 512$ and set $r = n$. We let $t = 1, q = 2^{80}$ and let $\ell = \gamma^*$. This assumes that we get at least one bit of entropy from each block. We would need $\gamma^* \approx 163$ to get 80 bits of security.
- **IT PRNG:** We again let $n = 512$. If we let $r = 256$, then we get (we also approximate $1/(1 - \rho) \leq 2$, very generously)

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{2^{129-\gamma^*} + \frac{\ell \cdot d'(\ell)}{2^{127}}} + t \frac{q^2}{2^{128}},$$

We let $\ell = \gamma^*$. Then, if we set for example $q = 2^{80}$. We would need the entropy loss, i.e., $\gamma^* - r = 162$ for 80 bits of security.

5.3 PRNGs from Sponges

Let $n \in \mathbb{N}$ and $n = r + c$. In the following, for an n -bit string s , let $s = s^{(r)} || s^{(c)}$ be decomposition of s into an r -bit and c -bit string. A PRNG using the Sponge paradigm can be obtained from a permutation π as follows (cf. Figure 5):

Construction 6 (PRNG from Sponges). The Sponge-based PRNG construction $\text{Spg} = (\text{refresh}, \text{next})$ uses a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to absorb and produce r -bit inputs and outputs, respectively, and is defined as follows:




Figure 5: Procedures *refresh* (processing a single-block input x_i) and *next* of Merkle-Damgård PRNG constructions with compression function F . Left: Computationally secure Construction 2; right: IT candidate Construction 3.

- $\text{refresh}^\pi(s, x) = \pi(s \oplus x \| 0^c)$, and
- $\text{next}^\pi(s) = (\pi(s) \oplus 0^r \| s^{(c)}, s^{(r)})$.

The next function design is due to Hutchinson [27], who simplified a proposal by Gazi and Tesaro [25]. Recall that the Merkle-Damgård constructions have a “parallel” next function in order to produce r/n blocks of random output with $r/n + 1$ calls to the ideal primitive, where the additional call is used to produce a new state. Were it not for this optimization, in order to obtain r bits of output, one would have to apply the next function r/n times in a row, which would result in twice the number of ideal-primitive calls.

The next function for Sponges, on the other hand, only makes a single call to the ideal primitive to produce both a new state and the random output. Therefore, no parallel next function is provided for the Sponge-based PRNG.

The security of Construction 6 is proved in the π -model, where π is a uniformly random permutation, which can be queried in both the forward and backward direction. The proof of the following theorem is deferred to Section 6.6.

Theorem 9 (Robustness of Sponge PRNGs). *Construction 6 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -robust PRNG in the π -model for*

$$\varepsilon_{\text{rob}} \leq 4t \cdot \left(\frac{\tilde{q}^2 + \tilde{q}\ell + \ell^2}{2^n} + \frac{\tilde{q}}{2^{\gamma^*}} + \frac{\tilde{q}^2}{2^c} \right),$$

where $\tilde{q} = q + r/n + 1$.

Observe that the bound in Theorem 9 is only reasonable when c is large enough, which matches the fact that CBC-based PRNGs—which correspond to the case $c = 0$, are not secure.

In the IT variant of the Sponge construction, *refresh* remains the same, but *next* will truncate the input state to r bits, which it outputs, and then zero out the state.

Construction 7 (IT-PRNG from Sponges). The Sponge-based PRNG construction $\text{Spg}_r = (\text{refresh}, \text{next})$ uses a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to absorb and produce r -bit inputs and outputs, respectively, and is defined as follows:

- $\text{refresh}^\pi(s, x) = \pi(s \oplus x \| 0^c)$, and
- $\text{next}^\pi(s) = (0^n, s[1..r])$.

An IT-PRNG based on the Sponge paradigm is a modification of the computational variant where the output of next is merely the truncated state to the first r bits along with 0^n as the new state.

Theorem 10 (IT-Robustness of Sponge PRNGs). *Construction 7 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust PRNG in the π -model for*

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \frac{\ell \cdot (\ell + \tilde{q})}{2^{c-1}}} + t\rho,$$

for $\rho = \frac{\tilde{q}^2}{2^c}$ where $\tilde{q} = q + t\ell$

Parameter choices. In terms of concrete parameters, observe the following for the PRNG constructions from Sponges above: above:

- **Computational PRNG:** SHA-3 like parameters have $n = 1600$ and $c = 1024$. We let $t = 1, q = 2^{80}$ and let $\ell = \gamma^*$. This assumes that we get at least one bit of entropy from each block. We would need $\gamma^* \approx 163$ to get 80 bits of security.
- **IT PRNG:** We let $n = 1600$ and $c = 1024$. In addition, we let $t = 1$ and $q = 2^{80}$. We also let $\ell = \gamma^*$. Therefore, we incur an entropy loss of 160 bits to get 80 bits of security.

6 Security Proofs for Computational Constructions

This section analyzes all previously presented computationally secure constructions. The main technique in all proofs is the so-called H-coefficient technique, which is discussed first, in Section 6.1. Section 6.2 then shows the security of the monolithic extractor from Section 3.2. While rather straight-forward, the corresponding proof is quite instructive for the security proofs of the PRNG constructions from Section 5. For these proofs, it is convenient to define and establish a number of intermediate properties, which are discussed in Section 6.3. Finally, Sections 6.4 to 6.6 establish the robustness (cf. Section 4.2) of the PRNG constructions.

6.1 The H-Coefficient Technique

When considering an adaptive distinguisher \mathcal{D} that tries to tell apart two different worlds, usually termed *real* and *ideal* experiments, the H-coefficient technique [40] is a handy tool for analyzing the distinguishing advantage of \mathcal{D} .

The H-coefficient technique considers transcripts between distinguisher \mathcal{D} and the challenger. These transcripts are partitioned into two groups: the *good* transcripts and the *bad* transcripts. For good transcripts τ , an H-coefficient proof will commonly derive a lower bound on the ratio of the probability of τ occurring in the real world and that of τ occurring in the ideal world. For bad transcripts, which are normally defined as the transcripts for which said lower bound cannot be derived, one upper bounds the probability that they occur. This latter bound can be proved in the *ideal* world, which usually greatly simplifies the derivation.

Transcripts. The interaction of \mathcal{D} with either the real or the ideal experiment produces a transcript T that contains the queries made by \mathcal{D} and the corresponding answers given by the challenger. For a fixed transcript τ , denote by $\mathbf{p}_0(\tau)$ and $\mathbf{p}_1(\tau)$ the probabilities that the real and ideal experiments, respectively, the challenger produces the answers in τ if asked the queries in τ .¹⁵ Similarly,

¹⁵Observe that $\mathbf{p}_0(\tau)$ and $\mathbf{p}_1(\tau)$ depend only on the corresponding experiment and are independent of \mathcal{D} .

the behavior of a distinguisher \mathcal{D} is described by a function $\mathbf{p}_{\mathcal{D}}(\tau)$ that assigns to τ the probability that \mathcal{D} produces the queries in τ if given the answers in τ . Observe that, therefore, the probability of a particular transcript τ occurring in an interaction of \mathcal{D} with the real experiment is $\mathbf{P}[T_0 = \tau] = \mathbf{p}_{\mathcal{D}}(\tau) \cdot \mathbf{p}_0(\tau)$ and, similarly, $\mathbf{P}[T_1 = \tau] = \mathbf{p}_{\mathcal{D}}(\tau) \cdot \mathbf{p}_1(\tau)$ for the ideal world. In the following, denote by \mathcal{T} the set of all transcripts τ .

Bounding the distinguishing advantage. The distinguishing advantage of \mathcal{D} is upper bounded by the statistical distance

$$\begin{aligned}
\text{SD}(T_0, T_1) &= \sum_{\tau \in \mathcal{T}} \max\{0, \mathbf{P}[T_1 = \tau] - \mathbf{P}[T_0 = \tau]\} \\
&= \sum_{\tau \in \mathcal{T}} \max\{0, \mathbf{p}_{\mathcal{D}}(\tau) \cdot \mathbf{p}_1(\tau) - \mathbf{p}_{\mathcal{D}}(\tau) \cdot \mathbf{p}_0(\tau)\} \\
&= \sum_{\tau \in \mathcal{T}} \mathbf{p}_{\mathcal{D}}(\tau) \cdot \mathbf{p}_1(\tau) \left(1 - \frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)}\right) \\
&= \sum_{\tau \in \mathcal{T}} \mathbf{P}[T_1 = \tau] \left(1 - \frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)}\right). \tag{2}
\end{aligned}$$

Suppose that for some set $\Gamma \subseteq \mathcal{T}$ of *good* transcripts, a lower bound

$$\frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)} \geq 1 - \varepsilon$$

is known for all $\tau \in \Gamma$ and some $\varepsilon \geq 0$. Then, (2) becomes

$$\begin{aligned}
\sum_{\tau \in \mathcal{T}} \mathbf{P}[T_1 = \tau] \left(1 - \frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)}\right) &\leq \sum_{\tau \in \Gamma} \left(1 - \frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)}\right) + \sum_{\tau \in \mathcal{T} \setminus \Gamma} \mathbf{P}[T_1 = \tau] \\
&\leq \varepsilon + \mathbf{P}[T_1 \in \mathcal{T} \setminus \Gamma],
\end{aligned}$$

where transcripts $\tau \in \mathcal{T} \setminus \Gamma$ are commonly referred to as *bad* transcripts. Given the above, applying the H-coefficient technique entails defining a set of good transcripts, bounding the fraction above, and showing that bad transcripts are unlikely. Note that the latter can be done in the ideal experiment, which is usually considerably easier than doing so in the real experiment.

Theorem 11 (H-coefficient method). *For two experiments described by $\mathbf{p}_0(\cdot)$ and $\mathbf{p}_1(\cdot)$, respectively, if there exists a set $\Gamma \subseteq \mathcal{T}$ and $\varepsilon, \delta \geq 0$ satisfying*

1. (*ratio analysis*) $\mathbf{p}_0(\tau)/\mathbf{p}_1(\tau) \geq 1 - \varepsilon$ for all $\tau \in \Gamma$ and
2. (*bad event analysis*) $\mathbf{P}[T_1 \notin \Gamma] \leq \delta$,

then the distinguishing advantage of any distinguisher \mathcal{D} is bounded by $\varepsilon + \delta$.

6.2 Monolithic Extractor

From Section 3.2, recall the monolithic extractor mono defined to work with a random oracle $F : \{0, 1\}^m \rightarrow \{0, 1\}^n$:

Construction 1 (Monolithic extractor). The monolithic seedless extractor $\text{mono}^G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ using a random oracle $G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined by

$$\text{mono}^G(x) := G(x).$$

Construction 1 is a seedless extractor according to Definition 3.

Theorem 1 (Monolithic seedless extraction). *Construction mono is a $(\gamma^*, q, \varepsilon)$ -extractor in the G -model for*

$$\varepsilon \leq \frac{q}{2^{\gamma^*}} .$$

The proof of Theorem 1 is a straight-forward application of the H-coefficient technique. The idea is to first show that unless \mathcal{A}_1 or \mathcal{A}_2 queries the input x provided by \mathcal{A}_1 , the real and ideal worlds (i.e., the cases where $b = 0$ and $b = 1$, respectively) are indistinguishable. That is, the corresponding ratio of transcript probabilities is 1. Transcripts where x is in the query list are defined to be *bad* transcripts, and the second part of the proof shows that bad transcripts are unlikely to occur due to the legitimacy of \mathcal{A} . The latter proof crucially relies on the fact that the H-coefficient technique enables performing the bad-event analysis in the *ideal* world.

Proof. Consider a transcript of the interaction between an attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and the challenger of the extraction game (as defined in Section 3.1). It consists of

- the input x provided by \mathcal{A}_1 ,
- the value y^* output by the game (which is either the output y_0 of the extractor on x or a uniformly random value from $\{0, 1\}^n$), and
- the query/answer list $L = L_1 \cup L_2$ of \mathcal{A}_1 's and \mathcal{A}_2 's interaction with F .

That is $\tau = (y^*, x, L)$. A *bad* transcript occurs when \mathcal{A}_1 or \mathcal{A}_2 queries F at input x , i.e., when L contains a pair of the form $(x, *)$. In order to apply Theorem 11, one merely needs to bound the probability ratio for good transcripts (Lemma 12) and the probability of a bad transcript occurring in the ideal world, i.e., for $b = 1$ (Lemma 13). \square

Lemma 12 (Ratio analysis). *For all good transcripts τ ,*

$$\frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)} = 1 .$$

Proof. Fix a good transcript τ and consider first $\mathbf{p}_1(\tau)$. Since in the ideal world y^* is sampled uniformly,

$$\mathbf{p}_1(\tau) = p_L \cdot 2^{-n} ,$$

where p_L denotes the probability that a uniform random function is consistent with the queries in L . In the real world,

$$\mathbf{p}_0(\tau) = p_L \cdot q_\tau ,$$

where q_τ is the probability that $F_L(x) = y^*$ over a function F_L that is sampled uniformly at random conditioned on being consistent with L . Since τ is a good transcript, F_L is not constrained by L at coordinate x , and, hence, $q_\tau = 2^{-n}$. \square

Remark 1. In the above proof, note that p_L does not include the probability that x appears in the transcript. Recall from Section 6.1 that the behaviors $\mathbf{p}(\tau)$ are the probabilities that the experiment produces the answers in τ *when given* the queries in τ (and the value x is a query by the distinguisher).

Lemma 13 (Bad event analysis). *For the set \mathcal{B} of bad transcripts (as defined above),*

$$\mathbb{P}[T_1 \in \mathcal{B}] \leq \frac{q}{2^{\gamma^*}} .$$

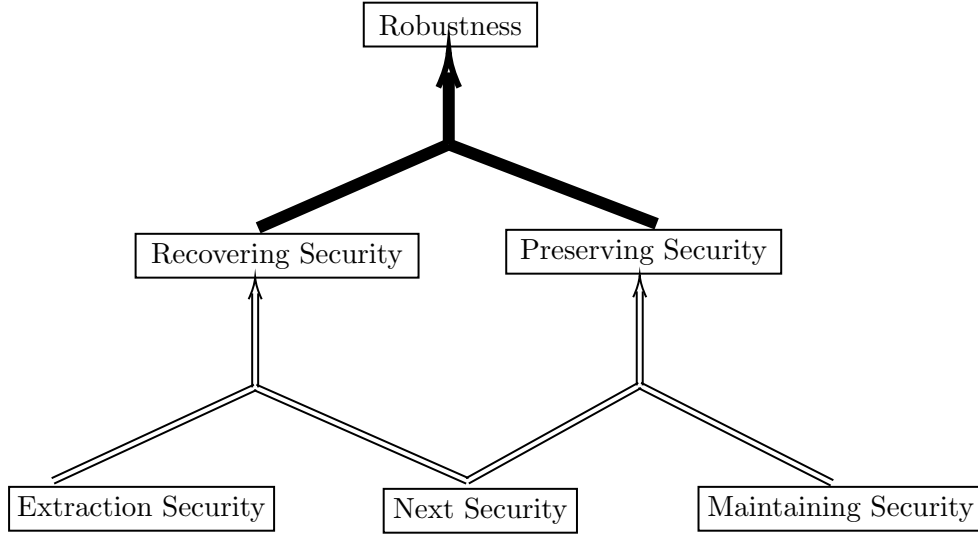


Figure 6: This figure represents the implication relations between the different intermediate notions of security. The filled arrows stand for a generic proof, while the unfilled arrows represent a construction-specific proof.

Proof. Recall that by the γ^* -legitimacy of $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$,

$$H_\infty(X|\Sigma\mathcal{L}_1) \geq \gamma^* .$$

Observe that in the ideal world, the output of the extraction game is a uniformly random value Y^* , which is independent of the input X produced by \mathcal{A}_1 . The sampling order of the ideal experiment can therefore be changed to be the following:

1. Sample F uniformly at random.
2. Run \mathcal{A}_1 until it outputs σ and x , thereby also generating the list of queries L_1 .
3. Choose y^* uniformly at random.
4. Run \mathcal{A}_2 on input (σ, y^*) , letting it make additional queries L_2 .
5. Resample the input X conditioned on $(\Sigma, \mathcal{L}_1) = (\sigma, L_1)$.

Note that since the conditioning includes \mathcal{L}_1 , \mathcal{A}_1 makes the same queries, L_1 , during the first run and the resampling process. Moreover, since conditioned on the values of (Σ, \mathcal{L}_1) , X and \mathcal{L}_2 are independent, the min-entropy condition holds for $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$, the list of all queries made by \mathcal{A} during the experiment, as well. That is,

$$H_\infty(X|\Sigma\mathcal{L}) \geq \gamma^* .$$

Thus, the probability that the resampled input X is contained in any query in the list L is at most $q \cdot 2^{-\gamma^*}$. \square

The PRNG Recovering Game

<pre> init $b \leftarrow \{0, 1\}$ </pre>	<pre> chall(s_0, x_1, \dots, x_ℓ) for $i = 1, \dots, \ell$ $s_i \leftarrow \text{refresh}(s_{i-1}, x_i)$ if $b = 0$ return $\text{next}(s_\ell)$ else $(s, y) \leftarrow \{0, 1\}^{n+r}$ return (s, y) </pre>
---	---

Figure 7: Oracles for PRNG recovering game.

6.3 Intermediate PRNG Security Notions

In keeping with tradition in PRNG literature, it is useful to define two simple properties called *recovering* and *preserving*. Recovering security requires that that if after an entropy drain, sufficient amount of entropy has been absorbed into the PRNG state, the output of the next function look random. Preserving security asks that after absorbing adversarially chosen inputs, a high-entropy state not become compromised, and the output of the next function after the absorption look random.

Recovering and preserving security can be shown to generically (i.e., for any PRNG construction) imply robustness. In order to establish these two properties themselves, it helps to introduce three further properties called *extraction*, *maintaining*, and *next security*:

- **Extraction security:** show that the PRNG *state* is indistinguishable from a uniform one after sufficient amounts of entropy have been absorbed;
- **Maintaining security:** show that the PRNG *state* is indistinguishable from a uniform one if it is random initially and arbitrary inputs are absorbed;
- **Next security:** show that the output of the next function is indistinguishable from a random value if it is called on a random input.

For each PRNG construction considered in this work, extraction and next security imply recovering security, and maintaining and next security imply preserving security. These proofs, however, are *not* generic and must be repeated for each PRNG construction. Figure 6 illustrates these implications.

6.3.1 Recovering and Preserving Security

As stated above, it is useful to define two simple properties called *recovering* and *preserving*, which together generically imply robustness via a hybrid argument.

Recovering security. The intuition behind Recovering security is that if after an entropy drain, sufficient amount of entropy (from the perspective of the attacker) has been absorbed into the state, then the output of *next* is indistinguishable from a uniformly random value in $\{0, 1\}^{n+r}$.

The corresponding game is depicted in Figure 7. It lets the attacker specify an initial state s_0 and a vector of inputs x_1, \dots, x_ℓ ; the inputs are then absorbed one-by-one, and *next* is called on the resulting state. The game returns the output of *next* if $b = 0$ and a uniform value if $b = 1$. The advantage of an attacker \mathcal{A} in the recovering game is denoted by $\text{Adv}_{\text{PRNG}}^{\text{rec}, P}(\mathcal{A})$.

The PRNG Preserving Game

init $s_0 \leftarrow \{0, 1\}^n$ $b \leftarrow \{0, 1\}$	chall (x_1, \dots, x_ℓ) for $i = 1, \dots, \ell$ $\quad s_i \leftarrow \text{refresh}(s_{i-1}, x_i)$ if $b = 0$ $\quad \text{return next}(s_\ell)$ else $\quad (s, y) \leftarrow \{0, 1\}^{n+r}$ $\quad \text{return } (s, y)$
---	--

Figure 8: Oracles for PRNG preserving game.

Similarly to the robustness game, an attacker has to satisfy a legitimacy condition. In particular, \mathcal{A} is γ^* -legitimate if

$$H_\infty(X_1, \dots, X_\ell | \Sigma \mathcal{L} S_0) \geq \gamma^*,$$

where

- Σ is the state of \mathcal{A} just before the call to **chall**,
- \mathcal{L} is the list of query and answers \mathcal{A} has made to P up to the call to **chall**, and
- S_0 is the initial state that the adversary provides.

For the recovering game, q again denotes the maximum number of P -queries that \mathcal{A} makes, and ℓ is the maximum number of blocks with which it calls oracle **chall**; a corresponding attacker is referred to as (q, ℓ) -attacker.

Definition 13. A PRNG construction $\text{PRNG} = (\text{refresh}, \text{next})$ is said to be $(\gamma^*, q, \ell, \varepsilon)$ -recovering in the P -model if for every γ^* -legitimate (q, ℓ) -attacker,

$$\text{Adv}_{\text{PRNG}}^{\text{rec}, P}(\mathcal{A}) \leq \varepsilon.$$

Preserving security. At a high level, preserving security requires that by absorbing adversarially chosen inputs, a high-entropy state cannot become compromised, and the output of **next** after the absorption is indistinguishable from a uniformly random value in $\{0, 1\}^{n+r}$.

The corresponding game is depicted in Figure 8. It lets the attacker specify a vector of inputs x_1, \dots, x_ℓ ; the inputs are then absorbed into a *randomly chosen* state one-by-one, and **next** is called on the resulting state. The game returns the output of **next** if $b = 0$ and a uniform value if $b = 1$. The advantage of an attacker \mathcal{A} in the preserving game is denoted by $\text{Adv}_{\text{PRNG}}^{\text{pre}, P}(\mathcal{A})$. There is no legitimacy constraint on \mathcal{A} ; the parameters q and ℓ are defined as before for a (q, ℓ) -attacker.

Definition 14. A PRNG construction $\text{PRNG} = (\text{refresh}, \text{next})$ is said to be (q, ℓ, ε) -preserving in the P -model if for every (q, ℓ) -attacker,

$$\text{Adv}_{\text{PRNG}}^{\text{pre}, P}(\mathcal{A}) \leq \varepsilon.$$

Recovering and preserving imply robustness. As mentioned above, in order to establish robustness of a PRNG construction, it suffices to prove that it is both *recovering* and *preserving*.

Theorem 14. Consider a PRNG construction $\text{PRNG} = (\text{refresh}, \text{next})$ for which **refresh** makes α P -calls and **next** makes β P -calls. Furthermore, assume PRNG is both

- $(\gamma^*, q, \ell, \varepsilon_{\text{rec}})$ -recovering and
- $(q, \ell, \varepsilon_{\text{pre}})$ -preserving

in the P -model. Then, PRNG is also $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -robust in the P -model, where

$$\varepsilon_{\text{rob}} \leq t \cdot (\varepsilon_{\text{rec}} + \varepsilon_{\text{pre}}) .$$

The proof of Theorem 14 is provided in Appendix C.1.

6.3.2 Extraction, Maintaining, and Next Security

Extraction security. Extraction security requires that after absorbing blocks with high joint entropy, the state of the PRNG be indistinguishable from a uniformly random one. The corresponding game is a variant of the game for recovering security (cf. Figure 7) in which next is not applied to s_ℓ ; instead, s_ℓ or a uniformly random value is output, depending on whether $b = 0$ or $b = 1$.

The legitimacy of an attacker \mathcal{A} as well as parameters q and ℓ are defined identically to the recovering game (cf. Section 4.2); the advantage of \mathcal{A} against extraction security of PRNG in the P -model is denoted by $\text{Adv}_{\text{PRNG}}^{\text{ext},P}(\mathcal{A})$.

Maintaining security. Maintaining security is a variant of the preserving game (cf. Figure 7) in which next is not applied to s_ℓ ; instead, s_ℓ or a uniformly random value is output, depending on whether $b = 0$ or $b = 1$. The parameters q and ℓ are defined identically to the preserving game (cf. Section 4.2); the advantage of \mathcal{A} against maintaining security of PRNG in the P -model is denoted by $\text{Adv}_{\text{PRNG}}^{\text{mtn},P}(\mathcal{A})$.

Next security. Next security requires that the output of the next function on a uniformly random state be indistinguishable from a uniformly random string. That is, an attacker \mathcal{A} , making at most q queries to the ideal primitive P , tries to distinguish $\text{next}^P(S)$ from U_{n+r} for a uniformly random $S \in \{0, 1\}^n$. Denote by $\text{Adv}_{\text{PRNG}}^{\text{next},P}(\mathcal{A})$ the advantage of \mathcal{A} .

6.4 PRNGs from Merkle-Damgård

This section establishes the robustness of the PRNG construction $\text{MD} = (\text{refresh}, \text{next})$ based on Merkle-Damgård with a random compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$. Recall that the construction is defined as follows (cf. Figure 3):

Construction 2 (PRNG from Merkle-Damgård). The (m, r) -PRNG construction $\text{MD} = (\text{refresh}, \text{next})$ based on Merkle-Damgård with a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined as follows:¹⁶

- $\text{refresh}^F(s, x) = F(s, x)$, and
- $\text{next}^F(s) = (F(s, 0), F(s, 1) \parallel \dots \parallel F(s, r/n))$.

Note that iteratively absorbing some input blocks x_1, \dots, x_ℓ via **refresh**, starting with a state s_0 is identical to applying the Merkle-Damgård construction to the input with initialization vector (IV) s_0 , which is denoted by $\text{MD}_{s_0}^F(x_1, \dots, x_\ell)$ in the remainder of this section.

¹⁶The integer arguments to the compression function are to be naturally mapped to $\{0, 1\}^n$.

Theorem 5 (Robustness of Merkle-Damgård PRNGs). *Construction 2 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -robust PRNG in the F -model for*

$$\varepsilon_{\text{rob}} \leq 2t \cdot \left(\frac{\tilde{q}^2 + \tilde{q}\ell + \ell^2}{2^n} + \frac{\tilde{q}}{2^{\gamma^*}} \right),$$

where $\tilde{q} = q + r/n + 1$.

As outlined in Section 6.3, the idea of the proof is to first establish extraction, maintaining, as well as next security, and to then show that extraction and next security imply recovering security and that maintaining and next security imply preserving security. More precisely, the proof proceeds as follows:

- **Extraction security:** In order to establish extraction security using the H-coefficient method (cf. Theorem 11 in Section 6.1), one first defines *bad* transcripts as transcripts for which the attacker has queried F on all coordinates needed to evaluate the PRNG on the inputs it provided.

For good transcripts, there is at least one missing F query to obtain the output of the PRNG. The probability ratio is lower bounded by arguing that the output of the PRNG looks random to the attacker unless there is a collision among the remaining F queries.

The probability of bad transcripts is upper bounded in the *ideal* world ($b = 1$). The proof uses the legitimacy of the attacker in a resampling argument to argue that it is unlikely that the attacker makes all the queries necessary to evaluate the PRNG on the inputs it provides. Special care has to be taken to handle collisions in the F queries.

- **Maintaining security:** The corresponding proof also employs the H-coefficient method, but is considerably more straightforward. A *bad* transcript is a transcript for which the initial (random) state of the PRNG appears in the F -queries. The ratio is bounded analogously to the extraction proof, and bounding the probability of a bad transcript is trivial.
- **Next security:** This is again a simple H-coefficient proof, which amounts to showing that unless the attacker queries F on an input $(s, *)$, where s is the (random) state to which next is applied, the output of next looks random.
- **Recovering security:** Showing that the Merkle-Damgård construction achieves recovering security is a simple hybrid argument: First, one uses extraction security to argue that the state after absorbing the inputs can be replaced by a random value. Second, using next security, one argues that the output of next on said random value looks random.
- **Preserving security:** Proving preserving security is also a simple hybrid argument: First, one uses maintaining security to argue that the state after absorbing the inputs can be replaced by a random value. Second, using next security, one argues that the output of next on said random value looks random.

The final bound in Theorem 5 follows by applying Theorem 14.

6.4.1 Extraction Security

Lemma 15 (Extraction security). *The advantage of any γ^* -legitimate (q, ℓ) -attacker \mathcal{A} against extraction security of MD is bounded by*

$$\text{Adv}_{\text{MD}}^{\text{ext}, F}(\mathcal{A}) \leq \varepsilon_{\text{MD}}^{\text{ext}}(\gamma^*, q, \ell) := \frac{q^2 + q\ell + \ell^2}{2^n} + \frac{2q}{2^{\gamma^*}}.$$

Proof. In order to prove Lemma 15 using the H-coefficient method, consider a transcript¹⁷

$$\tau = (s^*, s_0, x_1, \dots, x_\ell, L)$$

of the interaction between the \mathcal{A} and the extraction game, where s^* is the value returned by the game, L is the set of queries to the oracle made by the adversary \mathcal{A} , and where s_0 is the initial state and x_1, \dots, x_ℓ the inputs provided by \mathcal{A} . Let $\ell' \geq 0$ be maximal such that

$$((y_{i-1}, x_i), y_i) \in L$$

for some values $y_0, y_1, \dots, y_{\ell'}$ with $y_0 = s_0$. A transcript τ is called a *bad transcript* if $\ell' = \ell$; otherwise, τ is called *good*.

In order to apply Theorem 11, one merely needs to bound the probability ratio for good transcripts (Lemma 16) and the probability of a bad transcript occurring in the ideal world, i.e, for $b = 1$ (Lemma 17). \square

Lemma 16 (Ratio analysis). *For all good transcripts τ ,*

$$\frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)} \geq 1 - \frac{q\ell + \ell^2}{2^n}.$$

Proof. Fix a good transcript τ and consider first $\mathbf{p}_1(\tau)$. Since in the ideal world s^* is sampled uniformly,

$$\mathbf{p}_1(\tau) = p_L \cdot 2^{-n},$$

where p_L denotes the probability that a uniform random function is consistent with the queries in L . In the real world,

$$\mathbf{p}_0(\tau) = p_L \cdot q_\tau,$$

where q_τ is the probability that $\text{MD}_{s_0}^{F_L}(x_1, x_2, \dots, x_\ell) = s^*$ over a function F_L that is sampled uniformly at random conditioned on being consistent with L .

It remains to derive a lower bound on q_τ . To that end, observe that due to τ being a good transcript, $\ell' < \ell$. Hence, q_τ is the probability (over F_L) that

$$Y_\ell := \text{MD}_{y_{\ell'}}^{F_L}(x_{\ell'+1}, \dots, x_\ell) = s^*.$$

Consider the intermediate chaining values $Y_{\ell'+1}, \dots, Y_{\ell-1}$ and, for $i = \ell' + 1, \dots, \ell - 1$, define the event FRESH_i that Y_i is *fresh*, i.e., there is no query $((Y_i, *), *) \in L$ and $Y_i \neq Y_j$ for $j < i$. Let,

$$\text{FRESH} := \bigcap_{i=\ell'+1}^{\ell-1} \text{FRESH}_i.$$

Then,

$$\mathbf{P}[Y_\ell = s^* | \text{FRESH}] = 2^{-n}$$

since the conditioning implies that $F(Y_{\ell-1}, x_\ell)$ is freshly sampled and, hence, Y_ℓ is a uniformly random value. Moreover,

$$\mathbf{P} \left[\overline{\text{FRESH}}_i \mid \bigcap_{k=\ell'+1}^{i-1} \text{FRESH}_k \right] \leq \frac{q + \ell}{2^n}$$

¹⁷In order to keep notation simple, ℓ —here and in the following—is the number of inputs in a particular τ , not the upper bound from Lemma 15.

as there are at most $q + \ell$ non-fresh values by the time Y_i is sampled. Here, the notation \bar{A} denotes the complement of an event A . Therefore,

$$\mathbb{P}[\text{FRESH}] \geq 1 - \frac{q\ell + \ell^2}{2^n},$$

and, finally,

$$\begin{aligned} q_\tau &\geq \mathbb{P}[\text{FRESH}] \cdot \mathbb{P}[Y_\ell = s^* | \text{FRESH}] \\ &\geq \left(1 - \frac{q\ell + \ell^2}{2^n}\right) \cdot 2^{-n}, \end{aligned}$$

which implies

$$\frac{\mathfrak{p}_0(\tau)}{\mathfrak{p}_1(\tau)} \geq 1 - \frac{q\ell + \ell^2}{2^n}.$$

□

Lemma 17 (Bad event analysis). *For the set \mathcal{B} of bad transcripts (as defined above),*

$$\mathbb{P}[T_1 \in \mathcal{B}] \leq \frac{2q}{2^{\gamma^*}} + \frac{q^2}{2^n}.$$

Proof. Observe that in the ideal world, the output of the extraction game is a uniformly random value s^* , which is independent of the initial state S_0 and the inputs X_1, \dots, X_ℓ . The sampling order of the ideal experiment can therefore be changed to be the following:

1. Sample F uniformly at random.
2. Run \mathcal{A} until it outputs $(\tilde{s}_0, \tilde{x}_1, \dots, \tilde{x}_{\tilde{\ell}})$, thereby also generating the state σ and the list of queries L_0 before the challenge.
3. Choose s^* uniformly at random.
4. Continue running \mathcal{A} on σ and s^* , letting it make additional queries L_1 .
5. Resample the inputs (X_1, \dots, X_ℓ) conditioned on $(\Sigma, \mathcal{L}_0, S_0) = (\sigma, L_0, s_0)$.

Remember that \mathcal{L} is the random variable for the set of queries made while L is a particular value that the random variable takes. Observe that since the conditioning includes \mathcal{L}_0 , \mathcal{A} makes the same queries, L_0 , during the first run and the resampling process. Moreover, since conditioned on $(\Sigma, \mathcal{L}_0, S_0)$, (X_1, \dots, X_ℓ) and \mathcal{L}_1 are independent, the min-entropy condition holds for $\mathcal{L} = \mathcal{L}_0 \cup \mathcal{L}_1$, the list of all queries made by \mathcal{A} during the experiment, as well. That is,

$$\mathbb{H}_\infty(X_1, \dots, X_\ell | \Sigma \mathcal{L} S_0) \geq \gamma^*.$$

Let \mathcal{E} be the event that there exists a collision in \mathcal{L} , i.e., two or more queries have the same answer. Note that

$$\mathbb{P}[T_1 \in \mathcal{B}] \leq \mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}] + \mathbb{P}[\bar{\mathcal{E}}] \leq \mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}] + \frac{q^2}{2^n}.$$

Towards bounding $\mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}]$, consider now a particular triple $z = (\sigma, L, s_0) \in \mathcal{E}$, which is shorthand for L being collision-free. Define a *potential chain* as values y_0, y_1, \dots, y_ℓ for *some* ℓ such that $y_0 = s_0$ and

$$((y_{i-1}, v_i), y_i) \in L.$$

for $i = 1, \dots, \ell$ and some values v_1, \dots, v_ℓ . Observe that without collisions, L can contain at most q potential chains. Clearly, conditioned on $Z = z$, $T_1 \in \mathcal{B}$ if and only if for some potential chain, $X_i = v_i$ for all $i = 1, \dots, \ell$. Hence, by the legitimacy of \mathcal{A} ,

$$\mathbb{P}[T_1 \in \mathcal{B} | Z = z] \leq q \cdot p_z ,$$

where $p_z := \text{Pred}(\overline{X} | Z = z)$. In expectation,

$$\begin{aligned} \mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}] &= \sum_{z \in \mathcal{E}} \mathbb{P}[Z = z | \mathcal{E}] \cdot \mathbb{P}[T_1 \in \mathcal{B} | Z = z] \\ &\leq \sum_{z \in \mathcal{E}} \mathbb{P}[Z = z | \mathcal{E}] \cdot q \cdot p_z \\ &= q \cdot \text{Pred}(\overline{X} | Z \mathcal{E}) \\ &\leq q \cdot \frac{\text{Pred}(\overline{X} | Z)}{1 - q^2/2^n} \\ &\leq \frac{q}{(1 - q^2/2^n) \cdot 2^{\gamma^*}} \leq \frac{2q}{2^{\gamma^*}} , \end{aligned}$$

using that¹⁸ $q^2/2^n \leq 1/2$ and where the penultimate inequality is due to

$$\begin{aligned} \text{Pred}(X | Z) &\geq \sum_{z \in \mathcal{E}} \mathbb{P}[Z = z] \cdot p_z \\ &= \mathbb{P}[\mathcal{E}] \cdot \sum_{z \in \mathcal{E}} \frac{\mathbb{P}[Z = z]}{\mathbb{P}[\mathcal{E}]} \cdot p_z \\ &= \mathbb{P}[\mathcal{E}] \cdot \text{Pred}(\overline{X} | Z \mathcal{E}) . \end{aligned}$$

□

6.4.2 Maintaining Security

Lemma 18 (Maintaining security). *The advantage of any (q, ℓ) -attacker \mathcal{A} against maintaining security is bounded by*

$$\text{Adv}_{\text{MD}}^{\text{mtn}, F}(\mathcal{A}) \leq \varepsilon_{\text{MD}}^{\text{mtn}}(q, \ell) := \frac{q\ell + \ell^2}{2^n} + \frac{q}{2^n} .$$

Proof. To bound the advantage of an attacker \mathcal{A} at guessing b via an H-coefficient proof, consider a transcript

$$\tau = (s^*, s_0, x_1, \dots, x_\ell, L)$$

between \mathcal{A} and the maintaining game, where s^* is the output of the game and L are the queries made by \mathcal{A} . A transcript is *bad* if there is a query of the type $((s_0, *), *) \in L$ and *good* otherwise. The probability of a bad transcript occurring in the $b = 1$ case is clearly at most $|L|/2^n$. Moreover, for good transcripts τ ,

$$p_1(\tau) = 2^{-2n} \cdot p_L ,$$

where p_L denotes the probability that a uniform random function is consistent with the queries in L . Furthermore, by an argument similar to that in the proof of Lemma 16,

$$p_0(\tau) \geq \left(1 - \frac{q\ell + \ell^2}{2^n}\right) \cdot 2^{-2n} \cdot p_L .$$

The lemma follows by applying Theorem 11. □

¹⁸This can always be assumed as the bound would otherwise be vacuous.

6.4.3 Next Security

Lemma 19 (Next security). *The advantage of any q -attacker \mathcal{A} against next security is bounded by*

$$\text{Adv}_{\text{MD}}^{\text{next},F}(\mathcal{A}) \leq \varepsilon_{\text{MD}}^{\text{next}}(q) := \frac{q}{2^n}.$$

Proof. For a straight-forward H-coefficient proof, consider the transcript

$$\tau = (s, s^*, L),$$

where s is the initial state, s^* is the input given to \mathcal{A} , and L are the queries \mathcal{A} makes to F . A transcript is *bad* if L contains a query of the form $((s, *), *)$ and *good* otherwise. It is easily seen that for good transcripts, the probability ratio is 1, whereas, in the ideal world, where \mathcal{A} 's view is completely independent of S , the probability of a bad event is at most $q/2^n$. \square

6.4.4 Recovering Security

In the following, let $\varepsilon_{\text{MD}}^{\text{ext}}(\gamma^*, q, \ell)$ and $\varepsilon_{\text{MD}}^{\text{next}}(q)$ be as in Lemmas 15 and 19. Extraction and next security together imply recovering security:

Lemma 20 (Recovering Security). *For every γ^* -legitimate (q, ℓ) -attacker \mathcal{A} ,*

$$\text{Adv}_{\text{MD}}^{\text{rec},F}(\mathcal{A}) \leq \varepsilon_{\text{MD}}^{\text{ext}}(\gamma^*, q + r/n + 1, \ell) + \varepsilon_{\text{MD}}^{\text{next}}(q + r/n + 1).$$

Proof. For $b \in \{0, 1\}$, denote by H_b the recovering experiment conditioned on the secret bit having the value b . Moreover, define a hybrid experiment $H_{\frac{1}{2}}$ in which the challenge oracle returns $\text{next}^F(U_n)$ to \mathcal{A} . By the triangle inequality, to prove the lemma, it suffices to bound the distance between experiments H_0 and $H_{\frac{1}{2}}$ and $H_{\frac{1}{2}}$ and H_1 .

- Towards bounding the distance between H_0 and $H_{\frac{1}{2}}$, consider the following attacker \mathcal{A}_{ext} against extraction security of MD: \mathcal{A}_{ext} runs \mathcal{A} answering its oracle queries by passing queries to F . At some point, \mathcal{A} outputs $(s_0, x_1, \dots, x_\ell)$. \mathcal{A}_{ext} forwards $(s_0, x_1, x_2, \dots, x_\ell)$ to the challenger. \mathcal{A}_{ext} receives s^* as response from the challenger. \mathcal{A}_{ext} now computes next on the input s^* , by making $r/n + 1$ additional queries to the primitive, on the inputs (s^*, i) for $i = 0, 1, \dots, r/n$. It forwards to \mathcal{A} the values $(F(s^*, 0), F(s^*, 1) \parallel \dots \parallel F(s^*, r/n))$. It proceeds to respond to \mathcal{A} 's oracle calls as before and waits for \mathcal{A} 's guess bit. It merely forwards the same bit as its guess to the challenger.

When the challenger's bit is 0, \mathcal{A}_{ext} perfectly simulates towards \mathcal{A} the distribution H_0 . When the challenger's bit is 1, \mathcal{A} is given $\text{next}(U_n)$. This corresponds to the hybrid distribution $H_{\frac{1}{2}}$. Moreover, it is easily seen that if \mathcal{A} is γ^* -legitimate, so is \mathcal{A}_{ext} . Thus, the advantage of \mathcal{A} in distinguishing hybrids H_0 and $H_{\frac{1}{2}}$ is upper-bounded by the advantage of \mathcal{A}_{ext} in the extraction game which is $\varepsilon_{\text{MD}}^{\text{ext}}(\gamma^*, q + r/n + 1, \ell)$.

- Towards bounding the distance between $H_{\frac{1}{2}}$ and H_1 , consider the following attacker $\mathcal{A}_{\text{next}}$ against next security of MD: $\mathcal{A}_{\text{next}}$ runs \mathcal{A} answering its oracle queries by passing queries to F . When \mathcal{A} outputs $(s_0, x_1, \dots, x_\ell)$, $\mathcal{A}_{\text{next}}$ returns its own distinguishing challenge, continuing to answer oracle queries for \mathcal{A} . In the end, it outputs \mathcal{A} 's guess bit.

It is straight-forward to verify that $\mathcal{A}_{\text{next}}$ perfectly simulates $H_{\frac{1}{2}}$ to \mathcal{A} when given $\text{next}^F(U_n)$ and H_1 when given U_{n+r} . Thus, the advantage of \mathcal{A} in distinguishing hybrids $H_{\frac{1}{2}}$ and H_1 is upper-bounded by the advantage of $\mathcal{A}_{\text{next}}$ in the next game, which is $\varepsilon_{\text{MD}}^{\text{next}}(q + r/n + 1)$. \square

6.4.5 Preserving Security

In the following, let $\varepsilon_{\text{MD}}^{\text{mtn}}(q, \ell)$ and $\varepsilon_{\text{MD}}^{\text{next}}(q)$ be as in Lemmas 18 and 19. Maintaining and next security together imply preserving security:

Lemma 21 (Preserving Security). *For every (q, ℓ) -attacker \mathcal{A} ,*

$$\text{Adv}_{\text{MD}}^{\text{pre}, F}(\mathcal{A}) \leq \varepsilon_{\text{MD}}^{\text{mtn}}(q + r/n + 1, \ell) + \varepsilon_{\text{MD}}^{\text{next}}(q + r/n + 1).$$

Proof. For $b \in \{0, 1\}$, denote by H_b the preserving experiment conditioned on the secret bit having the value b . Moreover, define a hybrid experiment $H_{\frac{1}{2}}$ in which the challenge oracle returns $\text{next}^F(U_n)$ to \mathcal{A} . By the triangle inequality, to prove the lemma, it suffices to bound the distance between experiments H_0 and $H_{\frac{1}{2}}$ and $H_{\frac{1}{2}}$ and H_1 .

- Towards bounding the distance between H_0 and $H_{\frac{1}{2}}$, consider the following attacker \mathcal{A}_{mtn} against maintaining security of MD: \mathcal{A}_{mtn} runs \mathcal{A} answering its oracle queries by passing queries to F . At some point, \mathcal{A} outputs (x_1, \dots, x_ℓ) . \mathcal{A}_{mtn} forwards $(x_1, x_2, \dots, x_\ell)$ to the challenger. \mathcal{A}_{mtn} receives s^* as response from the challenger. \mathcal{A}_{mtn} now computes next on the input s^* , by making $r/n + 1$ additional queries to the primitive, on the inputs (s^*, i) for $i = 0, 1, \dots, r/n$. It forwards to \mathcal{A} the values $(F(s^*, 0), F(s^*, 1) \parallel \dots \parallel F(s^*, r/n))$. It proceeds to respond to \mathcal{A} 's oracle calls as before and waits for \mathcal{A} 's guess bit. It merely forwards the same bit as its guess to the challenger.

When the challenger's bit is 0, the \mathcal{A}_{mtn} perfectly simulates towards \mathcal{A} the distribution H_0 . When the challenger's bit is 1, \mathcal{A} is given $\text{next}(U_n)$. This corresponds to the hybrid distribution $H_{\frac{1}{2}}$. Thus, the advantage of \mathcal{A} in distinguishing hybrids H_0 and $H_{\frac{1}{2}}$ is upper-bounded by the advantage of \mathcal{A}_{mtn} in the extraction game which is $\varepsilon_{\text{MD}}^{\text{mtn}}(q + r/n + 1, \ell)$.

- Towards bounding the distance between $H_{\frac{1}{2}}$ and H_1 , consider the following attacker $\mathcal{A}_{\text{next}}$ against next security of MD: $\mathcal{A}_{\text{next}}$ runs \mathcal{A} answering its oracle queries by passing queries to F . When \mathcal{A} outputs (x_1, \dots, x_ℓ) , $\mathcal{A}_{\text{next}}$ returns its own distinguishing challenge, continuing to answer oracle queries for \mathcal{A} . In the end, it outputs \mathcal{A} 's guess bit.

It is straight-forward to verify that $\mathcal{A}_{\text{next}}$ perfectly simulates $H_{\frac{1}{2}}$ to \mathcal{A} when given $\text{next}^F(U_n)$ and H_1 when given U_{n+r} . Thus, the advantage of \mathcal{A} in distinguishing hybrids $H_0, H_{\frac{1}{2}}$ is upper-bounded by the advantage of $\mathcal{A}_{\text{next}}$ in the next game which is $\varepsilon_{\text{MD}}^{\text{next}}(q + r/n + 1)$.

□

6.5 PRNGs from Merkle-Damgård with Davies-Meyer

This section establishes the robustness of the PRNG construction $\text{DM} = (\text{refresh}, \text{next})$ based on Merkle-Damgård with the Davies-Meyer compression function, which is analyzed using an ideal cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. Recall that the construction is defined as follows (cf. Figure 4):

Construction 4 (PRNG from MD-DM). The (n, r) -PRNG construction $\text{DM} = (\text{refresh}, \text{next})$ based on Merkle-Damgård with Davies-Meyer (MD-DM) uses a cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and is defined as follows:¹⁹

¹⁹The integer arguments to the cipher are to be naturally mapped to $\{0, 1\}^n$.

- $\text{refresh}^E(s, x) = E(x, s) \oplus s$, and
- $\text{next}^E(s) = (E(0, s) \oplus s, E(1, s) \oplus s \parallel \dots \parallel E(r/n, s) \oplus s)$.

Note that iteratively absorbing some input blocks x_1, \dots, x_ℓ via **refresh**, starting with a state s_0 is identical to applying the MD-DM construction to the input with initialization vector (IV) s_0 , which is denoted by $\text{MD-DM}_{s_0}^E(x_1, \dots, x_\ell)$ in the remainder of this paper.

Theorem 7 (Robustness of MD-DM PRNGs). *Construction 4 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -robust PRNG in the E -model for*

$$\varepsilon_{\text{rob}} \leq 4t \cdot \left(\frac{\tilde{q}^2 + \tilde{q}\ell + \ell^2}{2^n} + \frac{\tilde{q}}{2^{\gamma^*}} \right),$$

where $\tilde{q} = q + r/n + 1$.

The robustness of the MD-DM PRNG construction is proved along similar lines as that of the MD construction with a random compression function (cf. Section 6.5). It is highly recommended to read that proof first. The proof again establishes extraction, maintaining, and next security, before showing that these imply recovering and preserving security. The final bound in Theorem 5 follows by applying Theorem 14.

6.5.1 Extraction Security

Lemma 22 (Extraction security). *The advantage of any γ^* -legitimate (q, ℓ) -attacker \mathcal{A} against extraction security of DM is bounded by*

$$\text{Adv}_{\text{DM}}^{\text{ext}, E}(\mathcal{A}) \leq \varepsilon_{\text{MD}}^{\text{ext}}(\gamma^*, q, \ell) := \frac{q^2 + 2(q\ell + \ell^2)}{2^n} + \frac{4q}{2^{\gamma^*}}.$$

Proof. The transcript has the identical format as previously, i.e.,

$$\tau = (s^*, s_0, x_1, \dots, x_\ell, L)$$

except that L is now the set of E -queries. To define bad transcripts, let $\ell' \geq 0$ be maximal such that

$$((x_i, y_{i-1}), y_i \oplus y_{i-1}) \in L$$

for some values $y_0, y_1, \dots, y_{\ell'}$ with $y_0 = s_0$. A transcript τ is a bad transcript if

- $\ell' = \ell$ or
- $\ell' = \ell - 1$ and there exists a query $((x_\ell, *), y_{\ell-1} \oplus s^*) \in L$.

As before, to apply Theorem 11, one merely needs to bound the probability ratio for good transcripts (Lemma 23) and the probability of a bad transcript occurring in the ideal world, i.e, for $b = 1$ (Lemma 24). \square

Lemma 23 (Ratio analysis). *For all good transcripts τ ,*

$$\frac{\mathbf{p}_0(\tau)}{\mathbf{p}_1(\tau)} \geq 1 - \frac{2(q\ell + \ell^2)}{2^n}.$$

Proof. As in Section 5.1, for a good transcript,

$$p_1(\tau) = p_L \cdot 2^{-n} \quad \text{and} \quad p_0(\tau) = p_L \cdot q_\tau ,$$

where p_L denotes the probability that a uniform random cipher is consistent with the queries in L and where q_τ is the probability that

$$\text{MD-DM}_{s_0}^{E_L}(x_1, x_2, \dots, x_\ell) = s^*$$

over a cipher E_L that is sampled uniformly at random conditioned on being consistent with L .

To derive a lower bound on q_τ , due to τ being a good transcript, it suffices to consider the two cases

- (1) $\ell' \leq \ell - 2$ and
- (2) $\ell' = \ell - 1$.

For (1), q_τ is the probability (over E_L) that

$$Y_\ell := \text{DM}_{y_{\ell'}}^{E_L}(x_{\ell'+1}, \dots, x_\ell) = s^* .$$

Consider the intermediate chaining values $Y_{\ell'} = y_{\ell'}, Y_{\ell'+1}, \dots, Y_{\ell-1}$ (defined via evaluations of E). Moreover, let $L_{\ell'}, L_{\ell'+1}, \dots, L_{\ell-2}$ be the set of points at which E is defined after evaluating these intermediate values, i.e., $L_{\ell'} := L$ and

$$L_i := L_{i-1} \cup \{(x_i, Y_{i-1}), Y_i \oplus Y_{i-1}\}$$

for $i = \ell' + 1, \dots, \ell - 2$. Furthermore, for a set \tilde{L} of E -queries define the sets $\text{Free-In}_{\tilde{L}}$ of *free inputs* and $\text{Free-Out}_{\tilde{L}}$ of *free outputs*, i.e.,

$$y \in \text{Free-In}_{\tilde{L}} : \iff ((*, y), *) \notin \tilde{L} \quad \text{and} \quad y \in \text{Free-Out}_{\tilde{L}} : \iff ((*, *), y) \notin \tilde{L} .$$

Finally, for $i = \ell' + 1, \dots, \ell - 2$, define the event FRESH_i that Y_i is a *fresh input*, i.e.,

$$Y_i \in \text{Free-In}_{L_{i-1}} .$$

However, for $Y_{\ell-1}$, let $\text{FRESH}_{\ell-1}$ be the event *not only* that $Y_{\ell-1} \in \text{Free-In}_{L_{\ell-2}}$ but also that $Y_{\ell-1} \oplus s^*$ is a *free output*, i.e.,

$$Y_{\ell-1} \oplus s^* \in \text{Free-Out}_{L_{\ell-2}} .$$

Let,

$$\text{FRESH} := \bigcap_{i=\ell'+1}^{\ell-1} \text{FRESH}_i .$$

Then, on the one hand,

$$\begin{aligned} \text{P}[Y_\ell = s^* | \text{FRESH}] &= \text{P}[E(x_\ell, Y_{\ell-1}) = Y_{\ell-1} \oplus s^* | \text{FRESH}] \\ &= \frac{1}{|\text{Free-Out}_{L_{\ell-2}}|} \geq \frac{1}{2^n} \end{aligned}$$

since the conditioning implies that $(x_\ell, Y_{\ell-1})$ is a fresh input to E and that $Y_{\ell-1} \oplus s^*$ actually is in $\text{Free-Out}_{L_{\ell-2}}$. On the other hand, in order to bound

$$\text{P} \left[\overline{\text{FRESH}_{\ell-1}} \mid \bigcap_{k=\ell'+1}^{\ell-2} \text{FRESH}_k \right] ,$$

observe that, for any values $y_{\ell'+1}, \dots, y_{\ell-2}$ consistent with the conditioning, $\text{FRESH}_{\ell-1}$ is violated if²⁰

$$E_L(x_{\ell-1}, y_{\ell-2}) \in \overline{\text{Free-In}}_{L_{\ell-2}} \oplus y_{\ell-2}$$

or

$$E_L(x_{\ell-1}, y_{\ell-2}) \in \overline{\text{Free-Out}}_{L_{\ell-2}} \oplus y_{\ell-2} \oplus s^* .$$

The probability that the former condition is violated is at most

$$\begin{aligned} \frac{|\overline{\text{Free-Out}}_{L_{\ell-2}} \cap \overline{\text{Free-In}}_{L_{\ell-2}} \oplus y_{\ell-2}|}{|\overline{\text{Free-Out}}_{L_{\ell-2}}|} &\leq \frac{|\overline{\text{Free-In}}_{L_{\ell-2}} \oplus y_{\ell-2}|}{|\overline{\text{Free-Out}}_{L_{\ell-2}}|} \\ &\leq \frac{q + \ell}{2^n - (q + \ell)} . \end{aligned}$$

The same bound via a similar argument is obtained for the latter condition as well as

$$\mathbb{P} \left[\overline{\text{FRESH}}_i \mid \bigcap_{k=\ell'+1}^{i-1} \text{FRESH}_k \right] \leq \frac{q + \ell}{2^n - (q + \ell)}$$

for $i = \ell' + 1, \dots, \ell - 2$. Therefore,

$$\mathbb{P}[\text{FRESH}] \geq 1 - \frac{q\ell + \ell^2}{2^n - (q + \ell)} ,$$

and, finally,

$$\begin{aligned} q_\tau &\geq \mathbb{P}[\text{FRESH}] \cdot \mathbb{P}[Y_\ell = s^* | \text{FRESH}] \\ &\geq \left(1 - \frac{q\ell + \ell^2}{2^n - (q + \ell)} \right) \cdot 2^{-n} , \end{aligned}$$

for case (1). For case (2), note that due to $\ell' = \ell - 1$, $(x_\ell, y_{\ell-1})$ is a fresh input to E , and, furthermore, $y_{\ell-1} \oplus s^* \in \text{Free-Out}_L$. Thus, in this case the probability that $Y_\ell = s^*$ is at least 2^{-n} .

Combining both cases, case (1) dominating, one obtains

$$\frac{p_0(\tau)}{p_1(\tau)} \geq 1 - \frac{q\ell + \ell^2}{2^n - (q + \ell)} \geq 1 - \frac{2(q\ell + \ell^2)}{2^n} ,$$

using that $q + \ell \leq 2^{n-1}$, an assumption one may always make since the bound in the lemma is vacuous otherwise. \square

Lemma 24 (Bad event analysis). *For the set \mathcal{B} of bad transcripts (as defined above),*

$$\mathbb{P}[T_1 \in \mathcal{B}] \leq \frac{4q}{2^{\gamma^*}} + \frac{q^2}{2^n} .$$

Proof. The same resampling approach as in the proof of Lemma 17 applies here as well. However, the collisions require additional care. Two queries E queries $((k, u), v)$ and $((k', u'), v')$ are said to *collide* if

$$v \oplus u = v' \oplus u' .$$

²⁰Using the notation $A \oplus b = \{x \oplus b \mid x \in A\}$.

It is easily verified that the probability, over E , that any two such queries collide is at most $(2^n - 1)^{-1}$. Let \mathcal{E} be the event that there exists such a collision in \mathcal{L} . Note that

$$\mathbb{P}[T_1 \in \mathcal{B}] \leq \mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}] + \mathbb{P}[\overline{\mathcal{E}}] \leq \mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}] + \frac{q^2}{2^n}.$$

Towards bounding $\mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}]$, consider a triple $z = (\sigma, L, s_0) \in \mathcal{E}$, which, once more, is shorthand for L being collision-free. Define a *potential chain* as ℓ , for *some* ℓ , values y_0, y_1, \dots, y_ℓ such that $y_0 = s_0$ and, for some values k_1, \dots, k_ℓ ,

- (a) $((k_i, y_{i-1}), y_i \oplus y_{i-1}) \in L$ for $i = 1, \dots, \ell$ or
- (b) $((k_i, y_{i-1}), y_i \oplus y_{i-1}) \in L$ for $i = 1, \dots, \ell - 1$ and $((k_\ell, y_\ell), y_{\ell-1} \oplus s^*) \in L$.²¹

Without collisions, L can contain at most $2q$ potential chains; this can be proved by induction: Consider a set collision-free set L' of E -queries and assume that the number of potential chains is at most $2|L'|$. Consider an additional query $((k, u), v)$ that does not cause a collision; it may only (but need not) create a new potential chain in two ways:

- $u = v' \oplus u'$: this corresponds to extending in the sense of (a) above and can only be true for a single previous query $((k', u'), v')$ if L is collision-free;
- $v = v' \oplus u' \oplus s^*$: this corresponds to creating a chain of type (b) and, again, can only hold for one query $((k', u'), v')$ if L is collision-free.

Summarizing, the new query creates at most two new potential chains.

Clearly, conditioned on $Z = z$, $T_1 \in \mathcal{B}$ if and only if for some potential chain, $X_i = k_i$ for all $i = 1, \dots, \ell$. Hence, by the legitimacy of \mathcal{A} ,

$$\mathbb{P}[T_1 \in \mathcal{B} | Z = z] \leq 2q \cdot p_z,$$

where $p_z := \text{Pred}(\overline{X} | Z = z)$. The remainder of the proof proceeds in the exact same fashion as the proof of Lemma 17. \square

6.5.2 Maintaining Security

The maintaining security of the Davies-Meyer PRNG construction DM is proved along similar lines as that of the MD construction with a random compression function. This section discusses the few differences.

Lemma 25 (Maintaining security). *The advantage of any (q, ℓ) -attacker \mathcal{A} against maintaining security is bounded by*

$$\text{Adv}_{\text{DM}}^{\text{mtn}, E}(\mathcal{A}) \leq \varepsilon_{\text{DM}}^{\text{mtn}}(q, \ell) := \frac{2(q\ell + \ell^2)}{2^n} + \frac{q}{2^n}.$$

Proof. Similarly to Lemma 18, maintaining security is shown via an H-coefficient proof. Once more, one considers transcripts

$$\tau = (s^*, s_0, x_1, \dots, x_\ell, L),$$

with the difference that L refers to E -queries here. A *bad* transcript contains a query of the type $((*, s_0), *) \in L$. Again, the probability of a bad transcript occurring in the $b = 1$ case is at most $|L|/2^n$, and for good transcripts,

$$p_1(\tau) = 2^{-2n} \cdot p_L \quad \text{and} \quad p_0(\tau) \geq \left(1 - \frac{2(q\ell + \ell^2)}{2^n}\right) \cdot 2^{-2n} \cdot p_L,$$

where the latter follows via an argument similar to that in the proof of Lemma 23. \square

²¹Observe that in case (b) the “intuitive chain” goes up to $y_{\ell-1}$ but y_ℓ is not part of it.

6.5.3 Next Security

Next security of DM is defined and proved in a fashion analogous to the case with a random compression function.

Lemma 26 (Next security). *The advantage of any q -attacker \mathcal{A} against next security is bounded by*

$$\text{Adv}_{\text{DM}}^{\text{next},E}(\mathcal{A}) \leq \varepsilon_{\text{DM}}^{\text{next}}(q) := \frac{q}{2^n}.$$

Proof. For a straight-forward H-coefficient proof, consider the transcript

$$\tau = (s, s_0^*, s_1^*, \dots, s_{r/n}^*, L),$$

where s is the initial state, the values s_i^* are the input to \mathcal{A} , and L are the queries \mathcal{A} makes to F . A transcript is *bad* if L contains a query of the form $((i, s), *)$ or $((i, *), s \oplus s_i^*)$ for some i ; otherwise, τ is called *good*. It is easily seen that for good transcripts, the probability ratio is *at least* 1, and, in the ideal world, where \mathcal{A} 's view is completely independent of S , the probability of a bad event is at most $q/2^n$. \square

6.5.4 Recovering Security

In the following, let $\varepsilon_{\text{DM}}^{\text{ext}}(\gamma^*, q, \ell)$ and $\varepsilon_{\text{DM}}^{\text{next}}(q)$ be as in Lemmas 22 and 26. Once more, extraction and next security together imply recovering security:

Lemma 27 (Recovering Security). *For every γ^* -legitimate (q, ℓ) -attacker \mathcal{A} ,*

$$\text{Adv}_{\text{DM}}^{\text{rec},E}(\mathcal{A}) \leq \varepsilon_{\text{DM}}^{\text{ext}}(\gamma^*, \ell, q + r/n + 1) + \varepsilon_{\text{DM}}^{\text{next}}(q + r/n + 1).$$

The proof of the lemma is completely analogous to that of Lemma 20 and is omitted.

6.5.5 Preserving Security

In the following, let $\varepsilon_{\text{MD}}^{\text{mtn}}(q, \ell)$ and $\varepsilon_{\text{MD}}^{\text{next}}(q)$ be as in Lemmas 18 and 19. Maintaining and next security together imply preserving security:

Lemma 28 (Preserving Security). *For every adversary \mathcal{A} ,*

$$\text{Adv}_{\text{DM}}^{\text{pre},E}(\mathcal{A}) \leq \varepsilon_{\text{MD}}^{\text{mtn}}(q + r/n + 1, \ell) + \varepsilon_{\text{MD}}^{\text{next}}(q + r/n + 1).$$

The proof of the lemma is completely analogous to that of Lemma 21 and is omitted.

6.6 PRNGs from Sponges

This section establishes the robustness of the PRNG construction $\text{Spg} = (\text{refresh}, \text{next})$ based on the Sponge paradigm with a random permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$. Recall that the construction is defined as follows (cf. Figure 5):

Construction 6 (PRNG from Sponges). The Sponge-based PRNG construction $\text{Spg} = (\text{refresh}, \text{next})$ uses a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to absorb and produce r -bit inputs and outputs, respectively, and is defined as follows:

- $\text{refresh}^\pi(s, x) = \pi(s \oplus x \| 0^c)$, and

- $\text{next}^\pi(s) = (\pi(s) \oplus 0^r \| s^{(c)}, s^{(r)})$.

Note that iteratively absorbing some input blocks x_1, \dots, x_ℓ via **refresh**, starting with a state s_0 is identical to applying the Sponge construction to the input with initialization vector (IV) s_0 , which is denoted by $\text{Sponge}_{s_0}^\pi(x_1, \dots, x_\ell)$ in the remainder of this section.

Theorem 9 (Robustness of Sponge PRNGs). *Construction 6 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -robust PRNG in the π -model for*

$$\varepsilon_{\text{rob}} \leq 4t \cdot \left(\frac{\tilde{q}^2 + \tilde{q}\ell + \ell^2}{2^n} + \frac{\tilde{q}}{2^{\gamma^*}} + \frac{\tilde{q}^2}{2^c} \right),$$

where $\tilde{q} = q + r/n + 1$.

The proof of the robustness of the Sponge PRNG construction follows the same outline as that of the MD construction (cf. Section 6.4). It is highly recommended to read that proof first. One crucial difference between the Merkle-Damgård constructions and Sponges is that Sponges do *not* satisfy extraction/maintaining security (with good parameters). For example, given the state of a Sponge PRNG after absorbing a single (possibly high-entropy) input, a simple inverse query to π results in a value of the form $a \| 0^c$, which is unlikely to happen in the ideal world ($b = 1$). This is handled by explicitly introducing a “hit” probability, i.e., the probability that the attacker queries π^{-1} on the final state of the Sponge. Recovering and preserving security are then established by arguing that the hit probability is low when **next** is applied to the state. The final bound in Theorem 9 follows by applying Theorem 14.

6.6.1 Extraction Security

The extraction security of the Sponge PRNG construction is defined and proved along similar lines as that of the previous constructions. This section discusses the differences. Denote by $\text{Adv}_{\text{SpG}}^{\text{inv}, \pi}(\mathcal{A})$ the probability that \mathcal{A} queries π^{-1} on the value s^* returned by the challenger.

Lemma 29 (Extraction security). *The advantage of any γ^* -legitimate (q, ℓ) -attacker \mathcal{A} against extraction security of DM is bounded by*

$$\text{Adv}_{\text{SpG}}^{\text{ext}, \pi}(\mathcal{A}) \leq \varepsilon_{\text{SpG}}^{\text{ext}}(\gamma^*, \ell, q) := \frac{q + 2(q\ell + \ell^2)}{2^n} + \frac{2q}{2^{\gamma^*}} + \frac{q^2}{2^c} + \text{Adv}_{\text{SpG}}^{\text{inv}, \pi}(\mathcal{A}).$$

In the following, for convenience, let

$$\delta_{\text{SpG}}^{\text{ext}}(\gamma^*, \ell, q) := \frac{q + 2(q\ell + \ell^2)}{2^n} + \frac{2q}{2^{\gamma^*}} + \frac{q^2}{2^c}.$$

Proof. The transcript has the identical format as previously, i.e.,

$$\tau = (s^*, s_0, x_1, \dots, x_\ell, L)$$

except that L is now the set of π -queries. To define bad transcripts let $\ell' \geq 0$ be maximal such that there exist $(u_i, v_i) \in L$ with $u_1 = s_0 \oplus x_1 \| 0^c$ and

$$u_i = v_{i-1} \oplus x_i \| 0^c$$

for $i = 2, \dots, \ell - 1$. A transcript τ is a bad transcript if

- **(hit)** $(*, s^*) \in L$ or

- **(chain)** $\ell' = \ell$.

As before, to apply Theorem 11, one merely needs to bound the probability ratio for good transcripts (Lemma 30) and the probability of a bad transcript occurring in the ideal world, i.e, for $b = 1$ (Lemma 31). \square

Lemma 30 (Ratio analysis). *For all good transcripts τ ,*

$$\frac{p_0(\tau)}{p_1(\tau)} \geq \left(1 - \frac{2(q\ell + \ell^2)}{2^n}\right).$$

Proof. As in Section 5.1, for a good transcript,

$$p_1(\tau) = p_L \cdot 2^{-n} \quad \text{and} \quad p_0(\tau) = p_L \cdot q_\tau,$$

where p_L denotes the probability that a uniform random permutation is consistent with the queries in L and where q_τ is the probability that

$$\text{Sponge}_{s_0}^{\pi_L}(x_1, x_2, \dots, x_\ell) = s^*$$

over a permutation π_L that is sampled uniformly at random conditioned on being consistent with L .

To derive a lower bound on q_τ , note that due to τ being a good transcript $\ell' < \ell$. Hence, q_τ is the probability (over π_L) that

$$V_\ell := \text{Sponge}_{v_{\ell'}}^{\pi_L}(x_{\ell'+1}, \dots, x_\ell) = s^*.$$

Consider the intermediate values $U_{\ell'+1}, V_{\ell'+1}, U_{\ell'+2}, V_{\ell'+2}, \dots, V_{\ell-1}, U_\ell$, where

$$U_{\ell'+1} = v_{\ell'} \oplus x_{\ell'+1} \parallel 0^c$$

and

$$U_i = V_{i-1} \oplus x_i \parallel 0^c$$

for $i = \ell' + 2, \dots, \ell$ as well as

$$V_i = \pi(U_i)$$

for $i = \ell' + 1, \dots, \ell - 1$. Define, for $i = \ell' + 1, \dots, \ell$, the event FRESH_i that

- U_i is *fresh*, i.e., there is no query of type $(U_i, *) \in L$ and $U_i \neq U_j$ for $j < i$, and
- V_{i-1} is not a *hit*, i.e., $V_i \neq s^*$;

observe that $\text{FRESH}_{\ell'+1}$ is always true due to the maximality of ℓ' and the fact that τ is a good transcript. Let,

$$\text{FRESH} := \bigcap_{i=\ell'+1}^{\ell} \text{FRESH}_i.$$

Then,

$$\mathbb{P}[V_\ell = s^* | \text{FRESH}] \geq 2^{-n}$$

since the conditioning implies that U_ℓ is a fresh input and therefore V_ℓ is chosen uniformly from a set of size at most 2^n , which contains s^* due to τ being a good transcript and no hits occurring while evaluating the intermediate values.

Moreover, observe that if U_{i-1} is a fresh input, the probability that V_{i-1} hits is at most $(2^n - (q + \ell))^{-1}$, and the probability that U_i is not fresh is at most $(q + \ell)(2^n - (q - \ell))^{-1}$ as there are at most $q + \ell$ non-fresh values when U_i is sampled uniformly from a set of size at least $2^n - (q + \ell)$. Hence,

$$\mathbb{P} \left[\overline{\text{FRESH}}_i \mid \bigcap_{k=\ell'+1}^{i-1} \text{FRESH}_k \right] \leq \frac{q + \ell + 1}{2^n - (q + \ell)},$$

and

$$\mathbb{P}[\text{FRESH}] \geq 1 - \frac{q\ell + \ell^2 + \ell}{2^n - (q - \ell)},$$

and, finally,

$$\begin{aligned} q_\tau &\geq \mathbb{P}[\text{FRESH}] \cdot \mathbb{P}[Y_\ell = s^* | \text{FRESH}] \\ &\geq \left(1 - \frac{q\ell + \ell^2 + \ell}{2^n - (q - \ell)} \right) \cdot 2^{-n}, \end{aligned}$$

which implies

$$\frac{\mathfrak{p}_0(\tau)}{\mathfrak{p}_1(\tau)} \geq 1 - \frac{q\ell + \ell^2}{2^n - (q - \ell)} \geq 1 - \frac{2(q\ell + \ell^2)}{2^n}.$$

□

Lemma 31 (Bad event analysis). *For the set \mathcal{B} of bad transcripts (as defined above),*

$$\mathbb{P}[T_1 \in \mathcal{B}] \leq \frac{2q}{2^{\gamma^*}} + \frac{q}{2^n} + \frac{q^2}{2^c} + \text{Adv}_{\text{SpG}}^{\text{inv}, \pi}(\mathcal{A}).$$

Proof. The same resampling approach as in the proof of Lemma 17 applies here as well. First, observe that a hit occurs if either one of the forward queries returns s^* or if the attacker makes a backward query on s^* . Hence, the probability of a hit is at most

$$\frac{q}{2^n} + \text{Adv}_{\text{SpG}}^{\text{inv}, \pi}(\mathcal{A}).$$

Consider the following directed graph $G = (V, E)$ based on the query set L :

- the nodes are the capacity parts that appear in L , i.e.,

$$V = \{v^{(c)} \mid (v, *) \in L \vee (*, v) \in L\};$$

- two nodes are connected by a labeled edge if a corresponding query has been made, i.e.,

$$E = \{(u^{(c)}, v^{(c)}, l) \mid (u, v) \in L \wedge l = u^{(r)} \oplus v^{(r)}\}.$$

L is called *collision-free* if there is at most one path *with unique labels* from $s_0^{(c)}$ to every other node in G . Let \mathcal{E} be the event that L is *not* collision-free. Note that

$$\mathbb{P}[T_1 \in \mathcal{B}] \leq \mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}] + \mathbb{P}[\overline{\mathcal{E}}] \leq \mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}] + \frac{q^2}{2^c}.$$

Towards bounding $\mathbb{P}[T_1 \in \mathcal{B} | \mathcal{E}]$, consider a triple $z = (\sigma, L, s_0) \in \mathcal{E}$, which, once more, is shorthand for L being collision-free. Define a *potential chain* to be, for *some* ℓ , any sequence $(u_1, v_1), \dots, (u_\ell, v_\ell) \in L$ such that $u_1 = s_0 \oplus \mu_1 || 0^c$ and

$$u_i = v_{i-1} \oplus \mu_i || 0^c$$

for $i = 2, \dots, \ell - 1$ and some values μ_1, \dots, μ_ℓ . Note that $(u_1^{(c)}, v_1^{(c)}), \dots, (u_\ell^{(c)}, v_\ell^{(c)})$ describe a path from $s_0^{(c)}$ to $v_\ell^{(c)}$ with labels μ_1, \dots, μ_ℓ . Hence, that for a collision-free L , there are at most q potential chains.

Clearly, conditioned on $Z = z$, $T_1 \in \mathcal{B}$ if and only if for some potential chain, $X_i = \mu_i$ for all $i = 1, \dots, \ell$. Hence, by the legitimacy of \mathcal{A} ,

$$\mathbb{P}[T_1 \in \mathcal{B} | Z = z] \leq q \cdot p_z ,$$

where $p_z := \text{Pred}(\bar{X} | Z = z)$. The remainder of the proof proceeds in the exact same fashion as the proof of Lemma 17. \square

6.6.2 Maintaining Security

The maintaining security of the Sponge PRNG construction Spg is proved along similar lines as that of the previous constructions. This section discusses the differences.

Lemma 32 (Maintaining security). *The advantage of any (q, ℓ) -attacker \mathcal{A} against maintaining security is bounded by*

$$\text{Adv}_{\text{Spg}}^{\text{mtn}, \pi}(\mathcal{A}) \leq \varepsilon_{\text{Spg}}^{\text{mtn}}(q, \ell) := \frac{2(q\ell + \ell^2)}{2^n} + \frac{2q}{2^n} + \text{Adv}_{\text{Spg}}^{\text{inv}, \pi}(\mathcal{A}) .$$

In the following, for convenience, let

$$\delta_{\text{Spg}}^{\text{mtn}}(\gamma^*, \ell, q) := \frac{2(q\ell + \ell^2)}{2^n} + \frac{2q}{2^n} .$$

Proof. Similarly to the preceding maintaining proofs, maintaining security of Spg is shown via an H-coefficient proof. Once more, one considers transcripts

$$\tau = (s^*, s_0, x_1, \dots, x_\ell, L) ,$$

where L refers to π -queries here. In a *bad* transcript, L contains a query of the type

- $(s_0, *)$, which happens (with in the case $b = 1$) with probability at most $|L|/2^n$ since s_0 is completely independent of \mathcal{A} 's view, or
- $(*, s^*)$, which happens (in the case $b = 1$) with probability $|L|/2^n$ via a forward query or with probability $\text{Adv}_{\text{Spg}}^{\text{inv}, \pi}(\mathcal{A})$ via a backward query.

As for good transcripts

$$p_1(\tau) = 2^{-2n} \cdot p_L \quad \text{and} \quad p_0(\tau) \geq \left(1 - \frac{2(q\ell + \ell^2)}{2^n}\right) \cdot 2^{-2n} \cdot p_L ,$$

where the latter follows via an argument similar to that in the proof of Lemma 30. \square

6.6.3 Next Security

Recall that the next function next of the Sponge construction computes, on an input state s_0 ,

$$(s, y) = \text{next}^\pi(s_0) = (\pi(s_0) \oplus 0^r \| s_0^{(c)}, s_0^{(r)}) ,$$

where s is the new state and y is the output. Next security demands that if s_0 is chosen uniformly at random, then the output of next be indistinguishable from U_{n+r} to an attacker \mathcal{A} making at most to q queries to π . Denote by $\text{Adv}_{\text{Spg}}^{\text{next}, \pi}(\mathcal{A})$ the advantage of \mathcal{A} .

Lemma 33 (Next security). *The advantage of any q -attacker \mathcal{A} against next security is bounded by*

$$\text{Adv}_{\text{SpG}}^{\text{next},\pi}(\mathcal{A}) \leq \varepsilon_{\text{SpG}}^{\text{next}}(q) := \frac{2q}{2^c}.$$

Proof. For a simple H-coefficient proof, consider a transcript

$$\tau = (s_0^{(c)}, s^*, y^*, L),$$

where s_0 is the initial state, s^* is the new state, y^* is the output value, and L are the queries to π . A bad transcript is a transcript with a query of the type

- $(y^* \| s_0^{(c)}, *) \in L$ or
- $(*, s^* \oplus 0^r \| s_0^{(c)}) \in L$.

Since the view of \mathcal{A} in the ideal world, where \mathcal{A} 's view is independent of $s_0^{(c)}$, the probability of a bad transcript is at most $2q/2^c$.

For a good transcript, observe that

$$p_1(\tau) = 2^{-c} \cdot 2^{-r} \cdot 2^{-n} \cdot p_L,$$

where p_L denotes the probability that a uniform random permutation is consistent with the queries in L . Moreover,

$$p_0(\tau) = 2^{-c} \cdot 2^{-r} \cdot q_\tau \cdot p_L.$$

Note that since τ is a good transcript, $(y^*, s_0^{(c)})$ is a fresh input to π and $s \oplus 0^r \| s_0^{(c)}$ is still available; hence $q_\tau \geq 2^{-n}$. \square

6.6.4 Recovering Security

In the following, let $\varepsilon_{\text{SpG}}^{\text{ext}}(\gamma^*, q, \ell)$ and $\varepsilon_{\text{SpG}}^{\text{next}}(q)$ be as in Lemmas 29 and 33. Once more, extraction and next security together imply recovering security:

Lemma 34 (Recovering Security). *For every γ^* -legitimate (q, ℓ) -attacker \mathcal{A} ,*

$$\text{Adv}_{\text{SpG}}^{\text{rec},\pi}(\mathcal{A}) \leq \delta_{\text{SpG}}^{\text{ext}}(\gamma^*, \ell, q + r/n + 1) + \frac{q + r/n + 1}{2^n} + 2 \cdot \varepsilon_{\text{SpG}}^{\text{next}}(q + r/n + 1).$$

Proof. As in the proof of Lemma 20 consider, for $b \in \{0, 1\}$, the recovering experiment H_b conditioned on the secret bit having the value b . Moreover, again define a hybrid experiment $H_{\frac{1}{2}}$ in which the challenge oracle returns $\text{next}^\pi(U_n)$ to \mathcal{A} .

- The distance between H_0 and $H_{\frac{1}{2}}$ is bounded by a similar reduction \mathcal{A}_{ext} to extraction security as in Lemma 20. Hence, it is at most

$$\varepsilon_{\text{SpG}}^{\text{ext}}(\gamma^*, \ell, q + r/n + 1) \leq \delta_{\text{SpG}}^{\text{ext}}(\gamma^*, \ell, q + r/n + 1) + \text{Adv}_{\text{SpG}}^{\text{inv},\pi}(\mathcal{A}_{\text{ext}}).$$

In order to bound the probability of a hit in the ideal world of extraction security, one analyses this event in a hybrid world similar to H_1 . By next security and the fact that hit occurs with probability at most $(q + r/n + 1) \cdot 2^{-n}$ in H_1 in that hybrid,

$$\text{Adv}_{\text{SpG}}^{\text{inv},\pi}(\mathcal{A}_{\text{ext}}) \leq \frac{q + r/n + 1}{2^n} + \varepsilon_{\text{SpG}}^{\text{next}}(q + r/n + 1).$$

- By an argument similar to that in Lemma 20, one uses next security to show that the advantage of \mathcal{A} in distinguishing hybrids $H_{\frac{1}{2}}$ and H_1 is upper-bounded by $\varepsilon_{\text{MD}}^{\text{next}}(q)$.

\square

6.6.5 Preserving Security

In the following, let $\varepsilon_{\text{SpG}}^{\text{mtn}}(q, \ell)$ and $\varepsilon_{\text{SpG}}^{\text{next}}(q)$ be as in Lemmas 29 and 33. Once more, maintaining and next security together imply preserving security:

Lemma 35 (Preserving Security). *For every (q, ℓ) -attacker \mathcal{A} ,*

$$\text{Adv}_{\text{SpG}}^{\text{pre}, \pi}(\mathcal{A}) \leq \delta_{\text{SpG}}^{\text{mtn}}(\ell, q + r/n + 1) + \frac{q + r/n + 1}{2^n} + 2 \cdot \varepsilon_{\text{SpG}}^{\text{next}}(q + r/n + 1) .$$

The proof proceeds similarly to that of Lemma 34—except that reductions to maintaining security are made instead of extraction security—and is omitted.

7 Security Proofs for IT Constructions

This section analyzes all previously presented IT secure constructions. Used by all proofs in this section are the three propositions stated and proved in Section 7.1. Section 7.2 then shows the security of the monolithic extractor from Section 3.2. While rather straight-forward, the corresponding proof is quite instructive for the security proofs of the PRNG constructions from Section 5. For these proofs, it is convenient to define and establish an intermediate property, which is discussed in Section 7.3. Finally, Sections 7.4 and 7.5 establish the robustness (cf. Section 4.2) of the PRNG constructions.

7.1 Information-Theoretic Preliminaries

The *collision probability* of a random variable X is defined simply as $\text{Coll}(X) := \sum_x \mathbb{P}[X = x]^2$. Moreover, let $\text{Coll}(X|y) := \sum_x \mathbb{P}[X = x|Y = y]^2$, and define the conditional collision probability

$$\text{Coll}(X|Y) := \mathbf{E}_{y \leftarrow Y}[\text{Coll}(X|y)] .$$

The following two propositions relate the statistical distance from uniform to the collision probability. The first is well known and at the core of the proof of the leftover hash lemma [29]; the latter is proved for self-containment.

Proposition 1. *For any random variable X with size- N range, and a uniformly distributed U with the same range,*

$$\text{SD}(X, U) \leq \frac{1}{2} \sqrt{N \cdot \text{Coll}(X) - 1} .$$

Proposition 2. *Let F be chosen uniformly at random from a set \mathcal{F} . Then, for any random variable X with size- N range (arbitrarily correlated with \mathcal{F}), and a uniformly distributed U with the same range, independent of F ,*

$$\text{SD}((X, F), (U, F)) \leq \frac{1}{2} \sqrt{N \cdot \text{Coll}(X|F) - 1} .$$

Proof. By Proposition 1,

$$\text{SD}((X, F), (U, F)) \leq \frac{1}{2} \sqrt{N|\mathcal{F}| \cdot \text{Coll}(X, F) - 1} .$$

Moreover,

$$\begin{aligned}
\text{Coll}(X, F) &= \sum_{x, f} \mathbb{P}[(X, F) = (x, f)]^2 \\
&= \frac{1}{|\mathcal{F}|} \sum_f \mathbb{P}[F = f] \sum_x \mathbb{P}[X = x | F = f]^2 \\
&= \frac{\text{Coll}(X|F)}{|\mathcal{F}|},
\end{aligned}$$

from which the proposition follows. \square

The following proposition will also be useful.

Proposition 3. *Consider two random variables X and Y with identical range and let \mathcal{E} and \mathcal{E}' be events on their respective probability spaces. Assume $\mathbb{P}[\mathcal{E}] = \mathbb{P}[\mathcal{E}']$, then*

$$\text{SD}(X, Y) \leq \text{SD}(X|\mathcal{E}, Y|\mathcal{E}') + \mathbb{P}[\bar{\mathcal{E}}].$$

Proof. Observe that

$$\begin{aligned}
\text{SD}(X, Y) &= \frac{1}{2} \sum_x |P_X(x) - P_Y(x)| \\
&= \frac{1}{2} \sum_x \left| P_{X|\mathcal{E}}(x) \mathbb{P}[\mathcal{E}] + P_{X|\bar{\mathcal{E}}}(x) \mathbb{P}[\bar{\mathcal{E}}] - P_{Y|\mathcal{E}'}(x) \mathbb{P}[\mathcal{E}'] - P_{Y|\bar{\mathcal{E}'}}(x) \mathbb{P}[\bar{\mathcal{E}'}] \right| \\
&\leq \mathbb{P}[\mathcal{E}] \cdot \frac{1}{2} \sum_x |P_{X|\mathcal{E}}(x) - P_{Y|\mathcal{E}'}(x)| + \mathbb{P}[\bar{\mathcal{E}}] \cdot \frac{1}{2} \sum_x |P_{X|\bar{\mathcal{E}}}(x) - P_{Y|\bar{\mathcal{E}'}}(x)| \\
&\leq \text{SD}(X|\mathcal{E}, Y|\mathcal{E}') + \mathbb{P}[\bar{\mathcal{E}}].
\end{aligned}$$

\square

7.2 Monolithic Extractor

From Section 3.2, recall the monolithic extractor `mono` defined to work with a random oracle $F : \{0, 1\}^m \rightarrow \{0, 1\}^n$:

Construction 1 (Monolithic extractor). The monolithic seedless extractor $\text{mono}^G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ using a random oracle $G : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined by

$$\text{mono}^G(x) := G(x).$$

Construction 1 is an IT seedless extractor according to Definition 3.

Theorem 2 (Monolithic seedless IT-extraction). *Construction `mono` is a $(\gamma^*, q, \varepsilon)$ -IT-extractor in the G -model for*

$$\varepsilon \leq \frac{1}{2} \sqrt{\frac{2^{-(\gamma^* - n)}}{1 - \rho}} + \rho,$$

where $\rho = q/2^{\gamma^*}$.

The proof of Theorem 2 proceeds by bounding the statistical distance of \mathcal{A}_2 's views in the real and ideal experiments via the corresponding collision probabilities (as done in the proof of the left-over hash lemma). In the proofs of the actual PRNG constructions in the following sections, bounding said collision probabilities constitutes the bulk of the proof and is quite involved.

Proof. Consider the extraction game corresponding to Definition 3 (cf. Section 3.1). Since \mathcal{A}_2 gets to make an unbounded number of queries to the random oracle G , one may equivalently consider the game where \mathcal{A}_2 simply gets the entire function table of G as input. Therefore, the distinguishing advantage of \mathcal{A}_2 is upper bounded by

$$\text{SD}((\Sigma, Y_0, \mathcal{L}_1, G), (\Sigma, Y_1, \mathcal{L}_1, G)) ,$$

where Σ is (the random variable corresponding to) the state information output by \mathcal{A}_1 , $Y_0 = G(X)$, Y_1 a uniform random string, and \mathcal{L}_1 the query/answer list by \mathcal{A}_1 .

Let \mathcal{E} be the event that \mathcal{A}_1 does not query the value X it provides as input to the extractor. Observe that this event can be defined in both the real ($b = 0$) and ideal ($b = 1$) experiments. In the ideal experiment, the probability of \mathcal{E} not occurring is easy to bound: Since Y_1 is uniformly random and independent of X , consider the following equivalent way of sampling a tuple $(\Sigma, X, Y_1, \mathcal{L}_1, G)$:

1. Sample Y_1 and G uniformly at random.
2. Run \mathcal{A}_1^G to produce Σ , \tilde{X} , and \mathcal{L}_1 .
3. Rerun \mathcal{A}_1^G with fresh randomness, but conditioned on the state information being Σ and the queries being \mathcal{L}_1 . This results in a new value X .
4. Output $(\Sigma, X, Y_1, \mathcal{L}_1, G)$.

It is easily seen that the distribution produced via resampling is identical to that of the actual experiment.

Note that for particular values σ and L_1 ,

$$\mathbb{P}[\bar{\mathcal{E}}|\sigma L_1] \leq q \cdot \text{Pred}(X|\sigma L_1) ,$$

which can be easily seen due to the alternative sampling above. Hence, taking expectations,

$$\mathbb{P}[\bar{\mathcal{E}}] \leq q \cdot \text{Pred}(X|\Sigma \mathcal{L}_1) \leq q \cdot 2^{-\gamma^*} ,$$

using the γ^* -legitimacy of \mathcal{A} in the last step.

Using Proposition 3, one obtains

$$\text{SD}((\Sigma, Y_0, \mathcal{L}_1, G), (\Sigma, Y_1, \mathcal{L}_1, G)) \leq \text{SD}((\Sigma, Y_0, \mathcal{L}_1, G)|\mathcal{E}, (\Sigma, Y_1, \mathcal{L}_1, G)|\mathcal{E}) + q \cdot 2^{-\gamma^*} .$$

In order to bound the statistical distance conditioned on \mathcal{E} , condition additionally on arbitrary values $z = (\sigma, L_1)$ (with non-zero probability given \mathcal{E}). Observe that under such conditioning, G is chosen uniformly at random from all functions consistent with G . Using Proposition 2, one bounds the desired statistical distance as

$$\text{SD}((Y_0, G)|z\mathcal{E}, (Y_1, G)|z\mathcal{E}) \leq \frac{1}{2} \sqrt{2^n \cdot \text{Coll}(Y_0|Gz\mathcal{E}) - 1} .$$

To bound the collision probability $\text{Coll}(Y|Gz\mathcal{E})$, consider the following experiment:²²

²²Observe that, in general, $\text{Coll}(U|V)$ is equal to the probability that $U = U'$ in the experiment where one jointly samples U and V , and then resamples U' conditioned on the value of V .

1. Sample G uniformly at random consistent with L_1 .
2. Sample inputs $X \leftarrow \mathcal{A}_1^G$ and $X' \leftarrow \mathcal{A}_1^G$ independently but conditioned on $Z = z$ and \mathcal{E} .
3. Compute $Y_0 \leftarrow \text{mono}^G(X)$ and $Y'_0 \leftarrow \text{mono}^G(X')$ respectively.

The collision probability $\text{Coll}(Y|Gz)$ is therefore equal to

$$\text{P}[Y_0 = Y'_0] \leq \text{P}[X = X'] + \text{P}[Y = Y'|X \neq X']$$

in the experiment above. The former term is at most $p_z := \text{Pred}(X|Z = z, \mathcal{E})$. For the latter term, consider arbitrary $x \neq x'$. Since neither x nor x' is covered by L_1 , $Y_0 = Y'_0$ occurs with probability 2^{-n} . Hence,

$$\text{SD}((Y_0, G)|z\mathcal{E}, (Y_1, G)|z\mathcal{E}) \leq \frac{1}{2} \sqrt{2^n(p_z + 2^{-n}) - 1} \leq \frac{1}{2} \sqrt{2^n p_z}.$$

Using Jensen's inequality, one obtains

$$\begin{aligned} \text{SD}((Y_0, G)|z\mathcal{E}, (Y_1, G)|z\mathcal{E}) &\leq \frac{1}{2} \sqrt{2^n \cdot \text{Pred}(X|Z\mathcal{E})} \\ &\leq \frac{1}{2} \sqrt{\frac{2^{-(\gamma^* - n)}}{(1 - \rho)}}, \end{aligned}$$

where the last inequality follows from

$$\begin{aligned} \text{Pred}(X|Z) &\geq \sum_z \text{P}[Z = z \wedge \mathcal{E}] \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \sum_{z \in \mathcal{E}} \frac{\text{P}[Z = z \wedge \mathcal{E}]}{\text{P}[\mathcal{E}]} \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \text{Pred}(\bar{X}|Z\mathcal{E}). \end{aligned}$$

and $\text{Pred}(X|Z) \leq 2^{-\gamma^*}$, by the legitimacy of \mathcal{A} . □

7.3 Intermediate IT PRNG Security Notions

Similarly to the computational case, it is useful to consider a simplified security notion, *recovering* security, for IT PRNGs. Recall that for the computational case, recovering security requires that the output of the next function next look random once sufficient amounts of entropy have been accumulated in the PRNG's state. In the IT case, the requirement is relaxed by only requiring that the output be indistinguishable from $(0^n, U_r)$. That is, call to `next` always resets the PRNG state to 0^n . This is without loss, as the definition of legitimacy considers every call to `next` an entropy drain.

The IT recovering game is represented in Figure 9. The advantage of an attacker \mathcal{A} in this game is denoted by $\text{Adv}_{\text{PRNG}}^{\text{rec-IT}, P}(\mathcal{A})$. An attacker \mathcal{A} in this game is γ^* -legitimate if

$$\text{H}_\infty(X_1, \dots, X_\ell | \Sigma \mathcal{L} S_0) \geq \gamma^*,$$

where

- Σ is the state of \mathcal{A} just before the call to **chall**,

Oracles of PRNG Recovering-IT Game

<pre> init $b \leftarrow \{0, 1\}$ </pre>	<pre> chall(s_0, x_1, \dots, x_ℓ) for $i = 1, \dots, \ell$ $s_i \leftarrow \text{refresh}(s_{i-1}, x_i)$ if $b = 0$ return $\text{next}(s_\ell)$ else $y \leftarrow \{0, 1\}^r$ return $(0^n, y)$ </pre>
---	--

Figure 9: Oracles for PRNG Recovering-IT game.

- \mathcal{L} is the list of query and answers \mathcal{A} has made to P up to the call to **chall**, and
- S_0 is the initial state that the adversary provides.

A (q, ℓ) -IT-attacker here is one that can make at most q queries to P before its challenge (and arbitrarily many afterwards), and such that the input to **chall** consists of at most ℓ blocks.

Definition 15. A PRNG construction $\text{PRNG} = (\text{refresh}, \text{next})$ is said to be $(\gamma^*, q, \ell, \varepsilon)$ -IT-recovering in the P -model if for every γ^* -IT-legitimate (q, ℓ) -IT-attacker,

$$\text{Adv}_{\text{PRNG}}^{\text{rec-IT}, P}(\mathcal{A}) \leq \varepsilon .$$

Observe that the definition of recovering security for IT PRNGs is—up to syntactical differences—equivalent to the security of the PRNG as an extractor, which is obtained by replacing the PRNG `next` algorithm by an algorithm `finalize` that simply discards the state output by `next`. In other words, an IT-robust PRNG can be viewed as an IT-secure online extractor. In particular, the following theorem is in principle little more than just a union bound.

Theorem 36. Let $\text{PRNG} = (\text{refresh}, \text{next})$ be a PRNG for which `refresh` makes α P -calls and `next` makes β P -calls. Let the PRNG be $(\gamma^*, q, \ell, \varepsilon_{\text{rec}})$ -IT-recovering in the P -model. Then, PRNG is also $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust in the P -model, where

$$\varepsilon_{\text{rob}} \leq t \cdot \varepsilon_{\text{rec}} .$$

The proof can be found in Appendix C.2.

7.4 IT PRNGs from Merkle-Damgård

This section establishes the robustness of the PRNG construction $\text{MD}_r = (\text{refresh}, \text{next})$ based on Merkle-Damgård with a random compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$. Recall that the construction is defined as follows (cf. Figure 3):

Construction 3 (IT-PRNG from Merkle-Damgård). The (m, r) -PRNG construction $\text{MD}_r = (\text{refresh}, \text{next})$ based on Merkle-Damgård with a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined as follows:

- $\text{refresh}^F(s, x) = F(s, x)$, and
- $\text{next}^F(s) = (0^n, s[1..r])$.

Construction 3 achieves the following security:

Theorem 6 (IT-Robustness of Merkle-Damgård PRNGs). *Construction 3 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust PRNG in the F -model, where*

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 \cdot \frac{\tilde{q}^2 2^r}{2^{2n}}} + t\rho ,$$

for $\rho = \frac{\tilde{q}^2}{2^r}$ where $\tilde{q} = q + t\ell$.

The theorem will be a direct consequence of the following lemma, and Theorem 36 above, which adds a multiplicative factor t to obtain the bound in the theorem.

Lemma 37. *For every γ^* -IT-legitimate (q, ℓ) -attacker in the ideal compression function model,*

$$\text{Adv}_{\text{MD}_r}^{\text{rec-IT}, \gamma^*}(\mathcal{A}) \leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 \frac{q^2 2^r}{2^{2n}}} + \rho ,$$

where $\rho = \frac{q^2}{2^r}$.

Before we turn to a proof, we note the challenges here. In particular, the core of the proof will be to show a bound on the collision probability of the output of `next` for a random compression function F , which in turn will reduce to outputting the truncation of the output of the MD construction. This problem resembles that studied by [17] and by [24]. However, the difficulty here is that we need to consider the set of queries previously done by the adversary \mathcal{A} . This will significantly complicate the proof—we also modify slightly the graph-theoretic formalization adopted by these previous works to something more amenable to our more complex setting.

Proof (of Lemma 37). We define a few random variables which we will be using in our proofs.

- F a randomly chosen compression function, to which the adversary is given access. We use F both for the oracle itself, as well as for the random variable describing the entire function table.
- ℓ : Number of blocks input to the challenge oracle (which is a random variable itself, we overload notation here, using the same letter we use in the bound on ℓ in the lemma statement).
- $\bar{X} = (\bar{X}_1, \dots, \bar{X}_\ell)$: the blocks input to the challenge oracle.
- \tilde{Y}_ℓ : output of MD_r . Let us remind ourselves that this is s truncated to r bits (the output of `next`);
- $Z = (\Sigma, \mathcal{L}, S_0)$: “side information” where Σ is the attacker state before challenge, \mathcal{L} is the attacker query/answers to F before challenge, S_0 is the initial PRNG state provided by \mathcal{A} ;
- U_r : uniform r -bit string.

The advantage of the adversary \mathcal{A} in the recovering game is bounded by $\text{SD}((\tilde{Y}_\ell, Z, F), (U_r, Z, F))$. Therefore, it is sufficient to upper-bound that. The reasoning is quite simple. Note that Z contains the state of the attacker Σ just before it makes the challenge query. This means that \mathcal{A} cannot tell apart real from random.

We also define an event \mathcal{E} where the answers to the F -queries by \mathcal{A} are distinct *when truncated to the first r bits*. Note that there are a maximum of q queries in this list. We would like to point

out that \mathcal{E} has the same probability of occurring in either experiment, since the experiments are identical up to the point when this even is defined. Therefore, by Proposition 3,

$$\text{SD}((\tilde{Y}_\ell, Z, F), (U_r, Z, F)) \leq \text{SD}((\tilde{Y}_\ell, Z, F)|\mathcal{E}, (U_r, Z, F)|\mathcal{E}) + \frac{q^2}{2^r};$$

For convenience we let $\rho := q^2/2^r$ for the remainder of the proof. In order to bound the statistical distance conditioned on \mathcal{E} , we can rewrite the same as as

$$\text{SD}((\tilde{Y}_\ell, Z, F)|\mathcal{E}, (U_r, Z, F)|\mathcal{E}) = \sum_{z \in \mathcal{E}} \text{P}[Z = z|\mathcal{E}] \cdot \text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z), \quad (3)$$

where $z \in \mathcal{E}$ is to denote that the sum is taken over all side informations $Z = z$, satisfying \mathcal{E} .²³ Define $p_z := \text{Pred}(\bar{X}|Z = z)$, and observe that

$$\mathbf{E}_z[p_z] = \text{Pred}(\bar{X}|Z) \leq 2^{-\gamma^*},$$

where the latter inequality follows from the assumption $\text{H}_\infty(\bar{X}|Z) \geq \gamma^*$. Moreover,

$$\text{H}_\infty(\bar{X}|Z\mathcal{E}) \geq \gamma^* - \log(1 - \rho)^{-1},$$

which is due to

$$\begin{aligned} \text{Pred}(\bar{X}|Z) &\geq \sum_{z \in \mathcal{E}} \text{P}[Z = z] \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \sum_{z \in \mathcal{E}} \frac{\text{P}[Z = z]}{\text{P}[\mathcal{E}]} \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \text{Pred}(\bar{X}|Z\mathcal{E}). \end{aligned}$$

From Lemma 38 we prove below, we will get,

$$\text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}}}.$$

Using Jensen's inequality, (3) becomes, for $\alpha = 2^r$ and $\beta = \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}}$,

$$\begin{aligned} \text{SD}((\tilde{Y}_\ell, F)|\mathcal{E}, (U_r, F)|\mathcal{E}) &\leq \frac{1}{2} \sqrt{\alpha \text{Pred}(\bar{X}|Z\mathcal{E}) + \beta} \\ &\leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \frac{2^{2r}}{2^{2n}}}. \end{aligned}$$

□

Lemma 38. For $z \in \mathcal{E}$ and the random variables as defined earlier,

$$\text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}}}.$$

Proof. Fix $z = (\sigma, L, s_0)$. Observe that, conditioned on z , F is distributed uniformly over the set of all functions that agree with L . Thus, by Proposition 2,

$$\text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z) \leq \frac{1}{2} \sqrt{2^r \cdot \text{Coll}(\tilde{Y}_\ell|Fz) - 1}. \quad (4)$$

To bound the collision probability, we consider the following experiment:

²³Therefore \mathcal{E} can be omitted in the conditioning of the statistical distance.

- choose F uniformly consistent with L
- sample inputs $\bar{X} = (\bar{X}_1, \dots, \bar{X}_\ell)$ and $\bar{X}' = (\bar{X}'_1, \dots, \bar{X}'_{\ell'})$ independently but conditioned on $Z = z$.
- compute \tilde{Y}_ℓ and $\tilde{Y}'_{\ell'}$ as the truncated MD evaluations with F of \bar{X} and \bar{X}'

Then, we bound the probability that $\tilde{Y}_\ell = \tilde{Y}'_{\ell'}$ in this experiment as

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \mathbb{P}[\bar{X} = \bar{X}'] + \mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'} | \bar{X} \neq \bar{X}']. \quad (5)$$

Clearly, the former is at most p_z . To bound the latter, we fix arbitrary inputs $\bar{x} \neq \bar{x}'$ of lengths ℓ and ℓ' , respectively. We also assume, wlog, that the evaluation of \bar{x}' is not completely covered by L ; due to the collision-freeness of L . Let \bar{x}_{k+1} be the first block of \bar{x} not covered by \mathcal{L} and similarly \bar{x}'_{k+1} for \bar{x}' . We let $k = \ell$ if all blocks are covered.

The structure graph. We will now model the evaluation producing \tilde{Y}_ℓ and $\tilde{Y}'_{\ell'}$ as a (labeled) graph process. In particular, let us define for convenience

$$x^{(i)} := \begin{cases} \bar{x}_i & i \leq \ell \\ \bar{x}'_{(i-\ell)} & \text{otherwise} \end{cases}$$

for all $i = 1, \dots, \ell + \ell'$, and we let $\tilde{\ell} = \ell + \ell'$.

For a given compression function F (consistent with L), first define a labeled (multi-)graph $H_F(\bar{x}, \bar{x}') = (\mathcal{V}, \mathcal{E}, \mathcal{L})$ vertex set $\mathcal{V} \subseteq \{0, 1\}^n$. Each edge will have a label $\mathcal{L}(e) = (x, c)$ – where $x \in \{0, 1\}^m$ and $c \in \{\text{red}, \text{blue}\}$. Now we add edges as follows, allowing replications of edges (we will then explain below how to remove duplicates). We will implicitly define \mathcal{V} as the set of n -bit strings which are endpoint to an edge:

Blue edges. For every $((x, y), y')$ in L we add an edge (y, y') with label (x, blue) , and refer to such edges as the *blue* edges.

Red edges. Define

$$s_i := \begin{cases} 0^n & i = 0 \\ F(s_{i-1}, \bar{x}_i) & 1 \leq i \leq \ell \\ F(0^n, \bar{x}'_1) & i = \ell + 1 \\ F(s_{i-1}, \bar{x}'_{(i-\ell)}) & \ell + 2 \leq i \leq \tilde{\ell}. \end{cases}$$

Now for $i = 1, \dots, \ell$, we add edge (s_{i-1}, s_i) with label (\bar{x}_i, red) , unless the edge (s_{i-1}, s_i) is already present with label (\bar{x}_i, c) for $c \in \{\text{blue}, \text{red}\}$. We add the edge $(0^n, s_{\ell+1})$ with label (\bar{x}'_1, red) , unless the edge (s_{i-1}, s_i) is already present with label (\bar{x}'_1, c) for $c \in \{\text{blue}, \text{red}\}$. Finally, we add each (s_{i-1}, s_i) for $i = \ell + 2, \dots, \tilde{\ell}$ with label (\bar{x}'_i, red) , unless the edge (s_{i-1}, s_i) is already present with label (\bar{x}'_i, c) for $c \in \{\text{blue}, \text{red}\}$. We refer to these edges we added as the *red* edges.

Note in particular that we may have two identical edges e_1 and e_2 , but in this case they will have labels (x_1, c_1) and (x_2, c_2) with $x_1 \neq x_2$, and moreover, at least one of them is red by our assumption on L .

Definition 16. The structure graph of $G_F = G_F(\bar{x}, \bar{x}') = (\mathcal{V}, \mathcal{E}, \mathcal{L})$ is the graph obtained from H_F as follows. We first look at all isomorphic graphs to G_F with vertices $\{0, \dots, |\mathcal{V}| - 1\}$ such that 0^n is mapped to 0. We then pick the lexicographically first such graph.²⁴

Let $\mathcal{G}_L(\bar{x}, \bar{x}')$ be the set of all $G_F(\bar{x}, \bar{x}')$ for an F compatible with L . Note that G_F will have two (possibly overlapping) paths starting from 0 with edges labeled by \bar{x} and \bar{x}' , containing blue and red edges. We say that G_F is *colliding* if these two paths end in the same vertex, and let $\text{Coll}_L(\bar{x}, \bar{x}') \subseteq \mathcal{G}_L(\bar{x}, \bar{x}')$ be the set of colliding structure graphs.

Definition 17 (Accidents). Let now \mathcal{B} be the set of vertices of the sub-graph of G_F induced by the blue edges, plus 0. We now traverse the path induced by \bar{x} , and then the path induced by \bar{x}' . Each time we encounter a vertex which is not in \mathcal{B} , we add it to it. Each time we encounter a vertex which is already in \mathcal{B} , we say that an accident has occurred. We let $\text{Acc}(G_F)$ be the number of accidents in G_F .

We also let $\mathcal{G}_L^a(\bar{x}, \bar{x}')$ to be the set of $H \in \mathcal{G}_L(\bar{x}, \bar{x}')$ with $\text{Acc}(H) = a$. We state the following lemmas.

Lemma 39. Let F be sampled randomly consistent with L , and $\bar{x} \neq \bar{x}'$. Then, let \tilde{Y}_ℓ and \tilde{Y}'_ℓ be values obtained after truncating at the end of the MD evaluations on inputs \bar{x} and \bar{x}' respectively. Then,

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell] \leq \mathbb{P}[G_F(\bar{x}, \bar{x}') \in \text{Coll}_L(\bar{x}, \bar{x}')] + \frac{1}{2^r}.$$

Proof. We rewrite $\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell]$ as:

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell] \leq \mathbb{P}[G_F(\bar{x}, \bar{x}') \in \text{Coll}_L(\bar{x}, \bar{x}')] + \mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell | G_F(\bar{x}, \bar{x}') \notin \text{Coll}_L(\bar{x}, \bar{x}')]$$

We take a closer look at the latter term. Note that the graph is a fixed graph and it has no collision at the end nodes. We proceed to assign random, yet distinct values to the vertices. These are chosen from $\{0, 1\}^n$. Note that it is sufficient to look at the two output vertices *locally* without looking at the global state. There are $2^n(2^n - 1)$ pairs of values for these output vertices such that these values are distinct. However, of these, $2^r 2^{n-r}(2^{n-r} - 1)$ have values which are equal in the first r bits and yet are distinct n -bit strings. Therefore,

$$\begin{aligned} \mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell | G_F(\bar{x}, \bar{x}') \notin \text{Coll}_L(\bar{x}, \bar{x}')] &\leq \frac{2^r 2^{n-r}(2^{n-r} - 1)}{2^n(2^n - 1)} \\ &= \frac{2^{n-r} - 1}{2^n - 1} \\ &\leq \frac{2^{n-r}}{2^n} = \frac{1}{2^r}. \end{aligned}$$

Putting it together, we have:

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell] \leq \mathbb{P}[G_F(\bar{x}, \bar{x}') \in \text{Coll}_L(\bar{x}, \bar{x}')] + \frac{1}{2^r}. \quad (6)$$

□

Lemma 40. Let F be sampled randomly consistent with L , and $\bar{x} \neq \bar{x}'$. Let $H \in \mathcal{G}_L(\bar{x}, \bar{x}')$. Then,

$$\mathbb{P}[G_F(\bar{x}, \bar{x}') = H] = \frac{1}{2^{n \cdot \text{Acc}(H)}}.$$

²⁴Indeed, the actual labeling of graphs will not matter below.

Proof. We have two messages $\bar{x} \neq \bar{x}'$ and let $x^{(i)}$ be as defined before. We have a randomly sampled F which is consistent with L . Let us assume that the values $S_1, \dots, S_{\tilde{\ell}}$ are revealed to us stepwise. S_i is the random variable representing the vertex after block i . Let $G_F = G_F(\bar{x}, \bar{x}')$. We define a consistency notion as follows: G_F is consistent with a given graph H after step $i \leq \tilde{\ell}$, denoted by Cons_i if the structure graphs $G_F^{(i)}$ and $H^{(i)}$ are equal as triples $(\mathcal{V}, \mathcal{E}, \mathcal{L})$ where $G_F^{(i)}$ is the graph G_F obtained after the first i blocks are processed. We define $H^{(i)}$ similarly. We assume that Cons_i is true for some i . We now bound $\text{P}[\text{Cons}_{i+1} | \text{Cons}_i]$. We look at the step $i+1$ in H and there are three possibilities on how the edge for message block $x^{(i+1)}$ looks:

- *Fresh:* It arrives at a new vertex which is not present in $H^{(i)}$.
- *Determined:* It follows an existing edge, i.e there exists a label for the edge of the form $(x^{(i+1)}, \cdot)$
- *Accident:* It causes an accident. In this case, $G_F^{(i+1)}$ will only be consistent if the edge corresponding to $x^{(i+1)}$ lands on the same vertex as in $H^{(i+1)}$. Note that this accident is a fresh-evaluation, i.e the output is not determined in the first i steps and is therefore chosen randomly from 2^n values. In other words, $\text{P}[\text{Cons}_{i+1} | \text{Cons}_i] = \frac{1}{2^n}$ in this case.

Note that the third case happens $\text{Acc}(H)$ times. Therefore, we have:

$$\text{P}[G_F(\bar{x}, \bar{x}') = H] = \text{P}[\text{Cons}_{\tilde{\ell}}] \leq \frac{1}{2^{n \cdot \text{Acc}(H)}}.$$

The latter inequality arises by upper-bounding $\text{P}[\text{Cons}_{i+1} | \text{Cons}_i]$ with 1 in the case of Fresh and Determined edges. \square

Lemma 41. *Let F be sampled randomly consistent with L , and $\bar{x} \neq \bar{x}'$. Then,*

$$\text{P}[\text{Acc}(G_F) \geq 2] \leq \frac{64\ell^4}{2^{2n}} + \frac{16\ell^2 q^2}{2^{2n}}.$$

Proof. We define \mathcal{G}_L^a , as before, to be the set of all structure graphs containing exactly a accidents.

$$\begin{aligned} \text{P}[\text{Acc}(G_F) \geq 2] &= \sum_{a=2}^{\infty} \text{P}[\text{Acc}(G_F) \geq 2] \\ &= \sum_{a=2}^{\infty} \sum_{g \in \mathcal{G}_L^a} \text{P}[G_F = g] \\ &\leq \sum_{a=2}^{\infty} \frac{|\mathcal{G}_L^a|}{2^{n \cdot a}}. \end{aligned}$$

The last step follows from Lemma 40. Observe that for fixed base graph and fixed messages, entire graph is defined by list of accidents. Each accident is defined by an edge (i, j) , where there are $\tilde{\ell} := \ell + \ell'$ choices for i and $\tilde{\ell} + q$ for j . Thus, there are at most $(\tilde{\ell}(\tilde{\ell} + q))^c$ graphs with a accidents.

$$\begin{aligned} \text{P}[\text{Acc}(G_F) \geq 2] &\leq \sum_{a=2}^{\infty} \frac{|\mathcal{G}_L^a|}{2^{n \cdot a}} \\ &\leq \sum_{a=2}^{\infty} \left(\frac{\tilde{\ell}(\tilde{\ell} + q)}{2^n} \right)^a \\ &\leq \left(\frac{\tilde{\ell}(\tilde{\ell} + q)}{2^n} \right)^2 \leq \frac{4\tilde{\ell}^4}{2^{2n}} + \frac{4\tilde{\ell}^2 q^2}{2^{2n}} \leq \frac{64\ell^4}{2^{2n}} + \frac{16\ell^2 q^2}{2^{2n}}. \end{aligned}$$

where we assume $\tilde{\ell} \leq 2\ell$ □

Note that $H \in \text{Coll}_L = \text{Coll}_L(\bar{x}, \bar{x}')$ necessarily implies $\text{Acc}(H) \geq 1$. We can then say, for $G_F = G_F(\bar{x}, \bar{x}')$,

$$\begin{aligned} \text{P}[G_F \in \text{Coll}_L] &\leq \text{P}[G_F \in \text{Coll}_L \wedge \text{Acc}(G_F) = 1] + \text{P}[\text{Acc}(G_F) \geq 2] \\ &\leq \frac{|\text{Coll}_L \cap \mathcal{G}_L^1|}{2^n} + \frac{64\ell^4}{2^{2n}} + \frac{16\ell^2 q^2}{2^{2n}}. \end{aligned} \quad (7)$$

by Lemmas 40 and 41. We are going to upper bound the number of graphs in $\text{Coll}_L \cap \mathcal{G}_L^1$.

We reduce handling the case of a general L now to the simpler case where $L = \emptyset$, so that we can resort to the counting argument of [24].

Now, let $G \in \text{Coll}_L \cap \mathcal{G}_L^1(\bar{x}, \bar{x}')$. Let \mathcal{E}_B be the set of all blue edges which are on the \bar{x} - and the \bar{x}' -paths. Moreover, let $\mathcal{E}'_B \subseteq \mathcal{E}_B$ be the set of all *initial* blue edges, i.e., the set of blue edges on these paths which are not preceded by any red edge. We prove the following lemma.

Lemma 42. *If $G \in \text{Coll}_L \cap \mathcal{G}_L^1(\bar{x}, \bar{x}')$, then $\mathcal{E}'_B = \mathcal{E}_B$*

Proof. We prove by contradiction by assuming that this is not true. In other words, there exists at least one edge $e \in \mathcal{E}_B \setminus \mathcal{E}'_B$. We now proceed to show that this necessarily leads to at least two collisions, contradicting $G \in \text{Coll}_L \cap \mathcal{G}_L^1(\bar{x}, \bar{x}')$. Wlog, assume that this lies on the \bar{x} path, and let $e_i = (u_i, v_i)$ be the first such edge. Then, this edge must be preceded by a red-edge, e_{i-1} . Then, clearly this implies that the graph contains at least one accident.

Now, let us take a look at the \bar{x}' path. We have, by our definition of G (it is in the set Coll_L), that there should exist a path from e_i to the end point of the \bar{x}' path. We will sketch a proof to show that this is not possible without having a second collision. We take a look at this path from e_i to the end point of \bar{x}' , V' .

- The path is of entirely blue edges.

We look at the edge into V' on the \bar{x}' path. Note that this cannot be a new blue edge as it would violate the no-collision constraint on L . If this was a red edge we have a second collision, by our definition. The only option is to take the same blue edge as on the e_i -path. In other words, the penultimate vertex on e_i -path and \bar{x}' path are the same. We can argue similarly for the penultimate vertex by showing that the edges into them should be the same blue edge on both the paths. Working backwards we arrive at the point u_i . Thus, we have shown that the message blocks corresponding to the edges along the path from e_i to V' is the same on both \bar{x} and \bar{x}' . We have also assumed that e_i is the first non-initial blue edge. In other words, the path before e_i should be all red. Since $\bar{x} \neq \bar{x}'$, there should be a block before e_i which would differ in the message and hence this would constitute a collision on the red edges. Therefore, in this case it is impossible to have the final outputs to be the same without causing a second collision.

- The path is of entirely red edges.

The argument is similar to the previous case. The only difference is that the \bar{x}' -path could take a blue edge into the e_i -path. However, this would still constitute a new collision. The same reasoning follows giving us the same conclusion.

- The path is mixed red and blue edges.

We have already shown that if a blue edge follows a red edge on the path then it is a collision. We can therefore assume that the set of blue edges occur together, followed by the set of red edges. The reasoning is pretty similar to the previous two cases.

In other words, a graph G cannot be in Coll_L and $\mathcal{G}_L^1(\bar{x}, \bar{x}')$ when there exist non-initial blue edges. \square

The direct consequence of Lemma 42 is that the number of colliding structure graphs with one accident is the same even if we remove all the non-initial blue edges. In addition, removing the initial blue edges will also not decrease the number of colliding graphs with one accident. This is true because the initial blue edges can be replaced by red edges without increasing the number of accidents by the properties of L . In other words, if we define $L' \subseteq L$ to be the set of queries which define the initial blue edges, we have

$$|\text{Coll}_L \cap \mathcal{G}_L^1(\bar{x}, \bar{x}')| = |\text{Coll}_{L'} \cap \mathcal{G}_{L'}^1(\bar{x}, \bar{x}')| \leq |\text{Coll}_\emptyset \cap \mathcal{G}_\emptyset^1(\bar{x}, \bar{x}')| \leq \ell d'(\ell). \quad (8)$$

where the last inequality is from [24]. Here $d'(n)$ is defined as follows:

$$d'(n) := \max_{n' \in \{1, \dots, n\}} |\{d \in \mathbb{N} : n' \bmod d = 0\}|.$$

Combining Equations 15, 7 and 8 we get,

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell] \leq \frac{1}{2^r} + \frac{\ell \cdot d'(\ell)}{2^n} + \frac{64\ell^4}{2^{2n}} + \frac{16\ell^2 q^2}{2^{2n}}$$

Substituting the above value in Equations 17 and 18, we get:

$$\text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}}}.$$

This concludes the proof of Lemma 38 \square

7.5 IT PRNGs from Merkle-Damgård with Davies-Meyer

This section establishes the robustness of the PRNG construction $\text{DM}_r = (\text{refresh}, \text{next})$ based on Merkle-Damgård with the Davies-Meyer compression function in the ideal-cipher model. Recall that the construction is defined as follows (cf. Figure 4):

Construction 5 (IT-PRNG from MD-DM). The (n, r) -PRNG construction $\text{DM}_r = (\text{refresh}, \text{next})$ using Merkle-Damgård with Davies-Meyer (MD-DM) uses a block cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and is defined as follows:

- $\text{refresh}^E(s, x) = E(x, s) \oplus s$, and
- $\text{next}^E(s) = (0^n, s[1..r])$.

Construction 5 achieves the following security:

Theorem 8 (IT-Robustness of MD-DM PRNGs). *Construction 5 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust PRNG in the E -model, where*

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + 64\ell^4 \cdot \frac{2^r}{2^{2n-2}} + 16\ell^2 \tilde{q}^2 \cdot \frac{2^r}{2^{2n-2}}} + t\rho,$$

for $\rho = \frac{\tilde{q}^2}{2^r}$ where $\tilde{q} = q + t\ell$

Our robustness proof for this construction uses the analysis from Section 7.4. The intuition here is that the structure graphs were defined for a compressing function F and we instantiate it with a Davies-Meyer construction and extend the arguments. However, there is a subtle difference because of the underlying primitive. In Davies-Meyer the primitive is an ideal cipher. Thus, we state and prove Lemma 45 which is the counterpart of Lemma 40 for the Davies-Meyer instantiation. We then apply the results of Lemma 45 in the proofs of Lemma 41 and results from [24] to prove Theorem 7.

Proof. The proof follows from Lemma 43 and Theorem 36. \square

Lemma 43. *For every γ^* -IT-legitimate (q, ℓ) -attacker, in the ideal cipher model,*

$$\text{Adv}_{\text{DM}}^{\text{rec-IT}, \gamma^*}(\mathcal{A}) \leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + \frac{64\ell^4 2^r}{2^{2n-2}} + \frac{16\ell^2 q^2 2^r}{2^{2n-2}}} + \rho.$$

where $\rho = \frac{q^2}{2^r}$

Proof. The arguments for this proof is similar to the proof of Lemma 37. As before we define the random variables $\tilde{Y}_\ell, Z, \Sigma, U_r$. The difference arises in the definition of $Z = (\Sigma, \mathcal{L}, S_0)$. We define the \mathcal{L} as follows: If \mathcal{L}' is the set of queries made by \mathcal{A} to the ideal cipher E , then for every $((x, y), z) \in \mathcal{L}'$, add $((x, y), z \oplus y)$ to \mathcal{L} . This is to ensure that the \mathcal{L} is consistent with the evaluation of Davies-Meyer on the input (x, y) . We now upper-bound the adversary \mathcal{A} in the IT-recovering game by showing an upperbound for $\text{SD}((\tilde{Y}_\ell, Z, E), (U_r, Z, E))$.

We also define the event \mathcal{E} which denotes the event when \mathcal{L} entries are distinct *when truncated to the first r bits*. Note that there are a maximum of q queries in this list. We would like to point out that \mathcal{E} has the same probability of occurring in either experiment. Therefore, by Proposition 3,

$$\text{SD}((\tilde{Y}_\ell, Z, E), (U_r, Z, E)) \leq \text{SD}((\tilde{Y}_\ell, Z, E)|\mathcal{E}, (U_r, Z, E)|\mathcal{E}) + \frac{q^2}{2^r};$$

We let $\rho := q^2/2^r$ for the remainder of the proof. In order to bound the statistical distance conditioned on \mathcal{E} , we can rewrite the same as

$$\text{SD}((\tilde{Y}_\ell, Z, E)|\mathcal{E}, (U_r, Z, E)|\mathcal{E}) = \sum_{z \in \mathcal{E}} \text{P}[Z = z|\mathcal{E}] \cdot \text{SD}((\tilde{Y}_\ell, E)|z, (U_r, E)|z), \quad (9)$$

where $z \in \mathcal{E}$ is to denote that the sum is taken over all side informations $Z = z$, satisfying \mathcal{E} .²⁵ We define $p_z := \text{Pred}(\bar{X}|Z = z)$, and as shown in Lemma 37, we have

$$\text{H}_\infty(\bar{X}|Z\mathcal{E}) \geq \gamma^* - \log(1 - \rho)^{-1},$$

From Lemma 44 we get,

$$\text{SD}((\tilde{Y}_\ell, E)|z, (U_r, E)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + 64\ell^4 \frac{2^r}{2^{2n-2}} + 16\ell^2 q^2 \frac{2^r}{2^{2n-2}}}$$

Using Jensen's inequality, (9) becomes, for $\alpha = 2^r$ and $\beta = \ell \cdot d'(\ell) \cdot \frac{2^r}{2^{n-1}} + 64\ell^4 \cdot \frac{2^r}{2^{2n-2}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n-2}}$,

$$\begin{aligned} \text{SD}((\tilde{Y}_\ell, E)|\mathcal{E}, (U_r, E)|\mathcal{E}) &\leq \frac{1}{2} \sqrt{\alpha \text{Pred}(\bar{X}|\mathcal{L}\mathcal{E}) + \beta} \\ &\leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + 64\ell^4 \frac{2^r}{2^{2n-2}} + 16\ell^2 \frac{q^2 2^r}{2^{2n-2}}}. \end{aligned}$$

\square

²⁵Therefore \mathcal{E} can be omitted in the conditioning of the statistical distance.

Lemma 44. For $z \in \mathcal{E}$ and the random variables as defined earlier,

$$\text{SD}((\tilde{Y}_\ell, E)|z, (U_r, E)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + 64\ell^4 \frac{2^r}{2^{2n-2}} + 16\ell^2 q^2 \frac{2^r}{2^{2n-2}}}$$

Proof. We fix $z = (\sigma, L, s_0)$. We have that z, E is distributed uniformly over the set of all ideal ciphers that agree with L . Thus, by Proposition 2,

$$\text{SD}((\tilde{Y}_\ell, E)|z, (U_r, E)|z) \leq \frac{1}{2} \sqrt{2^r \cdot \text{Coll}(\tilde{Y}_\ell|Ez) - 1}. \quad (10)$$

To bound the collision probability, we consider the following experiment:

- choose E uniformly consistent with L
- sample inputs $\bar{X} = (\bar{X}_1, \dots, \bar{X}_\ell)$ and $\bar{X}' = (\bar{X}'_1, \dots, \bar{X}'_{\ell'})$ independently but conditioned on $Z = z$.
- compute \tilde{Y}_ℓ and $\tilde{Y}'_{\ell'}$ as the truncated MD evaluations with E of \bar{X} and \bar{X}'

Now, we have that:

$$\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \text{P}[\bar{X} = \bar{X}'] + \text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'} | \bar{X} \neq \bar{X}']. \quad (11)$$

Clearly, the former is at most p_z . To get a bound on the second term, we proceed to fix arbitrary inputs $\bar{x} \neq \bar{x}'$ of lengths ℓ and ℓ' , respectively. This is similar to the process that is adopted in the proof of Lemma 38. We construct a similar Structure Graph. Note that in Lemma 38, we assumed that the primitive was a compressing function. In particular, Davies Meyer is also a compressing function. Therefore, we can define similar structure graph for this construction by viewing $F(y, x) := E(x, y) \oplus y$. From Lemma 39 we get the following result:

$$\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \text{P}[G_F(\bar{x}, \bar{x}') \in \text{Coll}_L(\bar{x}, \bar{x}')] + \frac{1}{2^r}.$$

However, the proof of Lemma 40 changes. We prove a corresponding Lemma as follows:

Lemma 45. For an E sampled randomly consistent with L with F defined as above, and $\bar{x} \neq \bar{x}'$. Let $H \in \mathcal{G}_L(\bar{x}, \bar{x}')$. Then,

$$\text{P}[G_F(\bar{x}, \bar{x}') = H] = \frac{1}{2^{(n-1) \cdot \text{Acc}(H)}}.$$

Proof. We have two messages $\bar{x} \neq \bar{x}'$ and let $x^{(i)}$ be as defined before. We have a randomly sampled E which is consistent with L . The proof is similar to Lemma 40. However, the proof differs in the case of a colliding edge, i.e the edge causes an accident. In this case, $G_F^{(i+1)}$ will only be consistent if the edge corresponding to $x^{(i+1)}$ lands on the same vertex as in $H^{(i+1)}$. However, the evaluation is not a random function F which chooses a value, uniformly at random, from 2^n values. The evaluation is dependent on the definition of the ideal cipher E . The output in this case is chosen, from a minimum of $2^n - q - i$ values. In other words, $\text{P}[\text{Cons}_{i+1} | \text{Cons}_i] \leq \frac{1}{2^{n-q-i}} \leq \frac{1}{2^{(n-1)}}$ in this case. Note that the third case happens $\text{Acc}(H)$ times. Therefore, we have:

$$\text{P}[G_F(\bar{x}, \bar{x}') = H] = \text{P}[\text{Cons}_\ell] \leq \frac{1}{2^{(n-1) \cdot \text{Acc}(H)}}.$$

□

Applying this Lemma to the proof of Lemma 41 and the result from [24], we get:

$$\text{SD}((\tilde{Y}_\ell, E)|z, (U_r, E)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + 64\ell^4 \frac{2^r}{2^{2n-2}} + 16\ell^2 q^2 \frac{2^r}{2^{2n-2}}}.$$

□

7.6 IT PRNGs from Sponge

This section establishes the robustness of the PRNG construction $\text{Spg}_r = (\text{refresh}, \text{next})$ based on the Sponge Construction in the ideal permutation model. Recall that the construction is defined as follows (cf. Figure 5):

Construction 7 (IT-PRNG from Sponges). The Sponge-based PRNG construction $\text{Spg}_r = (\text{refresh}, \text{next})$ uses a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to absorb and produce r -bit inputs and outputs, respectively, and is defined as follows:

- $\text{refresh}^\pi(s, x) = \pi(s \oplus x \| 0^c)$, and
- $\text{next}^\pi(s) = (0^n, s[1..r])$.

Construction 7 achieves the following security:

Theorem 10 (IT-Robustness of Sponge PRNGs). *Construction 7 is a $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust PRNG in the π -model for*

$$\varepsilon_{\text{rob-it}} \leq \frac{t}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \frac{\ell \cdot (\ell + \tilde{q})}{2^{c-1}}} + t\rho,$$

for $\rho = \frac{\tilde{q}^2}{2^c}$ where $\tilde{q} = q + t\ell$

In this section, we take a look at the proof of robustness for the IT-PRNG based on the Sponge construction.

Lemma 46. *For every γ^* -IT-legitimate (q, ℓ) -attacker, in the ideal permutation model,*

$$\text{SD}((\tilde{Y}_\ell, \pi)|\mathcal{E}, (U_r, \pi)|\mathcal{E}) \leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \frac{\ell \cdot (\ell + q)}{2^{c-1}}} + \rho.$$

where $\rho = \frac{q^2}{2^c}$

Proof (of Lemma 46). We define a few random variables which we will be using in our proofs.

- π : a randomly chosen permutation, to which the adversary is given access. We use π both for the oracle itself, as well as for the random variable describing the entire function table.
- ℓ : Number of blocks input to the challenge oracle (which is a random variable itself, we overload notation here, using the same letter we use in the bound on ℓ in the lemma statement).
- $\bar{X} = (\bar{X}_1, \dots, \bar{X}_\ell)$: the blocks input to the challenge oracle.
- \tilde{Y}_ℓ : output of MD_r . Let us remind ourselves that this is s truncated to r bits (the output of next);
- $Z = (\Sigma, \mathcal{L}, S_0)$: “side information” where Σ is the attacker state before challenge, \mathcal{L} is the attacker query/answers to π before challenge, S_0 is the initial PRNG state provided by \mathcal{A} ;
- U_r : uniform r -bit string.

The advantage of the adversary \mathcal{A} in the recovering game is bounded by $\text{SD}((\tilde{Y}_\ell, Z, \pi), (U_r, Z, \pi))$. Therefore, it is sufficient to upper-bound that. The reasoning is quite simple. Note that Z contains the state of the attacker Σ just before it makes the challenge query. This means that \mathcal{A} cannot tell apart real from random. For ease of this discussion we define the idea of supernode. State nodes have values in $\{0, 1\}^n$. They are clustered to supernodes where a super node with label in $\{0, 1\}^c$ contains all state nodes having that same values in the last c bits. Therefore, there are 2^c supernodes with each having 2^r state nodes.

We also define an event \mathcal{E} where the queries in L land do not collide on the last c bits of other elements. This is true even for inverse queries, i.e $\pi^{-1}(v)$ should result in a new supernode. We would like to point out that \mathcal{E} has the same probability of occurring in either experiment, since the experiments are identical up to the point when this event is defined. Therefore, by Proposition 3,

$$\text{SD}((\tilde{Y}_\ell, Z, \pi), (U_r, Z, \pi)) \leq \text{SD}((\tilde{Y}_\ell, Z, \pi)|\mathcal{E}, (U_r, Z, \pi)|\mathcal{E}) + \frac{q^2}{2^c};$$

For convenience we let $\rho := \frac{q^2}{2^c}$ for the remainder of the proof. In order to bound the statistical distance conditioned on \mathcal{E} , we can rewrite the same as as

$$\text{SD}((\tilde{Y}_\ell, Z, \pi)|\mathcal{E}, (U_r, Z, \pi)|\mathcal{E}) = \sum_{z \in \mathcal{E}} \text{P}[Z = z|\mathcal{E}] \cdot \text{SD}((\tilde{Y}_\ell, \pi)|z, (U_r, \pi)|z), \quad (12)$$

where $z \in \mathcal{E}$ is to denote that the sum is taken over all side informations $Z = z$, satisfying \mathcal{E} .²⁶ Define $p_z := \text{Pred}(\bar{X}|Z = z)$, and observe that

$$\mathbf{E}_z[p_z] = \text{Pred}(\bar{X}|Z) \leq 2^{-\gamma^*},$$

where the latter inequality follows from the assumption $H_\infty(\bar{X}|Z) \geq \gamma^*$. Moreover,

$$H_\infty(\bar{X}|Z\mathcal{E}) \geq \gamma^* - \log(1 - \rho)^{-1},$$

which is due to

$$\begin{aligned} \text{Pred}(\bar{X}|Z) &\geq \sum_{z \in \mathcal{E}} \text{P}[Z = z] \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \sum_{z \in \mathcal{E}} \frac{\text{P}[Z = z]}{\text{P}[\mathcal{E}]} \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \text{Pred}(\bar{X}|Z\mathcal{E}). \end{aligned}$$

From Lemma 47 we prove below, we will get,

$$\text{SD}((\tilde{Y}_\ell, \pi)|z, (U_r, \pi)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \frac{\ell \cdot (\ell + q)}{2^{c-1}}}.$$

Using Jensen's inequality, (12) becomes, for

$$\begin{aligned} \text{SD}((\tilde{Y}_\ell, \pi)|\mathcal{E}, (U_r, \pi)|\mathcal{E}) &\leq \frac{1}{2} \sqrt{2^r \text{Pred}(\bar{X}|Z\mathcal{E}) + \frac{\ell \cdot (\ell + q)}{2^{c-1}}} \\ &\leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \frac{\ell \cdot (\ell + q)}{2^{c-1}}}. \end{aligned}$$

□

²⁶Therefore \mathcal{E} can be omitted in the conditioning of the statistical distance.

Lemma 47. For $z \in \mathcal{E}$ and the random variables as defined earlier,

$$\text{SD}((\tilde{Y}_\ell, \pi)|z, (U_r, \pi)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \frac{\ell \cdot (\ell + q)}{2^{c-1}}}.$$

Proof. We fix $z = (\sigma, L, s_0)$. We have that z, π is distributed uniformly over the set of all ideal permutations that agree with L . Thus, by Proposition 2,

$$\text{SD}((\tilde{Y}_\ell, \pi)|z, (U_r, \pi)|z) \leq \frac{1}{2} \sqrt{2^r \cdot \text{Coll}(\tilde{Y}_\ell|\pi z) - 1}. \quad (13)$$

We consider the following experiment to bound the collision probability:

- choose π uniformly consistent with L
- sample inputs $\bar{X} = (\bar{X}_1, \dots, \bar{X}_\ell)$ and $\bar{X}' = (\bar{X}'_1, \dots, \bar{X}'_{\ell'})$ independently but conditioned on $Z = z$.
- compute \tilde{Y}_ℓ and $\tilde{Y}'_{\ell'}$ as the Sponge evaluation of \bar{X} and \bar{X}'

Now, we have that:

$$\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \text{P}[\bar{X} = \bar{X}'] + \text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'} | \bar{X} \neq \bar{X}']. \quad (14)$$

By definition, we have that the former term is p_z . We take a closer look at the latter term. We fix arbitrary inputs \bar{x} and \bar{x}' . Let the length of them be ℓ and ℓ' respectively. We also assume, wlog, that the evaluation of \bar{x}' is not completely covered by L ; due to the collision-freeness of L . Let \bar{x}_{k+1} be the first block of \bar{x} not covered by \mathcal{L} and similarly \bar{x}'_{k+1} for \bar{x}' . We let $k = \ell$ if all blocks are covered. We then apply Lemma 48 to conclude the proof. \square

Lemma 48. Let π be sampled randomly consistent with L , and $\bar{x} \neq \bar{x}'$. Then, let \tilde{Y}_ℓ and $\tilde{Y}'_{\ell'}$ be values obtained after truncating at the end of the Sponge evaluations on inputs \bar{x} and \bar{x}' respectively. Then,

$$\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \frac{\tilde{\ell}(q + \tilde{\ell})}{2^{n-1}} + \frac{1}{2^r}.$$

Proof. We denote by $\text{Coll}_L(\bar{x}, \bar{x}')$ the event that the sponge evaluation on inputs \bar{x}, \bar{x}' collide at the final state node, for a permutation π sampled consistently with L . We rewrite $\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}]$ as:

$$\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \text{P}[\text{Coll}_L(\bar{x}, \bar{x}')] + \text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'} | \overline{\text{Coll}_L(\bar{x}, \bar{x}')}]$$

We take a closer look at the latter term. Note that the evaluation of the inputs is fixed and it has no collision at the end nodes. We proceed to assign random, yet distinct values to the final state nodes. These are chosen from $\{0, 1\}^n$. Note that it is sufficient to look at the two output vertices *locally* without looking at the global state. There are $2^n(2^n - 1)$ pairs of values for these output vertices such that these values are distinct. However, of these, $2^r 2^{n-r}(2^{n-r} - 1)$ have values which are equal in the first r bits and yet are distinct n -bit strings. Therefore,

$$\begin{aligned} \text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'} | \overline{\text{Coll}_L(\bar{x}, \bar{x}')}] &\leq \frac{2^r 2^{n-r}(2^{n-r} - 1)}{2^n(2^n - 1)} \\ &= \frac{2^{n-r} - 1}{2^n - 1} \\ &\leq \frac{2^{n-r}}{2^n} = \frac{1}{2^r}. \end{aligned}$$

Putting it together, we have:

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell] \leq \mathbb{P}[\text{Coll}_L(\bar{x}, \bar{x}')] + \frac{1}{2^r} . \quad (15)$$

Now we bound the value of $\mathbb{P}[\text{Coll}_L(\bar{x}, \bar{x}')] .$ Consider the intermediate evaluations starting from $k + 1$. These are the inputs not covered by the list L . Define them as: $U_{k+1}, V_{k+1}, \dots, U_{\tilde{\ell}}, V_{\tilde{\ell}}$ such that:

$$U_{k+1} = s_k \oplus \bar{x}_{k+1} || 0^c$$

and for $i = k + 2, \dots, \tilde{\ell}$,

$$U_i = V_{i-1} \oplus x_i || 0^c .$$

In addition, for $i = k + 1, \dots, \tilde{\ell}$

$$V_i = \pi(U_i) .$$

For $i = k + 1, \dots, \tilde{\ell}$, denote by FRESH_i be the event U_i is fresh, i.e, $\pi(U_i)$ has not yet been defined. It is clear that FRESH_{k+1} is true. Let

$$\text{FRESH} = \bigcap_{i=k+1}^{\tilde{\ell}-1} \text{FRESH}_i$$

Then,

$$\mathbb{P}[\text{Coll}_L(\bar{x}, \bar{x}')] \leq \mathbb{P}[\text{Coll}_L(\bar{x}, \bar{x}') | \text{FRESH}] + \mathbb{P}[\overline{\text{FRESH}}]$$

Clearly the former term is,

$$\mathbb{P}[\text{Coll}_L(\bar{x}, \bar{x}') | \text{FRESH}] \leq \frac{1}{2^n - \tilde{\ell} - q}$$

Given that FRESH_{i-1} is true, this implies that $V_{i-1} = \pi(U_{i-1})$ was freshly chosen. Therefore, for FRESH_i to be false, V_{i-1} must have been chosen such that $U_i = V_{i-1} \oplus x_i$ is not fresh. Remember that, V_{i-1} would have been uniformly sampled from at least $2^n - q - \tilde{\ell}$ values. Therefore,

$$\mathbb{P}[\overline{\text{FRESH}}_i | \bigcap_{j=k+1}^{i-1} \text{FRESH}_j] \leq \frac{(q + i)}{2^n - \tilde{\ell} - q} .$$

In other words, $\mathbb{P}[\overline{\text{FRESH}}] \leq \frac{(\tilde{\ell}-1)(q+\tilde{\ell}-1)}{2^{n-1}}$. □

References

- [1] Boaz Barak and Shai Halevi. A model and architecture for pseudo-random generation with applications to /dev/random. In Vijayalakshmi Atluri, Catherine Meadows, and Ari Juels, editors, *ACM CCS 05*, pages 203–212, Alexandria, Virginia, USA, November 7–11, 2005. ACM Press.
- [2] Boaz Barak, Russell Impagliazzo, and Avi Wigderson. Extracting randomness using few independent sources. In *45th FOCS*, pages 384–393, Rome, Italy, October 17–19, 2004. IEEE Computer Society Press.
- [3] Boaz Barak, Ronen Shaltiel, and Eran Tromer. True random number generators secure in a changing environment. In Colin D. Walter, Çetin Kaya Koç, and Christof Paar, editors, *CHES 2003*, volume 2779 of *LNCS*, pages 166–180, Cologne, Germany, September 8–10, 2003. Springer, Heidelberg, Germany.

- [4] Elaine Barker and John Kelsey. NIST Special Publication 800-90A (A revision of SP 800-90) Recommendation for random number generation using deterministic random bit generators. <https://csrc.nist.gov/publications/detail/sp/800-90a/rev-1/final>, 2012.
- [5] Elaine Barker and John Kelsey. Recommendation for random number generation using deterministic random bit generators. NIST Special Publication 800-90A, 2012.
- [6] Mihir Bellare, Viet Tung Hoang, and Sriram Keelveedhi. Instantiating random oracles via UCEs. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 398–415, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [7] Mihir Bellare, Krzysztof Pietrzak, and Phillip Rogaway. Improved security analyses for CBC MACs. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 527–545, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Heidelberg, Germany.
- [8] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the indifferentiability of the sponge construction. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 181–197, Istanbul, Turkey, April 13–17, 2008. Springer, Heidelberg, Germany.
- [9] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Sponge-based pseudo-random number generators. In Stefan Mangard and François-Xavier Standaert, editors, *CHES 2010*, volume 6225 of *LNCS*, pages 33–47, Santa Barbara, CA, USA, August 17–20, 2010. Springer, Heidelberg, Germany.
- [10] Manuel Blum. Independent unbiased coin flips from a correlated biased source—a finite state markov chain. *Combinatorica*, 6(2):97–108, 1986.
- [11] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 670–683, Cambridge, MA, USA, June 18–21, 2016. ACM Press.
- [12] Stephen Checkoway, Ruben Niederhagen, Adam Everspaugh, Matthew Green, Tanja Lange, Thomas Ristenpart, Daniel J. Bernstein, Jake Maskiewicz, Hovav Shacham, and Matthew Fredrikson. On the practical exploitability of dual EC in TLS implementations. In *Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, August 20-22, 2014.*, pages 319–335, 2014.
- [13] Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
- [14] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity (extended abstract). In *26th FOCS*, pages 429–442, Portland, Oregon, October 21–23, 1985. IEEE Computer Society Press.
- [15] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.
- [16] Benny Chor, Oded Goldreich, Johan Håstad, Joel Friedman, Steven Rudich, and Roman Smolensky. The bit extraction problem of t-resilient functions (preliminary version). In *26th FOCS*, pages 396–407, Portland, Oregon, October 21–23, 1985. IEEE Computer Society Press.

- [17] Yevgeniy Dodis, Rosario Gennaro, Johan Håstad, Hugo Krawczyk, and Tal Rabin. Randomness extraction and key derivation using the CBC, cascade and HMAC modes. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 494–510, Santa Barbara, CA, USA, August 15–19, 2004. Springer, Heidelberg, Germany.
- [18] Yevgeniy Dodis, David Pointcheval, Sylvain Ruhault, Damien Vergnaud, and Daniel Wichs. Security analysis of pseudo-random number generators with input: /dev/random is not robust. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 647–658, Berlin, Germany, November 4–8, 2013. ACM Press.
- [19] Yevgeniy Dodis, Thomas Ristenpart, and Salil P. Vadhan. Randomness condensers for efficiently samplable, seed-dependent sources. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 618–635, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Heidelberg, Germany.
- [20] Yevgeniy Dodis, Adi Shamir, Noah Stephens-Davidowitz, and Daniel Wichs. How to eat your entropy and have it too - optimal recovery strategies for compromised RNGs. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 37–54, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [21] D. Eastlake, J. Schiller, and S. Crocker. *RFC 4086 - Randomness Requirements for Security*, June 2005.
- [22] Niels Ferguson. Private communication, 2013.
- [23] Niels Ferguson and Bruce Schneier. *Practical Cryptography*. John Wiley & Sons, Inc., New York, NY, USA, 1 edition, 2003.
- [24] Peter Gazi, Krzysztof Pietrzak, and Michal Rybár. The exact PRF-security of NMAC and HMAC. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 113–130, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [25] Peter Gazi and Stefano Tessaro. Provably robust sponge-based PRNGs and KDFs. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 87–116, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- [26] Danny Harnik and Moni Naor. On everlasting security in the hybrid bounded storage model. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006, Part II*, volume 4052 of *LNCS*, pages 192–203, Venice, Italy, July 10–14, 2006. Springer, Heidelberg, Germany.
- [27] Daniel Hutchinson. A robust and sponge-like PRNG with improved efficiency. In Roberto Avanzi and Howard M. Heys, editors, *SAC 2016*, volume 10532 of *LNCS*, pages 381–398, St. John’s, NL, Canada, August 10–12, 2016. Springer, Heidelberg, Germany.
- [28] Daniel Hutchinson. A robust and sponge-like PRNG with improved efficiency. Cryptology ePrint Archive, Report 2016/886, 2016. <http://eprint.iacr.org/2016/886>.
- [29] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24, Seattle, WA, USA, May 15–17, 1989. ACM Press.

- [30] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th FOCS*, pages 230–235, Research Triangle Park, North Carolina, October 30 – November 1, 1989. IEEE Computer Society Press.
- [31] Information technology - Security techniques - Random bit generation. ISO/IEC18031:2011, 2011.
- [32] Jesse Kamp, Anup Rao, Salil P. Vadhan, and David Zuckerman. Deterministic extractors for small-space sources. *J. Comput. Syst. Sci.*, 77(1):191–220, 2011.
- [33] John Kelsey, Bruce Schneier, and Niels Ferguson. Yarrow-160: Notes on the design and analysis of the yarrow cryptographic pseudorandom number generator. In *In Sixth Annual Workshop on Selected Areas in Cryptography*, pages 13–33. Springer, 1999.
- [34] Killmann, W. and Schindler, W. A proposal for: Functionality classes for random number generators. AIS 20 / AIS31, 2011.
- [35] Hugo Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 631–648, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Heidelberg, Germany.
- [36] David Lichtenstein, Nathan Linial, and Michael E. Saks. Some extremal problems arising from discrete control processes. *Combinatorica*, 9(3):269–287, 1989.
- [37] John M. Intel digital random number generator (DRNG) software implementation guide. <https://software.intel.com/en-us/articles/intel-digital-random-number-generator-drng-software-implementation-guide>, 2014.
- [38] Noam Nisan and David Zuckerman. More deterministic simulation in logspace. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 235–244, 1993.
- [39] Noam Nisan and David Zuckerman. Randomness is linear in space. *J. Comput. Syst. Sci.*, 52(1):43–52, 1996.
- [40] Jacques Patarin. The “coefficients H” technique (invited talk). In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *SAC 2008*, volume 5381 of *LNCS*, pages 328–345, Sackville, New Brunswick, Canada, August 14–15, 2009. Springer, Heidelberg, Germany.
- [41] Berry Schoenmakers and Andrey Sidorenko. Cryptanalysis of the dual elliptic curve pseudorandom generator. Cryptology ePrint Archive, Report 2006/190, 2006. <http://eprint.iacr.org/2006/190>.
- [42] Thomas Shrimpton and R. Seth Terashima. A provable-security analysis of Intel’s secure key RNG. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 77–100, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [43] Dan Shumow and Niels Ferguson. On the possibility of a back door in the nist sp800-90 dual ec prng. *CRYPTO Rump Session*, 2007.
- [44] Pratik Soni and Stefano Tessaro. Public-seed pseudorandom permutations. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 412–441, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.

- [45] Luca Trevisan and Salil P. Vadhan. Extracting randomness from samplable distributions. In *41st FOCS*, pages 32–42, Redondo Beach, CA, USA, November 12–14, 2000. IEEE Computer Society Press.
- [46] John von Neumann. Various techniques used in connection with random digits. In A.S. Householder, G.E. Forsythe, and H.H. Germond, editors, *Monte Carlo Method*, pages 36–38. National Bureau of Standards Applied Mathematics Series, 12, Washington, D.C.: U.S. Government Printing Office, 1951.
- [47] Wikipedia. /dev/random. <http://en.wikipedia.org/wiki//dev/random>, 2004. [Online; accessed 09-February-2014].
- [48] Joanne Woodage and Dan Shumow. An analysis of the NIST SP 800-90A standard. In *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19 - May 23, 2019*.

A Comparison to Previous PRNG Security Notions

Overview. The first robustness notion for PRNGs was proposed by Barak and Halevi [1] and required that PRNGs be:

- *resilient*, i.e., the attacker cannot predict the output of the PRNG even if he can influence the inputs;
- *forward secure*, i.e., upon state compromise, previous outputs of the PRNG remain random; and
- *post-compromise secure*, i.e., if sufficient entropy is absorbed into the PRNG state, the outputs regain security.

The definition was developed further in a crucial way by Dodis *et al.* [18], who considered the setting where entropy only *gradually* accumulates in the state (as opposed to the model in [1], where inputs are required to have high entropy in order to recover from compromise. In order to analyse constructions such as sponges, Gazi and Tessaro [25] extended the robustness definition further to a setting with a public random permutation.

Distribution samplers. In order to capture distributions for which the PRNG is expected to be robust, Dodis *et al.* [18] considered distribution samplers Sam , which with every input x_i for the PRNG would also provide an entropy estimate γ_i and a leakage value z_i . A sampler was considered *legitimate* if every input x_i has the promised min entropy γ_i even conditioned on all *other* inputs (i.e., past and future) as well as on *all* values γ_j and z_j . A very important consequence of this definition is that (it can be shown that) extraction with such a sampler Sam is impossible without a randomly chosen seed about which the sampler has no information. Dodis *et al.* provided a construction that uses such a seed and whose refresh function actually is information-theoretically secure. However, security completely breaks down as soon as the sampler can depend on the seed, which is why said construction was deemed unsuitable by practitioners since seed-independence of the input may not be guaranteed in practice.

In an effort to analyze more practical constructions Gazi and Tessaro [25] extended the robustness notion to incorporate ideal primitives. In particular, for their analysis of Sponge, they considered samplers Sam^π with oracle access to a random permutation π and required an entropy notion related to that above, but required, informally, that inputs be unpredictable even if in addition to the above, one also conditions on the queries by the sampler (albeit the sampler is allowed *private* queries for every input it generates). While the extension to ideal models allows to analyze constructions used in practice, the work by Gazi and Tessaro does not solve the need for a seed (about which Sam^π has no information), without which the aforementioned impossibility still holds.

Main differences between existing notions and the new one. The new definition put forth in this work differs from previous ones in several ways: First, the legitimacy notion is such that there is no need for a seed. Second, because of that, the sampler and the distinguisher are merged into a single attacker, which would not be possible with the previous notions due to the need to hide the seed from Sam . Third, the new legitimacy notion measures not the sum of entropies of particular blocks conditioned on each other, but the entropy of an entire sequence of inputs *as a whole*. Finally, the new notion dispenses with the need for entropy estimates by simply only considering attackers that provide inputs for which the average min-entropy at certain points is high enough conditioned on the state and the queries of the attacker. Observe that the entropy estimates

were originally introduced to circumvent the impossibility of randomness estimation, which would have been needed in reductions to computationally secure primitives (such as using a PRG for the next function). Hence, entropy estimates were an unnatural artifact of the definition anyway.

B Constructions of Online Extractors

This section presents four simple, intuitive, and—most importantly—practical seedless online-extractor constructions:

- a construction based on the *Merkle-Damgård paradigm* using a public *fixed-length compression function*;
- a construction based on the *Merkle-Damgård paradigm* using the *Davies-Meyer compression function* (as in SHA-2), which is built from any public block cipher; and
- a construction based on the *Sponge paradigm* (as in SHA-3), which uses a public permutation.
- a construction based on the HMAC function, which is a variant of the Merkle-Damgård paradigm, using a public *fixed-length compression function*.

The four resulting online extractor constructions are pictorially represented in Figure 1. For all these 4 variants we actually present two constructions: one achieving computational security, and one achieving information-theoretic (IT) security. Except for the sponges, where the computational and IT variants are syntactically identical, for the other three setting the IT variant simply truncates the computational variant to the first r bits.

B.1 Extractors from Merkle-Damgård

An online extractor can be obtained from a compression function F as follows:²⁷

Construction 8 (Online extractor from Merkle-Damgård). The (m, n) -online extractor construction $\text{MD} = (\text{refresh}, \text{finalize})$ based on Merkle-Damgård with a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined as follows:

- $\text{refresh}^F(s, x) = F(s, x)$, and
- $\text{finalize}^F(s) = s$.

The security of Construction 8 is proved in the F -model, where F is a uniformly random function. The theorem below is identical to Lemma 15.

Theorem 49 (Online extractor from Merkle-Damgård). *Construction 8 is a $(\gamma^*, q, \varepsilon)$ -online extractor in the F -model for*

$$\varepsilon \leq \frac{q^2 + q\ell + \ell^2}{2^n} + \frac{2q}{2^{\gamma^*}}.$$

An IT-secure online extractor based on Merkle-Damgård can be obtained if the finalize function simply truncates the state:

²⁷To reduce notational clutter, the algorithms `refresh` and `finalize` of the extractor constructions are not “branded” with the design name. There will be no ambiguity as to which construction is meant in any place in this paper.

Construction 9 (IT online extractor from Merkle-Damgård). The (m, r) -IT-online extractor construction $\text{MD}_r = (\text{refresh}, \text{finalize})$ based on Merkle-Damgård with a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined as follows:

- $\text{refresh}^F(s, x) = F(s, x)$, and
- $\text{finalize}^F(s) = s[1..r]$.

The security of Construction 9 is proved in the F -model, where F is a uniformly random function. To state the theorem for the IT construction, for an integer ℓ , let

$$d'(\ell) = \max_{\ell' \in \{1, \dots, \ell\}} |\{d \in \mathbb{N} : d|\ell'\}|.$$

Observe that, asymptotically, $d'(\ell)$ grows very slowly, i.e., as $\ell^{o(1)}$. Furthermore, let F be a random compression function. The following theorem is equivalent to Lemma 37.

Theorem 50 (IT online extractor from Merkle-Damgård). *Construction 9 is a $(\gamma^*, q, \varepsilon)$ -IT-online extractor in the F -model, where*

$$\varepsilon \leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 \frac{q^2 2^r}{2^{2n}}} + \rho,$$

where $\rho = \frac{q^2}{2^r}$.

B.2 Extractors from Merkle-Damgård with Davies-Meyer

The Davies-Meyer compression function maps two inputs $a \in \{0, 1\}^m$ and $b \in \{0, 1\}^n$ to an n -bit string

$$E(b, a) \oplus a,$$

where E is an arbitrary block cipher (where b is the key and a the input).²⁸ Correspondingly, an online extractor can be obtained from E as follows:

Construction 10 (Online extractor from MD-DM). The (k, n) -online extractor construction $\text{DM} = (\text{refresh}, \text{finalize})$ based on Merkle-Damgård with Davies-Meyer (MD-DM) uses a cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and is defined as follows:

- $\text{refresh}^E(s, x) = E(x, s) \oplus s$, and
- $\text{finalize}^E(s) = s$.

The security of Construction 10 is proved in the E -model, where E is a cipher chosen uniformly at random from the set of all ciphers and can be queried in both the forward and backward direction. The theorem below is identical to Lemma 22

Theorem 51 (Online extractor from MD-DM). *Construction 10 is a $(\gamma^*, q, \varepsilon_{\text{rob}})$ -robust online extractor in the E -model for*

$$\varepsilon_{\text{rob}} \leq \frac{q^2 + 2(q\ell + \ell^2)}{2^n} + \frac{4q}{2^{\gamma^*}}.$$

²⁸A (block) cipher is an efficiently computable and invertible permutation $E(k, \cdot) : \{0, 1\}^n \rightarrow \{0, 1\}^n$ for every key $k \in \{0, 1\}^k$.

An IT-secure online extractor based on MD-DM can be obtained if the finalize function simply truncates the state:

Construction 11 (IT online extractor from MD-DM). The (k, r) -IT-online construction $\text{DM}_r = (\text{refresh}, \text{finalize})$ using Merkle-Damgård with Davies-Meyer (MD-DM) uses a block cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and is defined as follows:

- $\text{refresh}^E(s, x) = E(x, s) \oplus s$, and
- $\text{finalize}^E(s) = s[1..r]$.

The security of Construction 11 is proved in the E -model, where E is a cipher chosen uniformly at random from the set of all ciphers and can be queried in both the forward and backward direction. Let $d'(\ell)$ be defined as in Section 5.1. The following theorem is equivalent to Lemma 43.

Theorem 52 (IT online extractor from MD-DM). *Construction 11 is a $(\gamma^*, q, \varepsilon_{\text{rob}})$ -IT-online extractor in the E -model, where*

$$\text{Adv}_{\text{DM}}^{\text{rec-IT}, \gamma^*}(\mathcal{A}) \leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \frac{2^r}{2^{n-1}} + \frac{64\ell^4 2^r}{2^{2n-2}} + \frac{16\ell^2 q^2 2^r}{2^{2n-2}} + \rho}.$$

where $\rho = \frac{q^2}{2^r}$.

B.3 Extractors from Sponges

Let $n \in \mathbb{N}$ and $n = r + c$. In the following, for an n -bit string s , let $s = s^{(r)} \| s^{(c)}$ be decomposition of s into an r -bit and c -bit string. An online extractor using the Sponge paradigm can be obtained from a permutation π as follows:

Construction 12 (Online extractor from Sponges). The Sponge-based online-extractor construction $\text{Spg} = (\text{refresh}, \text{finalize})$ uses a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to absorb and produce r -bit inputs and outputs, respectively, and is defined as follows:

- $\text{refresh}^\pi(s, x) = \pi(s \oplus x \| 0^c)$, and
- $\text{finalize}^\pi(s) = s[1..r]$.

Observe that for Sponge-based extractors, even the computational variant needs to truncate the state, otherwise the output of the extractor could be inverted by the attacker, which renders the constructions insecure.

The security of Construction 12 is proved in the π -model, where π is a uniformly random permutation, which can be queried in both the forward and backward direction. The proof of the following theorem follows along similar lines as that of Lemma 34.

Theorem 53 (Online extractor from Sponges). *Construction 12 is a $(\gamma^*, q, \ell, \varepsilon)$ -online extractor in the π -model for*

$$\varepsilon \leq 2 \cdot \left(\frac{q + q\ell + \ell^2}{2^n} + \frac{q^2}{2^c} + \frac{q}{2^{\gamma^*}} \right).$$

Observe that the bound in Theorem 9 is only reasonable when c is large enough, which matches the fact that CBC-based online extractors—which correspond to the case $c = 0$, are not secure.

As pointed out before, the IT-secure online extractor based on Sponge would be similar to the computational variant. We reproduce the definition for completeness.

Construction 13 (IT Online extractor from Sponges). The Sponge-based IT-online-extractor construction $\text{Spg} = (\text{refresh}, \text{finalize})$ uses a permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to absorb and produce r -bit inputs and outputs, respectively, and is defined as follows:

- $\text{refresh}^\pi(s, x) = \pi(s \oplus x \| 0^c)$, and
- $\text{finalize}^\pi(s) = s[1..r]$.

The security of Construction 13 is proved in the π -model, where π is an ideal permutation chosen uniformly at random from the set of all permutations and can be queried in both the forward and backward direction. The following theorem is equivalent to Lemma 46.

Theorem 54 (IT online extractor from Sponges). *Construction 13 is a $(\gamma^*, q, \varepsilon_{\text{rob}})$ -IT-online extractor in the π -model, where*

$$\text{Adv}_{\text{Spg}}^{\text{rec-IT}, \gamma^*}(\mathcal{A}) \leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \frac{\ell \cdot (\ell + q)}{2^{c-1}}} + \rho .$$

where $\rho = \frac{q^2}{2^c}$.

B.4 Extractors from HMAC

In practice, uniformly random key material is often derived from high-entropy inputs (resulting, e.g., from a key-agreement protocol) using a *key-derivation function (KDF)*. A common paradigm, suggested by Krawczyk [35], to construct KDFs is to combine an extractor with a variable-length pseudorandom function (VL-PRF). The most widely used KDF is *HKDF*, which uses the HMAC mode of operation for compression functions (CF) to instantiate both the extractor and the VL-PRF. This section considers the security of HMAC as a *seedless extractor* w.r.t. the new legitimacy condition put forth by this work. Together with a VL-PRF, the seedless HMAC extractor can then be used to justify security for the seedless variant of HKDF. The latter extension is discussed in Section 3.5.

HMAC. HMAC is similar to Merkle-Damgård, but requires additional CF calls, designed to prevent extension attacks when HMAC is used as a PRF. Concretely, for a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$, the HMAC construction takes a key $k \in \{0, 1\}^m$ as well as inputs $x_1, \dots, x_\ell \in \{0, 1\}^m$ and outputs²⁹

$$\text{HMAC}(k, x_1, \dots, x_\ell) := \text{MD}^F(k \oplus \text{opad}, \text{MD}^F(k \oplus \text{ipad}, x_1, \dots, x_\ell)) ,$$

where $\text{ipad} \neq \text{opad} \in \{0, 1\}^m$ are arbitrary constants. HMAC can be used as seedless extractor by fixing, say, $k := 0$. That is, the online extractor based on HMAC would be defined as follows:

Construction 14 (Online extractor from HMAC.). The (m, n) -online extractor construction $\text{HMAC} = (\text{refresh}, \text{finalize})$ based on HMAC with a compression function $F : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ is defined as follows:

- the initial state is $s_0 = F(0, \text{ipad})$;
- $\text{refresh}^F(s, x) = F(s, x)$; and

²⁹The IV to the Merkle-Damgård construction is 0.

- $\text{finalize}^F(s) = F(F(0, \text{opad}), s)$.

The security of Construction 8 is proved in the F -model, where F is a uniformly random function.

Theorem 55 (Online extractor from HMAC). *Construction 14 is a $(\gamma^*, q, \varepsilon)$ -online extractor in the F -model for*

$$\varepsilon \leq \frac{q^2 + q + q\ell + \ell^2}{2^n} + \frac{2q}{2^{\gamma^*}}.$$

Proof (sketch). To prove the security of HMAC as an extractor according to Definition 3, one first considers a hybrid experiment, in which the output of the extractor is computed as

$$F(F(0, \text{opad}), U),$$

for a value chosen uniformly at random. The indistinguishability of the original extraction game and the hybrid experiment is established via a simple hybrid argument: essentially, the reduction (to the security of MD as an extractor) simply prepends a block ipad to the inputs x_1, \dots, x_ℓ , and upon receiving a value y , it additionally computes $F(F(0, \text{opad}), y)$. It is easily seen that adding ipad to the input does not affect the conditional entropy of the input—and therefore the legitimacy of the reduction.

Finally, it is easily seen that in the hybrid experiment, the advantage of any attacker is zero unless it queries $F(F(0, \text{opad}), U)$, which happens with probability at most $q/2^n$. \square

An IT-secure online extractor based on Merkle-Damgård can be obtained if the finalize function additionally truncates the output:

Construction 15 (IT online extractor from HMAC.). The (m, r) -IT-online extractor construction $\text{HMAC}' = (\text{refresh}, \text{finalize})$ based on HMAC with a compression function $F : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is defined as follows:

- the initial state is $s_0 = F(0, \text{ipad})$;
- $\text{refresh}^F(s, x) = F(s, x)$; and
- $\text{finalize}^F(s) = F(F(0, \text{opad}), s)[1..r]$.

The security of Construction 15 is proved in the F -model, where F is a uniformly random function. To state the theorem for the IT construction, for an integer ℓ , let

$$d'(\ell) = \max_{\ell' \in \{1, \dots, \ell\}} |\{d \in \mathbb{N} : d|\ell'\}|.$$

Observe that, asymptotically, $d'(\ell)$ grows very slowly, i.e., as $\ell^{o(1)}$. Furthermore, let F be a random compression function. The following theorem is equivalent to Lemma 37.

Theorem 56 (IT online extractor from HMAC). *Construction 15 is a $(\gamma^*, q, \varepsilon)$ -IT-online extractor in the F -model, where*

$$\varepsilon \leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}} + (q + \ell + 2) \cdot \frac{2^r}{2^n} + \rho},$$

where $\rho = \frac{q^2}{2^n}$.

Proof. We define a few random variables which we will be using in our proofs.

- F a randomly chosen compression function, to which the adversary is given access. We use F both for the oracle itself, as well as for the random variable describing the entire function table.
- ℓ : Number of blocks input to the challenge oracle (which is a random variable itself, we overload notation here, using the same letter we use in the bound on ℓ in the lemma statement).
- $\bar{X} = (\bar{X}_1, \dots, \bar{X}_\ell)$: the blocks input to the challenge oracle.
- \tilde{Y}_ℓ : output of HMAC'. Let us remind ourselves that this is s truncated to r bits (the output of finalize);
- $Z = (\Sigma, \mathcal{L}, S_0)$: “side information” where Σ is the attacker state before challenge, \mathcal{L} is the attacker query/answers to F before challenge, S_0 is the initial state provided by \mathcal{A} ;
- U_r : uniform r -bit string.
- S : the state of the extractor. For the purposes of this proof, we will be using S to indicate the state of the extractor that is the input to finalize.

The advantage of the adversary \mathcal{A} in the online extractor game is bounded by $\text{SD}((\tilde{Y}_\ell, Z, F), (U_r, Z, F))$. Therefore, it is sufficient to upper-bound just that. The reason follows from the fact that Z contains the state of the attacker Σ just before it makes the challenge query. This means that \mathcal{A} cannot tell apart real from random.

Much like the earlier proofs, we define an event \mathcal{E} where the answers to the F -queries by \mathcal{A} are distinct *when truncated to the first r bits*. Note that there are a maximum of q queries in this list. As pointed out earlier \mathcal{E} has the same probability of occurring in either experiment, since the experiments are identical up to the point when this event is defined. Therefore, by Proposition 3,

$$\text{SD}((\tilde{Y}_\ell, Z, F), (U_r, Z, F)) \leq \text{SD}((\tilde{Y}_\ell, Z, F)|\mathcal{E}, (U_r, Z, F)|\mathcal{E}) + \frac{q^2}{2^r};$$

For convenience we let $\rho := q^2/2^r$ for the remainder of the proof. In order to bound the statistical distance conditioned on \mathcal{E} , we can rewrite the same as

$$\text{SD}((\tilde{Y}_\ell, Z, F)|\mathcal{E}, (U_r, Z, F)|\mathcal{E}) = \sum_{z \in \mathcal{E}} \text{P}[Z = z|\mathcal{E}] \cdot \text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z), \quad (16)$$

where $z \in \mathcal{E}$ is to denote that the sum is taken over all side informations $Z = z$, satisfying \mathcal{E} .³⁰ Define $p_z := \text{Pred}(\bar{X}|Z = z)$, and observe that

$$\mathbf{E}_z[p_z] = \text{Pred}(\bar{X}|Z) \leq 2^{-\gamma^*},$$

where the latter inequality follows from the assumption $\text{H}_\infty(\bar{X}|Z) \geq \gamma^*$. Moreover,

$$\text{H}_\infty(\bar{X}|Z\mathcal{E}) \geq \gamma^* - \log(1 - \rho)^{-1},$$

which is due to

$$\begin{aligned} \text{Pred}(\bar{X}|Z) &\geq \sum_{z \in \mathcal{E}} \text{P}[Z = z] \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \sum_{z \in \mathcal{E}} \frac{\text{P}[Z = z]}{\text{P}[\mathcal{E}]} \cdot p_z \\ &= \text{P}[\mathcal{E}] \cdot \text{Pred}(\bar{X}|Z\mathcal{E}). \end{aligned}$$

³⁰Therefore \mathcal{E} can be omitted in the conditioning of the statistical distance.

From Lemma 57 we prove below, we will get,

$$\text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}} + (q + \ell + 2) \cdot \frac{2^r}{2^n}} .$$

Using Jensen's inequality, (16) becomes, for $\alpha = 2^r$ and $\beta = \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}} + (q + \ell + 1) \cdot \frac{2^r}{2^n}$,

$$\begin{aligned} \text{SD}((\tilde{Y}_\ell, F)|\mathcal{E}, (U_r, F)|\mathcal{E}) &\leq \frac{1}{2} \sqrt{\alpha \text{Pred}(\bar{X}|\mathcal{L}\mathcal{E}) + \beta} \\ &\leq \frac{1}{2} \sqrt{\frac{2^{r-\gamma^*}}{(1-\rho)} + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}} + (q + \ell + 2) \cdot \frac{2^r}{2^n}} . \end{aligned}$$

□

Lemma 57. For $z \in \mathcal{E}$ and the random variables as defined earlier,

$$\text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z) \leq \frac{1}{2} \sqrt{2^r p_z + \ell \cdot d'(\ell) \cdot \frac{2^r}{2^n} + 64\ell^4 \cdot \frac{2^r}{2^{2n}} + 16\ell^2 q^2 \cdot \frac{2^r}{2^{2n}} + (q + \ell + 2) \cdot \frac{2^r}{2^n}} .$$

Proof. Fix $z = (\sigma, L, s_0)$. When we condition on z , it becomes evident that F is uniformly distributed over all functions that agree with L . We now use Proposition 2 to get that:

$$\text{SD}((\tilde{Y}_\ell, F)|z, (U_r, F)|z) \leq \frac{1}{2} \sqrt{2^r \cdot \text{Coll}(\tilde{Y}_\ell|Fz) - 1} . \quad (17)$$

To bound the collision probability, we consider the following experiment:

- choose F uniformly consistent with L
- sample inputs $\bar{X} = (\bar{X}_1, \dots, \bar{X}_\ell)$ and $\bar{X}' = (\bar{X}'_1, \dots, \bar{X}'_{\ell'})$ independently but conditioned on $Z = z$.
- compute S, S' as the refresh evaluations with F of inputs \bar{X} and \bar{X}' respectively.
- compute \tilde{Y}_ℓ and $\tilde{Y}'_{\ell'}$ as the truncated HMAC' evaluations with F of \bar{X} and \bar{X}' , i.e $\tilde{Y}_\ell = F(F(0, \text{opad}), S)[1..r]$ and $\tilde{Y}'_{\ell'} = F(F(0, \text{opad}), S')[1..r]$

We first condition on the event \mathcal{E} that $F(0, \text{opad}) \neq F(0, \text{ipad})$. Then, we bound the probability that $\tilde{Y}_\ell = \tilde{Y}'_{\ell'}$ in this modified experiment which is conditioned on \mathcal{E} as

$$\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \frac{1}{2^n} + \text{P}[\bar{X} = \bar{X}'] + \text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'} | \bar{X} \neq \bar{X}']^{31} . \quad (18)$$

Now, the first term reduces to at most p_z . Towards bounding the second term, we fix arbitrary inputs $\bar{x} \neq \bar{x}'$ of lengths ℓ and ℓ' , respectively. We also assume, wlog, that the evaluation of \bar{x}' is not completely covered by L ; otherwise it would violate the collision-freeness of L . Let \bar{x}_{k+1} be the first block of \bar{x} not covered by L and similarly \bar{x}'_{k+1} for \bar{x}' . We let $k = \ell$ if all blocks are covered. We now use the results of Lemma 58 to upperbound $\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}]$. This concludes the proof. □

Lemma 58. For fixed inputs $\bar{x} \neq \bar{x}'$ and the random variables as defined earlier,

$$\text{P}[\tilde{Y}_\ell = \tilde{Y}'_{\ell'}] \leq \frac{\ell \cdot d'(\ell)}{2^n} + \frac{64\ell^4}{2^{2n}} + \frac{16\ell^2 q^2}{2^{2n}} + \frac{q + \ell + 1}{2^n} + \frac{1}{2^r} .$$

³¹The conditioning on \mathcal{E} is dropped from the latter two terms.

Proof. Upon fixing it to be arbitrary inputs, we can drop the conditioning on $\bar{X} \neq \bar{X}'$. In this setting, we can condition on the equality of S and S' as follows:

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell] \leq \mathbb{P}[S = S'] + \mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell | S \neq S'] . \quad (19)$$

Let us take a look at the first term. Notice that this is the collision probability of S when run on two fixed inputs \bar{x} and \bar{x}' . In other words, this proof is similar to the bounding of collision probability of truncated outputs of MD when run on \bar{X} with side information z . We now proceed to construct the structure graph as done in the Proof of Lemma 38. The key difference here is that we need the structure graph to be colliding, i.e $\mathbb{P}[S = S'] = \mathbb{P}[G_F(\bar{x}, \bar{x}') \in \text{Coll}_L(\bar{x}, \bar{x}')] where $\text{Coll}_L(\bar{x}, \bar{x}')$ are the set of colliding structure graphs. We now use the results of Lemmas 40, 41 and 42 to conclude that:$

$$\mathbb{P}[S = S'] \leq \frac{\ell \cdot d'(\ell)}{2^n} + \frac{64\ell^4}{2^{2n}} + \frac{16\ell^2 q^2}{2^{2n}} . \quad (20)$$

Again, we look at the second term. We can fix the values to be arbitrary $s \neq s'$. Note that conditioning could impact the randomness of the function F . However, when we run on fixed inputs \bar{x}, \bar{x}' . This would mean that F is uniformly random conditioned on the set of queries in L and the set of evaluations of \bar{x}, \bar{x}' resulting in states s, s' respectively. In this setting, define a bad event as an event \mathcal{E} such that the final output are not the same. Therefore, we can bound:

$$\mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell] \leq \mathbb{P}[\tilde{Y}_\ell = \tilde{Y}'_\ell | \mathcal{E}] + \mathbb{P}[\neg \mathcal{E}] \leq \frac{1}{2^r} + \frac{q + \ell + 1}{2^n}$$

where the last inequality is a result of Lemma 39.

We can look at $\mathbb{P}[\neg \mathcal{E}]$ as conditioned on an event that the output of $F(0, \text{opad})$ collides with one of the other at most $q + \ell$ values corresponding to the L and the evaluations of the inputs. When this output does not collide, then the final output is freshly sampled and by the randomness of F we get that the final outputs can collide with probability at most $\frac{1}{2^n}$. Furthermore, the probability that output of $F(0, \text{opad})$ collides is $\frac{q+\ell}{2^n}$. This concludes the proof. \square

C Hybrid Proofs

C.1 Recovering and Preserving Imply Robustness

Theorem 14. *Consider a PRNG construction $\text{PRNG} = (\text{refresh}, \text{next})$ for which refresh makes α P -calls and next makes β P -calls. Furthermore, assume PRNG is both*

- $(\gamma^*, q, \ell, \varepsilon_{\text{rec}})$ -recovering and
- $(q, \ell, \varepsilon_{\text{pre}})$ -preserving

in the P -model. Then, PRNG is also $(\gamma^, q, t, \ell, \varepsilon_{\text{rob}})$ -robust in the P -model, where*

$$\varepsilon_{\text{rob}} \leq t \cdot (\varepsilon_{\text{rec}} + \varepsilon_{\text{pre}}) .$$

For the proof of Theorem 14, consider a γ^* -legitimate (canonical) (q, t, ℓ) -attacker \mathcal{A} . The proof is through a series of hybrid games, where the advantage of \mathcal{A} in the final game is easily seen to be 0. In order to define the hybrids, let a *nice* next query be either **get-next** or **next-ror**. A nice next query is

- *recovering* if it is the first such query after the MRED and
- *preserving* otherwise.

Define the hybrid rob_i to be the robustness game for PRNG where for first i nice next queries, the output of `next` is replaced by a uniform random string of length $n + r$. Moreover, consider an intermediate hybrid $\text{rob}_{i+\frac{1}{2}}$, where the challenger also replaces the output of `next` with a random string if the $(i + 1)^{\text{st}}$ query is *preserving*.

Denote by η_i the probability that \mathcal{A} outputs 1 when interacting with hybrid rob_i . In the following, preserving security will be used to argue the indistinguishability of hybrids rob_i and $\text{rob}_{i+\frac{1}{2}}$ (Claim 59), and recovering security (Claim 60) for $\text{rob}_{i+\frac{1}{2}}$ and rob_i .

Claim 59. For all $i = 0, \dots, t$, $|\eta_i - \eta_{i+\frac{1}{2}}| \leq \varepsilon_{\text{pre}}$.

Proof. The two hybrids only differ in the case when $(i + 1)^{\text{st}}$ next query is *preserving*. Hence, assume that the adversary \mathcal{A} ensures that the $(i + 1)^{\text{st}}$ query is indeed preserving, which serves to maximize its advantage. Consider the following attacker \mathcal{A}' against the preserving security of PRNG: Initially, \mathcal{A}' sets

- $b \leftarrow_{\$} \{0, 1\}$,
- $s \leftarrow 0^n$,
- $j \leftarrow 0$, and
- $\chi \leftarrow \lambda$, where λ is the empty string.

Then, \mathcal{A}' runs \mathcal{A} , simulating all oracle calls made by \mathcal{A} answering the queries from \mathcal{A} as follows: At all times, P -queries by \mathcal{A} are simply forwarded by \mathcal{A}' to its own P -oracle and back. Furthermore, while $j \leq i$:

- **adv-refresh**(x): \mathcal{A}' simply ignores the query.
- **set-state**(s'): \mathcal{A}' just sets $s \leftarrow s'$.
- **get-state**: \mathcal{A}' returns s .
- **get-next** or **next-ror**: \mathcal{A}' chooses $(s, y) \leftarrow_{\$} \{0, 1\}^{(n+r)}$ uniformly at random, increments $j \leftarrow j + 1$, and returns y .

Once $j = i + 1$, \mathcal{A}' simulates the oracles as follows:

- **adv-refresh**(x): \mathcal{A}' appends x to χ , i.e., $\chi \leftarrow \chi || x$.
- **set-state**(s'): \mathcal{A}' sets $s \leftarrow s'$.
- **get-state**: \mathcal{A}' returns s .
- **get-next**: \mathcal{A}' calls its challenge oracle to obtain $(s, y) \leftarrow \mathbf{chall}(s, \chi)$, increments $j \leftarrow j + 1$, and returns y .
- **next-ror**: \mathcal{A}' calls its challenge oracle to obtain $(s, y_0) \leftarrow \mathbf{chall}(s, \chi)$, chooses $y_1 \leftarrow_{\$} \{0, 1\}^r$, increments $j \leftarrow j + 1$, and returns y_b .

Subsequently, i.e., once $j > i + 1$, \mathcal{A}' uses the state s returned by **chall** and its P -access to keep simulating the oracles consistent with rob_{i+1} . In the end, \mathcal{A}' outputs whatever \mathcal{A} outputs.

Let the challenge bit of the challenger for \mathcal{A}' be \tilde{b} . Note that when $\tilde{b} = 0$, \mathcal{A}' perfectly simulates hybrid $\text{rob}_{i+\frac{1}{2}}$ for \mathcal{A} . Similarly, if $\tilde{b} = 1$, \mathcal{A}' perfectly simulates hybrid rob_{i+1} for \mathcal{A} . (Recall that \mathcal{A} is canonical.) \square

The next step is to show that the hybrids $\text{rob}_{i+\frac{1}{2}}$ and rob_{i+1} are indistinguishable from each other.

Claim 60. For all $i = 0, \dots, t - 1$, $|\eta_{i+\frac{1}{2}} - \eta_{i+1}| \leq \varepsilon_{\text{rec}}$.

Proof. The two hybrids only differ in the case when $(i + 1)^{\text{st}}$ next query is *recovering*. Hence, assume that the adversary \mathcal{A} ensures that the $(i + 1)^{\text{st}}$ query is indeed recovering, which serves to maximize its advantage. Consider the following attacker \mathcal{A}' against the recovering security of PRNG: Initially, \mathcal{A}' sets

- $b \leftarrow_{\$} \{0, 1\}$,
- $s \leftarrow 0^n$,
- $j \leftarrow 0$, and
- $\chi \leftarrow \lambda$, where λ is the empty string.

Then, \mathcal{A}' runs \mathcal{A} , simulating all oracle calls made by \mathcal{A} answering the queries from \mathcal{A} as follows: At all times, P -queries by \mathcal{A} are simply forwarded by \mathcal{A}' to its own P -oracle and back. Furthermore, while $j \leq i$:

- **adv-refresh**(x): \mathcal{A}' simply ignores the query.
- **set-state**(s'): \mathcal{A}' just sets $s \leftarrow s'$.
- **get-state**: \mathcal{A}' returns s .
- **get-next** or **next-ror**: \mathcal{A}' chooses $(s, y) \leftarrow_{\$} \{0, 1\}^{(n+r)}$ uniformly at random, increments $j \leftarrow j + 1$, and returns y .

Once $j = i + 1$, \mathcal{A}' simulates the oracles as follows:

- **adv-refresh**(x): \mathcal{A}' appends x to χ , i.e., $\chi \leftarrow \chi || x$.
- **set-state**(s'): \mathcal{A}' sets $s \leftarrow s'$.
- **get-state**: \mathcal{A}' returns s .
- **get-next**: \mathcal{A}' calls its challenge oracle to obtain $(s, y) \leftarrow \mathbf{chall}(s, \chi)$, increments $j \leftarrow j + 1$, and returns y .
- **next-ror**: \mathcal{A}' calls its challenge oracle to obtain $(s, y_0) \leftarrow \mathbf{chall}(s, \chi)$, chooses $y_1 \leftarrow_{\$} \{0, 1\}^r$, increments $j \leftarrow j + 1$, and returns y_b .

Subsequently, i.e., once $j > i + 1$, \mathcal{A}' uses the state s returned by **chall** and its P -access to keep simulating the oracles consistent with rob_{i+1} . In the end, \mathcal{A}' outputs whatever \mathcal{A} outputs.

Let the challenge bit of the challenger for \mathcal{A}' be \tilde{b} . Note that when $\tilde{b} = 0$, \mathcal{A}' perfectly simulates hybrid $\text{rob}_{i+\frac{1}{2}}$ for \mathcal{A} . Similarly, if $\tilde{b} = 1$, \mathcal{A}' perfectly simulates hybrid rob_{i+1} for \mathcal{A} . (Recall that \mathcal{A} is canonical.)

It remains to argue that \mathcal{A}' is γ^* -legitimate. To that end, recall from Section 4.2 that the legitimacy condition is defined in the *legitimacy* game. Observe in particular, that up to the time \mathcal{A}' makes its challenge queries, \mathcal{A} 's view is exactly as it would be in the legitimacy game, in which one considers the following random variables immediately before \mathcal{A} makes the i^{th} call to oracle **get-next** or **next-ror**:

- \mathcal{L}_i : the list of P -queries by \mathcal{A} and the corresponding answers;
- Σ_i : the state of \mathcal{A} ;
- \bar{X}_i : vector of inputs provided by \mathcal{A} since the *the most recent entropy drain (MRED)*; and
- S_i : the state of the PRNG immediately after the MRED.

In a similar fashion, consider the following random variables pertaining to \mathcal{A}' just before its call to the challenge oracle:

- X , the input vector \mathcal{A}' provides to **chall**;
- Σ , the state of \mathcal{A}' just before the call to **chall**;
- \mathcal{L} , the list of P -queries and answers by \mathcal{A}' before the call to **chall**; and
- S_0 ; the state \mathcal{A}' provides to **chall**.

That is, it needs to be established that

$$H_\infty(X|\Sigma\mathcal{L}S_0) \geq \gamma^* . \quad (21)$$

using the argument of the γ^* -legitimacy of \mathcal{A} ,

$$H_\infty(\bar{X}_i|\Sigma_i\mathcal{L}_iS_i) \geq \gamma^* .$$

First, it is easily verified that $S = S_i$. Second, observe that Σ only needs to contain, in addition to Σ_i , the values of b and j , which are clearly independent of X . Third, clearly $\mathcal{L} = \mathcal{L}_i$. Finally, note that $X = \bar{X}_i$. From the preceding, (21) follows. \square

C.2 IT-Recovering Implies IT-Robustness

Theorem 36. *Let $\text{PRNG} = (\text{refresh}, \text{next})$ be a PRNG for which **refresh** makes α P -calls and **next** makes β P -calls. Let the PRNG be $(\gamma^*, q, \ell, \varepsilon_{\text{rec}})$ -IT-recovering in the P -model. Then, PRNG is also $(\gamma^*, q, t, \ell, \varepsilon_{\text{rob}})$ -IT-robust in the P -model, where*

$$\varepsilon_{\text{rob}} \leq t \cdot \varepsilon_{\text{rec}} .$$

As seen in the proof of Theorem 14, define the hybrid rob_i to be the IT-robustness game for PRNG where for first i nice next queries, the output of **next** is replaced by a uniform random string of length $n + r$. We denote by η_i the probability that \mathcal{A} outputs 1 when interacting with hybrid rob_i .

Claim 61. *For all $i = 0, \dots, t - 1$, $|\eta_i - \eta_{i+1}| \leq \varepsilon_{\text{rec}}$.*

The proof of Claim 61 is completely analogous to that of Claim 60 and is therefore omitted.