

A Provable-Security Analysis of Intel’s Secure Key RNG

Thomas Shrimpton and R. Seth Terashima

Dept. of Computer Science, Portland State University
{teshrim, seth}@cs.pdx.edu

Abstract. We provide the first provable-security analysis of the Intel Secure Key hardware RNG (ISK-RNG), versions of which have appeared in Intel processors since late 2011. To model the ISK-RNG, we generalize the PRNG-with-inputs primitive, introduced by Dodis et al. at CCS’13 for their /dev/[u]random analysis. The concrete security bounds we uncover tell a mixed story. We find that ISK-RNG lacks backward-security altogether, and that the forward-security bound for the “truly random” bits fetched by the RDSEED instruction is potentially worrisome. On the other hand, we are able to prove stronger forward-security bounds for the pseudorandom bits fetched by the RDRAND instruction. En route to these results, our main technical efforts focus on the way in which ISK-RNG employs CBCMAC as an entropy extractor.

Keywords: random number generator, entropy extraction, provable security

A version of this paper appears in Eurocrypt 2015. This is the full version.

Table of Contents

1	Introduction	3
1.1	Security findings for the ISK-RNG	3
1.2	Improvements to the PWI model	5
2	Preliminaries	6
3	The ISK-RNG architecture	7
4	Analysis of the ISK-RNG extractor	9
5	Modeling the ISK-RNG as a PWI	11
5.1	The PWI model	11
5.2	Mapping ISK-RNG into the PWI model	12
6	PWI Security	12
6.1	Basic notions	12
6.2	Masking functions and updated security notions	13
6.3	PWI-Security Theorems	16
7	Security of the ISK-RNG as a PWI	16
7.1	Negative Results	17
7.2	Positive results	18
7.3	Discussion of results	21
A	PWIs with non-pseudo-uniformly random state	23
B	On proving robustness	24
C	On the read-only robustness of ISK-RNG	28
D	Towards a notion of PWI availability	31
E	Security Proofs	32

1 Introduction

In late 2011, Intel began production of Ivy Bridge processors, which introduced a new pseudorandom number generator (PRNG), fully implemented in hardware. Access to this PRNG is through the `RDRAND` instruction (pronounced “read rand”), and benchmarks demonstrate a throughput of over 500 MB/s on a quad-core Ivy Bridge processor [10]. The forthcoming Broadwell architecture will support an additional instruction, `RDSEED` (“read seed”), which delivers true random bits, as opposed to cryptographically pseudorandom ones. Both `RDRAND` and `RDSEED` fall under the Intel Secure Key umbrella, so we will refer to the new hardware as the ISK-RNG [11].

The ISK-RNG has received a third-party lab evaluation [8], commissioned by Intel, but has yet to receive an academic, provable-security treatment along the lines of that given the `/dev/[u]random` software RNGs by a line of papers [7,1,5,13]. We provide such a treatment.

Our abstract model for the ISK-RNG is that of a PRNG-with-input (PWI), established by Barak and Halevi [1] and extended by Dodis et al. [5]. To better capture important design features of the ISK-RNG we make several improvements to the PWI abstraction, which have significant knock-on effects for the associated security notions. Our results establish the security of the ISK-RNG relative to these notions. Our findings are mixed, suggesting that in some cases `RDSEED` may not be as secure as one might hope, but with stronger results for `RDRAND`.

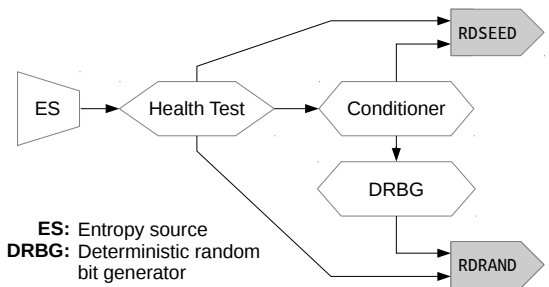


Fig. 1. Overview of Intel’s hardware PRNG.

The ISK-RNG architecture. A detailed description of the ISK-RNG can be found in Section 3, but we’ll provide a short sketch here. At a high-level, the ISK-RNG consists of four main components, as shown in Figure 1. At the heart is the hardware *entropy source*, which uses thermal noise to generate random bits and then writes them into a 256-bit raw-sample buffer. This buffer is subjected to a battery of heuristic *health tests*, which try to determine if the buffer contents are sufficiently random. The raw entropy bits are not assumed to be *uniformly* random — they may be biased or correlated. So a *conditioner* (i.e. an entropy extractor), repeatedly reads from this buffer, combining multiple 256-bit samples and compressing them into a single 128-bit string, hopefully one that is close to uniformly random.

These uniform bit strings then periodically reseed a deterministic PRNG (based on CTR-AES), providing a high-speed source of pseudorandom bits. Calls to the `RDRAND` instruction read from these bits, whereas calls to `RDSEED` will read directly from the conditioner output.

1.1 Security findings for the ISK-RNG

We consider security of the ISK-RNG relative to four PWI-security notions, adopted (with modifications) from Dodis, Pointcheval, Ruhault, Vergniaud and Wichs [5] (hereafter DPRVW): *resilience*, the apparent randomness of `RDRAND` and `RDSEED` outputs; *forward security*, the apparent randomness of previous `RDRAND` and `RDSEED` outputs once the PWI state is revealed; *backward security*, the apparent randomness of future `RDRAND` and `RDSEED` outputs from a corrupted PWI state; and *robustness*, the apparent randomness of `RDRAND` and `RDSEED` outputs when state observation and corruption may happen at arbitrary times.

Using estimates for the quality of the entropy source derived from the findings of [8], we are able to show the following results (in a random permutation model):

1. As far as the resilience of `RDRAND` and `RDSEED` is concerned, `RDRAND` delivers pseudorandom bits with a comfortable security margin. On the other hand, `RDSEED` delivers truly random bits but with a security margin that becomes worrisome if an adversary can see a large number of outputs from either interface. If he controls an unprivileged process on the same physical machine, this could happen very quickly.
2. For forward security, `RDRAND` and `RDSEED` also provide these respective security margins, as long as one is willing to make some reasonable assumptions about the adversary’s limitations.
3. The ISK-RNG does *not* provide backward security because the hardware indefinitely retains stale state when the ISK-RNG is not in active use. However, we are able to quantify the lifespan of this information when the ISK-RNG *is* in active use, thus proving backwards security and a read-only form of robustness against a class of “slow” adversaries.

Interpretation. In this context, forward security, backward security, and robustness are only relevant to those concerned about attackers who (1) are able to obtain physical access to the machine and (2) sophisticated enough to read or tamper with registers directly (the registers in question are not accessible through software, even by the operating system). Moreover, the window of opportunity for an attacker trying to compromise forward security (i.e., trying to reconstruct past random values given current access to the machine) is under a millisecond, barring pathological failures of the entropy source. Hence we suspect most practitioners will be concerned only with resilience.

As far as resilience, then, we prove `RDRAND` to be secure under a reasonable set of assumptions regarding the quality of the entropy source and a reasonable but heuristic assumption regarding AES-128: namely that it can be modeled as a random permutation when used with a specific fixed, publicly known key. We provide concrete, quantitative analysis in Section 7.3; the results are encouraging.

The situation with `RDSEED` is more complicated, because the security bounds become quantitatively quite weak in this context. We believe, but cannot prove, that this weakness does not correspond to a practical attack. Our suspicion is that an actual attack would require the adversary to have a precise physical model of the entropy source (the exact parameters of which appear to change from chip to chip [8]), and compute, by brute force, the distribution induced by processing streams from this entropy source using CBC-MAC under the previously mentioned AES key. Such an attack would clearly be computationally infeasible as long as the number of possible streams is large, but the relevant portion of the security bound is for computationally unbounded adversaries. (Recall that `RDSEED` is designed to provide truly random bits, rather than “merely” cryptographically pseudorandom ones.)

The stronger `RDRAND` results hold even if an attacker can access both interfaces.

Analyzing the ISK-RNG entropy extractor. The core technical results of the paper are concerned with analyzing the ISK-RNG entropy extractor, which employs CBC-MAC over AES-128, using the fixed string $\text{AES}_0(1)$ as the AES key. Although Intel documents [16] appeal to a CRYPTO’02 paper by Dodis, Gennaro, Håstad and Krawczyk [4] for support, this direct appeal is not well founded. There are significant technical obstacles to overcome before these CBC-MAC results can be applied. For example, because extractor-dependent state is maintained across extractions (including state revealed to the adversary by `RDSEED`), a crucial “seed independence” assumption is violated. The

CRI report [8], on the other hand, ignores the issue entirely by making an implicit assumption that applying CBC-MAC-AES to an arbitrary input with 128 bits of min-entropy will produce an output close to a uniformly random 128-bit string, an assumption known to be false with respect to *any* entropy extractor (not just CBC-MAC) [15]. We discuss and resolve these issues in Section 4.

1.2 Improvements to the PWI model

For our abstract model, we take the *pseudorandom number generator with input* (PWI) primitive, formalized by DPRVW as a model for `/dev/[u]random`. At a high level, a PWI surfaces three algorithms: one to initialize the internal state of the primitive, one that produces an output for use by calling applications (updating the state in the process), and one that updates the state as a function of an *externally* provided input. Exposing an external input captures the practical situation in which PRNG outputs may depend upon external sources of (assumed) entropy.

One contribution of this paper is to generalize the PWI abstraction in ways that better capture not only the ISK-RNG, but also, we hope, other real-world PWIs. These include allowances for:

1. Non-uniform PWI state: Real-world PRNGs like ISK-RNG have internal state that is not properly modeled as being uniformly random; for example, because they contain counters or fixed strings. DPRVW make progress in abandoning the restrictions of prior works, whose security notions mandate that the internal state be indistinguishable from a uniformly random one. However, they do not go far enough, as they still assume the PWI is *initialized* with a uniform random state. This leads to problems: one can construct a PWI that is provably “forward-secure” in the sense defined by DPRVW, but is clearly not forward-secure in the commonly accepted sense. Conversely, there are PWIs that seemingly should be considered forward-secure, but are not under the DPRVW definition. We provide examples in Appendix A.
2. Realistic initialization: Relatedly, the current PWI formalism cannot model setup procedures, such as those used by the ISK-RNG. Addressing this shortcoming is particularly important since weak initialization has led to attacks on real-world PWIs [9,6]; security definitions should expose these vulnerabilities.
3. Multiple external interfaces: the ISK-RNG API exposes two interfaces that depend on the same state, namely `RDRAND` and `RDSEED`. Hence the model should reflect this. We note that ISK-RNG isn’t unique in this regard: `/dev/random` and `/dev/urandom` also share some state [13], and so generalizing the PWI model to allow multiple interfaces is a necessary prerequisite to establishing any positive results there.
4. Blocking behavior: the ISK-RNG and other real-world RNGs (e.g. `/dev/random`) block when their tests suggest that they have not accumulated sufficient entropy to generate random numbers securely. Capturing this behavior allows us to ensure that this mechanism cannot inadvertently compromise security. We briefly discuss the related issue of availability in Appendix D.
5. Race conditions: As is common in cryptography, the passage of time in the PWI model is implicitly driven by the adversary’s oracle queries. However, the ISK-RNG state may change asynchronously with respect to queries made to oracles in the existing PWI model. We introduce a “dummy” `wait()` oracle that prompts the PWI to complete its current highest-priority atomic task.

To deal with non-uniform state, we introduce an analytical tool called a *masking function*. Loosely speaking, a masking function M is a tool for specifying what the “ideal” version $M(S)$ of any given PWI state S would be. This allows us to give general results about PWI security (e.g. what can be achieved when the state is ideal), yet admits per-scheme specification of what “ideal” means. We define masking functions, and incorporate them into the DPRVW’s security notions in such a way that their results can be quickly lifted to our setting. Masking functions also allow us to frame an appropriate definition for secure initialization: i.e. does the setup procedure produce a state S that is indistinguishable from $M(S)$?

2 Preliminaries

Notation. We denote the set of all n -bit strings as $\{0, 1\}^n$, and the set of all (finite) binary strings as $\{0, 1\}^*$. Given $x, y \in \{0, 1\}^*$, both xy and $x \parallel y$ denote their concatenation, and $|x|$ is the length of x . If $|x| = |y|$, $x \oplus y$ is the bitwise XOR of x and y . The symbol ε denotes the empty string. The set $\text{Perm}(n)$ denotes the set of permutations on $\{0, 1\}^n$.

When S is a finite set, we assume that it is equipped with the uniform distribution unless otherwise specified. For any distribution \mathcal{S} , the notation $X \stackrel{\$}{\leftarrow} \mathcal{S}$ indicates X is a random variable sampled from \mathcal{S} . Similarly, if F is a randomized algorithm, $X \stackrel{\$}{\leftarrow} F(x_1, \dots, x_n)$ means that X is sampled from the distribution induced by providing F with the indicated arguments. An adversary A is a randomized algorithm, and we adopt the shorthand $A \Rightarrow y$ to mean that when its execution halts, it outputs y . When an algorithm P is provided oracle (black-box, unit-time) access to an algorithm Q , we write P^Q .

Entropy and Sources. If X and X' are random variables, their statistical distance is $\Delta(X, X') = \frac{1}{2} \sum_x |\Pr[X = x] - \Pr[X' = x]|$, where the sum is over the union of the supports of X and X' . The min-entropy of X is $\mathbf{H}_\infty(X) = -\max_x (\log \Pr[X = x])$, and the worst-case min-entropy of X given X' is $\mathbf{H}_\infty(X | X') = -\log(\max_{x, x'} \Pr[X = x | X' = x'])$. When X is a random variable and \mathcal{E} is some event, we denote by $X|_{\mathcal{E}}$ the random variable X conditioned on \mathcal{E} ; i.e., for any x in the support of X , $\Pr[X|_{\mathcal{E}} = x] = \Pr[X = x | \mathcal{E}]$.

An *entropy source* \mathcal{D} is a randomized algorithm that, given a state string $\sigma \in \{0, 1\}^*$, samples a tuple $(\sigma', I, \gamma, z) \in \{0, 1\}^* \times \{0, 1\}^p \times \mathbb{R}_{\geq 0} \times \{0, 1\}^*$. Let $(\sigma_i, I_i, \gamma_i, z_i) \stackrel{\$}{\leftarrow} \mathcal{D}(\sigma_{i-1})$ be a sequence of samples, where $\sigma_0 = \varepsilon$, and $i = 1, \dots, q_{\mathcal{D}}$ for some integer $q_{\mathcal{D}}$. We say that entropy source \mathcal{D} is *legitimate* if $H_\infty(I_j | (I_i, z_i, \gamma_i)_{i \neq j}) \geq \gamma_j$. In this paper, we assume all entropy sources are legitimate.

In this definition, $\sigma, \sigma' \in \{0, 1\}^*$ represent the current and new states for \mathcal{D} , respectively. The string $I \in \{0, 1\}^p$ is what will be to be fed as input to the PWI, and should provide fresh entropy. The quantity $\gamma \in \mathbb{R}_{\geq 0}$ is an estimate for the amount of entropy contained in I . We note that γ is strictly a convenient book-keeping device in the PWI model, and is not intended to reflect an actual output of the entropy source being modeled. Our security notions will formalize attacker capabilities of interest, but we also allow for side-information (about I) that an attacker might obtain through means not otherwise explicit in the model (e.g. timing or power side-channels). This side information will be encoded in the string z .

Cryptographic building blocks. A blockcipher is a function $E : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that for each key $K \in \{0, 1\}^\kappa$, $E(K, \cdot)$, written $E_K(\cdot)$, is a permutation on $\{0, 1\}^n$. Given $\text{IV} \in \{0, 1\}^n$, $K \in \{0, 1\}^\kappa$, and $X_i \in \{0, 1\}^n$ for $i \in [0.. \nu]$, define

$$\text{CTR}_K^{\text{IV}}(X_0 \cdots X_\nu) = (X_0 \oplus E_K(\text{IV})) \parallel \cdots \parallel (X_\nu \oplus E_K(\text{IV} + \nu)).$$

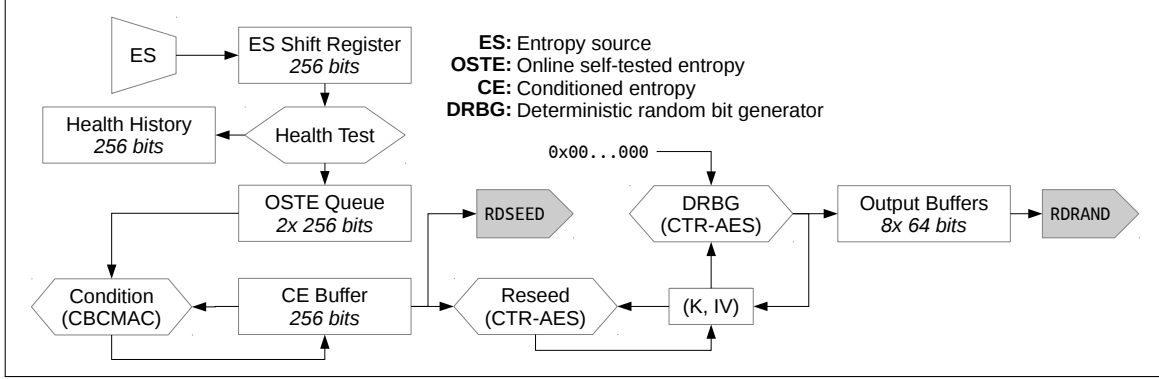


Fig. 2. Block diagram for Intel’s RDRAND implementation. The CBCMAC computation uses AES-128 with the fixed key $K' = \text{AES}_0(1)$. The DRBG runs AES-128 in counter mode to produce $\{0, 1\}^{128 \cdot 3}$ bits of output; the first 256 bits are used to update the key K and IV; the final 128 bits are sent to the output buffer, which is read by the RDRAND instruction.

(We define the $+$ operator on $\{0, 1\}^n$ as addition modulo 2^n on the unsigned integers encoded by the operands.) Further define

$$\text{CBCMAC}_K^{\text{IV}}(X_0 \cdots X_\nu) = \text{CBCMAC}_K^{E_K(\text{IV} \oplus X_0)}(X_1 \cdots X_\nu),$$

and $\text{CBCMAC}_K^{\text{IV}}(\varepsilon) = \text{IV}$. Describing the standard CBCMAC algorithm in this manner simplifies descriptions of programs that compute CBCMAC online. We omit an explicit IV from the notation when $\text{IV} = 0^n$. In this paper, the implicit blockcipher E will always be AES-128 ($\kappa = n = 128$).

The pseudorandom-permutation (PRP) advantage of an adversary A attacking a blockcipher $E : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is defined as $\text{Adv}_E^{\text{PRP}}(A) = \Pr [A^{E_K} \Rightarrow 1] - \Pr [A^\pi \Rightarrow 1]$, with probabilities over the coins of A and the random variables $K \xleftarrow{\$} \{0, 1\}^\kappa$ and $\pi \xleftarrow{\$} \text{Perm}(n)$.

3 The ISK-RNG architecture

This section describes the design of the ISK-RNG. Unless otherwise noted, this information comes from the CRI report [8]. The design can be divided roughly into three phases: entropy generation, entropy extraction, and expansion. Raw bits from the generation phase are fed into an entropy extractor, which is tasked with turning biased or correlated bits into uniform random strings. The expansion step uses these strings to seed a deterministic PRNG, which can produce cryptographically pseudorandom outputs at high speeds.

The design is shown in Figure 2. In this figure, rectangular boxes indicate values we consider part of the ISK-RNG state, hexagons indicate procedures that read and modify the state, and the shaded arrows indicate assembly instructions that allow (unprivileged) processes to read from the indicated buffer.

Entropy generation, Health Tests, and “Swelling”. The hardware entropy source (labeled ES) is a dual differential jamb latch with feedback; thermal noise resolves a latch formed by two cross-coupled inverters, generating a random bit before the system is reset. Bits from the entropy source are written into a 256-bit shift register.

Every 256 writes, the contents of the register are subjected to a series of health tests. These count how many times certain specified bit strings appear, and verify that the results are within normal limits. For example, the substring 010 may occur between 9 and 57 times, inclusive. These substrings and the corresponding numbers of allowable occurrences are intended to catch pathologically bad failures while keeping the false-positive rate low. (For reference, a uniformly random 256-bit string would be flagged as unhealthy approximately 1% of the time.) If the current ES register fails one of the tests, that 256-bit source-sample is flagged as *unhealthy*. We refer interested readers to the CRI report [8] for a more detailed description; for our purposes, it suffices to say there is some fixed set $\mathcal{H} \subseteq \{0, 1\}^{256}$ of strings that pass the health tests. The health-history register tracks how many of the last 256 samples passed the health test. This is a first-in first-out buffer, where a 1-bit means that a sample was deemed healthy, and a 0-bit mean that a sample was deemed unhealthy. The global health of the ISK-RNG is captured by a property call *swellness*.

Definition 1 (Swell ISK-RNG). *The ISK-RNG is said to be swell if at least 128 of the last 256 samples were healthy, i.e. if the health-history register contains at least 128 1s.* \square

Whether or not the current sample passes the health test, it is appended to the Online Self-Tested Entropy (OSTE) queue, and it is the OSTE queue that provides input to the extraction phase.

Extraction. Strings in the OSTE queue are not assumed to be uniformly random. Instead, each 256-bit entry is assumed to have a certain amount of min-entropy. The CBCMAC construction, over AES with key $K' = \text{AES}_0(1)$ [12], is employed as an entropy extractor, in order to turn strings in the OSTE queue into two 128-bit *conditioned entropy* (CE) strings. These are held in the CE buffer, which is initially all zeros, and are used to service RDSEED instructions and to reseed the DRBG. An important property of the CE buffer is its *availability*.

Definition 2 (CE buffer availability). *The CE buffer is available if (1) the ISK-RNG is swell, and (2) both 128-bit halves of the CE buffer (CE_0 and CE_1) have been updated using m healthy OSTE values since the most recent RDSEED call and the most recent DRBG reseeding. For Ivy Bridge chips, $m = 2$; for Broadwell chips, $m = 3$ [12].* \square

When the CE buffer is not available, the hardware will replenish the OSTE buffers with fresh entropy and feed them into a running CBCMAC calculation until a sufficient number of healthy samples have been conditioned. So if at some point $\text{CE}_0 = X$ and then the CE buffer is used to service a RDSEED instruction (making the CE buffer unavailable), the hardware will collect entropy strings $I_1, I_2, I_3, \dots \in \{0, 1\}^{256}$ and reassign $\text{CE}_0 \leftarrow \text{CBCMAC}_0(XI_1I_2I_3 \dots)$ online until there exist $i_1 < i_2 < \dots < i_m$ such that $I_{i_j} \in \mathcal{H}$ for $j \in [1..m]$ and the ISK-RNG is swell. Then the processes will repeat for CE_1 .

The particulars of the way CBCMAC is used in the ISK-RNG extractor, along with the notions of swellness and availability, play a large role in Section 4.

Expansion. To reseed the DRBG, the contents of CE_0 and CE_1 are used to generate a key and IV (respectively) for counter mode encryption over AES. This reseeding process only happens when the CE buffer is available. It takes the current key and IV, (K, IV) , and updates them by computing $K \parallel IV \leftarrow \text{CTR}_K^{\text{IV}}(\text{CE}_1 \parallel \text{CE}_2)$. Initially, $K = IV = 0^{128}$. However, using CTR with this non-random key is not a problem as long as the CE buffer is (close to) uniformly random: since the CE buffer is XORed into the CTR keystream, it can act as a one-time pad.

A pseudorandom value R is generated by computing $R \| K \| IV \leftarrow \text{CTR}_K^IV(0^{3 \cdot 128})$. (Note that this process also irreversibly updates K and IV , which helps provide forward security.) The ISK-RNG writes R to an output buffer, which is read by `RDRAND`. This FIFO output buffer [10] can contain up to eight 64-bit values. ISK-RNG allows a maximum of 511 64-bit values to be generated between reseeding operations; after this, it will only return 0s and will clear the carry bit to signal an error.

Setup. When the ISK-RNG powers on, the ISK-RNG performs a series of known-answer, built-in self-tests. Then the conditioned entropy (CE) buffer is cleared and the deterministic random bit generator (DRBG) is reseeded four times [12]. Each reseeding operation requires reconditioning the CE buffer until it is available. Finally, the system populates the eight output buffers using the DRBG.

Standards compliance. Intel states [14] that ISK-RNG is compliant with NIST’s SP800-90B & C draft standards. Whereas `RDRAND` can provide bit strings with “only” a 128-bit security level (since it uses AES-128 in CTR mode), `RDSEED` has no such limitation.

4 Analysis of the ISK-RNG extractor

As we will see, some of the PWI-security results for the ISK-RNG are not as strong as one might hope. Much of this is due to weak concrete bounds on its CBCMAC entropy extractor, which is tasked with turning the presumably biased and correlated bits from the entropy source into uniformly distributed strings. Let us explain.

Previous CBC-MAC results are not (directly) helpful. A paper by Dodis, Gennaro, Håstad, Krawczyk and Rabin [4] analyzes the security of CBC-MAC as an entropy extractor, and their results are cited by Intel documents [16] to support the ISK-RNG design. Because generic PRFs-as-entropy-extractors results [3] are too weak to be useful, the analysis of [4] takes place in the random permutation model. That is, instead of considering CBC-MAC over a blockcipher with a random key, they consider CBC-MAC over a random permutation. This model is a heuristic: even, say, AES equipped with a random key would not be a random permutation. In fact, CBC-MAC within the ISK-RNG uses AES with the fixed key $K' = \text{AES}_0(1)$ (on all chips). This fact may strike one as alarming. But we believe that a “nothing-up-my-sleeve” value for the extractor seed is a reasonable choice. (Generating the seed from the entropy source would be highly suspect from a theoretical perspective, because one requires that the extractor seed be “independent” of the entropy distribution.)

Anyway, our primary goal here is to identify what we *can* say about ISK, even if we’re forced to use a heuristic model. Dodis et al.[4] provide the following theorem:

Theorem 1 (CBCMAC entropy extractor [4]). *Fix positive integers k and L . Let $I \in \{0, 1\}^{Lk}$ be a random variable, $R \xleftarrow{\$} \{0, 1\}^k$ be a uniform random string, and let $\pi \xleftarrow{\$} \text{Perm}(k)$ be a random permutation. Then $\Delta((\pi, R), (\pi, \text{CBCMAC}_\pi(I))) \leq \frac{1}{2} \sqrt{2^{k - H_\infty(I)} + \frac{\mathcal{O}(L^2)}{2^k}}$.*

Unfortunately, one cannot simply apply this theorem to the CBC-MAC-based extractor used in ISK-RNG, without attending to the following two significant obstacles:

(1) As we noted in Section 3, the CBCMAC-based extractor uses its own previous output as the first block of its next input. Consequently, the CBCMAC inputs are not independent of the

seed. This pushes leftover-hash-lemma style results like Theorem 1 out of scope, and furthermore prevents us from employing a black-box hybrid argument to lift the results to the multiple-query setting.

(2) The $\mathcal{O}(L^2)$ term is problematic, contributing a $\mathcal{O}(L/2^{k/2})$ term to bound.¹ We note that this is significantly worse than the familiar $\mathcal{O}(L^2/2^k)$ “birthday bound” — although the two both become vacuous when $L \approx 2^{k/2}$, the former violates a desired security level $\epsilon \ll 1$ much sooner (hidden constants being equal). The weak bound is exacerbated by the fact that L may grow very quickly in the ISK-RNG during periods of time when the CE buffer is not available.

Analyzing the CBC-MAC extractor. In this section we present results that allow us to overcome these hurdles, bringing Theorem 1 into scope. In particular, the main technical result of this section is the following theorem. Loosely, it says that we can still obtain a hybrid-like bound, even though a black-box hybrid argument isn’t possible. Moreover, we can avoid the problem of “runaway” input strings (resulting in large L) by, in effect, only counting a fixed-length prefix of such strings.

Theorem 2. *Fix positive integers L, k, q and n with $q \leq n$. For $i \in [1..n]$, let $I_i \in \{0, 1\}^*$ be random variables with lengths divisible by k , and sample $R_i \xleftarrow{\$} \{0, 1\}^k$. Fix $\pi \xleftarrow{\$} \text{Perm}(k)$. Define I_i^L and I_i^R to be the unique strings such that $|I_i^L| = \min\{|I_i|, Lk\}$ and $I_i = I_i^L I_i^R$. Let $C_i = \text{CBCMAC}_\pi(C_{i-1} \| I_i)$, where C_0 is a random variable independent of π and each I_i and R_i . Then:*

$$\Delta((\pi, C_1, \dots, C_q, I_{>q}), (\pi, R_1, \dots, R_q, I_{>q})) \leq \frac{1}{2} \sum_{i=1}^q \sqrt{2^{k - \mathbf{H}_\infty(I_i^L | I_{>i}, I_i^R)} + \frac{\mathcal{O}((L+1)^2)}{2^k}},$$

where $I_{>m} = (I_{m+1}, \dots, I_n)$ for integer m .

Proof. Setting $R_0 = C_0$ (our applications will use either $C_0 \xleftarrow{\$} \{0, 1\}^k$ or $C_0 \xleftarrow{\$} \{0^k\}$), define:

$$\delta_i = \Delta((\pi, R_1, \dots, R_{i-1}, C_i^L, C_{i+1}^L, \dots, C_q^L), (\pi, R_1, \dots, R_{i-1}, R_i, C_{i+1}^L, \dots, C_q^L)).$$

where $C_j^i = \text{CBCMAC}_\pi(R_{i-1} \| I_i)$ if $j = i$, and $C_j^i = \text{CBCMAC}_\pi(C_{j-1}^i \| I_j)$ if $j > i$. As a consequence, we have $\Delta((\pi, C_1, \dots, C_q, I_{>q}), (\pi, R_1, \dots, R_q, I_{>q})) \leq \sum_{i=1}^q \delta_i$. Since each R_i is independent and uniform and C_j^i can be computed from π, R_{i-1} , and $I_{>i}^+ = (I_{>i}, I_i^R)$ whenever $j > i$,

$$\begin{aligned} \delta_i &\leq \Delta((\pi, R_{i-1}, R_i, I_{>i}^+), (\pi, R_{i-1}, C_i^L, I_{>i}^+)) \\ &\leq \sum_{r,s} \Delta((\pi, r, R_i, s), (\pi, r, C_i^L |_{(R_{i-1}, I_{>i}^+) = (r,s)}, s)) \cdot \Pr[R_{i-1} = r] \Pr[I_{>i}^+ = s]. \end{aligned}$$

Let \mathcal{P} be the event $(R_{i-1} = r) \wedge (I_{>i}^+ = s)$; then $C_i^L |_{\mathcal{P}} = \text{CBCMAC}_\pi(r \| I_i^L |_{\mathcal{P}} \| c^R)$, where c^R is the first component of s , corresponding to the (now fixed) value of I_i^R . Note that $I_i |_{\mathcal{P}}$ is distributed identically to $I_i |_{I_{>i}^+ = s}$. Define $R'_i = \text{CBCMAC}_\pi(r \| I_i^L |_{\mathcal{P}})$. This now brings the results of Theorem 1 into scope:

$$\Delta((\pi, r, R_i, s), (\pi, r, R'_i, s)) \leq \frac{1}{2} \sqrt{2^{k - \mathbf{H}_\infty(I_i^L | I_{>i})} + \frac{\mathcal{O}((L+1)^2)}{2^k}}.$$

¹ A set of slides published by Intel [16] claims a much stronger result based on Theorem 1. However, in addition to failing to account for point (1) above, the difference appears to stem from a mistake in translating notation. Specifically, the above theorem from [4] writes the second term under the radical as $K \cdot \epsilon(L, K)$, where $\epsilon(L, K) = \mathcal{O}(L^2/K^2)$ and in our notation $K = 2^k$. The Intel slides, however, appear to have mistranscribed this term as $L \cdot \epsilon(L, K)$ (in their notation, $L = b$ and $K = 2^n$). Since $L \ll K$ for values of interest, Intel’s claim significantly underestimates the concrete security bound.

Now, $C_i^i|_{\mathcal{P}} = \text{CBCMAC}_{\pi}^{R_i}(c^R)$. However, given that c^R is fixed, $\tau_{\pi, c^R}(\cdot) = \text{CBCMAC}_{\pi}(\cdot)(c^R)$ is a permutation for every possible π . Therefore since R_i is uniformly distributed,

$$\begin{aligned}
\Delta((\pi, r, R_i, s), (\pi, r, C_i^i|_{\mathcal{P}}, s)) &= \frac{1}{2} \sum_{(\rho, x)} |\Pr[(\pi, R_i) = (\rho, x)] - \Pr[(\pi, C_i^i|_{\mathcal{P}}) = (\rho, x)]| \\
&= \frac{1}{2} \sum_{(\rho, x)} \left| \frac{1}{2^k} - \Pr[C_i^i|_{\mathcal{P}} = x \mid \pi = \rho] \right| \Pr[\pi = \rho] \\
&= \frac{1}{2} \sum_{(\rho, x)} \left| \frac{1}{2^k} - \Pr[R_i' = \tau_{\rho, c^R}^{-1}(x) \mid \pi = \rho] \right| \Pr[\pi = \rho] \\
&= \frac{1}{2} \sum_{(\rho, x)} |\Pr[(\pi, R_i) = (\rho, x)] - \Pr[(\pi, R_i') = (\rho, x)]| \\
&= \Delta((\pi, r, R_i, s), (\pi, r, R_i', s)),
\end{aligned}$$

with summations over $(\rho, x) \in \text{Perm}(k) \times \{0, 1\}^k$. This completes the proof. \square

It remains to show that, with high probability, the (potentially) truncated extractor input contains sufficient min-entropy. Note that making reasonable min-entropy assumptions regarding the entropy source is not sufficient; for example, the approximate 1% false-positive rate of the health tests on uniformly random 256-bit strings implies that there are at least 2^{249} *unhealthy* strings. Therefore the entropy source could produce *only* unhealthy samples, resulting in unbounded L , and still have high min-entropy. In order to avoid such pathological behavior, we will later (in Section 7.2) need to introduce additional assumptions regarding the rate at which the entropy source produces healthy samples. Ultimately, we will choose L such that we have a high probability of never needing more than $L/2$ samples, but such that $L/2^{k/2}$ is small, as this term will dominate our security bounds.

5 Modeling the ISK-RNG as a PWI

Building upon DPRVW, here we define the syntax of a PWI. We give the syntax first, and then discuss what it captures, pointing out where our definition differs from DPRVW.

5.1 The PWI model

Definition 3 (PWI). *Let p , and ℓ be non-negative integers, and let $\text{IFace}, \text{Seed}, \text{State}$ be non-empty sets. A PRNG with input (PWI) with interface set IFace , seed space Seed , and state space State is a tuple of deterministic algorithms $\mathcal{G} = (\text{setup}, \text{refresh}, \text{next}, \text{tick})$, where*

- **setup** takes no input, and generates an initial PWI state $S_0 \in \text{State}$. Although **setup** itself is deterministic, it may be provided oracle access to an entropy source \mathcal{D} , in which case its output S_0 will be a random variable determined by the coins of \mathcal{D} .
- **refresh** : $\text{Seed} \times \text{State} \times \{0, 1\}^p \rightarrow \text{State}$ takes a seed $\text{seed} \in \text{Seed}$, the current PWI state $S \in \text{State}$, and string $I \in \{0, 1\}^p$ as input, and returns new state.
- **next** : $\text{Seed} \times \text{IFace} \times \text{State} \rightarrow \text{State} \times (\{0, 1\}^{\ell} \cup \{\perp\})$ takes a seed, the current state, and an interface label $m \in \text{IFace}$, and returns a new state, and either ℓ -bit output value or a distinguished, non-string symbol \perp .

- `tick` : $\text{Seed} \times \text{State} \rightarrow \text{State}$ takes a seed and the current state as input, and returns a new state. □

We will typically omit explicit mention of the the seed argument to `refresh`, `next` and `tick`, unless it is needed for clarity

The `setup` algorithm captures the initialization of the PWI, in particular its internal state. Unlike DPRVW, whose syntax requires `setup` to generate the PWI seed, we view the seed as something generated externally and provided to the PWI. Permitting an explicit setup procedure is necessary to correctly model ISK-RNG and, more generally, allows us to formulate an appropriate security definition for PWI initialization.

The `refresh` algorithm captures the incorporation of new entropy into the PWI state. Like DPRVW, we treat the entropy source as external. This provides a clean and general way to model the source as untrusted to provide consistent, high-entropy outputs.

Our `next` algorithm captures the interface exposed to (potentially adversarial) parties that request PWI outputs. By embellishing the DPRVW syntax for `next` with the interface set `interface`, we model APIs that expose multiple functionalities that access PWI state. This is certainly the case for the ISK-RNG, via the `RDRAND` and `RDSEED` instructions, as well as `/dev/[u]random`. We also model blocking by letting `next` return \perp .

The `tick` algorithm is entirely new, and requires some explanation. In the security notions formalized by DPRVW, the passage of “time” is implicitly driven by adversarial queries. (This is typical for security notions, in general.) But real PRNGs like the ISK-RNG may have behaviors that update the state in ways that are not cleanly captured by an execution model that is driven by entropy-input events (`refresh` calls), or output-request events (`next` calls). The `tick` algorithm handles this, while allowing our upcoming security notions to retain the tradition of being driven by adversarial queries: the adversary will be allowed to “query” the `tick` oracle, causing one unit of time to pass and state changes to occur.

5.2 Mapping ISK-RNG into the PWI model

We now turn our attention to mapping the ISK-RNG specification into the PWI model. Figure 3 summarizes the state that our model tracks. Figure 4 provides our model for the PWI `setup`, `refresh`, `next`, and `tick` oracles. Two additional procedures, `DRBG` and `reseed`, are used internally.

6 PWI Security

Having defined the syntax for PWIs, we can now introduce corresponding security notions. The basic notions are those of DPRVW, with a few notable alterations. To handle issues of non-uniform state and (more) realistic initialization procedures, we introduce a new technical tool, masking functions, that allows us to cleanly address these issues.

6.1 Basic notions

Here we define four PWI-security notions, in the game-playing framework [2]. In each there is a (potentially adversarial) entropy source \mathcal{D} , and an adversary A . The latter is provided access to the oracles detailed in Figure 5 (top), and what distinguishes the four notions are restrictions applied to the queries of the adversary A . In particular, we consider the following games:

Robustness (ROB): no restrictions on queries.

Variable	Bits	Description
ESSR	256	Entropy source shift register
window	8	Counts new bits in the ESSR
OSTE ₁	256	} Online self-tested entropy buffers
OSTE ₂	256	
CE ₀	128	} Conditioned entropy buffers
CE ₁	128	
ptr	1	Tracks CE buffer to condition next
health	256	Tracks health of last 256 ES samples
K	128	DRBG key (For AES-CTR)
IV	128	DRBG IV (For AES-CTR)
out _{1,...,8}	512	Eight 64-bit output buffers
outcount	≥ 4	Counts number of full output buffers
count	≥ 9	Counts DRBG calls since reseeding
CEfull	1	Set if CE buffers are available
block	1	Set if reseed has priority over RDSEED

Fig. 3. State variables of the ISK-RNG

Forward security (FWD): no queries to `set-state` are allowed; and a single query to `get-state` is allowed, and this must be the final query.

Backward security (BWD): no queries to `get-state` are allowed; a single query to `set-state` is allowed, and this must be the first query.

Resilience (RES): no queries to `get-state` or `set-state` are allowed.

See DPRVW for additional discussion. We note that all games share common `initialize` and `finalize` procedures, shown in Figure 5 (bottom). Thus, the robustness-advantage of A in attacking \mathcal{G} is defined to be $\mathbf{Adv}_{\mathcal{G},\mathcal{D}}^{\text{rob}}(A) = 2 \Pr[\text{ROB}_{\mathcal{G},\mathcal{D}}(A) = 1] - 1$. The forward security, backward security, and resilience advantages $\mathbf{Adv}_{\mathcal{G},\mathcal{D}}^{\text{fwd}}(A)$, $\mathbf{Adv}_{\mathcal{G},\mathcal{D}}^{\text{bwd}}(A)$, and $\mathbf{Adv}_{\mathcal{G},\mathcal{D}}^{\text{res}}(A)$ are similarly defined. It is clear that robustness implies forwards and backwards security, and both of these independently imply resilience.

We note that, because the PRNG cannot reasonably be expected to produce random-looking outputs without sufficient entropy or with a known or corrupted state, the various security experiments track (1) a boolean variable `corrupt` and (2) a value γ measuring the total entropy that has been fed into the PRNG since `corrupt` was last set. These serve as book-keeping devices to prevent trivial wins. The `corrupt` flag is cleared whenever γ exceeds some specified threshold γ^* .

6.2 Masking functions and updated security notions

As noted earlier, the DPRVW security definitions assume the PWI state is initially uniformly random. However, this does not realistically model the behavior of real-world PWIs, notably ISK-RNG, which do not attempt to reach a pseudorandom state; for example, they may maintain counters. (Indeed one can construct PWIs that would be deemed secure when starting from a uniformly random state, but that would not be secure in actuality; the reverse is also true. See Appendix A for examples.) Yet, clearly, some portion of the PWI state must be unpredictable to an attacker, as otherwise one cannot expect PWI outputs to look random.

To better capture real-world characteristics of PWI state, we introduce the idea of a *masking function*. A masking function M over state space `State` is a randomized algorithm from `State`

```

Oracle setup(ES):
01 for  $i = 1, 2, 3, 4$  do
02    $S.CE_0 \leftarrow \text{CBCMAC}_{K'}(S.CE_0)$ 
03   while  $S.ptr = 0$  do
04      $I \xleftarrow{\$} \text{ES}$ 
05      $S \leftarrow \text{refresh}(S, I)$ 
06      $S.CE_1 \leftarrow \text{CBCMAC}_{K'}(S.CE_1)$ 
07     while  $S.ptr = 1$  do
08        $I \xleftarrow{\$} \text{ES}$ 
09        $S \leftarrow \text{refresh}(S, I)$ 
10      $S \leftarrow \text{reseed}(S)$ 
11   for  $i = 1, 3, 5, 7$  do
12      $(S, R) \leftarrow \text{DRBG}(S)$ 
13      $S.out_i \parallel S.out_{i+1} \leftarrow R$ 
14    $S.outcount \leftarrow 8$ 
15   return  $S$ 

Oracle DRBG(S):
16  $S.IV \leftarrow S.IV + 1$ 
17  $R \leftarrow \text{CTR}_K^V(0^{128})$ 
18 if  $S.CEfull$  then
19    $S \leftarrow \text{reseed}(S)$ 
20 else if  $S.count < 512$ 
21    $S.K \parallel S.V \leftarrow \text{CTR}_{S.K}^{S.V+1}(0^{256})$ 
22    $S.count \leftarrow S.count + 1$ 
23 else
24   return  $(S, \perp)$ 
25   return  $(S, R)$ 

Oracle tick(S):
26 if  $S.CEfull$  and  $S.count > 0$  then
27    $S \leftarrow \text{reseed}(S)$ 
28   return  $S$ 
29 if  $S.count < 512$  then
30   if  $S.outcount < 8$  then
31      $S.outcount \leftarrow S.outcount + 1$ 
32      $(S, R) = \text{DRBG}(S)$ 
33      $S.out_{outcount} \leftarrow R$ 
34   return  $S$ 
35   return  $S$ 

Oracle refresh(S, I):
36 shift( $S.ESSR, I$ )
37  $S.window \leftarrow S.window + 1 \bmod 256$ 
38 if  $S.window = 0$  then
39   shift( $S.health, \text{isHealthy}(S.ESSR)$ )
40    $S.OSTE_2 \leftarrow S.OSTE_1$ 
41    $S.OSTE_1 \leftarrow S.ESSR$ 
42    $i \leftarrow S.ptr$ 
43    $I_j^i \leftarrow I_j^i \parallel S.OSTE_2$  // Record-keeping
44    $S.CE_i \leftarrow \text{CBCMAC}_{K'}^{S.CE_i}(OSTE_2)$ 
45   if  $\text{sum}(S.health) \geq 128$  then
46     if  $\text{isHealthy}(OSTE_2)$  then
47        $S.samples \leftarrow S.samples + 1$ 
48     if  $S.samples = m$  then
49        $S.samples \leftarrow 0$ 
50     if  $S.ptr = 0$  then
51        $S.ptr \leftarrow 1$ 
52     else
53        $S.ptr \leftarrow 0$ ;  $S.CEfull \leftarrow 1$ 
54      $C_j^0 \parallel C_j^1 \leftarrow S.CE$  // Record-keeping
55      $j \leftarrow j + 1$ ; // Record-keeping
56   return  $S$ 

Oracle reseed(S):
57  $S.K \parallel S.V \leftarrow \text{CTR}_K^{V+1}(S.CE)$ 
58  $S.CE_0 \leftarrow \text{CBCMAC}_{K'}(S.CE_0)$ 
59  $S.CE_1 \leftarrow \text{CBCMAC}_{K'}(S.CE_1)$ 
60  $S.count \leftarrow 0$ ;  $S.CEfull \leftarrow 0$ 
61  $S.ptr \leftarrow 0$ ;  $S.block \leftarrow 0$ 
62   return  $S$ 

Oracle next(interface, S):
63 if interface = RDRAND then
64   if  $S.outcount = 0$  then return  $(S, \perp)$ 
65    $R \leftarrow \text{LSB}_{64}(S.out_1)$ 
66   for  $i = 1, \dots, 7$  do
67      $S.out_i \leftarrow S.out_{i+1}$ 
68    $S.outcount \leftarrow S.outcount - 1$ 
69   return  $(S, R)$ 
70 else if interface = RDSEED
71   if  $S.CEfull = 0$  then
72     return  $(S, \perp)$ 
73   if  $S.block = 1, S.count > 0$  then
74     return  $(S, \perp)$ 
75    $R \leftarrow S.CE_0 \parallel S.CE_1$ 
76    $S.CEfull \leftarrow 0$ ;  $S.ptr \leftarrow 0$ 
77    $S.CE_0 \leftarrow \text{CBCMAC}_{K'}(S.CE_0)$ 
78    $S.CE_1 \leftarrow \text{CBCMAC}_{K'}(S.CE_1)$ 
79    $S.block \leftarrow 1$ 
80   return  $(S, R)$ 

```

Fig. 4. The above oracles describe the behavior of ISK-RNG from within the PWI model. See Table 3 for a description of the state variables $S.*$. All bits are initially zero. For Ivy Bridge chips, $m = 2$, and for Broadwell chips $m = 3$. The key $K' = \text{AES}_0(1)$ is fixed across all chips. The function $\text{shift}(x, y)$ sets value of x to the right-most $|x|$ bits of $x \parallel y$. Lines marked with a “Record-keeping” comment are there to aid in proofs and exposition.

<p>Oracle \mathcal{D}-refresh:</p> <pre> $(\sigma, I, \gamma, z) \stackrel{\\$}{\leftarrow} \mathcal{D}(\sigma)$ $S \leftarrow \text{refresh}(S, I)$ $c \leftarrow c + \gamma$ if $c \geq \gamma^*$ then $\text{corrupt} \leftarrow \text{false}$ return (γ, z) </pre>	<p>Oracle next-ror(m):</p> <pre> if corrupt then return \perp $(S, R_0) \leftarrow \text{next}(m, S)$ if $R_0 = \perp$ then $R_1 \leftarrow \perp$ else $R_1 \stackrel{\\$}{\leftarrow} \{0, 1\}^\ell$ return R_b </pre>	<p>Oracle get-next(m):</p> <pre> $(S, R) \leftarrow \text{next}(m, S)$ if corrupt then $c \leftarrow 0$ return R </pre> <p>Oracle wait:</p> <pre> $S \leftarrow \text{tick}(S)$ return ε </pre>	<p>Oracle get-state:</p> <pre> $c \leftarrow 0$ $\text{corrupt} \leftarrow \text{true}$ return S </pre> <p>Oracle set-state(S^*):</p> <pre> $c \leftarrow 0$ $\text{corrupt} \leftarrow \text{true}$ $S \leftarrow S^*$ </pre>
--	--	--	--

<p>Procedure initialize:</p> <pre> $\sigma \leftarrow 0$; $\text{seed} \stackrel{\\$}{\leftarrow} \text{Seed}$; $i \leftarrow 0$ $S \leftarrow \text{setup}^{\text{ES}}$ $c \leftarrow n$; $\text{corrupt} \leftarrow \text{false}$ $b \stackrel{\\$}{\leftarrow} \{0, 1\}$ return $(\text{seed}, (\gamma_j, z_j)_{j=1}^i)$ </pre>	<p>Oracle ES:</p> <pre> $i \leftarrow i + 1$ $(\sigma, I, \gamma_i, z_i) \stackrel{\\$}{\leftarrow} \mathcal{D}(\sigma)$ return I </pre>	<p>Procedure finalize(b):</p> <pre> if $b = b^*$ then return 1 else return 0 </pre>
--	---	---

Fig. 5. Top: Oracles for the PWI security games. **Bottom:** the shared initialize and finalize procedures for the PWI security games. Recall that the output of initialize is provided to adversary A as input, and the output of finalize is the output of the game.

to itself. As an example, if states consist of a counter c , a fixed identifier id , and a buffer B of (supposedly) entropic bits, then $M(c, \text{id}, B)$ might be defined to return (c, id, B') where B' is sampled by M from some distribution.

A masked state is meant to capture whatever characterizes a “good” state of a PWI, i.e. after it has accumulated a sufficient amount of externally provided entropy. Informally, for any state S , we want that (1) a PWI with state $M(S)$ should produce pseudorandom outputs, and (2) after the PWI has gathered sufficient entropy, its state S should be indistinguishable from $M(S)$.

To the second point, the initial PWI state S_0 is of particular importance. In the following definition, we characterize masking functions M such that the initial S_0 and $M(S_0)$ are indistinguishable.

Definition 4 (Honest-initialization masking functions). Let \mathcal{D} be an entropy source, $\mathcal{G} = (\text{setup}, \text{refresh}, \text{next})$ be a PWI with state space State , A be an adversary, and $M : \text{State} \rightarrow \text{State}$ be a masking function. Let (seed, Z) be the random variable returned by running the initialize() (Figure 5) using \mathcal{G} and \mathcal{D} , and let S_0 be the state produced by this procedure. Set $\text{Adv}_{\mathcal{G}, \mathcal{D}, M}^{\text{init}}(A) = \Pr[A(S_0, \text{seed}, Z) \Rightarrow 1] - \Pr[A(M(S_0), \text{seed}, Z) \Rightarrow 1]$. If $\text{Adv}_{\mathcal{G}, \mathcal{D}, M}^{\text{init}}(A) \leq \epsilon$ for any adversary A running in time t , then M is a $(\mathcal{G}, \mathcal{D}, t, \epsilon)$ -honest-initialization masking function. \square

Note that the above definition is made with respect to a specific \mathcal{D} . The assumptions required of \mathcal{D} (e.g., that it will provide a certain amount of entropy within a specified number of queries) will depend on the PWI in question, but should be as weak as possible.

We now define “bootstrapped” versions of the PWI security goals, which always begin from a masked state. This will allow us to reason about security when the PWI starts from an “ideal” state, i.e. what we expect after an secure initialization of the system.

Definition 5 (Bootstrapped security). Let \mathcal{G} be a PWI and M be a masking function. For $x \in \{\text{fwd}, \text{bwd}, \text{res}, \text{rob}\}$, let $\text{Adv}_{\mathcal{G}, \mathcal{D}}^{x/M}(A)$ be defined as $\text{Adv}_{\mathcal{G}, \mathcal{D}}^x(A)$, except with line 02 of the initialize procedure (Fig. 5) changed, to execute instead $S' \stackrel{\$}{\leftarrow} \text{setup}^{\text{ES}}; S \stackrel{\$}{\leftarrow} M(S')$. \square

6.3 PWI-Security Theorems

Bootstrapped security notions are useful, because they allow the analysis to begin with an idealized state. However, this comes at a cost: we need to ensure that the masking function is honest in the sense that it accurately reflects the result of running the `setup` procedure. The following theorem states the intuitive result that if the masking function is secure (and honest), then security when the PWI begins in a masked state $M(S)$ implies security when the PWI begins in state S . We omit the simple proof, which follows from a standard reduction argument.

Theorem 3. *Let \mathcal{G} be a PWI, \mathcal{D} be an entropy source, and M be a masking function. Suppose M is a $(\mathcal{G}, \mathcal{D}, t, \epsilon)$ -honest initialization mask. Then for any $x \in \{\text{fwd}, \text{bwd}, \text{res}, \text{rob}\}$ there exists some adversary $B(\cdot)$ such that for any adversary A , $\mathbf{Adv}_{\mathcal{G}, \mathcal{D}}^x(A) \leq \mathbf{Adv}_{\mathcal{G}, \mathcal{D}}^{x/M}(B(A)) + \epsilon$. Further, if it takes time t' to compute M , and A makes q queries and runs in time t , then $B(A)$ makes q queries and runs in time $\mathcal{O}(t) + t'$.*

For a second general result, we revisit a nice theorem by DPRVW and adapt it to our model. The theorem states that if a PWI possesses two weaker security properties — roughly, the ability to randomize a corrupted state after harvesting sufficient entropy and the ability to keep its state pseudorandom in the presence of adversarial entropy — then it is robust. These definitions, however, again assume that a state “should” appear uniformly random. We present modified definitions that instead use masking functions, and prove an analogous theorem. While the transition involves a couple subtleties — in particular, we require an idempotence property of the masking function — the proof is essentially identical to the one in [5]; therefore we make an informal statement here and defer the formal treatment to Appendix B.

Theorem 4 (Informal). *Let \mathcal{G} be a PWI. Suppose there exists a mask M such that: (1) When starting from an arbitrary initial state S of the adversary’s choosing, the final PWI state S' is indistinguishable from $M(S')$ provided the PWI obtains sufficient entropy; (2) When starting from an initial state $M(S)$ (for adversarially chosen S), the final PWI state S' is indistinguishable from $M(S')$, even if the adversary controls the intervening entropy input strings; (3) \mathcal{G} produces pseudorandom outputs when in a masked state. Then \mathcal{G} is robust.*

7 Security of the ISK-RNG as a PWI

We are now positioned to analyze the security of ISK-RNG. To begin, we demonstrate some simple attacks that violate both forwards and backwards security (hence robustness, too). Next, we show that by placing a few additional restrictions on adversaries — restrictions that are well-motivated by the hardware — we can recover forward security. As we said in our introduction, the concrete security bounds we prove are not as strong as one might hope, due to some limitations of CBCMAC’s effectiveness as an entropy extractor in the ISK-RNG. However, we are able to prove somewhat better results when legitimate parties use only the `RDRAND` interface, even when attackers also have access to `RDSEED`. This means that, e.g., a hostile process can’t use its access to `RDSEED` to learn information about `RDRAND` return values used by a would-be victim; the result also implies stronger results for Ivy Bridge chips, where `RDSEED` is not available.

For the remainder of Section 7, we fix the following constants: $p = 1$ is the length of each entropy input; $k = 128$ is the length of each CBCMAC input block (since ISK-RNG uses AES); $\text{IFace} = \{\text{RDSEED}, \text{RDRAND}\}$ are the ISK-RNG interfaces; $m = 2, 3$ is the number of healthy samples

required by Ivy Bridge and Broadwell, respectively, before the CE buffer is available; and $\ell = 64$ is the length of the PWI outputs. Although RDRAND also allows programs to request 16 or 32 bits, this is implemented by fetching then truncating a 64-bit output, and similarly with RDSEED [12]. Therefore we assume without loss of generality that the adversary only requests the full 64 bits.

Recall that in the PWI model, the entropy source leaks information γ about each input string. We assume that every 256th such string (each one a single bit, $p = 1$) leaks the health of the corresponding 256 bit string (as determined by the online health test). Hence the adversary will always know the health of the OSTE buffers and the value of the health buffer. This is not simply a convenience: because the CE buffer is not available until it has been reconditioned with m healthy samples, RDSEED may leak health information through a timing side channel.

When the CE buffer is available, it can be used to reseed the DRBG or to service a RDSEED instruction. Priority is given to whichever was not last used [12]. However, because the PWI model cannot describe pending RDSEED instructions, the adversary must explicitly use its wait oracle to yield when it has priority: a wait invocation uses the CE to reseed, while a RDSEED invocation returns its contents.

The adversary’s wait oracle also allows us to account for the fact that updating the eight 64-bit output buffers is not an atomic operation. By using the tick function (invoked by wait) to only fill one at a time, we conservatively allow the adversary to control if a reseeding operation intervenes. Note that tick will reseed rather than fill an output buffer if reseeding is desired ($S.\text{count} > 0$) and possible ($S.\text{CEfull} = 1$). This reflects the priorities of the hardware [12].

In order to save power, the entropy source goes to sleep if all the output buffers are full, the CE buffer is available, and no RDRAND instructions have been processed since the last reseed [12]. The PWI model, however, requires that we continue to provide \mathcal{D} -refresh access to the adversary. Our decision to leak health information to the adversary allows us to avoid any problems here: the adversary knows when the entropy source sleeps, so we can restrict the adversary to not make \mathcal{D} -refresh calls when it does.

To make this power-saving hardware constraint “work” with the PWI model, we assume that each healthy 256-bit block produced by the entropy source contains at least γ bits of min-entropy. Formally, define $(\sigma_i, b_i, \gamma_i, z_i) = \mathcal{D}(\sigma_{i-1})$ for $i \geq 1$ (where $\sigma_0 = \varepsilon$), and let $I_i = b_{256i}b_{256i+1} \cdots b_{256i+255}$. We assume $\mathbf{H}_\infty(I_i \mid (\sigma_j, I_j, \gamma_j, z_j)_{j \neq i}, I_i \in \mathcal{H}) \geq \gamma$, for some $\gamma > 0$, and require that $\sum_{j=256i}^{256i+255} \gamma_j \geq \gamma$ whenever $I_i \in \mathcal{H}$. We set $\gamma^* = m\gamma$ to demand, in effect, that ISK-RNG delivers on its implicit promise that m healthy entropy samples are sufficient. At the end of this section, we will draw from the CRI report’s analysis to find reasonable estimates for γ and discuss the implications.

7.1 Negative Results

We begin with some quick negative results, showing that the ISK-RNG achieves neither forward nor backwards security. This immediately rules out robustness, too. We again emphasize that these negative results will be followed by positive results for realistic classes of restricted adversaries; we present them primarily to motivate the coming restrictions.

Theorem 5 (ISK-RNG lacks forward security). *There exists an adversary A making one next- ror query and one get-state query such that for any entropy source \mathcal{D} , $\text{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}}(A) = 1 - 2^{-128}$.*

Theorem 6 (ISK-RNG lacks backward security). *There exists an adversary A making one next- ror query and one set-state query such that for any entropy source \mathcal{D} , $\text{Adv}_{\text{ISK}, \mathcal{D}}^{\text{bwd}}(A) = 1 - 2^{-128}$.*

In the case of backwards security, the adversary sets some initial state S with $S.\text{samples} = 0$, makes a sequence of \mathcal{D} -refresh calls to clear the corrupt flag (which, by our previously state assumptions, will happen as soon as the CE buffer becomes available), and finally assigns $X \leftarrow \text{next-ror}(\text{RDRAND})$. The adversary then checks if $X = S.\text{out}_1$, and outputs 0 if this is the case and 1 otherwise. For forward security, the adversary assigns $X \leftarrow \text{next-ror}(\text{RDSEED})$, then learns the resulting state S using $\text{get-state}()$. If $X = \text{AES}_0^{-1}(S.\text{CE}_0) \parallel \text{AES}_0^{-1}(S.\text{CE}_1)$, the adversary outputs 0; otherwise, the it outputs 1. (Here, $\mathbf{0} = 0^{128}$.)

However, these results are very conservative. In the case of forward security, the hardware will quickly recondition the CE buffer and refill the output buffers, effectively erasing all state that could be used to compute previous outputs. Backwards security is more complicated because not only do future outputs persist in the output buffer indefinitely, but future DRBG keys are leaked via the ESSR, OSTE, and CE buffers. Once the output buffers are flushed, though, these other buffers will quickly be overwritten with fresh entropy.

7.2 Positive results

We now turn our attention to restricted, but still conservative, classes of adversary in order to produce positive results.

Additional assumptions. We further assume that in the forward-security game, adversaries do not make their get-state query until they have allowed the output buffers to be refilled. This assumption is motivated by the speed with which the hardware will automatically accomplish this: at the reported RDRAND throughput of 500 MB/s, all eight 64-bit buffers can be refilled around 8 million times per second. Formally:

Definition 6 (Delayed adversaries). *An adversary A attacking ISK-RNG in the forward-security game is delayed if after making its last get-next and next-ror queries, A calls \mathcal{D} -refresh until the CE buffer is available, then calls wait nine times before making its get-state query. \square*

This will trigger a reseed and then refill any empty output buffers.

Moreover, we will assume there is some positive probability β such that each 256-bit block of bits from the entropy source is healthy with probability at least β . Formally (recall that $\mathcal{H} \subseteq \{0, 1\}^{256}$ is the set of strings deemed healthy by ISK-RNG’s online health tests):

Definition 7 (β -healthy). *Let \mathcal{D} be an entropy source and fix $\beta > 0$. Let $\mathcal{H} \subseteq \{0, 1\}^{256}$ be the set of strings deemed healthy by the ISK-RNG. For $i = 1, 2, 3, \dots$ define $(\sigma_i, b_i, \gamma_i, z_i) = \mathcal{D}(\sigma_{i-1})$ (where $\sigma_0 = \varepsilon$), and for $j = 0, 1, 2, \dots$, define $B_j = b_{256j} \parallel b_{256j+1} \parallel \dots \parallel b_{256j+255}$. Let $H_j = 1$ if $B_j \in \mathcal{H}$, and set $H_j = 0$, otherwise. Then \mathcal{D} is β -healthy if for all such j and all $H \in \{0, 1\}^{j-1}$, $\Pr[B_j \in \mathcal{H} \mid (H_\ell)_{\ell < j} = H] \geq \beta$. \square*

So for any positive integers ℓ and L_m , we can upper bound the probability that the sequence $(B_i)_{i=\ell}^{\ell+L_m-1}$ contains fewer than m healthy values using: $\Pr[|\{j : B_j \in \mathcal{H}, \ell \leq j < \ell + L_m\}| < m] \leq \sum_{i=0}^{m-1} \binom{L_m}{i} \beta^i (1 - \beta)^{L_m-i}$.

Remark 1. Our goal is to identify under what reasonable assumptions ISK-RNG could be deemed secure, and, as we argued at the end of Section 4, this requires making an assumption about the entropy source’s ability to produce “healthy” samples (a min-entropy assumption is too weak). We settled on the above β -healthy assumption because it is simple and fairly broad: we do not

k	CBCMAC blocksize (128 bits)
m	Number of “healthy” $2k$ -bit strings that need to be conditioned before the CE buffer becomes available ($m = 2, 3$ for Ivy Bridge and Broadwell chips, respectively).
L_m	Parameter we can freely choose to keep both $\hat{\epsilon}(L_m)$ and $\epsilon(L_m)$ small.
γ	An assumed lower bound on the conditional min-entropy of healthy strings.

Fig. 6. Summary of values used for theorem statements.

assume the probabilities of samples being healthy are constant or even independent, just that the conditional probabilities don’t dip below the β threshold. Moreover, we later show that the “unhealthy sample rate” could easily be fifty times the ideal 1% false-positive base rate without significantly damaging our bounds. Finally, even the β -healthy assumption is more than we need. We require an upperbound on the probability on the left-hand side of the above equation, and the β -healthy assumption provides a natural, concrete way to think about this probability.

Rigorously testing the β -healthy assumption without access to the entropy source is problematic. That being said, barring such access, we doubt it would be possible to do significantly better.

With these assumptions, we are ready to continue on to our positive results. Our first step is to define an appropriate masking function that describes an “ideal” state, and then to prove that setup creates such a state. This lets later proofs simply assume we begin in an idealized state (see Theorem 3).

ISK-RNG masking function. Fix the masking function $M : \{0, 1\}^n \rightarrow \{0, 1\}^n$ that on input S , overwrites $S.CE$, $S.K$, $S.IV$, and $S.out_{1,\dots,8}$ with independent, uniformly random strings of the appropriate lengths, leaves all other portions of the state untouched, and returns the result (refer back to Fig. 3 for a listing of the components of the ISK-RNG state S). This is the ISK-RNG masking function.

Recall the results of Theorem 2. For convenience, we define

$$\epsilon(L_m) = \mathcal{O}(L_m + 1)/2^{k/2} \quad \text{and} \quad \hat{\epsilon}(L_m) = \sum_{i=0}^{m-1} \binom{L_m}{i} \beta^i (1 - \beta)^{L_m - i},$$

where $\epsilon(L_m)$ is from Theorem 2 and $\hat{\epsilon}(L_m)$ is the above bound on the probability of obtaining fewer than m healthy samples from a β -healthy entropy source within L_m trials. Our theorem statements refer to various previously defined values, summarized in Figure 6

The following lemma says that if AES is a secure PRP (against adversaries making three queries) and each healthy sample from the entropy source has sufficiently large min-entropy, then the ISK-RNG masking function is honest. That is, that the ISK-RNG setup procedure successfully places the hardware in a state where (we will show) it can begin producing pseudorandom outputs.

Lemma 1 (ISK-RNG masking function is honest). *Fix positive integers k and m , and fix $0 < \beta \leq 1$. Let L_m be a positive integer. Let M be the ISK-RNG masking function. Let \mathcal{D} be a β -healthy entropy source. Then for any adversary A , there exists an adversary B running in the same time and making three queries such that $\mathbf{Adv}_{\text{ISK}, \mathcal{D}, M}^{\text{init}}(A) \leq 2^{(k-m\gamma)/2+2} + 4\epsilon(L_m) + 8\hat{\epsilon}(L_m) + 5(\mathbf{Adv}_{\text{AES}}^{\text{PRP}}(B) + \frac{3}{2^k})$.*

The proof is deferred to Appendix E. Using reasonable estimates for the big- \mathcal{O} constant and γ (discussed in Section 7.3) provides us with an upper bound of roughly 2^{-60} for the first three terms of the security bound for both $m = 2, 3$.

Remark 2. The PRP term may be problematic if one takes the view that RDSEED should offer information-theoretic security. That is, Lemma 2 says that the ISK-RNG initialization procedure yields state — which includes the CE buffers — that is only computationally indistinguishable from “ideal”. However, we observe that if one adjusts the masking function to leave the output buffers unchanged, and demands a post-setup reconditioning (which the hardware endeavors to provide, anyway), one could indeed use the result to prove information-theoretic RDSEED security. However, this would be at the expense of *not* being able to prove security of the RDRAND interface, a task which necessarily requires computational assumptions.

Forward security. Our exploration of forward security proceeds in two steps. To begin, we introduce a new game, M -RDRAND, which differs from M -FWD in that the next-ror oracle always returns the “real” value R_0 when queried on the RDSEED interface, but behaves normally during queries to the RDRAND interface. Define

$$\mathbf{Adv}_{\mathcal{G}, \mathcal{D}}^{\text{fwd-RDRAND}/M}(A) = 2 \Pr [M\text{-RDRAND}(A) \Rightarrow 1] - 1.$$

Proving the security of this game is not only a useful intermediate step in proving the security of M -FWD, but also can be interpreted as measuring the strength of RDRAND return values when an adversary also has access to the RDSEED instruction (which can be used to learn information about the ISK-RNG state, but that we do not require to return pseudorandom values). This distinction is valuable, because the concrete bounds on the M -FWD experiment are not as strong as one would hope.

Theorem 7 (M -RDRAND). *Let A be a delayed adversary making q queries to RDRAND and running in time t . Then there exists an adversary B making three queries and running in time $\mathcal{O}(t)$ such that $\mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd-RDRAND}/M}(A) \leq 2(q + 4) (\mathbf{Adv}_{\text{AES}}^{\text{PRP}}(B) + \frac{3}{2^k})$.*

The proof appears in Appendix E. Barring an efficient attack on AES (that only uses three queries!) this bound is quite strong. If q were to grow quite large, say on the order of $q \approx 2^{80}$, then the bound might begin to approach 2^{-40} , which seems a reasonable safety margin. However, even at the reported rate of around 500 MB/s, ISK-RNG would take over 70 years to reach this point. Moreover, the hybrid factor of q is likely a conservative artifact of the proof.

Note, however, that this bound applies to ISK-RNG when starting in an “ideal” masked state; one needs to add in the bound from Lemma 1 to account for initialization. As we mentioned earlier, reasonable estimates for the big- \mathcal{O} constant and γ (see Section 7.3) place this term at roughly 2^{-60} .

We now proceed to the “full” forward-security result, where both the RDRAND and the RDSEED interfaces are required to produce indistinguishable-from-random outputs. Since RDSEED reads directly from the CE buffer, this bound relies more heavily on the entropy source and CBCMAC extractor (and less on the computational security of AES). Again, see Appendix E for a proof.

Theorem 8 (ISK-RNG’s masked forward security). *Fix a positive integers k and m , and fix $0 < \beta \leq 1$. Let L_m be a positive integer. Let A be a delayed adversary making a combined q*

queries to `get-next` and `next-ror`. Then if \mathcal{D} is β -healthy, there exists some adversary B making three queries and running in the same time as A such that

$$\begin{aligned} \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}/M}(A) &\leq (q + 1) \left(2^{(k-m\gamma)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right) \\ &\quad + 2(q + 4) \left(\mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right). \end{aligned}$$

Corollary 1. *Let A be a delayed adversary making a combined q queries to its `get-next` and `next-ror` oracles. If \mathcal{D} is β -healthy, then there exists an adversary B making three queries and running in the same time as A such that*

$$\begin{aligned} \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}}(A) &\leq (q + 5) \left(2^{(k-m\gamma)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right) \\ &\quad + (2q + 13) \left(\mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right), \end{aligned}$$

where the remaining quantities are defined as in Theorem 8.

The corollary follows from applying Theorem 3 to Theorem 8 and Lemma 1. We defer our discussion of this bound to Section 7.3. First, we briefly turn our attention to the questions of backwards security and robustness.

Backwards security and Robustness. The issue with obtaining backwards security (and hence robustness) is that future outputs can linger in the output buffers indefinitely: the hardware will shutdown the entropy source after all the buffers are full and the CE buffer is available. Hence, state remains compromised until fresh entropy filters through the $\text{ESSR} \rightarrow \text{OSTE}_1 \rightarrow \text{OSTE}_2 \rightarrow \text{CE}$ buffers and is used to reseed the DRBG, without first being siphoned off by `RDSEED`.

Consider the worst-case scenario for Ivy Bridge chips, where only the `RDRAND` interface is available. Following a state compromise, the next eight outputs are revealed by the output buffers, the next 511 may be computed using the compromised DRBG seed, the next 511 may be computed using a DRBG seed determined by the compromised CE buffer, and the next 511 may be computed using a DRBG key determined by the compromised `OSTE` and `ESSR` buffers. This amounts to slightly more than 12 KB of outputs that an adversary could potentially predict.

However, we show in Appendix C that if one restricts the model to “read-only” adversaries (by denying adversaries access to `set-state` but permitting access to `get-state`) and one discounts wins based on the above attacks (by denying adversaries access to `next-ror` until after the “corrupted” values have already been replaced) then ISK-RNG is secure. The concrete bounds we obtain are essentially identical to those provided by Theorems 7 and 8, depending on whether or not one requires the `RDSEED` interface to be secure. See the appendix for further discussion of how these restrictions can be interpreted along with a formal theorem statement and proof.

7.3 Discussion of results

Let us examine the bound of Corollary 1 in detail. We specialize to the parameters used by Intel: $k = 128$ (a consequence of using AES), $m = 2$ for Ivy Bridge chips, and $m = 3$ for Broadwell chips.

To estimate γ , we turn to the CRI report [8]. Hamburg, Kocher, and Marson subjected raw entropy source bits (using data provided by Intel) to a battery of statistical tests. Using a Markov

model with 12 bits of state, they estimate the entropy source produces approximately 0.65 bits of min-entropy per bit of output. However, this was an average (some states of the Markov model resulted in more predictable bits), and a 12-bit state, though perhaps necessary to collect enough samples for a meaningful empirical analysis, is not enough for our purposes. Therefore let us suppose a more conservative rate of 0.5, leading to $\gamma = 128$.

This sets the $(q+5)2^{(k-m\gamma)/2}$ term of our bound to $(q+5)2^{-64}$ for Ivy Bridge (where $m = 2$) and $(q+5)2^{-128}$ for Broadwell (where $m = 3$). The latter bound is quite strong, but, given how quickly q can grow, the former may be worrisome if one wishes to maintain strong security guarantees (e.g., one wishes to cap an adversary’s advantage at 2^{-40}). But this is not the dominate term.

We next consider the term $(q+5)(\epsilon(L_m) + 2\hat{\epsilon}(L_m))$. If we set the big-O constant of ϵ to c (so $\epsilon(L_m) = cL/2^{64}$) then we can choose L_m to optimize this expression. Taking $\beta = 1/2$, $c = \sqrt{10}$, which we believe to be conservative,² gives an upper bound of $(q+5)2^{-56}$; a more generous $\beta = 0.99$, $c = 1$ improves the upper bound to about $(q+5)2^{-60}$. (These bounds are accurate for both $m = 2$ and $m = 3$, although the corresponding values for L_m differ considerably.)

At this point, limiting an adversary’s advantage to 2^{-40} is difficult—an adversarial process gathering random bits at the benchmarked rate of 500 MB/s could issue the maximum allowable number of queries in under one millisecond. Or at least, this is the case if we demand that RDSEED produces uniform random outputs. On the other hand, if one only needs RDRAND to be secure, then Theorem 7 suggests that limiting an adversary’s advantage to 2^{-40} is entirely reasonable; in this setting, we only pick up a single $4(\epsilon(L_m) + 2\hat{\epsilon}(L_m))$ term even after moving to the unmasked forward-security setting, with no troublesome multiplicative factor of q .

The remaining term, $(2q+13)(\mathbf{Adv}_{\text{AES}}^{\text{PRP}}(B) + 3/2^{128})$, is likely to be negligible (recall that B is permitted only three queries).

Our analysis does not point to any obvious, practical attacks (aside from the trivial ones that exploit the output buffers, though it seems a stretch to deem those practical). However, it exposes the CBCMAC extraction process as the likely weakest link, and quantifies the extent of that weakness. An actual attack would need to exploit how the specific output distribution of the entropy source interacts with CBCMAC under the fixed key K' .

Acknowledgements

The authors wish to thank DJ Johnston and Jesse Walker of Intel for answering our questions regarding the design and implementation details of ISK-RNG. Both Terashima and Shrimpton were supported by NSF grants CNS-0845610 and CNS-1319061.

References

1. Boaz Barak and Shai Halevi. A model and architecture for pseudo-random generation with applications to `/dev/random`. In *Proceedings of the 12th ACM conference on Computer and communications security*, pages 203–212. ACM, 2005.
2. Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *Advances in Cryptology—EUROCRYPT 2006*, pages 409–426. Springer, 2006.
3. Olivier Chevassut, Pierre-Alain Fouque, Pierrick Gaudry, and David Pointcheval. The twist-AUGmented technique for key exchange. In *Public Key Cryptography*, pages 410–426. Springer, 2006.

² An author of [4] assures us that the asymptotic constant is “certainly less than 10” (and our c is the square root of this constant). A perfect entropy source would give $\beta = 0.99$ since the health tests have a 0.01 false-positive rate.

4. Yevgeniy Dodis, Rosario Gennaro, Johan Håstad, Hugo Krawczyk, and Tal Rabin. Randomness extraction and key derivation using the CBC, cascade and HMAC modes. In *Advances in Cryptology—CRYPTO 2004*, pages 494–510. Springer, 2004.
5. Yevgeniy Dodis, David Pointcheval, Sylvain Ruhault, Damien Vergniaud, and Daniel Wichs. Security analysis of pseudo-random number generators with input: `/dev/random` is not robust. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 647–658. ACM, 2013.
6. Adam Everspaugh, Yan Zhai, Robert Jellinek, Thomas Ristenpart, and Michael Swift. Not-so-random numbers in virtualized linux and the Whirlwind RNG. *IEEE Symposium on Security And Privacy*, 2014.
7. Zvi Gutterman, Benny Pinkas, and Tzachy Reinman. Analysis of the Linux random number generator. In *Security and Privacy, 2006 IEEE Symposium on*, pages 15–pp. IEEE, 2006.
8. Mike Hamburg, Paul Kocher, and Mark E Marson. Analysis of Intel’s Ivy Bridge digital random number generator. *Online: http://www.cryptography.com/public/pdf/Intel_TRN_G_Report_20120312.pdf*, 2012.
9. Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J Alex Halderman. Mining your ps and qs: Detection of widespread weak keys in network devices. In *USENIX Security Symposium*, pages 205–220, 2012.
10. Gael Hofemeier. 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>, August 2012. Accessed May 2014.
11. Gael Hofemeier and Robert Chesebrough. Introduction to Intel AES-NI and Intel Secure Key instructions. <https://software.intel.com/en-us/articles/introduction-to-intel-aes-ni-and-intel-secure-key-instructions>, July 2012. Accessed May 2014.
12. JD Johnston (Intel). Personal communication, May 2014.
13. Patrick Lacharme, Andrea Röck, Vincent Strubel, and Marion Videau. The Linux pseudorandom number generator revisited. *IACR Cryptology ePrint Archive*, 2012:251, 2012.
14. John Mechalas. The difference between RDRAND and RDSEED. <https://software.intel.com/en-us/blogs/2012/11/17/the-difference-between-rdrand-and-rdseed>, November 2012. Accessed April 2014.
15. Jaikumar Radhakrishnan and Amnon Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13(1):2–24, 2000.
16. Jesse Walker. Conceptual foundations of the Ivy Bridge random number generator. http://www.ists.dartmouth.edu/docs/walker_ivy-bridge.pdf, November 2011.

A PWIs with non-pseudo-uniformly random state

We contend the security definitions of [5] are flawed. Recall that in [5], the experiments defining FWD, BWD, RES, and ROB security begin by sampling a uniform random state. The advantage of an adversary in, for example, the FWD security definition is equivalent to our definition for $\mathbf{Adv}^{\mathcal{S}\text{-fwd}}$, where $\mathcal{S}(\cdot)$ returns a fresh random string on each invocation.

Consider following PWI $\mathcal{G} = (\text{setup}, \text{refresh}, \text{next})$ with $2n$ bits of state. Define setup to return 0^{2n} , $\text{refresh}(S, I) = S$, and $\text{next}(S_0 \parallel S_1) = (S'_0 \parallel S_1, R)$, where $|S_0| = |S_1| = |S'_0| = |R| = n$, $S'_0 \parallel R = f_{S_0}(0^m) \parallel f_{S_0}(1^m)$, and f is a PRF. It isn’t hard to show that for any A making q queries, there exists some B making $2q$ queries such that $\mathbf{Adv}_{\mathcal{G}}^{\text{fwd}/\mathcal{S}}(A) \leq \mathbf{Adv}_f^{\text{prf}}(B)$. However, \mathcal{G} cannot reasonably be said to have forward security because using $\mathcal{S}(\cdot)$ to create the initial state does not reflect anything approximating the behavior of setup .

One could argue that this is simply an artifact of \mathcal{G} ’s lack of *backward* security. After all, it is the adversary’s knowledge concerning the initial state that is problematic here. But this flaw should not be surfaced in the definition of *forward* security. Granted, we do need forward security to start from a “healthy” state. Using an explicit setup procedure allows us to neatly avoid dictating what that state should be without requiring the security definitions themselves to handle the issue on an *ad-hoc* basis.

Further consider the (admittedly pathological) example $\mathcal{G}' = (\text{setup}', \text{refresh}, \text{next}')$, where

$$\text{next}'(S_0 \parallel S_1) = \begin{cases} \text{next}(S_0, S_1) & \text{If } S_1 = 0^n, \\ (f_{S_0}(0^m) \parallel S_0, f_{S_0}(1^m)) & \text{otherwise,} \end{cases}$$

and $\text{setup}'(\text{seed}, \mathcal{D})$ uses an entropy extractor (e.g., the one described in [5]) to produce a random value $R \in \{0, 1\}^n$, then outputs $R \parallel 0^n$. Now the state leaks the PRF key used to generate the most recent next value — *unless* the right n -bits of state are 0^n . Here, starting with a uniform random state can artificially “break” a PWI that would otherwise have forward security.

B On proving robustness

When Dodis et al. introduced their PWI model [5], they presented a theorem showing that any PWI possessing two relatively weak security properties, *preserving* and *recovering* security, is also robust. However, both of these security definitions assume that state should ideally be uniformly random. The theorem therefore isn’t useful for PWIs like ISK-RNG, for which this assumption is false (see Appendix A). In this section we present analogous definitions that instead use masking functions, and prove a corresponding theorem. Although this result cannot be applied to ISK-RNG, which lacks even backwards security, we offer it as a contribution towards developing a theory for PWIs.

First we define idempotent masking functions. This property will be a mathematical convenience in the coming proof. It allows an adversary in a reduction argument to blindly apply a masking function to a state that may or may not already be masked without worrying that doing so will cause the simulated environment and the “expected” environment to diverge.

Definition 8 (Idempotent masking functions). *A masking function $M : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is idempotent if for any state $S \in \{0, 1\}^n$, $M(S)$ and $M(M(S))$ are identically distributed random variables.* \square

Both the masking function we use for our analysis of the ISK-RNG (pg. 19) and the masking function implicitly used by DPRVW for their construction [5] (which returns a uniform random string) are idempotent. More generally, a masking function is idempotent if it “hides” some secret portion of the state using a random value from a distribution that depends only on the non-secret part of the state. That is, fix some masking function M and suppose the state space can be decomposed into a product $\text{Public} \times \text{Secret}$ such that for each $C \in \text{Public}$, there exists a distribution M_C with the property that when $(C', R') \stackrel{\$}{\leftarrow} M(C, R)$, $\Pr[C = C'] = 1$ and R' is distributed according to M_C . Then M is idempotent. This follows immediately from the definition.

The security experiments for Recovering security and Preserving security are shown in Figure 7. We refer to the two games as **Recover** and **Preserve**, respectively. Our definitions are equivalent to those of [5] if one specializes to a M that returns a string sampled from $\{0, 1\}^n$ and considers only non-blocking PWIs.

Roughly, a PWI has recovering security if even when starting in a compromised state, it will, after harvesting sufficient entropy, eventually return to a pseudo-random state (as described by the masking function) and begin producing pseudo-random outputs. It has preserving security if it can continue to maintain a pseudo-random state and produce pseudo-random outputs even if given arbitrary, adversarially controlled inputs from the entropy source. Intuitively, these properties should suffice to ensure robustness, and indeed, this is the case.

Definition 9 ((Witnessed) Recovering-Security). *A PWI $\mathcal{G} = (\text{setup}, \text{refresh}, \text{next}, \text{tick})$ has $(t, q_{\mathcal{D}}, \gamma^*, \epsilon)$ -recovering security, witnessed by the masking function M , if, for any attacker A and legitimate sampler \mathcal{D} , both running in time t , the recovering advantage*

$$\text{Adv}_{\mathcal{G}; M}^{\text{rec}}(A) = 2 \Pr[\text{Recover}_{\mathcal{G}}(A, \mathcal{D}, M) = 1] - 1$$

Experiment Recover($\mathcal{G}, A, \mathcal{D}, M$):

```
01 (setup, refresh, next)  $\leftarrow \mathcal{G}$ 
02 seed  $\xleftarrow{\$} \mathcal{S}$ 
03  $b \xleftarrow{\$} \{0, 1\}$ ;  $\sigma_0 \leftarrow 0$ ;  $\mu \leftarrow 0$ 
04 for  $k = 1, \dots, q_{\mathcal{D}}$  do
05    $(\sigma_k, I_k, \gamma_k, z_k) \leftarrow \mathcal{D}(\sigma_{k-1})$ 
06  $(S_0, d, \sigma') \xleftarrow{\$} A^{\text{get-refresh}}(\text{seed}, \gamma_1, \dots, \gamma_{q_{\mathcal{D}}}, z_1, \dots, z_{q_{\mathcal{D}}})$ 
07 if  $\mu + d > q_{\mathcal{D}}$  or  $\sum_{j=\mu+1}^{\mu+d} \gamma_j < \gamma^*$  then
08   return 0
09 for  $j = 1, \dots, d$  do
10    $S_j \leftarrow \text{refresh}(S_{j-1}, I_{\mu+j}, \text{seed})$ 
11  $(S_0^*, R_0^*) \leftarrow \text{next}(S_d)$ 
12  $S_1^* \xleftarrow{\$} M(S_0^*)$ 
13 if  $R_0^* = \perp$  then
14    $R_1^* \leftarrow \perp$ 
15 else
16    $R_1^* \xleftarrow{\$} \{0, 1\}^\ell$ 
17  $b^* \xleftarrow{\$} A(\sigma', S_b^*, R_b, I_{\mu+d+1}, \dots, I_{q_{\mathcal{D}}})$ 
18 if  $b^* = b$  then
19   return 1
20 else
21   return 0
```

Proc get-refresh():

```
22  $\mu \leftarrow \mu + 1$ 
23 return  $I_\mu$ 
```

Experiment Preserve(\mathcal{G}, A, M):

```
24 (setup, refresh, next)  $\leftarrow \mathcal{G}$ 
25 seed  $\xleftarrow{\$} \mathcal{S}$ ;  $b \xleftarrow{\$} \{0, 1\}$ 
26  $(S'_0, I_1, \dots, I_d, \sigma') \xleftarrow{\$} A(\text{seed})$ 
27  $S_0 \xleftarrow{\$} M(S'_0)$ 
28 for  $j = 1, \dots, d$  do
29    $S_j \leftarrow \text{refresh}(S_{j-1}, I_j, \text{seed})$ 
30  $(S_0^*, R_0^*) \xleftarrow{\$} \text{next}(S_d)$ 
31  $S_1^* \xleftarrow{\$} M(S_0^*)$ 
32 if  $R_0^* = \perp$  then
33    $R_1^* \leftarrow \perp$ 
34 else
35    $R_1^* \xleftarrow{\$} \{0, 1\}^\ell$ 
36  $b^* \xleftarrow{\$} A(\sigma', S_b^*, R_b^*)$ 
37 if  $b^* = b$  then
38   return 1
39 else
40   return 0
```

Fig. 7. Security experiments for recovering and preserving security, with respect to a witness masking function M .

is at most ϵ . \square

Definition 10 ((Witnessed) Preserving Security). A PWI \mathcal{G} has (t, ϵ) -preserving security witnessed by the (possibly randomized) function $M : \{0, 1\}^n \rightarrow \{0, 1\}^n$ if for any attacker A running in time t , the preserving advantage $\text{Adv}_{\mathcal{G}; M}^{\text{pres}}(A) = 2 \Pr[\text{Preserve}(\mathcal{G}, A, M) = 1] - 1$ is at most ϵ . \square

With this setup, we can state our PWI robustness security result. With the above definitions in place, the proof itself is essentially identical to that of Theorem 1 of DPRVW; our exposition closely follows theirs, and we include it for the sake of completeness. The only two subtleties are the need for the idempotence property of masking functions and the change to `Preserve` that places the initial (unmasked) state under the adversary’s control.

Theorem 9. [Formal statement of Thm. 4] Let q_N be a positive integer. If there is an idempotent masking function $M : \{0, 1\}^n \rightarrow \{0, 1\}^n$ that witnesses the (t, ϵ_r) -recovering and (t, ϵ_p) -preserving security of a PWI \mathcal{G} , and M is $(\mathcal{G}, \mathcal{D}, t, \epsilon_h)$ -honest, then \mathcal{G} is $(t', \epsilon_h + q_N(\epsilon_r + \epsilon_p))$ -robust.

Proof. We refer to queries to either the `get-next` or `next-ror` oracles as `next` queries. A `next` query is *uncompromised* if it is made when `corrupt = false`, and *compromised* otherwise. If at any point in the experiment `corrupt` is `true`, the next *uncompromised* `next` query is a *recovering* query. The remaining uncompromised `next` queries are *preserving* queries.

When an adversary makes a recovering query, we associate with it a *most recent entropy drain* (MRED) query: namely, the most recent query to `get-state`, `set-state`, or `get-next`. (We assume without loss of generality that the adversary never makes a query to `next-ror` when `corrupt = true`.) Note that between any recovering query and its associated MRED query (which set $c \leftarrow 0$), the adversary must have made a series of queries to `D-refresh`, generating a sequence of *recovering samples* $I = (I_i, I_{i+1}, \dots, I_{i+d})$ with the property that the corresponding entropy estimates $(\gamma_i, \gamma_{i+1}, \gamma_{i+d})$ sum to at least γ^* .

Let M -ROB be the robustness experiment were the post-setup state S is overwritten with $M(S)$. Define game G_i to be the same as M -ROB $_{\mathcal{G}, \mathcal{D}}$ except that for the first i `next` queries:

- If the query is to `next-ror`, the challenger sets $S \leftarrow M(S)$ after Line 9, and always returns R_1 (not R_b).
- If the query is to `get-next`, the challenger sets $S \leftarrow M(S)$ after Line 15, and then if $R \neq \perp$, overwrites $R \leftarrow^{\$} \{0, 1\}^\ell$.

So for the first i `next` queries, the state gets overwritten with a mask of the state, and the adversary receives random bits.

Further define $G_{i+1/2}$, which behaves like G_i when the $(i + 1)$ st `next` query is preserving, and like G_{i+1} otherwise. Note that $\Pr[G_0(A) = 1] = \Pr[M\text{-ROB}_{\mathcal{G}, \mathcal{D}}(A) = 1] + \epsilon_h$ and that $\Pr[G_{q_N}(A) = 1] = 1/2$. (The latter equality holds because in Game G_{q_N} , all of the oracle outputs are independent of b). Therefore:

$$\begin{aligned} \Pr[M\text{-ROB}_{\mathcal{G}, \mathcal{D}}(A) = 1] &\leq \sum_{i=0}^{q_N-1} (|\Pr[G_i(A, \mathcal{D}) = 1] - \Pr[G_{i+1/2}(A, \mathcal{D}) = 1]| \\ &\quad + |\Pr[G_{i+1/2}(A, \mathcal{D}) = 1] - \Pr[G_{i+1}(A, \mathcal{D}) = 1]|). \end{aligned}$$

We show in two following lemmata that the former absolute value is upper bounded by ϵ_p , while the latter is upper bounded by ϵ_r . Since $|\Pr[\text{ROB}_{\mathcal{G}, \mathcal{D}}(A) = 1] - \Pr[M\text{-ROB}_{\mathcal{G}, \mathcal{D}}(A) = 1]| \leq \epsilon_h$, this completes the proof. \square

Lemma 2. *Let Game G_i be defined as above, with respect to some PWI \mathcal{G} . Then if \mathcal{G} is (t', ϵ_p) -preserving secure, $|G_i(A, \mathcal{D}) - G_{i+1/2}(A, \mathcal{D})| \leq \epsilon_p$ for any adversary A running in time $t \approx t'$.*

Proof. Given A , we will construct an adversary A' for the recovery game. Note that we may assume without loss of generality that the $(i+1)$ st next query will be a preserving query, because otherwise G_i and $G_{i+1/2}$ are identical. A' obtains a seed from the challenger, and uses it along with the code for \mathcal{D} to simulate G_i for A until the $(i+1)$ st next query.

Let S_0 be the state that A' uses to simulate G_i for A immediately following the i th next query. When A makes its $(i+1)$ st next query, A' gives the challenger S_0 along with the entropy inputs I_1, \dots, I_d it generated between A 's i th and $(i+1)$ st next queries. The challenger replies with (S_b^*, R_b^*) .

Next, A' flips its own coin $b' \xleftarrow{\$} \{0, 1\}$. If $b' = 0$, A' replies to A 's $(i+1)$ st next query with R^{**} (if this is a `get-next` query) or with (S^{**}, R^{**}) (if it's a `next-ror` query), where if $b' = 0$ then $(S^{**}, R^{**}) = (S^*, R^*)$, and if $b' = 1$ then $S^{**} = M(S^*)$ and

$$R^{**} = \begin{cases} \perp & \text{if } R^* = \perp \\ R & \text{otherwise, for } R \xleftarrow{\$} \{0, 1\}^\ell. \end{cases}$$

Finally, A' resumes simulating the environment for A , starting from state S^{**} , and following the specification of G_{i+1} . When A outputs b^* , A' outputs $b^{*'}$, which is 1 if $b^* = b'$, and 0 otherwise.

Now, if the original challenge bit was $b = 0$, then A' has exactly simulated G_i for A : the first i next queries followed the specification of G_i , while the $(i+1)$ st next query returned the “real” state and output bits if $b' = 0$, and returned uniformly random bits and a mask of the state if $b' = 1$. Further, because M is idempotent, the fact that the challenger applies a mask to S_0 does not change its distribution— G_i calls for the state to be masked after every next query, and so S_0 is “already” masked.

On the other hand, if the challenge bit was $b = 1$, then A' has exactly simulated $G_{i+1/2}$ for A . It returned uniformly random bits and a mask of the state on the $(i+1)$ st query. Our idempotence assumption again comes into play: when $b = 1$, A' receives a state $S^* = M(S)$ that has already been masked, and returns $S^{**} = M(M(S))$ to A ; however, these random variables are identically distributed.

Consequently,

$$\begin{aligned} |\Pr [G_i(A, \mathcal{D}) = 1] - \Pr [G_{i+1/2}(A, \mathcal{D}) = 1]| &= |\Pr [b' = b^* \mid b' = 0] - \Pr [b' = b^* \mid b' = 1]| \\ &= \left| 2 \Pr [b^{*'} = b] - 1 \right| \leq \epsilon_p, \end{aligned}$$

which is what we wanted. □

Lemma 3. *Let Game G_i be defined as above, with respect to some PWI \mathcal{G} . Then if \mathcal{G} is (t', ϵ_r) -recovery secure, $|G_{i+1/2}(A, \mathcal{D}) - G_{i+1}(A, \mathcal{D})| \leq \epsilon_r$ for any adversary A running in time $t \approx t'$.*

Proof. Given A , we will construct an adversary A' for the recovery game. Note that we may assume without loss of generality that the $(i+1)$ st next query will be a recovering query, because otherwise $G_{i+1/2}$ and G_i are identical. A' obtains a seed and information $(\gamma_j, z_j)_{j=1}^{q_{\mathcal{D}}}$ from the challenger, chooses a challenge bit $b' \xleftarrow{\$} \{0, 1\}$, and generates an initial state $S' \xleftarrow{\$} \text{setup}$. Then A' simulates G_i for A , providing A with the seed and using the leaked information it obtained from the challenger and its `get-refresh` oracle to simulate the \mathcal{D} -refresh oracle for A' .

After the i th next query, however, A' no longer immediately uses its `get-refresh` oracle to update its state in response to a \mathcal{D} -refresh query from A . Instead, it simply returns the appropriate (γ, z) pair. Then on the $(i+1)$ st next query, A' invokes its `get-refresh` oracle to update its state to the point immediately following the corresponding MRED. Call this state S_0 . Next, A' counts the number d of \mathcal{D} -refresh calls A made after the MRED, and submits (S_0, d) to the challenger. A' receives the challenge (S^*, R^*) in reply, along with the entropy strings $I_\nu, I_{\nu+1}, \dots, I_{q_{\mathcal{D}}}$ it will later use to resume simulating an environment for A .

If $b' = 0$, A' replies to A 's $(i+1)$ st next query with R^{**} (if this is a `get-next` query) or with (S^{**}, R^{**}) (if it's a `next-ror` query), where if $b' = 0$ then $(S^{**}, R^{**}) = (S^*, R^*)$, and if $b' = 1$ then $S^{**} = M(S^*)$ and

$$R^{**} = \begin{cases} \perp & \text{if } R^* = \perp \\ R & \text{otherwise, for } R \stackrel{\$}{\leftarrow} \{0, 1\}^\ell \end{cases}$$

Finally, A' uses $I_{\nu+1}, \dots, I_{q_{\mathcal{D}}}$ to resume simulating the environment for A , starting from state S^{**} , and following the specification of G_{i+1} . When A outputs b^* , A' outputs $b^{*'}$, which is 1 if $b^* = b'$, and 0 otherwise.

Now, if the original challenge bit was $b = 0$, then A' has exactly simulated $G_{i+1/2}$ for A : the first i next queries followed the specification of G_i , while the $(i+1)$ st next query returned the “real” state and output bits if $b' = 0$, and returned uniformly random bits and a mask of the state if $b' = 1$.

On the other hand, if the challenge bit was $b = 1$, then A' has exactly simulated G_{i+1} for A . It returned uniformly random bits and a mask of the state on the $(i+1)$ st query. This is where our idempotence assumption comes into play: when $b = 1$, A' receives a state $M(S)$ that has already been masked, and returns $M(M(S))$ to A ; however, these random variables are identically distributed.

Consequently,

$$\begin{aligned} |\Pr [G_{i+1/2}(A, \mathcal{D}) = 1] - \Pr [G_{i+1}(A, \mathcal{D}) = 1]| &= |\Pr [b' = b^* \mid b' = 0] - \Pr [b' = b^* \mid b' = 1]| \\ &= \left| 2 \Pr [b^{*'} = b] - 1 \right| \leq \epsilon_r, \end{aligned}$$

which was what we wanted. □

C On the read-only robustness of ISK-RNG

This appendix contains our analysis of the robustness of ISK-RNG against a restricted adversary, in a somewhat restricted model. We discussed the necessity of limiting the adversary on pg. 21; in short, future output values persist in output buffers even if the ES produces large amounts of entropy following a state compromise. Therefore these outputs (and in some cases many more, as we discussed earlier) will necessarily be compromised. The best we can hope for is that ISK-RNG will eventually recover, where now “eventually” means after it produces a certain number of (unsafe) outputs, rather than after it accumulates γ^* bits of entropy.

A second change to the ROB experiment we need to make is that we will not permit the adversary to make `set-state` queries. He may, however, make `get-state` queries without restriction. That is, we consider adversaries who can learn information about the PWI state, but not adversaries that can tamper with it. We feel this is a reasonable assumption to make because, in addition to the practical

difficulties of tampering with hardware not accessible to software, an adversary possessing both the required technical expertise and opportunity to tamper with hardware will have numerous other attack vectors at his disposal, anyway; attempting to defend against him would be futile.

The distinction between “read-only” adversaries and those with write capabilities may seem rather fine, perhaps even artificial, in the context of hardware RNGs. Our intent here is not to speculate on the relative difficulty of such attacks, but rather to determine how damaging they might be. This being said, we acknowledge that the distinction is likely more important in software RNGs (for example, the recent Heartbleed vulnerability in OpenSSL permitted attackers to read, but not modify, the target’s memory); hence, a read-only robustness notion may prove useful in that domain, as well.

Returning to ISK-RNG, the reason for a change to “read-only” adversaries is that a **set-state** query would permit the adversary to replace otherwise entropic bits with seed-dependent values, preventing us from leveraging results from the entropy-extraction literature, specifically the Left-over Hash Lemma and Theorem 1. ISK-RNG would propagate this dependency through all subsequent reconditionings. We were able to overcome these obstacles in the proof of Theorem 2 by bootstrapping off of the near-uniform randomness of each conditioner output to argue for the randomness of the next, but this remedy is unavailable here.

We call this new game (arbitrary **get-state** queries, no **set-state** queries) **ro-ROB**, for read-only robustness.

We call an adversary A *slow* if after making a **get-state** or **set-state** call, A does not make a **next-ror** call until it performs the following sequence of steps:

1. Invokes \mathcal{D} -refresh until the CE buffer becomes available.
2. Empties the CE buffer (by using the RDSEED interface, or by using RDRAND and then allowing ISK-RNG to reseed) and makes it available again. At this point, the “compromised” OSTE and ESSR buffers have been conditioned into the CE buffer, which therefore remains compromised.
3. Empties the CE buffer and makes it available again. One can now hope that its contents are now information theoretically random and independent of any compromised value.
4. Causes the DRBG to reseed (by using the RDRAND interface, which empties one of the output buffers, and then invoking `wait`). If the previous step successfully produced a fresh, random value for the CE buffer, then this value will act as a one-time pad when generating a new DRBG key and IV.
5. Flushes the remaining output buffers by invoking RDRAND and invokes \mathcal{D} -refresh until the CE buffer becomes available. (This final step prevents A from invoking **get-state** to learn the “fresh” entropy used to reseed the DRBG.)
6. Invokes `wait()` until the eight output buffers are repopulated.

We note that this seemingly-elaborate sequence of steps will naturally occur once processes running on the machine dispatch a sufficient number of RDRAND or RDSEED instructions.

Theorem 10 (ISK-RNG is read-only robust against slow adversaries). *Let A be a slow adversary running in time t , making q_{gs} queries to its **get-state** oracle, and a combined q queries to its **get-next** and **next-ror** oracles. Let L_m be a positive integer. Suppose \mathcal{D} is β -healthy for some $\beta > 0$. Then there exists an adversary B making three queries and running in time $\mathcal{O}(t)$ such that*

$$\text{Adv}_{\text{ISK}, \mathcal{D}}^{\text{ro-rob}/M}(A) \leq (q + 3q_{\text{gs}}) \left(2^{(k-m\gamma)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right) + 2(q + 8q_{\text{gs}}) \left(\text{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right),$$

where $\epsilon(L_m) = \mathcal{O}(L_m + 1)/2^{k/2}$ and $\hat{\epsilon}(L_m) = \sum_{i=0}^{m-1} \binom{L_m}{i} \beta^i (1 - \beta)^{L_m-i}$.

Proof. For $i \in [0..q_{\text{gs}}]$, define G_i to be the game that behaves identically to $\text{ro-ROB}_{\text{ISK}, \mathcal{D}}$ with challenge bit $b = 0$ (i.e., the next-ror oracle returns random values) until after i queries to get-state , and then behaves as though $b = 1$. So

$$\mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{ro-rob}/M}(A) = \Pr [G_{q_{\text{gs}}+1}(A) \Rightarrow 1] - \Pr [G_0(A) \Rightarrow 1].$$

Let S_i be the state following the i th get-state query in G_i (so S_i is not masked). Let S_i^j be that same state after the slow adversary A completes Step j above. Then $S_i^2.\text{CE}_0$ is a permutation as a function of $S_{i-1}.\text{CE}_0$ for any fixed values of $S_{i-1}.\text{ESSR}$, $S_{i-1}.\text{OSTE}$, and intervening \mathcal{D} outputs, and similarly for $S_i^2.\text{CE}_1$ with respect to $S_{i-1}.\text{CE}_1$. Therefore $S_i^2.\text{CE}$ is uniformly distributed, and likewise independent of π and any entropy source outputs.

Define the game G'_i to behave identically to G_i , except that after the CE buffer becomes available during Steps 3 and 5 above, its contents are immediately rewritten with uniform random values. It follows (cf. the proof of Lemma 1) that

$$\Pr [G'_i(A) \Rightarrow 1] - \Pr [G_i(A) \Rightarrow 1] \leq 2 \left(2^{(k-m\gamma^*)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right).$$

Define the game G''_i to behave identically to G'_i , except that the output of each AES invocation during Step 6 is replaced with a uniform random string. Follow the logic of Lemma 1, we have that there is some adversary B making three queries and running in time t such that

$$\Pr [G''_i(A) \Rightarrow 1] - \Pr [G'_i(A) \Rightarrow 1] \leq 8 \left(\mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right).$$

Finally, let G'''_i be the game that behaves identically to G''_i , except that immediately after A completes Step 6 after making its i th get-state query, G'''_i overwrites the state S with $M(S)$. But since S is at this point identically distributed to $M(S)$,

$$\Pr [G'''_i(A) \Rightarrow 1] - \Pr [G''_i(A) \Rightarrow 1] = 0.$$

Let q_i be the combined number of queries that A makes to get-next and next-ror between its i th and $(i+1)$ st get-state query. It follows that there is some adversary B_i making (q_i+1) queries and running in time t such that

$$\Pr [G_{i+1}(A) \Rightarrow 1] - \Pr [G'''_i(A) \Rightarrow 1] \leq \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}/M}(B_i).$$

(There is a subtle issue here: the M -FWD experiment starts by masking some S generated by the setup procedure, which is not the case in this reduction; however, the proof of Theorem 8 follows through cleanly starting from $M(S)$ for arbitrary S .)

Putting everything together,

$$\begin{aligned}
\mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{ro-rob}/M}(A) &\leq \sum_{i=0}^{q_{\text{gs}}} (\Pr [G_{i+1}(A) \Rightarrow 1] - \Pr [G_i(A) \Rightarrow 1]) \\
&\leq \sum_{i=0}^{q_{\text{gs}}} \left[2 \left(2^{(k-m\gamma^*)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right) \right. \\
&\quad \left. + 8 \left(\mathbf{Adv}_{\text{AES}}^{\text{PRP}}(B) + \frac{3}{2^k} \right) + \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}/M}(B_i) \right] \\
&\leq \sum_{i=0}^{q_{\text{gs}}} \left[(q_i + 3) \left(2^{(k-m\gamma^*)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right) + 2(q_i + 8) \left(\mathbf{Adv}_{\text{AES}}^{\text{PRP}}(B) + \frac{3}{2^k} \right) \right] \\
&\leq (q + 3q_{\text{gs}}) \left(2^{(k-m\gamma^*)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right) + 2(q + 8q_{\text{gs}}) \left(\mathbf{Adv}_{\text{AES}}^{\text{PRP}}(B) + \frac{3}{2^k} \right).
\end{aligned}$$

This completes the proof.

Discussion. This bound is very similar to the one provided by Theorem 8. Additionally, we note that if one only requires `RDRAND` to be secure (as opposed to both `RDRAND` and `RDSEED`), then one could instead obtain a bound similar to that of Theorem 7 by replacing the $\mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}/M}(B_i)$ term in the proof with a $\mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd-RDRAND}/M}(B_i)$ term, upper bounding it with the value provided by that theorem.

Consequently, the discussion of Theorems 8 and 7 applies more-or-less intact to Theorem 10 and its `RDRAND`-only variant. The discussion can be found in Section 7.3 on pg. 21.

D Towards a notion of PWI availability

If a PWI can block, that immediately raises the question of availability. In general, however, any attempt to estimate the entropy contained in a sequence of input strings will necessarily be heuristic and subject to failure; imagine, for example, an entropy source that internally generates a random 128-bit AES key, and then begins outputting a CTR mode key stream. As long as AES is secure (as a PRP and hence a PRF, up until a large number of outputs have been produced), then no efficient test will be able to distinguish these outputs from random. Yet after the first couple outputs, the rest have essentially *zero* conditional min-entropy.

Still, entropy estimation and blocking may be reasonable ways to address specific, foreseeable failures with the entropy source (perhaps including certain types of hardware failure). Therefore we propose the following definition as a step toward capturing availability.

A PWI \mathcal{P} is $(\mathcal{D}, t, q, q^*, \gamma^*, \epsilon)$ -available if for all adversaries A making q queries and running in time t , the probability of A winning the above game (which is parameterized by q^* and γ^*) is at most ϵ . (Note that a definition that simply requires the PWI to not block would not be well-suited for PWIs that endeavor to provide truly random, rather than cryptographically random, bits; this is the case with the `RDSEED` interface, and, ostensibly, the `/dev/random` device in Linux.)

The basic idea is that the PWI should be available as long as it has gathered at least γ^* bits of entropy since q^* queries ago. However, we do not count any entropy gathered prior to the adversary tampering with state. (We do allow the adversary to *view* the state without penalty.) By making

the entropy source \mathcal{D} non-adversarial, we leave room to prove availability with specific (classes of) entropy sources.

The question of whether, say, ISK-RNG meets this definition of availability would require assumptions on \mathcal{D} beyond simply that it provide high min-entropy. See the discussion immediately following the proof of Theorem 2.

Oracle \mathcal{D} -refresh:

$(\sigma, I, \gamma, z) \stackrel{\$}{\leftarrow} \mathcal{D}(\sigma)$
 $\gamma_t \leftarrow \gamma_t + \gamma$
 $S \leftarrow \text{refresh}(S, I)$
return (γ, z)

Oracle get-next:

$(S, R) \leftarrow \text{next}(S)$
if $R = \perp$ **then**
 if $\gamma_\mu + \gamma_{\mu+1} + \dots + \gamma_t > \gamma^*$
then
 $b \leftarrow 1$
 $t \leftarrow t + 1$
 if $\mu < t - q^*$ **then**
 $\mu \leftarrow t - q^*$
return R

Oracle get-state:

return S

Oracle set-state(S^*):

$\mu \leftarrow t$
 $\gamma_t \leftarrow 0$
 $S \leftarrow S^*$

Proc initialize:

$\text{seed} \leftarrow \text{setup}$
 $\sigma \leftarrow 0; S \stackrel{\$}{\leftarrow} \{0, 1\}^n$
 $\mu \leftarrow 0; t \leftarrow 0$
 $\gamma_0 \leftarrow \gamma^*; b \leftarrow 0$

Proc finalize:

return b

E Security Proofs

This appendix contains our proofs of security for ISK-RNG.

Lemma 1 (ISK-RNG masking function is honest). *Fix positive integers k and m , and fix $0 < \beta \leq 1$. Let L_m be a positive integer. Let M be the ISK-RNG masking function. Let \mathcal{D} be a β -healthy entropy source. Then for any adversary A , there exists an adversary B running in the same time and making three queries such that $\text{Adv}_{\text{ISK}, \mathcal{D}, M}^{\text{init}}(A) \leq 2^{(k-m\gamma)/2+2} + 4\epsilon(L_m) + 8\hat{\epsilon}(L_m) + 5(\text{Adv}_{\text{AES}}^{\text{PTP}}(B) + \frac{3}{2^k})$.*

The following proof refers to a large body of pseudocode, found in Figure 4 (pg. 14).

Proof. Recall that $\text{Adv}_{\text{ISK}, \mathcal{D}}^{\text{init}}(A) = \Pr[A(S, \pi, z) \Rightarrow 1] - \Pr[A(M(S), \pi, z) \Rightarrow 1]$, where S is the state produced by the `setup` procedure and z contains any side-channel information leaked to the adversary during setup. The probabilities are over the coins of `initialize()`, π , \mathcal{D} , and A . During `setup`, ISK-RNG reconditions the CE buffer four times, and reseeds immediately afterwards each time. Let $C = (C_0, C_1, C_2, C_3)$ be the four CE buffer values used to perform these reseeding operations. Note

that $X = (S, \pi, z)$ is a deterministic function of C , π , and $J = (S.\text{OSTE}_1, S.\text{OSTE}_2, S.\text{ESSR}, z)$; we make this explicit by writing $X = f(C, \pi, J)$. Let $R = (R_0, R_1, R_2, R_3)$ be sampled from $\{0, 1\}^{4k}$. We have

$$\begin{aligned} \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{init}}(A) &= \Pr [A(f(C, \pi, J)) \Rightarrow 1] - \Pr [A(f(R, \pi, J)) \Rightarrow 1] \\ &\quad + \Pr [A(f(R, \pi, J)) \Rightarrow 1] - \Pr [A(M(S), \pi, z) \Rightarrow 1] \\ &\leq \Delta((\pi, C)|_J, (\pi, R)|_J) + (\Pr [A(f(R, \pi, J)) \Rightarrow 1] - \Pr [A(M(S), \pi, z) \Rightarrow 1]). \end{aligned}$$

For $j \in [0..3]$, $i = 0, 1$, let I_j^i be defined as in Figure 4; that is, I_j^i is the string of bits from the entropy source that gets used to update $S.\text{CE}_i$ between the j th time and the $(j+1)$ st time that buffer is available. So, for example, $C_j = C_j^0 \parallel C_j^1$, where $C_j^i = \text{CBCMAC}(C_{j-1}^i I_j^i)$ and $C_0^0 = C_0^1 = 0^{128}$. For each i, j , let $I_j^i = B_{j,1}^i B_{j,2}^i \cdots B_{j,\ell(j)}^i$, with each $|B_{j,j'}^i| = 2k$. Let \mathcal{E} be the event that for each such (i, j) , $\left| \left\{ j' : B_{j,j'}^i \in \mathcal{H}, j' < L_m \right\} \right| \geq m$. There are eight such (i, j) pairs, so $\Pr [\neg \mathcal{E}] \leq 8\hat{\epsilon}(L_m)$. (In cases where $\ell(i) < L_m$, I_i necessarily contains m ‘‘healthy blocks’’ because ISK-RNG will only stop accumulating entropy if this condition holds.) Since π is independent of both \mathcal{E} and J , Theorem 2 tells us that

$$\begin{aligned} \Delta((\pi, C)|_J, (\pi, R)|_J) &\leq \Delta((\pi, C|_{J,\mathcal{E}}), (\pi, R)) + \Pr [\neg \mathcal{E}] \\ &\leq 8 \left(\frac{1}{2} \sqrt{2^{k-m\gamma} + \epsilon(L_m)^2} \right) + 8\hat{\epsilon}(L_m) \\ &\leq 2^{(k-m\gamma)/2+2} + 4\epsilon(L_m) + 8\hat{\epsilon}(L_m). \end{aligned}$$

We claim that there exists some adversary B making three queries and running in time t such that

$$\Pr [A(f(R, \pi, J)) \Rightarrow 1] - \Pr [A(M(S), \pi, z) \Rightarrow 1] \leq 5 \left(\mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right).$$

Consider the experiment $A(f(R, \pi, J))$. During the third reseed, ISK-RNG updates $S.K \parallel S.IV$ by encrypting $S.\text{CE}$ in counter-mode. But at this point in time, $S.\text{CE} = R_3$ is uniform random bits, and thus so are the new key and IV (R_3 acts as a one-time pad on the counter-mode keystream). Let K_0 and V_0 be the values of $S.K$ and $S.IV$ immediately following this operation. The next reseed operation creates a new DRBG key and IV, K_1 and V_1 , by computing $K_1 \parallel V_1 \leftarrow \text{CTR}_{K_0}^{V_0}(R_4)$. Finally, for $i \in [0..3]$, the output buffers are filled by computing $S.\text{out}_{2i} \parallel S.\text{out}_{2i+1} \parallel K_{i+2} \parallel V_{i+2} \leftarrow \text{CTR}_{K_{i+1}}^{V_{i+1}}(0^{3 \cdot 128})$. The proof of our claim concludes with a standard hybrid argument where the outputs of these five CTR computations are replaced one-by-one with random bits; the $3/2^k$ term falls out of the PRF-PRP switching lemma. The final CE buffer in the $A(f(R, \pi, J))$ experiment is already uniformly random (R_4), and the final step in the hybrid argument completes the masking function’s task of replacing $S.\text{out}$, $S.K$, and $S.IV$ with uniform random bits. \square

Theorem 7 (*M-RDRAND*). *Let A be a delayed adversary making q queries to RDRAND and running in time t . Then there exists an adversary B making three queries and running in time $\mathcal{O}(t)$ such that $\mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd-RDRAND}/M}(A) \leq 2(q+4) \left(\mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right)$.*

Proof. During an execution of $M\text{-RDRAND}(A)$, the DRBG key $S.K$ is changed during reseeding or after a `get-next` call. Let $M\text{-RDRAND}^\nu(A)$ be the same as $M\text{-RDRAND}$, but where $\text{AES}_{S.K}$ is replaced with a random function the first ν times this happens. Then $\Pr [M\text{-RDRAND}^\nu(A) \Rightarrow 1] -$

$\Pr [M\text{-RDRAND}^{\nu+1}(A) \Rightarrow 1] \leq \mathbf{Adv}_{\text{AES}}^{\text{prp}}(B_\nu) + \frac{3}{2^k}$, for some adversary B making three queries and running in the same time as A . (B_ν simulates $M\text{-RDRAND}^\nu$ for A using random functions for the first ν keys, uses its oracle for key $(\nu + 1)$, and then uses AES for the remainder of the experiment; B_ν returns whatever value A does.)

Let $d_\nu = 2 \Pr [M\text{-RDRAND}^\nu(A) \Rightarrow 1] - 1$. Observe that in Game $M\text{-RDRAND}^\nu$, we can defer assigning b a value until after query ν ; until that point, all the coins of the experiment are independent of b . Then $d_{q+4} = 0$ because the final state, revealed by `get-state()`, is likewise independent of b . It follows that

$$d_0 = d_0 - d_{q+4} = \sum_{i=0}^{i+3} (d_i - d_{i+1}) \leq 2 \sum_{\nu=1}^{q+4} \left(\mathbf{Adv}_{\text{AES}}^{\text{prp}}(B_\nu) + \frac{3}{2^k} \right).$$

But $d_0 = \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd-RDRAND}/M}(A)$. Taking B to be the B_ν with maximal advantage completes the proof. \square

Theorem 8 (ISK-RNG’s masked forward security). *Fix a positive integers k and m , and fix $0 < \beta \leq 1$. Let L_m be a positive integer. Let A be a delayed adversary making a combined q queries to `get-next` and `next-ror`. Then if \mathcal{D} is β -healthy, there exists some adversary B making three queries and running in the same time as A such that*

$$\begin{aligned} \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}/M}(A) &\leq (q + 1) \left(2^{(k-m\gamma)/2} + \epsilon(L_m) + 2\hat{\epsilon}(L_m) \right) \\ &\quad + 2(q + 4) \left(\mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right). \end{aligned}$$

Proof. During an execution of $M\text{-FWD}(A)$, the ISK-RNG CE buffer becomes available at most $q + 1$ times. Since the outputs of \mathcal{D} do not depend on A , we can fix its output bits ahead of time.

For $j \in [1..q + 1]$, let $C_j = C_j^0 \parallel C_j^1$ be the resulting sequence of available conditioned entropy buffers produced over the course of the experiment $M\text{-FWD}$, as defined as in Figure 4. For $i = 1, 2$ let I_j^i also be defined as in this figure. Thus for $j > 0$, $C_j^i = \text{CBCMAC}(C_{j-1}^i I_j^i)$ while C_0^i is the (uniformly random) value generated by M .

Let $G(A)$ denote the (random) output of A during an execution of $M\text{-FWD}(A)$. Let $G_i(A)$ denote the output of A in a modified version of this game where the CE buffer is immediately overwritten with a uniformly random value the first i times it becomes available. Let b be the challenge bit chosen by these games (as in Figure 5). Let $C = (C_1, \dots, C_{q+1})$ and let $R = (R_1, \dots, R_{q+1}) \stackrel{\$}{\leftarrow} (\{0, 1\}^{2k})^{q+1}$ be a tuple of independent, uniformly random values.

We have:

$$\begin{aligned} \mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}/M}(A) &= \Pr [G(A) = 1 \mid b = 0] - \Pr [G(A) = 1 \mid b = 1] \\ &= \Pr [G_0(A) = 1 \mid b = 0] - \Pr [G(A) = 1 \mid b = 1] \\ &= \sum_{j=0}^q (\Pr [G_j(A) = 1 \mid b = 0] - \Pr [G_{j+1}(A) = 1 \mid b = 0]) \\ &\quad + \Pr [G_{q+1}(A) = 1 \mid b = 0] - \Pr [G(A) = 1 \mid b = 1] \end{aligned}$$

To upper-bound the terms of the telescoping sum, we consider an adversary’s advantage in distinguishing G_j from G_{j+1} (with b fixed to 0 in both cases). Without loss of generality, assume the

adversary is additionally supplied with $I_{>j+1} = (I_{j+2}^0, I_{j+2}^1, \dots, I_{q+1}^0, I_{q+1}^1)$ when it makes its `get-state` query. This allows us to make the further assumption, again without loss of generality, that the adversary causes the CE buffer to become available exactly $j + 1$ times: if this happens fewer times, the games are identical, while the information $I_{>j+1}$, when combined with the information from `get-state`, is sufficient for the adversary to simulate further queries. Hence,

$$\Pr [G_j(A) = 1 \mid b = 1] - \Pr [G_{j+1}(A) = 1 \mid b = 1] \leq \Delta(\mathcal{C}_j, \mathcal{R}_j),$$

where $\mathcal{C}_j = (\pi, R_j, C_{j+1}, I_{>j+1})$, $\mathcal{R}_j = (\pi, R_j, R_{j+1}, I_{>j+1})$, and $R_0 = C_0$ is the random CE buffer value generated by M . (The probabilities on the left-hand side of the above inequality are also over the coins $(R_0, R_1, \dots, R_{j-1})$, but these do not contribute to the statistical distance.)

Keeping j fixed, for $i = 0, 1$, let $I_j^i = B_{j,1}^i B_{j,2}^i \cdots B_{j,\ell(j)}^i$, with each $|B_{j,j'}^i| = 2k$. Let \mathcal{E} be the event that for each such (i, j) , $\left| \left\{ j' : B_{j,j'}^i \in \mathcal{H}, j' < L_m \right\} \right| \geq m$. There are two such (i, j) pairs, so $\Pr [\neg \mathcal{E}] \leq 2(q+1)\hat{\epsilon}(L_m)$. (In cases where $\ell(i) < L_m$, I_i necessarily contains m “healthy blocks” because ISK-RNG will only stop accumulating entropy if this condition holds.) By Theorem 2,

$$\Delta(\mathcal{R}_j, \mathcal{C}_j) \leq \Delta(\mathcal{C}_j|_{\mathcal{E}}, \mathcal{R}_j) + \Pr [\neg \mathcal{E}] \leq \frac{1}{2} \sqrt{2^{k-m\gamma} + \epsilon(L_m)^2} + 2\hat{\epsilon}(L_m).$$

At this point we apply the argument of Theorem 7 to show

$$\Pr [G_{q+1}(A) = 1 \mid b = 0] - \Pr [G(A) = 1 \mid b = 1] \leq 2(q+4) \left(\mathbf{Adv}_{\text{AES}}^{\text{DFP}}(B) + \frac{3}{2^k} \right),$$

this time allowing the RDSEED interface to return random bits when $b = 0$. This change, however, has no bearing on the argument, as the “real” values in G_{q+1} are random anyway. Combining the two previous inequalities with our upper bound for $\mathbf{Adv}_{\text{ISK}, \mathcal{D}}^{\text{fwd}/M}(A)$ will complete the proof.

The only subtlety in applying the argument of Theorem 7 is that the last value returned by RDSEED may not be independent of the final contents of the CE buffer. This is why we assume in our theorem statement that A is delayed; i.e., that it invokes `D-refresh` until the CE buffer is available. The bound accounts for the fact that the final CE buffer, $R_{q'}$, must be statistically close to uniform and independent of $(R_i)_{i < q'}$. Since the buffer can be refreshed at most once for each `get-next` and `next-ror` query, $q' \leq q + 1$. \square