XDRBG: A Proposed Deterministic Random Bit Generator Based on Any XOF

John Kelsey[1,2], Stefan Lucks[3] and Stephan Müller[4]

[1] National Institute of Standards and Technology (NIST), Gaithersburg, USA
[2] COSIC, KU Leuven, Leuven, Belgium
[3] Bauhaus-Universität, Weimar, Germany
[4] atsec information security corp, Austin, USA

Abstract. A deterministic random bit generator (DRBG) generates pseudorandom bits from an unpredictable seed, i.e., a seed drawn from any random source with sufficient entropy. The current paper formalizes a security notion for a DRBG, in which an attacker may make any legal sequence of requests to the DRBG and sometimes compromise the DRBG state, but should still not be able to distingush DRBG outputs from ideal random bits. The paper proposes XDRBG, a new DRBG based on any eXtendable Output Function (XOF) and proves the security of the XDRBG in the ideal-XOF model. The proven bounds are tight, as demonstrated by matching attacks. The paper also discusses the security of XDRBG against quantum attackers. Finally, the paper proposes concrete instantiations of XDRBG, employing either the SHAKE128 or the SHAKE256 XDRBG. Alternative instantiations suitable for lightweight applications can be based on ASCON.

Keywords: pseudorandom bit generation · forward security · backward security · extendable output function (XOF)

1 Introduction

Generating random bits is a critical function in almost any secure cryptographic system. Usually, the process for generating these random bits is broken into two parts: First, an *entropy source* provides some unpredictable input string as a *seed*. Second, a cryptographic algorithm called a deterministic random bit generator (DRBG) in [BK15] and this work, and called a PRNG (pseudorandom number generator), cryptographic PRNG, PRG[1] or DRNG (deterministic random number generator) elsewhere, produces the output bits. See [Mec18, MMHH23, Fer19, AMD23, KSF99] for widely-used examples of a DRBG.

Informally, a DRBG is *seeded* by providing it with an input string with some guaranteed amount of min-entropy. It then must must provide output bits on demand, which (if the DRBG was properly seeded) are computationally indistinguishable from ideal random bits. Further, because the internal state of a DRBG might be compromised, output bits generated before a compromise must remain indistinguishable from ideal random bits (called "backtracking resistance" in [BK15]), and the DRBG must recover from a state compromise once provided sufficient new entropy (called "prediction resistance" in [BK15]).

DRBGs are almost always constructed from other cryptographic primitives, such as hash functions, block ciphers, or stream ciphers[2].

---

[1]PRG is also sometimes used for a different primitive that expands a seed to a long string without providing other features expected of DRBGs.

[2]Thus, the DRBGs in [BK15] are based on AES [Nat01], SHA2 [Nat15a], or SHA3 [Nat15b], and the DRBG currently used in Linux is based on ChaCha20 [NL18].

In this paper, we propose a new DRBG that follows the interface and requirements of two widely-used standards: NIST's SP 800-90A [BK15] and BSI's AIS 20/31 [Kil11, Pet23].

In [BDPA07], *sponge functions* were proposed as a new way of constructing hash functions. These functions use a large fixed permutation or function to process data, but unlike traditional hash functions, they can generate arbitrary-length outputs. The resulting broad class of cryptographic primitive (which might also be realized using other constructions) was named a XOF *(eXtended Output Function)* in [Nat15b]. Informally, a XOF works like a cryptographic hash function, but allows for an arbitrary-length output string. The calls $\mathsf{XOF}(x, n)$ and $\mathsf{XOF}(x, n + u)$ will yield identical results in their first n bits. A XOF can reasonably be modeled as a random oracle with an extremely long output for each query, which is then truncated to the desired output length. See [BDPVA08] for a security analysis of XOFs based on sponge functions such as SHAKE128 and SHAKE256. In recent years, many new sponge-based hash functions have been proposed; typically these also support XOF functionality. For example, see [RS16, AMMS16, BKL+19, DEM+20, DEMS22].

In this work, we: (1) Describe XDRBG, a new DRBG based on any XOF. (2) Propose a security notion appropriate for DRBGs. (3) Prove XDRBG secure under this security notion. The proof treats the XOF at hand as a random oracle. (4) Demonstrate classical attacks matching our security bounds. (5) Describe the quantum security of XDRBG. (6) Propose concrete parameters and instantiations for XDRBG. (7) Discuss some remaining open questions.

Relationship with Prior Work.   Bertoni et al. describe a generic construction for cryptographic random bit generation based on a sponge construction in [BDPV10]. Additional proposals along these lines were made by Gaži and Tessaro[GT16a, GT16b], Hutchinson [Hut16a, Hut16b], and Coretti et al. [CDKT19]. While XDRBG is broadly similar to these designs (and has been strongly influenced by the designs of Bertoni et al. and Coretti et al.), our design differs from these earlier works in some important ways:

1. XDRBG supports the interface defined for DRBGs in [BK15]. Unlike the designs of Bertoni et al. and Coretti et al., XDRBG supports distinct INSTANTIATE and RESEED calls, variable-length outputs from GENERATE calls, and untrusted additional inputs provided by the caller.

2. XDRBG is intended to be usable with any XOF, with the DRBG making queries to the XOF using a standard interface, and without making any assumptions about its internal workings. Thus, our next DRBG state is part of the output from the XOF instead of remaining in the capacity of the underlying sponge, and is re-input in subsequent XOF queries by the DRBG. Similarly, we model the XOF as an ideal object, rather than considering its underlying structure.

3. XDRBG is targeted for use in the realm of cryptographic random bit generation under SP 800-90 [BK15, TBK+18, BKM+22] and AIS 20/31 [Kil11, Pet23], where entropy sources are *independently evaluated and validated*. Thus we assume the availability of known amounts of entropy on demand, with entropy sources that are non-adversarial and whose entropy distributions are oracle-independent. (I.e., we adopt the model of [DGH+04] and later [WS19], rather than the model of [CDKT19].) We suspect that the techniques of [CDKT19] could be used to show that the XOF in XDRBG works as a seedless extractor (XDRBG is quite similar to their sponge-based PRNG), but we leave this for future work.

2 Preliminaries

2.1 Entropy

A DRBG samples a seed from a random source. The mathematical model for a random source is a distribution, and, for the purpose of the current paper, the all-essential property of a distribution is its min-entropy.

Let $\mathcal{D}^h$ be a distribution over strings $\{0,1\}^*$. We write $S \leftarrow_\$ \mathcal{D}^h$ if the string S is chosen according to $\mathcal{D}^h$. In that context, the superscript h indicates a lower bound $h \leq H_{\min}(\mathcal{D}^h)$ for the min-entropy

$$H_{\min}(\mathcal{D}^h) = -\log_2 \left(\max_{S \leftarrow_\$ \mathcal{D}^h, T \in \{0,1\}^*} (\Pr[S = T]) \right),$$

rather than the more traditional Shannon entropy $-\sum_{T \in \{0,1\}^*} (\Pr[S = T] * \log_2(\Pr[S = T]))$. Firstly, the min-entropy of a distribution is always a lower bound for the Shannon entropy of that distribution. Thus, by requiring *at least h bits of min-entropy*, we will always get at least h bits of Shannon entropy. Secondly, in our context the min-entropy is more intuitive: it describes the upper bound for the attacker's chance to guess the seed.

Below, we will consider two thresholds for the min-entropy:

- H_{init}: Whenever the XDRBG is instantiated the seed is drawn from a source with at least H_{init} bits of min-entropy.

- H_{rsd}: When the reseed command is called, the seed is drawn from a source with at least H_{rsd} bits of min-entropy.

2.2 Interface for DRBG

Following [BK15], we define three DRBG operations:

1. $V \leftarrow \text{INSTANTIATE}(S, \alpha)$ creates a DRBG state V, using seed material S and a string for optional personalization and other optional data[3] α. The seed S must be drawn from an entropy source with min-entropy H_{init}.

2. $V \leftarrow \text{RESEED}(V', S, \alpha)$ creates a DRBG state V from a previous state V', the seed S and the string α. The seed S must be drawn from an entropy source with min-entropy H_{rsd}.

3. $(V, \Sigma) \leftarrow \text{GENERATE}(V', \ell, \alpha)$ generates a new DRBG state V and an ℓ-bit output string Σ from the old state V' and the string α. Σ is required to be indistinguishable from random bits.

In NIST SP 800-90A [BK15], DRBGs are defined with a particular interface which assumes the DRBG can draw bits from the entropy source directly. Each DRBG has a state of its own, identified with a state handle.

For clarity, we prefer a somewhat simpler interface. Instead of using state handles to keep track of DRBG states, we simply pass in the DRBG state (a bit string in XDRBG) to the DRBG function as a parameter. Each DRBG function returns an updated DRBG state. Additionally, when entropy is provided to the DRBG, we draw entropy from the source and pass it in to the DRBG function as a parameter. Note that the cryptographic object being described is unchanged–only the description is different.

At first glance, INSTANTIATE and RESEED may seem redundant, but there is actually an important difference between them. INSTANTIATE(S, α) discards the previous state

---

[3]Throughout the paper, we use α (or α_j, $\alpha_{d,i}$, etc.) for those strings.

and the new state only depends on the seed and the optional string α. RESEED(V, S, α) refreshes the state based on the previous state, the seed, and the optional string α. This is reflected by the operations we defined above: INSTANTIATE does not take a DRBG state as input, while RESEED does take a DRBG state as input; both return a resulting DRBG state. This distinction also matters for the security analysis of the DRBG, and the amount of entropy required by each function call, as discussed below.

2.3   Security Level

For typical instantiations of the XDRBG, we will discuss their classical and quantum security level, which specify approximately how much computation is needed to distinguish the outputs of the DRBG from random bits. More precisely, if the attacker makes Q XOF queries, which implies Q to be a lower bound for the attacker's workload, then we claim $\mathcal{L}$-bit classical security if the attacker's chance to decide the DRBG output from random is at most $\frac{1}{2} + \epsilon$ for small ϵ and $Q \ll 2^{\mathcal{L}}$, say, $\epsilon \leq 2^{\mathcal{L}/16}$ and $Q \leq 2^{7\mathcal{L}/8}$. For quantum security, we will assume a straightforward application of Grover's algorithm for an attack, in which case a classical security level of $\mathcal{L}$ bits implies a quantum security level of $\mathcal{L}/2$ bit, because Grover's algorithm must iterate at least about $\mathcal{L}/2$ times to succeed with high probability.

2.4   Forward and Backward Security

Informally, forward security (called *backtracking resistance* in [BK15]) requires that earlier DRBG outputs remain secure when a later DRBG state is compromised. This is a property of the DRBG algorithm. Likewise, backward security (called *prediction resistance* in [BK15]) requires a DRBG to be able to recover from a compromise of its state. This requires a reseed method (or calling instantiate again) to provide additional entropy. A DRBG without access to any new seed material (e.g., from an entropy source) simply cannot achieve backward security.

AIS 20/31 defines two similar properties: *enhanced backward secrecy* and *enhanced forward secrecy*. Enhanced backward secrecy is approximately equivalent to backtracking resistance (aka, forward security); enhanced forward secrecy is approximately equivalent to prediction resistance (aka, backward security).

2.5   Multitarget Attack on INSTANTIATE

Given many devices, each instatiated once, or a few or even one single device instantiated many times, the following multitarget attack becomes possible.

Suppose at each startup the DRBG is instantiated with a seed S containing k bits of entropy. Assume $S \in \{0, 1\}^k$ to be uniformly distributed. Assume that after each instantiation at least k bits of output are produced. After I such instantiations, an attacker can recover one DRBG state with a $2^k/I$ search. The attacker simply guesses $2^k/I$ possible values of S, and for each one instantiates the DRBG with S, generates an output, and checks it against the I output values.

This is within the normal security bounds of any k-bit scheme. Nevertheless, it can substantially weaken some applications. For example, consider a single device with a 256-bit ECDSA [Nat13] key, supporting 128-bit security. Suppose the device instantiates its DRBG with $k = 128$ random bits at each startup, and that it is restarted and produces an ECDSA signature $I = 2^{32}$ times in its lifetime. In spite of formally supporting 128-bit security, this device is actually vulnerable to a 2^{96}-time attack which will recover its signing key!

Let R_1 be an upper bound on the number of times the DRBG will be instantiated. As long as at least $k + \log_2(R_1)$ bits of min-entropy is provided for the seed when instantiating

the DRBG, the multitarget attack is blocked. As will turn out below, k bits of min-entropy suffice for reseeding.

The requirements for DRBG instantiation in [BK15] include a *nonce* along with the entropy input. The proposed new requirements for DRBG instantiation in [BKM+22] replace the nonce requirement with a requirement for additional entropy. Both the old and new requirements effectively block this multitarget attack in the case that a given user does not instantiate their DRBG more than 2^{64} times.

2.6 Extendable Output Functions (`XOF`s)

Formally, a `XOF` is a function $\{0,1\}^* \to \{0,1\}^*$, which we model as a random oracle. To avoid returning an infinite sequence, the `XOF` gets an integer as a second parameter[4]: $\text{XOF} : \{0,1\}^* \times \mathbf{N} \to \{0,1\}^*$. Now $\text{XOF}(x, \ell)$ returns the first ℓ bits from the infinite sequence. Thus, for every $x \in \{0,1\}^*$, the first $\min(\ell, \ell')$ bits of $\text{XOF}(x, \ell)$ and $\text{XOF}(x, \ell')$ are the same, and, for all $x' \neq x$, the sequences $\text{XOF}(x, \ell)$ and $\text{XOF}(x', \ell')$ are two independent random sequences of ℓ and ℓ' bits, respectively. The algorithm below describes how our ideal `XOF` might be implemented by lazy sampling.

Algorithm 1 `XOF` definition.

1: **function** INIT
2:     $T \leftarrow \{\}$
    (# T holds a map $\{0,1\}^* \to \{0,1\}^*$. Initially, T is empty. #)
3: **function** $\text{XOF}(x, \ell)$
4:     **if** $x \in T$ **then**
5:         **if** $|T[x]| < \ell$ **then**
6:             $s' \leftarrow_\$ \{0,1\}^{\ell - |T[x]|}$
7:             $T[x] \leftarrow T[x] \parallel s'$
8:     **else**
9:         $T[x] \leftarrow_\$ \{0,1\}^\ell$
    (# Now $x \in T$, and $|T[x]| \geq \ell$ #)
10:     **return**($T[x]$ truncated to ℓ bits)

Of course, any real `XOF` does not provide unlimited ideal security. A typical `XOF` will allow an attacker to make a sequence of queries, and then the attacker will try to differentiate the `XOF` from an ideal `XOF`. Write W for the sum of the lengths of all inputs and outputs (in bit) made by the attacker. We say, a `XOF` supports k-bit security, if for every attacker, the advantage in distinguishing the `XOF` from random is at most $W/2^k$.

FIPS 202 [Nat15b] defines two sponge-based `XOF`s: `SHAKE128` and `SHAKE256`. Both employ a cryptographic 1600-bit permutation; `SHAKE128` employs an internal state (the *capacity*) of 256 bits and `SHAKE256` 512 bits.

A sponge-based `XOF` with a capacity of c bits can maintain about $c/2$ bit security against classical attackers [BDPVA08] and about $c/3$ bit security against quantum attackers [Cza21].[5] In that context, the term *security* has to be understood as *indifferentiability from a random oracle*, when the underlying cryptographic permutation is modelled as a random permutation.

Accordingly, `SHAKE128` can claim 128 bit security against classical and 85 bit security against quantum attackers, and `SHAKE256` can claim 256 bit security against classical and

---

[4]In some applications of a `XOF`, the output length is not known at the time its inputs are provided; these must support a somewhat more complicated interface.

[5]Cautionary note: To the best of our knowledge, the claimed $c/3$ bits of quantum security for a sponge with c bit capacity has so far only been published at an eprint server [Cza21], but not yet at a peer-reviewed conference or journal. In this paper, we assume the claim to be correct.

171 bit security against quantum attackers.

3 XDRBG **Definition**

XDRBG is a DRBG based on an underlying XOF. One could also view XDRBG as a *mode of operation* for a XOF to realize a DRBG. Note that unlike most prior designs of this kind, XDRBG does not assume anything about the internal structure of the XOF.

Conventions and Notation.

1. When we encode an integer as a bitstring, we always assume network byte order, and write X_n to represent encoding X as an n-bit integer.

2. For each instantiation of the XDRBG we will claim two approximate security levels: one with respect to classical attackers, and a second one with respect to quantum attackers employing Grover's algorithm. A security level of k bits implies that the given attacker, when restricted to time $t < 2^k$, succeeds with at most $t/2^k$ probability.

   The classical (quantum) security level of an instantiation of the XDRBG, employing a given XOF, is the minimum of the classical (quantum) security level of the XDRBG in the ideal-XOF model and the classical (quantum) security level of the XOF at hand.

3. The internal state of XDRBG is a single bitstring, V of fixed size $|V|$.

4. All XDRBG functions support an additional input α to personalize the DRBG; α can be empty. We assume a function ENCODE, which encodes its sequence of inputs into a string in an unambiguous way, so that there can be no $(S, \alpha, n) \neq (S', \alpha', n')$ with $\text{ENCODE}(S, \alpha, n) = \text{ENCODE}(S', \alpha', n')$. We write $|\text{ENCODE}| = |\text{ENCODE}(S, \alpha, n)| - |S| - |\alpha|$ for the stretch of the encoding. See appendix B for our recommended encoding function.

   Within our proofs, we also assume a function[6] PARSE. If S can be written as $S = \text{ENCODE}(s_1, s_2, i)$, the function returns $\text{PARSE}(S, 1) = s_1$, $\text{PARSE}(S, 2) = s_2$, and $\text{PARSE}(S, 3) = i$. Further, we use the convention that $\text{PARSE}(S, -1)$ gives the last entry in the tuple that was ENCODEd. When there is no tth element in the tuple that was encoded to construct the string S, $\text{PARSE}(S, t)$ returns a failure symbol $\perp$.

5. Output lengths are specified in bits. As discussed in Section 7, we recommend an upper limit on the output from each GENERATE call, called maxout.

6. INSTANTIATE and RESEED calls require entropy to result in a secure DRBG state. We thus define two parameters specifying how much min-entropy must be provided. Each INSTANTIATE requires at least H_{init} bits of min-entropy; each RESEED requires at least H_{rsd} bits of min-entropy.

The security analysis of XDRBG can continue without defining the size of the state ($|V|$) or the entropy bounds H_{rsd} and H_{init}. Concrete recommendations for these parameters – and a possible maxout limit for generate requests – appear in Section 7.

---

[6]This function is not implemented as part of XDRBG, but its existence is required for XDRBG to be secure–the encoded inputs must be unambiguously parseable.

Algorithm 2 XDRBG Definition

The function ENCODE : $\{0,1\}^* \times \{0,1\}^* \times \{0,1,2\} \rightarrow \{0,1\}^*$ takes two strings and an integer and returns a string. We require ENCODE$(s_1, s_2, i) \neq$ ENCODE(s_1', s_2', i') for $(s_1, s_2, i) \neq (s_1', s_2', i') \in \{0,1\}^* \times \{0,1\}^* \times \{0,1,2\}$.

1: **function** INSTANTIATE(seed, α)
    (# Returns $|V|$-bit state; source for seed: $\geq H_{init}$ bits min-entropy. #)
2:     $V \leftarrow$ XOF$($ ENCODE(seed, $\alpha, 0$), $|V|$ $)$
3:     **return**(V)

4: **function** RESEED(V', seed, α)
    (# Returns $|V|$-bit state; source for seed: $\geq H_{rsd}$ bits min-entropy. #)
5:     $V \leftarrow$ XOF$($ ENCODE$((V' \parallel$ seed $), \alpha, 1), |V|$ $)$
6:     **return**(V)

7: **function** GENERATE(V', ℓ, α)
    (# Returns $|V|$-bit state <u>and</u> ℓ-bit string Σ. #)
8:     $T \leftarrow$ XOF$($ ENCODE$(V', \alpha, 2), \ell + |V|$ $)$
9:     $V \leftarrow$ first $|V|$ bits of T
10:    $\Sigma \leftarrow$ last ℓ bits of T
11:    **return**(V, Σ)

Design Rationale. The goal of XDRBG is to provide an efficient and comprehensible DRBG based on any XOF. A simple, comprehensible design makes implementation, cryptanalysis, proving security, and checking the proof all easier. These considerations led us to define XDRBG so that: (1) Each call to a DRBG function (INSTANTIATE, RESEED, or GENERATE) results in a single XOF query, and (2) The first $|V|$ bits of the XOF output always become the new state V. (In GENERATE calls, the remaining bits of the XOF output become the DRBG output, Σ.) XDRBG assumes only access to the normal XOF interface, not access to any internal state of the XOF.

4   The DRBG Security Game

In order to reason about the security of our DRBG, we first need to define what it means for a DRBG to be secure. Informally, our security goals can be summarized as follows: *The attacker must not be able to distinguish the outputs from the DRBG from perfect random bits.* This must be true for any sequence of instantiate, generate and reseed calls to different devices. The only exception is output generated in the time following a state compromise and before the next intake of fresh entropy (i.e., before either reseeding or instantiating the DRBG).

For clarity of explanation, we will use the following terms in discussing the game:

- When the challenger or attacker interact with the XOF, this is a *query*.
  (Example: The challenger makes a XOF query.)

- When the challenger interacts with a DRBG instance, this is a *call*.
  (Example: The challenger makes an INSTANTIATE call.)

- When the attacker interacts with the challenger, this is a *request*.
  (Example: The attacker makes a R_OUT request.)

Thus, when the attacker makes a R_OUT request, this causes the challenger to make a GENERATE call to the DRBG, which then causes the challenger to make a XOF query.

4.1 Intuition for the Game

Broadly spoken, the goal of the security game is to capture all the legal ways an attacker might interact with devices implementing SP 800-90A type DRBGs. Previous security analyses of DRBGs have not captured the full range of these interactions, including attacker choice of additional inputs, and forcing many instantiations on the same device.

In the security game, the attacker can request any device to generate outputs, to reseed, or to instantiate. The challenger simulates all devices using access to the XOF and responds to each of the attacker's requests. The only constraint is that the first request to each device must be instantiate. The attacker is free to choose the *additional input* (α) for the requests. After every request, the attacker can learn the DRBG state of a device by *compromising* it. A device must recover from a compromise when an instantiate or reseed call is made.

The attacker can force a given device to repeatedly be instantiated. This captures the multitarget instantiation attack on a single device described above, as well as any other weaknesses in the instantiation process the DRBG might suffer from. By allowing the attacker to repeatedly compromise the state of the device and to reseed it, we capture any flaws in either the reseed process or in the backtracking or prediction resistance of the DRBG. Similarly, the attacker can provide additional inputs to different devices in an attempt to cause different devices' DRBG states to collide, or to try to force a device's future state to collide with one of its past states.

Consider the set of all DRBG outputs state, from all devices. The attacker's goal is to distinguish that set of outputs from a set of independent uniformly distributed random bits. Whenever a device state is compromised, i.e., has become known to the attacker, the attacker can compute the device's DRBG outputs on its own, and thus trivially distinguish them from random bits. Our game handles this by having compromised devices *always* provide the outputs from the DRBG, while the bits from an uncompromised device will be ideal random bits or DRBG bits depending on the value of the random bit b. We can thus exclude bits from a compromised device from the distinguishing goal, since they will be identical whether $b = 0$ or $b = 1$.

4.2 Game Definition and Rationale

To capture the above intuition formally, we specify a security game, see Algorithm 3. It can be seen as an extension of the proof model used in [WS19]. Our game incorporates a wider range of possible attacks, and closely tracks with the assumptions and requirements of the SP 800-90 series and AIS 20/31:

1. At the beginning of the game, the challenger generates a random bit b. If $b = 0$, all R_OUT requests will be answered by outputs from the DRBG; if $b = 1$, some of those requests will be answered with ideal random bits, instead. All other requests are answered in exactly the same way regardless of b.

2. We assume D devices the adversary can make requests to. The sum of all requests sent to all devices is exactly R. (An adversary making $R' < R$ requests can always be modelled by an adversary making exactly R requests while ignoring any responses after the first R' requests.)

3. An adversarial request is a quintuple $(\texttt{req}, d, i, \alpha_{d,i}, \texttt{leak})$:

   - $\texttt{req} \in \{\texttt{R\_INST}, \texttt{R\_RESEED}, \texttt{R\_OUT}(\ell)\}$ defines the call the device needs to make to handle the request.
   - $d \in \{1, \ldots, D\}$ is the index of the device the request goes to.
   - i is a counter: $(\texttt{req}, d, i, \alpha_{d,i}, \texttt{leak})$ is the i-th request to device d.

Algorithm 3 DRBG Security Game:
The attacker can access up to D devices and makes R requests in total. The challenger responds to valid requests. The attacker can directly query the XOF up to Q times at any time during the game, not counting the challenger's own XOF queries, made when addressing the R requests. (That is, the attacker get adaptive access to the XOF.) The attacker wins if the final message from the challenger is 1.

1: **function** CHALLENGER(D, Q, R)
   (# Start the game: randomly choose the secret bit b. #)
2:     $b \leftarrow\!\!\$ \{0, 1\}$
   (# Initiate tracker for requests. #)
3:     **for** $d \in \{1, \ldots, D\}, i \in \{1, \ldots R\}$ **do**
4:         $\text{Done}(d, i) = \mathsf{false}$

   (# Attacker first commits to distributions and then is granted access to XOF. #)
5:     **for** $i \in \{1, \ldots, D\}, j \in \{1, \ldots R\}, h \in \{H_{\text{init}}, H_{\text{rsd}}\}$ **do**
6:         Attacker chooses distributions $\mathcal{D}_{i,j}^h$, as elaborated in the text.

7:     Challenger grants direct access to XOF for attacker, for up to Q queries.
   (# Attacker makes R requests. Handle valid requests, ignore invalid ones. #)
8:     **for** step $\in \{1, \ldots, R\}$ **do**
9:         Attacker chooses request $(\mathtt{req}_{d,i}, d, i, \alpha_{d,i}, \mathtt{leak})$, as elaborated in the text.
10:        **if** $((\overline{\text{Done}(d, i)}) \wedge (i = 1) \wedge (\mathtt{req}_{d,i} = \mathtt{R\_INST})) \vee ((i > 1) \wedge \text{Done}(d, i - 1))$ **then**
11:            $V_{d,i} \leftarrow \text{HANDLE\_REQUEST}(V_{d,i-1}, b, \mathtt{req}_{d,i}, d, i, \alpha_{d,i}, \mathcal{D}_{d,i}^{H_{\text{init}}}, \mathcal{D}_{d,i}^{H_{\text{rsd}}})$ *(# see Alg. 4 #)*
12:            $\text{Done}(d, i) \leftarrow \mathsf{true}$
13:            **if** $\mathtt{leak}$ **then** *(# Compromise current state of device d. #)*
14:                Send DRBG state $V_{d,i}$ to attacker.
15:                $\mathtt{corrupt}_d \leftarrow \mathsf{true}$

   (# Finish the game: the attacker wins if it correctly guesses the secret bit b. #)
16:     Receive $\hat{b}$ from attacker.
17:     **if** $b = \hat{b}$ **then** send 1 to attacker **else** send 0 to attacker.

Algorithm 4 Subroutine Handle_Request from DRBG Security Game.

1: **function** HANDLE_REQUEST($V_{d,i-1}, b, \mathtt{req}_{d,i}, d, i, \alpha_{d,i}, \mathcal{D}_{d,i}^{H_{\text{init}}}, \mathcal{D}_{d,i}^{H_{\text{rsd}}}$)
2:     **if** $\mathtt{req}_{d,i}$ is $\mathtt{R\_OUT}(\ell)$ **then**
3:         $(V_{d,i}, Z) \leftarrow \text{GENERATE}(V_{d,i-1}, \ell, \alpha_{d,i})$
4:         **if** $b = 1$ AND NOT $\mathtt{corrupt}_d$ **then**
5:            $Z \leftarrow\!\!\$ \{0, 1\}^\ell$
6:         Send Z to attacker.
7:     **else if** $\mathtt{req}_{d,i}$ is $\mathtt{R\_RESEED}$ **then**
8:         $S_{d,i} \leftarrow\!\!\$ \mathcal{D}_{d,i}^{H_{\text{rsd}}}$
9:         $V_{d,i} \leftarrow \text{RESEED}(V_{d,i-1}, S_{d,i}, \alpha_{d,i})$
10:        $\mathtt{corrupt}_d \leftarrow \mathsf{false}$
11:     **else if** $\mathtt{req}_{d,i}$ is $\mathtt{R\_INST}$ **then**
12:         $S_{d,i} \leftarrow\!\!\$ \mathcal{D}_{d,i}^{H_{\text{init}}}$
13:         $V_{d,i} \leftarrow \text{INSTANTIATE}(S_{d,i}, \alpha_{d,i})$
14:         $\mathtt{corrupt}_d \leftarrow \mathsf{false}$
15:     **return** $(V_{d,i})$

- $\alpha_{d,i}$ denotes the additional information for the call.
- $\texttt{leak} \in \{\texttt{true}, \texttt{false}\}$ indicates a state compromise: if $\texttt{true}$, the adversary learns $V_{d,i}$, i.e., the state of device d at the end of handling the request.

4. A request $(\texttt{req}, d, i, \alpha_{d,i}, \texttt{leak})$ is *valid*, if and only if

- there has been no i-th valid request to device d before, and
- either this is the first valid request to device d (i.e., $i = 1$) and the device is requested to instantiate ($\texttt{req} = \texttt{R\_INST}$) or this is not the first valid request to device d (i.e., $i > 1$) and the $(i-1)$-th valid request to device d has already been handled.

The challenger only reacts to valid requests, cf. lines 10–15 from Algorithm 3. Note the usage of the array $\text{Done}(\cdot, \cdot)$ to keep track of valid requests.

This serves two purposes: It prevents requests to uninstantiated devices. And it enforces a consistent notation: Requests to device d are numbered by $i = 1$, $i = 2$, $i = 3$, ..., in that order and without skipping or repeating an integer.

5. Each valid request results in a single DRBG call made by the challenger (cf. Alg. 4). Each DRBG call leads to one single $\texttt{XOF}$ query (cf. Alg. 2). In parallel to the requests, the attacker can also query the $\texttt{XOF}$ up to Q times (cf. Line 7 from Alg. 3).

6. For each device, we assume the availability of a properly designed and tested entropy source (known as an NTG.1, PTG.2 or PTG.3 in AIS 20/31) with a specified min-entropy. As pointed out above, we assume lower bounds H_{init} and H_{rsd} for the min-entropy of our entropy sources, namely H_{init} when the DRBG is instantiated, and H_{rsd} when it is reseeded.

SP 800-90C requires a security parameter k and fixes the entropy bounds by $H_{\text{init}} = 3k/2$, and $H_{\text{rsd}} = k$. AIS 20/31 requires[7] $H_{\text{init}} = H_{\text{rsd}} = 240$.

Assuming the specified min-entropy, we claim the validity of our results for *all realistic* entropy sources.

- At the beginning of the game, the attacker will specify a sequence of distributions $\mathcal{D}_{d,1}^h, \ldots, \mathcal{D}_{d,R}^h$ with min-entropy $h \in \{H_{\text{rsd}}, H_{\text{init}}\}$ for each device d. As the distributions are formally specified by the attacker, they cover *all* realistic entropy sources with the required min-entropy.

  Obviously, the entropy sources from different devices can be distributed differently. Furthermore, a realistic entropy source may change its distribution over time, e.g., due to heating up or cooling down. Thus, the attacker must choose a distribution $\mathcal{D}_{d,i}^h$ for each possible triple (d, i, h).

- To restrict the entropy sources to *realistic* ones, we require the attacker to commit to all distributions $\mathcal{D}_{d,i}^h$ at the very beginning of the game, in advance of all queries to the $\texttt{XOF}$, either directly by the attacker querying the $\texttt{XOF}$, or indirectly from the attacker's requests (cf. lines 5–7 of Algorithm 3).[8]

---

[7]The current draft of AIS 20/31 [Pet23] stems from 2023. It requires either 240 bits of min-entropy or 250 bits of Shannon entropy. Older versions of AIS 20/31 [Kil11] did allow smaller amounts of entropy. AIS 20/31 also imposes additional requirements on seeding a DRNG [Pet23], which we disregard here.

[8]In reality, the distributions stem from random sources with the required amount of entropy. Giving the attacker the ability to choose those distributions $\mathcal{D}_{d,i}^h$ is not realistic, but matches the all-quantification: I.e., if we claim (as we actually do) security for all combinations of distributions with the required min-entropy, we can, as well, model this by the adversary choosing the combination of distributions it likes. On the other hand, if the attacker had access to the $\texttt{XOF}$ before committing to $\mathcal{D}_i^h$, it might choose some $\mathcal{D}_i^h$ with $\Pr[\texttt{XOF}(u, \ell) = \texttt{XOF}(u', \ell)] > 2^{-h} + 2^{-\ell}$ for $\ell \geq 1$ and $u, u' \leftarrow_\$ \mathcal{D}_i^h$, even though $\Pr[u = u'] \leq 2^{-h}$. This is not a realistic attack setting for a realistic entropy source.

This will allow us to apply Lemma 1 below: one cannot choose $u \neq u'$ with $\Pr[\texttt{XOF}(u, \ell) = \texttt{XOF}(u', \ell)] > 2^{-\ell}$, without first querying the XOF.

7. In some situations, the attacker may realistically be able to *compromise* a device i.e., to learn its internal state $V_{d,i}$. This is indicated by `leak`, see above.

   For each device d, the algorithm maintains a flag $\texttt{corrupt}_d$. The flag is set when the device's state is compromised, and the flag is cleared when the state is advanced randomly using fresh entropy, i.e., after each reseed and instantiate request. (There is no need to initialize $\texttt{corrupt}_d$, because the first valid request to any device d is always instantiate, which sets $\texttt{corrupt}_d$ without prior reading.)

   A stronger attack model might allow the attacker to *set* the device's state to a chosen or known value. Our game does not support this, because we consider such attacks unrealistic. But we briefly discuss their potential impact in section 8.3.

8. At the end of the game, the attacker tries to guess the random bit b chosen at the beginning of the game.

   In order for XDRBG to be acceptably secure, the probability of the attacker to win the game must be no greater than $1/2 + \epsilon$ for some very small ϵ.

5   Security Analysis

5.1   The Main Result and some Corollaries

Let (d, i) denote the i-th request made to device d. Our results depend on the maximum number λ_1 of requests (d, i) with the same $\alpha_{d,i}$, and on the total number λ_2 of unordered pairs $(d, i) \neq (d', i')$ with $\alpha_{d,i} = \alpha_{d',i'}$:

$$\lambda_1 = \max_a \left(|\{(d, i) : \alpha_{d,i} = a\}| \right) \quad \text{and} \quad \lambda_2 = \sum_a \binom{|\{(d, i) : \alpha_{d,i} = a\}|}{2}. \tag{1}$$

We assume $\binom{0}{2} = 0 = \binom{1}{2}$, and in general $\binom{N}{2} = \frac{N(N-1)}{2}$.

Theorem 1. *Let H_{init} and H_{rsd} be the min-entropy for `R_INST` and `R_RESEED` requests, respectively. Let $|V| \geq H_{init}$ be the state size of the DRBG. Let the attacker make Q queries and R requests and let λ_1 and λ_2 be defined as in Equation 1.*

The attacker's probability to win the DRBG game is at most $1/2 + \epsilon$, with

$$\epsilon \leq Q \left(\frac{\lambda_1}{2^{H_{init}} - Q - R} + \frac{1}{2^{H_{rsd}} - Q} + \frac{Q}{2 \times 2^{|V|}} \right) + \frac{\lambda_2}{2^{H_{init}}} + \frac{R^2}{2 \times 2^{|V|}} \tag{2}$$

Below, we describe the security we guarantee, depending on different policies for the choice of the $\alpha_{d,i}$. We assume $\log_2(Q) \ll H_{\text{rsd}} \leq H_{\text{init}} \leq |V|$. Because $|V|$ is always at least twice the security level of the DRBG, we also assume $\log_2(Q) \leq |V|/2$.

The first policy imposes *no restrictions* on the choice of the $\alpha_{d,i}$. The $\alpha_{d,i}$ may even be empty, or set to another constant string. With at most R adversarial requests, $\lambda_1 \leq R$ and $\lambda_2 \leq \binom{R}{2} \leq \frac{R^2}{2}$, which gives the following bound:

Corollary 1. *Let H_{init} and H_{rsd} be the min-entropy for `R_INST` and `R_RESEED` requests, respectively. Let $|V| \geq H_{init}$ be the state size of the DRBG.*

The attacker's probability to win the DRBG game, making Q queries and R requests, is at most $1/2 + \epsilon$, with

$$\epsilon \leq Q \times \left(\frac{R}{2^{H_{init}} - Q - R} + \frac{1}{2^{H_{rsd}} - Q} + \frac{Q}{2 \times 2^{|V|}} \right) + \frac{R^2}{2} \times \left(\frac{1}{2^{H_{init}}} + \frac{1}{2^{|V|}} \right).$$

The above corollary incentivizes to choose $H_{\text{rsd}} < H_{\text{init}}$, actually: $H_{\text{rsd}} \approx H_{\text{init}} - \log_2(R)$. However, we get improved bounds by requiring disjoint α. The second policy thus requires *unique* $\alpha_{d,i}$, i.e., if $\alpha_{d,i} \neq \alpha_{d',i'}$ for $(d,i) \neq (d',i')$. In this case, we have $\lambda_1 = 1$ and $\lambda_2 = 0$ and an improved bound:

Corollary 2. *Let H_{init} and H_{rsd} be the min-entropy for* R_INST *and* R_RESEED *requests, respectively. Let $|V| \geq H_{init}$ be the state size of the DRBG. Let all additional inputs $\alpha_{d,i}$ be unique. Let the attacker make Q queries and R requests.*

The attacker's probability to win the DRBG game is at most $1/2 + \epsilon$, with

$$\epsilon \leq Q \times \left(\frac{1}{2^{H_{init}} - Q - R} + \frac{1}{2^{H_{rsd}} - Q} + \frac{Q}{2 \times 2^{|V|}} \right) + \frac{R^2}{2 \times 2^{|V|}}.$$

I.e., *unique* additional data $\alpha_{d,i}$ would incentivise $H_{\text{init}} = H_{\text{rsd}}$. On the other hand, the $\alpha_{d,i}$ are de-facto nonces, and handling nonces may not always be desirable for a DRBG. Actually, the current draft of SP 800-90 [BKM+22] abandons the requirement to use a nonce, which has been imposed so far [BK15]. Instead, [BKM+22] explicitly requires $H_{\text{init}} > H_{\text{rsd}}$.

So finally, we study a policy *in between* no rules and strict uniqueness for all $\alpha_{d,i}$. Each device d has a unique name id_d, i.e., $\text{id}_d \neq \text{id}_{d'}$ for $d \neq d'$. The adversary makes at most R_{DEV} requests to each device, and there are at most D devices. I.e., the number of requests in total is at most $R = D \times R_{\text{DEV}}$. Set $\alpha_{d,i} = \text{id}_d$, i.e., all requests to a device just use the device's unique name as the additional input. We refer to this as *personalization*. Then $\lambda_1 \leq R_{\text{DEV}}$ and $\lambda_2 \leq D \times \binom{R_{\text{DEV}}}{2} \leq D \times \frac{R_{\text{DEV}}^2}{2} \leq \frac{R \times R_{\text{DEV}}}{2}$, and thus:

Corollary 3. *Let H_{init} and H_{rsd} be the min-entropy for* R_INST *and* R_RESEED *requests, respectively, and $|V| \geq H_{init}$ the state size of the DRBG. Assume D devices, each responding to at most R_{DEV} requests, id_d being a unique name for device $d \in \{1, \ldots, D\}$, and $\alpha_{d,i} = id_d$. Let the attacker make Q queries and at most $R = D \times R_{DEV}$ requests.*

The attacker's probability to win the DRBG game is at most $1/2 + \epsilon$, with $\epsilon \leq$

$$Q \times \left(\frac{R_{DEV}}{2^{H_{init}} - Q - R} + \frac{1}{2^{H_{rsd}} - Q} + \frac{Q}{2 \times 2^{|V|}} \right) + \frac{R \times R_{DEV}}{2} \times \left(\frac{1}{2^{H_{init}}} + \frac{1}{2^{|V|}} \right). \quad (3)$$

Once again, this incentivizes $H_{\text{rsd}} < H_{\text{init}}$, but now $H_{\text{rsd}} \approx H_{\text{init}} - \log_2(R_{\text{DEV}})$.

Remark. Sometimes, devices $d \neq d'$ with identical names: $\text{id}_d = \text{id}_{d'}$ may exist. For example, device IDs could be chosen at random, or a manufacturing error might cause two devices to have the same serial number. In this case, the claimed security bound still holds if we treat all devices with the same name as a single device. Namely, if we redefine R_{DEV} such that all devices $d_1, d_2, \ldots$ with $\text{id}_{d_1} = \text{id}_{d_2} = \cdots$ together do not respond to more than R_{DEV} queries, then equation 3 still applies.

5.2   The Proof of the Main Result

Notation for the Proof.   As before, we use (d,i) as a shorthand for requests $(b, \texttt{req}_{d,i}, d, i, \alpha_{d,i}, \texttt{leak})$. A generate request (d,i) is either *corrupted* or *faithful.*. Namely, it is corrupted, if, and only if, at the point of time Handle_Request is called, $\texttt{corrupt}_d = \texttt{true}$. Instantiate and reseed requests are always faithful.[9] We refer to the inputs to the attacker's XOF queries as $W_1, \ldots, W_Q$. Furthermore, recall that $V_{d,i}$ denotes the i-th state on device d. We write $U_{d,i}$ for the matching input from the challenger's matching XOF query, i.e., $V_{d,i} = \texttt{XOF}(U_{d,i}, |V|)$:

---

[9]Note that the security game *first* calls Handle_Request and *only then* considers the $\texttt{leak}$ parameter. Thus, a request $(d,i) = (b, \texttt{req}_{d,i}, d, i, \alpha_{d,i}, \texttt{leak})$ can be corrupted or faithful, regardless of $\texttt{leak}$. But if $\texttt{leak} = \texttt{true}$ then all subsequent generate requests $(d, i+1)$, $(d, i+2)$, $\ldots$ to the same device will be corrupted, until the first reseed or instantiate request is made.

- $U_{d,i} = \text{ENCODE}(S_{d,i}, \alpha_{d,i}, 0)$ for instantiate,

- $U_{d,i} = \text{ENCODE}((V_{d,i-1} \| S_{d,i}), \alpha_{d,i}, 1)$ for reseed, and

- $U_{d,i} = \text{ENCODE}(V_{d,i-1}, \alpha_{d,i}, 2)$ for generate.

Also, we say $U_{d,i}$ is *corrupted* or *faithful*, if (d,i) is so.

If W_j is a XOF query made by the attacker, write Q_n for the number of queries in $W_1, \ldots, W_Q$ with $\text{PARSE}(W_j, 3) = n$. Clearly, $Q_1 + Q_2 + Q_3 \leq Q$.

Independent XOF queries. Our proof consists of a sequence of lemmas. The first one is almost trivial.

Lemma 1. *For any $u \neq u'$ chosen independently from the XOF and any $\ell \geq 1$*

$$\Pr[\text{XOF}(u, \ell) = \text{XOF}(u', \ell)] = 2^{-\ell}.$$

Proof. Recall Algorithm 1. Since $u \neq u'$, the values $\text{XOF}(u, \ell)$ and $\text{XOF}(u', \ell)$ are chosen as two independent random ℓ-bit values. Their probability to collide is $2^{-\ell}$. □

Collisions. A part of our proof is based on bounding the probability of an input W_j for a XOF query made by the attacker to collide with a faithful $U_{d,i}$, i.e., the event $W_j = U_{d,i}$. We refer to the triple (j, d, i) as a *collision*. A collision can only occur if the input strings to the XOF are identical, and thus for every t, $\text{PARSE}(W_j, t) = \text{PARSE}(U_{d,i}, t)$. Specifically, this requires that $\text{PARSE}(W_j, -1) = \text{PARSE}(U_{d,i}, -1) = \alpha_{d,i}$, so that a given attacker XOF query can only collide with a XOF query from one kind of XDRBG call. Note that for any given string a, there are at most λ_1 requests (d, i) with $\alpha_{d,i} = a$.

Instantiate Collisions. We start with collisions (j, d, i), where (d, i) is an instantiate request, i.e., $\text{PARSE}(W_i, -1) = \text{PARSE}(U_{d,i}, -1) = 0$.

Lemma 2. *Let BAD_{UW}^{INIT} denote the event that during the attack game a XOF query j and an instantiate request (d, i) is made with $U_{d,i} = W_j$. Then*

$$\Pr[BAD_{UW}^{INIT}] \leq \frac{Q_1 \times \lambda_1}{2^{H_{init}} - Q}.$$

Proof. An instantiate collision (j, d, i) implies that the attacker actually did make both the XOF query W_j and the instantiate request (d, i) during the DRBG security game. If the query is made before request i, the attacker is trying to guess a random value with H_{init} bits of min-entropy, which will be chosen later (i.e., when the request is made). In this case, $\Pr[U_{d,i} = W_j] \leq 2^{-H_{\text{init}}}$. If request j is made before the query, and the event BAD_{UW}^{INIT} did not already occur before the request, then the attacker is trying to guess a fixed unknown value with H_{init} bits of min-entropy. But in this case, the attacker may already know at most $j - 1$ relationships $W_1 \neq U_{d,i}, \ldots, W_{j-1} \neq U_{d,i}$. Thus, the attacker's chance to guess $U_{d,i}$ when making up to Q_1 queries W_j with $\text{PARSE}(W_j, -1) = 1$ is $\Pr[U_{d,i} = W_j] \leq 1/(2^{H_{\text{init}}} - Q_1)$.

As there are at most $Q_1 \times \lambda_1$ triples (j, d, i), with (d, i) being an instantiate query and $\alpha_{d,i} = \text{PARSE}(W_j, 2)$, the probability of any instantiate collision (j, d, i) is

$$\Pr[\text{BAD}_{\text{UW}}^{\text{INIT}}] \leq \frac{Q_1 \times \lambda_1}{2^{H_{\text{init}}} - Q_1} \leq \frac{Q_1 \times \lambda_1}{2^{H_{\text{init}}} - Q}.$$

□

Repeating States.   Another event, which can be beneficial for the attacker, is the case that the challenger generates the same state $V_{d,i} = V_{d',i'}$ from two different requests (d,i) and (d',i').

Lemma 3. *Let BAD_{UW}^{INIT} be defined as in lemma 2. Let BAD_{VV} denote the event that some of the challenger's states during the attack game repeat. I.e., for two requests $(d,i) \neq (d',i')$, the i-th request to device d and the i'-th request to device d', it holds that $V_{d,i} = V_{d',i'}$. Then*

$$\Pr[\mathrm{BAD}_{VV} | \overline{BAD_{UW}^{INIT}}] \leq \frac{\lambda_2}{2^{H_{init}}} + \frac{R^2}{2 \times 2^{|V|}} + \frac{Q^2}{2 \times 2^{|V|}}$$

Proof. If $U_{d,i} \neq U_{d',i'}$,

$$\Pr[V_{d,i} = V_{d',i'} | U_{d,i} \neq U_{d',i'}] = 1/2^{|V|},$$

cf. lemma 1. As there exist at most $\binom{R}{2}$ unordered pairs $U_{d,i} \neq U_{d',i'}$

$$\Pr[\exists d,i,d',i' : (U_{d,i} \neq U_{d',i'} \wedge V_{d,i} = V_{d',i'})] \leq \binom{R}{2} \frac{1}{|V|} \leq \frac{R^2}{2 \times 2^{|V|}}. \tag{4}$$

We still need to discuss $\Pr[U_{d,i} = U_{d',i'}]$. W.l.o.g., (d,i) is made before (d',i'). If $i,i' > 1$, then $V_{d,i-1} \neq V_{d',i'-1}$ (otherwise, the event BAD_{VV} has already occurred).

If at least one of the requests (d,i), (d',i') is not an instantiate request, then $U_{d,i} \neq U_{d',i'}$ follows from $V_{d,i-1} \neq V_{d',j'-1}$. If $\alpha_{d,i} \neq \alpha_{d',i'}$, then $U_{d,i} \neq U_{d',i'}$.

To bound $\Pr[U_{d,i} = U_{d',i'}]$, we can thus assume $(\alpha_{d,i} = \alpha_{d',i'})$ and both (d,i) and (d',i') are instantiate requests. As the seeds $S_{d,i}$ and $S_{d',i'}$ are independently drawn, and the min-entropy of each seed is at most H_{init}, we get

$$\Pr[U_{d,i} = U_{d',i'}] \leq \Pr[S_{d,i} = S_{d',i'}] \leq \frac{1}{2^{H_{\mathrm{init}}}}.$$

As there exist at most λ_2 pairs $(d,i) \neq (d',i)$ with $\alpha_{d,i} = \alpha_{d',i'}$, we thus have

$$\Pr[\exists d,i,d',i' : U_{d,i} = U_{d',i'}] \leq \frac{\lambda_2}{2^{H_{\mathrm{init}}}}. \tag{5}$$

Finally, the attacker might try to *cause* a collision between two XDRBG states, using his control over the α and his large number of allowed XOF queries. For example, the attacker might compromise two different devices' internal states, V, V', and then search for an α, α' such that $\mathtt{XOF}(\mathrm{ENCODE}(V, \alpha, 2), |V|) = \mathtt{XOF}(\mathrm{ENCODE}(V', \alpha', 2), |V|)$, but perhaps there could be other ways to structure this collision search. We bound the probability of this happening by bounding the attacker's probability of finding any pair of inputs to the XOF that might lead to a colliding pair of DRBG states. Let $W[i], W[j]$ be XOF queries made by the attacker. The attacker's probability of finding a collision on the first $|V|$ bits of *any* XOF output is bounded by

$$\Pr[\mathtt{XOF}(W[i], |V|) = \mathtt{XOF}(W[j], |V|)] \leq \binom{Q}{2} \times 2^{-|V|} \leq \frac{Q^2}{2 \times 2^{|V|}} \tag{6}$$

The lemma follows from combining equations 4, 5, and 6.   □

Reseed Collisions.   How likely is a collision (j, d, i), if (d,i) is a reseed request?

Lemma 4. *Let BAD_{VV} be as defined in lemma 3. Let BAD_{UW}^{RSD} denote the event that (j, d, i) exist, where (d,i) is a reseed request and $U_{d,i} = W_j$. Then*

$$\Pr[BAD_{UW}^{RSD} | \overline{BAD_{VV}}] \leq \frac{Q_2}{2^{H_{rsd}} - Q}.$$

Proof. Assume (d, i) to be a reseed request, and $\mathrm{BAD_{UW}^{RSD}}$ did not alreaday occur, i.e., $U_{d,i} \notin \{W_1, \ldots, W_{j-1}\}$.

First, consider query j being made before reseed request (d, i). Thus, W_j is fixed, and $\Pr[U_{d,i} = W_j] \le \frac{1}{2^{H_{\mathrm{rsd}}}}$, since the min-entropy of $U_{d,i}$ is at least H_{rsd}.

Now consider reseed request (d, i) has been made first. Even if the attacker knows $V_{d,j-1}$, maybe due to a previous compromise, it still has to guess the current seed $S_{d,i}$. The probability to guess $S_{d,i}$ without prior knowledge is $\frac{1}{2^{H_{\mathrm{rsd}}}}$. If previously there had been Q_2 failed attempts to guess $S_{d,i}$, the probability to guess $U_{d,i}$ when choosing W_j is still $\Pr[W_j = U_{d,i}] \le \frac{1}{2^{H_{\mathrm{rsd}} - Q_2}}$.

Note that, even if the attacker knows $V_{d,i-1}$, we assume no repeating states (i.e., $\overline{\mathrm{BAD_{VV}}}$), so for any given W_j, there can be at most one request (d, i) such that (j, d, i) could possibly collide, namely the unique (d, i) with $\mathrm{PARSE}(W_j, 1) = V_{d,i-1}$. So in total, the attacker's chance to find a collision within Q queries is

$$\Pr[\mathrm{BAD_{UW}^{RSD}}|\overline{\mathrm{BAD_{VV}}}] \le \frac{Q_2}{2^{H_{\mathrm{rsd}}} - Q_2} \le \frac{Q_2}{2^{H_{\mathrm{rsd}}} - Q}$$

$\square$

Faithful Generate Requests. In the case of a generate request, observe that if (d, i) is compromised, i.e., if the attacker knows $V_{d,i-1}$, then the attacker may easily match a challenger $\mathtt{XOF}$ query, leading to a collision between W_j and $U_{d,i}$. But such matches are useless for the attacker, who is trying to distinguish the output from faithful generate requests from random values. So in our context, we consider only faithful generate requests.

Lemma 5. *Let BAD_{VV}, BAD_{UW}^{INIT}, and BAD_{UW}^{RSD} be as defined above. Let BAD_{UW}^{FGEN} denote the event that (j, d, i) exist, where (d, i) is a faithful generate query and $U_{d,i} = W_j$. Then*

$$\Pr[BAD_{UW}^{FGEN}|\overline{BAD_{VV}} \wedge \overline{BAD_{UW}^{INIT}} \wedge \overline{BAD_{UW}^{RSD}}] \le \frac{Q_3 \times \lambda_1}{2^{|V|} - Q_3 - R} \le \frac{Q_3 \times \lambda_1}{2^{|V|} - Q - R}$$

Proof. Assume (d, i) to be a faithful generate request, and $U_{d,i} \notin \{W_1, \ldots, W_{j-1}\}$. Consider the probability of the attacker to choose W_j with $W_j = U_{d,i}$. Essentially, the attacker has to guess $V_{d,j-1}$.

First, consider query j being made before request (d, i). The event $U_{d,i} = W_j$ only occurs if the $\mathtt{XOF}$ generates the right $|V|$-bit output. Thus $\Pr[U_{d,i} = W_j] \le \frac{1}{2^{|V|}}$.

Now consider query j being made after request (d, i). In this case, the attacker is guessing an unknown $|V|$ bit output from a challenger's $\mathtt{XOF}$ query. Without other information, $\Pr[W_j = U_{d,i}]$ would be $\le \frac{1}{2^{|V|}}$. But the attacker knows up to Q_3 values W_j which are not $U_{d,i}$, and there can be up to R previously-disclosed DRBG states, which the attacker can avoid to guess. So the probability of the attacker matching the $\mathtt{XOF}$ query from a single faithful GENERATE request is $\Pr[W_j = U_{d,i}] \le \frac{1}{2^{|V|} - Q_3 - R}$.

The claimed bound stems from the fact that there may be up to λ_1 faithful generate requests (d, i) with the same additional input $\alpha_{d,i} = \mathrm{PARSE}(U_{d,i}, 2)$ and up to Q_3 queries W_j with $\mathrm{PARSE}(W_j) = 3$. $\square$

No Bad Events. The last lemma we need for the main result is describes the adversarial advantage in the absence of bad events.

Lemma 6.

$$\Pr\left[\hat{b} = b \mid (\overline{BAD_{UW}^{INIT}} \wedge \overline{BAD_{VV}} \wedge \overline{BAD_{UW}^{RSD}} \wedge \overline{BAD_{UW}^{FGEN}})\right] = \frac{1}{2}$$

Proof. By the definition of the attack game (cf. alg. 3), all those outputs sent to the attacker which depend on b, stem from faithful generate requests.

The event $\overline{\text{BAD}_{\text{VV}}}$ implies that, since all the challenger's states $V_{d,i}$ are different, all the XOF inputs $U_{d,i}$ to answer generate requests are different. Thus, all the output bits generated in line 3 of the attack game stem from calls $\text{XOF}(U_{d,i},\ldots)$ with different inputs $U_{d,i}$. Depending on b, a faithful generate request will return either of the following values to the attacker:

- By the definition of the XOF, cf. Algorithm 1, all the output bits from a XOF query are uniformly distributed random bits. I.e., if $b = 0$, the challenger will compute $T \leftarrow\!\!\$\ \{0,1\}^{\ell+|V|}$ in the XOF query and the attacker will see the rightmost ℓ bits of T.

- If $b = 1$, the attacker is given $Z \leftarrow\!\!\$\ \{0,1\}^{\ell}$.

Regardless of b, if $\overline{\text{BAD}_{\text{VV}}}$ the distribution of responses to the attacker's faithful generate requests is the uniform distribution. To distinguish $b = 0$ from $b = 1$, the attacker must thus make a query $\text{XOF}(U_{d,i},\ldots)$, which matches a faithful request. Without such a matching query, i.e., when $\overline{\text{BAD}_{\text{UW}}^{\text{INIT}}} \wedge \overline{\text{BAD}_{\text{UW}}^{\text{RSD}}} \wedge \overline{\text{BAD}_{\text{UW}}^{\text{FGEN}}}$, all the answers to the attacker's queries are uniform random values, independent from the responses to faithful requests.   □

Proof of theorem 1.   According to lemma 6,

$$\Pr\left[\hat{b} = b \mid \left(\overline{\text{BAD}_{\text{UW}}^{\text{INIT}}} \wedge \overline{\text{BAD}_{\text{VV}}} \wedge \overline{\text{BAD}_{\text{UW}}^{\text{RSD}}} \wedge \overline{\text{BAD}_{\text{UW}}^{\text{FGEN}}}\right)\right] = \frac{1}{2}.$$

I.e., our advantage ϵ in distinguishing $b = 0$ from $b = 1$ is at most the probability to trigger one of the bad events. Hence ϵ is at most the sum of the the bounds from lemmas 2, 3, 4, and 5:

$$\epsilon \leq \frac{Q_1 \times \lambda_1}{2^{H_{\text{init}}} - Q} + \frac{\lambda_2}{2^{H_{\text{init}}}} + \frac{R^2}{2 \times 2^{|V|}} + \frac{Q^2}{2 \times 2^{|V|}} + \frac{Q_2}{2^{H_{\text{rsd}}} - Q} + \frac{Q_3 \times \lambda_1}{2^{|V|} - Q - R}.$$

We can simplify this to the claimed bound

$$\epsilon \leq Q\left(\frac{\lambda_1}{2^{H_{\text{init}}} - Q - R} + \frac{1}{2^{H_{\text{rsd}}} - Q} + \frac{Q}{2 \times 2^{|V|}}\right) + \frac{\lambda_2}{2^{H_{\text{init}}}} + \frac{R^2}{2 \times 2^{|V|}}$$

by applying $\frac{Q_1 \times \lambda_1}{2^{H_{\text{init}}} - Q} + \frac{Q_3 \times \lambda_1}{2^{|V|} - Q - R} \leq \frac{Q \times \lambda_1}{2^{H_{\text{init}}} - Q - R}$ (since $H_{\text{init}} \leq |V|$ and $Q_1 + Q_3 \leq Q$) and $\frac{Q_2}{2^{H_{\text{rsd}}} - Q} \leq \frac{Q}{2^{H_{\text{rsd}}} - Q}$ (since $Q_2 \leq Q$).   □

6   Matching Attacks

In this section, we will sketch attacks which closely match our claimed security bounds. For the case of simplicity, we do not put any constraints on the usage of the additional input – the attacker is even free to choose all the $a_{d,i}$ as the empty string. This is the case of corollary 1.

We assume each $\mathcal{D}_{d,i}^h$ to be the uniform distribution of h-bit values. This tightly matches the claimed min-entropy h for the $\mathcal{D}_{d,i}^h$.

6.1   Classical Attacks

Attack 1 ($R \approx 2 * 2^{H_{\text{init}}/2}$):   The core idea for this attack is to maximise the probability of the event BAD_{VV} and then to exploit it. Note that $H_{\text{init}} \leq |V|$.

The attacker makes $R/2$ instantiate-requests. After each instantiate-request, it makes one request to generate $\ell \gg H_{\text{init}}/2$ output bits.

If $b = 0$, the attacker can expect two of the ℓ-bit outputs to be identical: The seeds for the instantiate-requests are uniformly distributed H_{init}-bit values, with probability $> 1/2$ two of the R seeds will be identical. Thus, these two instantiate-requests will initiate the same output state, which will then be used by the subsequent generate requests to generate the same ℓ-bit output.

If $b = 1$, then, since $\ell \gg H_{\text{init}}/2$, the probability for any two independent ℓ-bit values to be identical is negligible.

Attack 2 ($Q \approx 2 * \max(2^{H_{\text{init}}}/R, 2^{H_{\text{rsd}}})$): The core idea for this attack is to try to guess one of the challenger's $\texttt{XOF}$ inputs $U_{d,i}$. We split the attack into two subcases:

Attack 2a ($2^{H_{\text{init}}}/R > 2^{H_{\text{rsd}}}$, $Q \approx 2 * 2^{H_{\text{init}}}/R$): Similarly to attack 1, the attacker makes $R/2$ instantiate requests, interleaved with $R/2$ requests to generate $\ell \gg H_{\text{rsd}}$ output bits each. In contrast to attack 1, the attacker now chooses $Q/2$ random seeds $S_1, \ldots, S_{Q/2} \in \{0,1\}^{H_{\text{rsd}}}$ and picks the ℓ rightmost bits from each of the Q strings $\texttt{XOF}(\texttt{XOF}(S_j, |V|), |V| + \ell)$. With significant probability, one of the attacker's random states S_j will collide with one of the challenger's states $S_{d,i}$.

If $b = 0$, the collision of the attacker's S_j with one of the challenger's states implies the same ℓ-bit output.

If $b = 1$, the probability for one of the attacker's $Q/2$ ℓ-bit output strings with one of the challengers $R/2$ ℓ-bit output strings is negligible, since $\ell > H_{\text{rsd}}$.

Attack 2b ($2^{H_{\text{init}}}/R < 2^{H_{\text{rsd}}}$, $Q \approx 2 * 2^{H_{\text{rsd}}}$): Consider a sequence of a corruption followed by a reseed request and then a request to generate $\ell \gg H_{\text{rsd}}$ output bits. Thanks to the corruption, the attacker knows the input state $V_{d,i}$ for reseed. The input state $V_{d,i+1}$ for generate is computed by $V_{d,i+1} = \texttt{XOF}((V_{d,i} \parallel S_{d,i} \parallel \alpha_{d,i}), |V|)$ from the unknown seed $S_{d,i} \leftarrow_\$ \mathcal{D}_i^{H_{\text{rsd}}}$.

If $b = 0$, the visible output consists of the ℓ rightmost bits from $\texttt{XOF}(V_{d,i+1}, |V| + \ell)$. By trying out all $2^{H_{\text{rsd}}}$ choices for $S_{d,i}$, the attacker can find $V_{d,i+1}$ with matching output bits.

If $b = 1$, the visible output consists of ℓ random bits. Since $\ell \gg H_{\text{rsd}}$, the probability for the existence of any $S \in \{0,1\}^{H_{\text{rsd}}}$, such that the ℓ rightmost bits from $\texttt{XOF}(\texttt{XOF}(V_{d,i} \parallel S, |V|), |V| + \ell)$ match ℓ random bits is negligible.

6.2 Quantum Security: Applying Grover's Algorithm

What happens if the attacker can use a quantum computer? In the current paper, we always assume the DRBG to run on a classical computer. By implication, the challenger is classical, and attack 1 still applies. As the attacker can use a quantum computer, it can make $\texttt{XOF}$ calls *in superposition*. This allows the quantum attacker to use Grover's algorithm to speed-up attack 2a from $Q \approx 2^{H_{\text{init}}}/R$ $\texttt{XOF}$ calls down to $Q \approx 2^{H_{\text{init}}/2}/\sqrt{R}$ and attack 2b from $Q \approx 2^{H_{\text{rsd}}}$ calls down to $Q \approx 2^{H_{\text{rsd}}/2}$. *This is a serious issue for quantum secure DRBGs.*

On the other hand, Grover's algorithm doesn't parallelize well. An implementation of attack 2a or 2b, running c *classical* cores in parallel, speeds up by a factor of c. The speed-up of Grover's algorithm from running c quantum cores is only $\sqrt{c}$.

A concrete example: If the classical attack takes time 2^{85} on a single classical core, then 2^{15} classical cores running in parallel suffice to reduce the wall-clock time for the attack to the equivalent of $2^{85}/2^{15} = 2^{70}$ sequential $\texttt{XOF}$ calls. If Grover's algorithm takes the same 2^{85} units of time on a single quantum core, we'd need 2^{30} quantum cores to mount the attack in time $2^{85}/\sqrt{2^{30}} = 2^{70}$. Given the same number of 2^{30} cores, but classical ones, and 2^{70} units of wall-clock time, we could classically exhaust a 100-bit search space. So

let us assume the attacker not to be willing to wait for more than the wall-clock time 2^{70} cryptographic operations would take,[10] either on a classical or on a quantum computer. In this case, 85-bit quantum security is as good as 100-bit classical security. In general, a classical security level of $70 + 2t$ bits is equivalent a quantum security level of only $70 + t$ bits.[11]

7   Concrete Proposals

7.1   Numerical Examples

Table 1 provides some numerical examples, derived from corollary 1 and 3, respectively. The security levels depend on the state size, the entropy bounds, on the number R of requests and, in the case of the bounds derived from corollary 3, also on the maximum number R_{DEV} of requests from a single device. E.g., if we assume $|V| \geq 256$, $H_{\text{init}} \geq 192$, $H_{\text{rsd}} \geq 128$, $R \leq 2^{64}$, and apply corollary 1, we can guarantee a classical security level of 128 bits, and a quantum security level of 64 bits. But we can increase the bound on R to $R \leq 2^{128}$, we can maintain the same approximate security bound if we assume personalized devices (i.e., each device is given a unique name as its additional input), restrict each single device to at most $R_{\text{DEV}} \leq 2^{56}$ queries and apply corollary 3.

Table 1: Approximate security levels for different instantiations of the XDRBG, assuming an ideal XOF and derived by applying corollary 1 or 3, respectively. The quantum security levels assume a straightforward application of Grover's algorithm. Each row describes lower bounds for each of $|V|$, H_{init}, and H_{rsd}, and an upper bound for R (and for R_{DEV} in the lower part), to achieve a given level $\mathcal{L}$ of classical and quantum security. Note that if $R = R_{\text{DEV}}$, the bounds from corollary 1 and 3 are the same. The connection of Q and $\mathcal{L}$ is explained in section 2.3.

			bounds derived from corollary 1		
$|V|$	H_{init}	H_{rsd}	$\log_2(R)$	approx. security level $\mathcal{L}$	
$\geq$	$\geq$	$\geq$	$\leq$	classical	quantum
256	192	128	64	128	64
512	240	240	56	192	96
512	384	256	128	256	128

			bounds derived from corollary 3			
$|V|$	H_{init}	H_{rsd}	$\log_2(R)$	$\log_2(R_{\text{DEV}})$	approx. security level $\mathcal{L}$	
$\geq$	$\geq$	$\geq$	$\leq$	$\leq$	classical	quantum
256	192	128	128	56	128	64
512	240	240	128	56	192	96
512	384	256	128	128	256	128

7.2   Recommendation on Personalization of XDRBG Instances

We recommend to personalize all implementations of the XDRBG *at least* when the bound for the number R of is less than 2^{64}. If the XOF itself provides a personalization option, as, e.g., cSHAKE does for the SHAKE XOF, one could make use of that option and then leave the α empty.

---

[10]We argue that this assumption is realistic. No attacker cares about 1 000 or 10 000 years to mount an attack on sequential hardware. The attacker cares about the amount of parallel hardware needed to finish the attack in a given amount of time. The threshold size 2^{70} is, of course, open for debate.

[11]Reality may be even worse for the quantum attacker. Evaluating a cryptographic primitive on quantum circuits should be slower, in practice, than evaluating the same primitive on classical hardware.

In Appendix A.2 we briefly discuss the approach of going beyond personalization or randomization by actually providing additional entropy for the additional input.

7.3   Proposed `XDRBG` Parameters Based on `SHAKE128` and `SHAKE256`

We propose three `XDRBG` instances in Table 4. The first one employs `SHAKE128` with a capacity of 256 bits and provides 128 bits of classical security and 64 bits of quantum security. This matches category one (the lowest category) from the NIST post-quantum security criteria, see Appendix D. The second and third instance employ `SHAKE256`, where we claim security category three for the second instance and category five, the top category, for the third instance.

Table 2: Three proposals for DRBG standards and their approximate security levels. The first assumes `SHAKE128` with its 256-bit capacity, the second and third `SHAKE256` with its 512-bit capacity. The first two require personalized devices, with the bound $R_{\mathrm{DEV}} \leq 2^{56}$ on the number of queries from a single device, the third one can achieve its claimed security without such a constraint. We set $|V| =$ capacity. The promised classical security levels stem from corollary 3, though for `XDRBG`-256 corollary 1 would suffice to derive exactly the same bound. The quantum security levels assume the application of Grover's algorithm. The *category* refers to the NIST post-quantum criteria, cf. Appendix D.

	capacity	H_{init}	H_{rsd}	$\log_2$ (R)	$\log_2$ (R_{DEV})	promised security level $\mathcal{L}$		
						classi- cal	quantum (Grover)	cate- gory
XDRBG-128	256	192	128	128	56	128	64	1
XDRBG-192	512	240	240	128	56	192	96	3
XDRBG-256	512	384	256	128	128	256	128	5

These proposals are inspired by existing or drafted standards. In SP 800-90, a DRBG has a security level, $k \in \{128, 192, 256\}$.[12] The min-entropy needed for instantiation is defined in terms of the security level: $H_{\mathrm{init}} \geq 3k/2, H_{\mathrm{rsd}} \geq k$. I.e., `XDRBG`-128 and `XDRBG`-256 match the cases $k = 128$ and $k = 256$, respectively. `XDRBG`-192 is inspired by the revised version of AIS 20/31 [Pet23], (still a draft as of this writing), which defines lower-limits on effective internal state size and entropy provided to the DRNG[13]. The current draft of AIS 20/31 requires 240 bits of min-entropy for instantiation, so $H_{\mathrm{init}} \geq 240, H_{\mathrm{rsd}} \geq 240$. An implementation meant to comply with both would simply choose the maximum of the two required values. It is always allowable to incorporate *more* entropy than required, or to assume/require a *smaller* number R of requests. In this way, an implementation can be compatible with both standards.

7.4   Alternative Proposals Based on the ASCON Permutation

Recently, [Nat23], NIST has announced plans to standardize ASCON [DEMS22], a lightweight family of authenticated encryption and hashing algorithms based on a 320-bit permutation. We anticipate corresponding standards for ASCON-based `XOF`s.

A typical constraint for *lightweight* primitives is the number of input and output bits – e.g., $b = 320$ bit for the ASCON permutation, in contrast to 1600 bit for the `SHAKE` permutation. As pointed out above, any sponge-based `XOF` with a capacity of c bit can only provide $c/2$ bit of classical and $c/3$ bit of quantum security. Note that the rate $r = b - c$ determines the maximum number of input bits to be absorbed or output bits to be squeezed from each time the permutation is called. I.e., the performance of a `XOF` is roughly proportional to the rate.

---

[12]Currently, SP 800-90 also supports $k = 112$, but we expect that to be removed from the next version.
[13]Recall that DRNG is the term used in AIS 20/31 for what we refer to as a DRBG.

Table 3: Two proposals for lightweight DRBG standards and their promised security levels, which again stem from corollary 3 in the classical case, and from Grover's algorithm in the quantum case. The *category* refers to the NIST post-quantum criteria, cf. Appendix D.

	capacity	H_{init}	H_{rsd}	$\log_2 (R)$	$\log_2 (R_{\text{DEV}})$	promised security level $\mathcal{L}$		
						classi-cal	quantum (Grover)	cate-gory
XDRBG-L-128	256	192	128	128	56	128	64	1
XDRBG-L-170	308	240	240	128	64	170	85	2

Table 3 proposes two ASCON-based lightweight variants of the XDRBG. The first one, XDRBG-L-128, is a lightweight alternative to XDRBG-128 with the same security claims. Note that the capacity is $c = 256$ bit, so the rate is still $r = 320 - c = 64$ bit. The second one, XDRBG-L-170 aims at improved security, namely category-2 quantum security. Performance-wise, the improved security comes at a steep price. Since the capacity is $c = 308$ bit, i.e., almost the full permutation size of 320 bit (and can't be much smaller, if one wants to claim category-2 security based on the bounds we have proven), this leaves only $r = 12 = 320 - c$ bits for the rate.

7.5 Limiting the Damage from State Compromise: `maxout`

XDRBG, like the DRBGs in [BK15], promises backtracking resistance *between* GENERATE calls, but not *within* a GENERATE call. Thus, a state compromise during a very long GENERATE output could expose previously-generated outputs from the GENERATE call to compromise. This risk was pointed out in [WS19].

This threat is partly mitigated[14] in SP 800-90A DRBGs by defining a limit, `maxout`, on the maximum output length from one single GENERATE call. That is, each GENERATE call must return no more than `maxout` bits of output. To generate $X > \texttt{maxout}$ bits of output, one has to call the generate function $\lceil X/\texttt{maxout} \rceil$ times.

While the choice of `maxout` has no effect on our security bounds (we promise backtracking resistance only between GENERATE calls, and our security definition reflects this), a not-too-large `maxout` limit is a low-cost defense against the impact of a state compromise. Although we cannot offer a rigorous analysis in defense of any particular choice, we believe a `maxout` of around 2048 bits for XDRBG provides a reasonable performance/security tradeoff. However, the precise value of `maxout` should be tailored to the underlying XOF, to maximize performance. For example, SHAKE128 has a rate of 1344 bits, and XDRBG uses the first 256 bits as its new DRBG state, so it would be wasteful to set the `maxout` of XDRBG-128 to 2048 bits–setting `maxout` = 2432 gets the full benefit of the bits produced by two permutation calls. Table 4 shows recommended `maxout` values for each of our proposed versions of XDRBG. These values are derived by finding the smallest number of complete permutation calls that give an output of at least 2048 bits after accounting for the generation of the next DRBG state. However, we emphasize that the specific choice of `maxout` is not based on the results of a rigorous security analysis, and could reasonably be changed when the requirements of the application using the DRBG require it.

7.6 Performance and Usefulness

While performance is usually not the *most* important feature of a DRBG, it still matters. Table 5 gives a preliminary performance comparison on several different platforms, including both low-end and high-end processors across three different architectures. In our data,

---

[14]In [BKM+22], an additional requirement is imposed: the full GENERATE output must be produced before any of the output is used.

Table 4: Recommended values for `maxout` for different proposed `XDRBG` parameters. In each case, `maxout` is chosen to be the smallest output length of at least 2048 bits that can be produced from an integer number of permutation calls, after accounting for the need to generate a new DRBG state. Note that `XDRBG-192` and `XDRBG-256` are identical other than the amount of entropy they require for instantiation and reseeding.

	Based on	Rate	Recommended `maxout`
XDRBG-128	SHAKE128	1344	2432
XDRBG-192	SHAKE256	1088	2752
XDRBG-256	SHAKE128	1088	2752
XDRBG-L-128	ASCON_XOF	64	2048
XDRBG-L-170	ASCON_XOF	12	2052

`XDRBG` based on `SHAKE256` is always competitive with HashDRBG or HMAC based on SHA256, and is usually faster.

Table 5: Comparing XDRBG to other hash-based DRBGs on a variety of platforms. All DRBGs claim 256 bits of classical security. Times given are in seconds required to fill a buffer with 1 GiB of output from the DRBG (smaller numbers mean better performance). The "Vec. Instr." column shows performance using whatever vector instructions were available on the given platform.

	`XDRBG-256` SHAKE256 Vec. Instr.	`XDRBG-256` SHAKE256	HashDRBG SHA256	HMACDRBG SHA256
AMD Ryzen 5950X	4.62	5.06	6.63	27.33
Intel 11th Gen i7-1195G7	2.64	5.28	5.81	23.11
Intel 12th Gen i7-1280P	3.96	4.15	4.68	19.97
Apple M2	2.18	2.89	4.88	20.54
ARM Cortex-A76 r4p1	6.17	6.52	9.59	41.15
ARM Cortex-A72 r0p3	12.26	12.33	18.15	78.20
ARM Cortex-A8 r2p5	62.68	185.45	186.71	784.66
ARM Cortex-A7 r0p5	81.81	249.85	242.79	1015.23
Si Five (RISC-V)	104.73	104.80	72.11	309.03

`XDRBG` is designed to use the normal `XOF` interface, rather than having access to any internal values or state, or making any assumptions about the underlying `XOF`'s inner workings other than its security strength. This makes the DRBG somewhat less efficient, but with the benefit that the DRBG can be implemented pretty efficiently with normal access to a `XOF` primitive such as `SHAKE256`, and will work as well for future `XOF`s–even ones not based on a sponge construction.

`XDRBG` is designed to keep the inputs to the `XOF` as small as possible, given its security requirements. Consider the number of permutation calls to handle the different types of requests:

- Handling a reseed request with a seed of length σ and additional input of length $|\alpha|$ needs to make $\lceil (|V| + \sigma + |\alpha| + |\text{ENCODE}|)/r \rceil$ permutation calls to absorb its entire input and to emit the first r bits of the new state. If $|V| > r$, it needs an additional $\lceil |V|/r \rceil - 1$ permutation calls to generate the new state in full. In total, this makes

$$\left\lceil \frac{|V| + \sigma + |\alpha| + |\text{ENCODE}|}{r} \right\rceil + \left\lceil \frac{|V|}{r} \right\rceil - 1 \quad \text{permutation calls.}$$

- Handling an instantiate request is almost the same as handling a reseed request, except that the input does not expect a seed of any size. Thus, it requires

$$\left\lceil \frac{\sigma + |\alpha| + |\text{ENCODE}|}{r} \right\rceil + \left\lceil \frac{|V|}{r} \right\rceil - 1 \quad \text{permutation calls.}$$

- Similarly, handling a generate request for ℓ bits of output takes

$$\left\lceil \frac{|V| + |\alpha| + |\text{ENCODE}|}{r} \right\rceil + \left\lceil \frac{|V| + \ell}{r} \right\rceil - 1 \quad \text{permutation calls.}$$

As a concrete (arguably quite typical) example, consider $\sigma = 512$, $|\text{ENCODE}| = 8$ (cf. appendix B) and $\ell = 256$. We assume the additional input not to be too long, say, $|\alpha| = 128$. The performance of our proposed XDRBG variants is as follows:

- XDRBG-128 has rate $r = 1344$ and $|V| = 256$. Regardless of the request, given our choice of parameters the XDRBG-128 calls the permutation only once.

- XDRBG-192 and XDRBG-256 have rate $r = 1088$ and $|V| = 512$. Both XDRBG-192 and XDRBG-256 make two permutation calls for reseed and one call for each of instantiate and generate.

- The rate and state size of XDRBG-L-128 are $r = 64$ and $|V| = 256$. A reseed thus takes $\lceil \frac{256+512+128+8}{64} \rceil + \lceil \frac{256}{64} \rceil - 1 = 18$ calls, instantiate $\lceil \frac{512+128+8}{64} \rceil + \lceil \frac{256}{64} \rceil - 1 = 14$ calls, and generate also $\lceil \frac{256+128+8}{64} \rceil + \lceil \frac{256+256}{64} \rceil - 1 = 14$ calls.

- For XDRBG-L-170 the rate is $r = 12$ and the state size is $|V| = 308$. This implies 101 calls for reseed, 89 calls for instantiate, and 73 calls for generate.

The above comparison is simplistic. E.g., we assumed a constant seed size $\sigma = 512$. In practice, σ will likely be proportional to the required entropy. I.e., we neglect a certain benefit for the less secure variants of the XDRBG. Nevertheless, we believe the above comparison still gives a reasonable idea for typical application scenarios.

8   Alternative Approaches

We considered and rejected a number of design alternatives and alternative security models for XDRBG.

8.1   Sponge-based vs XOF-based

The first design choice we had was whether to base the DRBG on a sponge function or a XOF. (Recall that a XOF provides a particular kind of functionality that can be implemented by sponge function, but might also be implemented in some other way.) As discussed above, several prior works, have proposed cryptographic PRNGs based on a sponge construction (specifically using a large ideal permutation), and it would have been relatively easy to adapt those to the requirements of SP 800-90A. However, we believe that basing a DRBG on a XOF provides more flexibility. XDRBG can be based on any XOF, regardless of its underlying structure or assumptions. A future XOF whose security does not rely on an ideal permutation assumption or even use a sponge construction will still work with XDRBG.

8.2   Reseed Interval

SP 800-90A defines a *reseed interval* for its DRBGs. This requires that the DRBG be reseeded after a certain number of GENERATE calls. A relatively small reseed interval provides a defense against the impact of a DRBG state compromise. However, a reseed interval short enough to provide substantial protection from state compromise would also make the DRBG algorithm unworkable in many environments. For example, SP 800-90C defines the RBG1 construction, which has access to live entropy only for instantiation; a similar construction is permitted as a DRG.3 under AIS 20/31. If XDRBG required reseeding

every ten or even one hundred GENERATE calls, it would be unworkable for use in these constructions. On the other hand, the huge reseed intervals (requiring a RESEED every 2^{32} or 2^{48} GENERATE calls) in [BK15] are not too costly, but their security benefit is negligible. For these reasons, we elected not to define a reseed interval as part of XDRBG.

However, we recommend that in any application where it is practical, XDRBG (or any other DRBG) should be reseeded periodically. As described in [KSWH98, CDK+22a], a reseed *must* wait until sufficient entropy is available to avoid the iterative guessing attack / premature next condition.

A conditional reseed may serve as an easy alternative to a fixed reseed interval without forcing the application to wait (**if** the min-entropy of the seed is H_{rsd} **then** reseed before continuing **else** continue without reseeding). Alas, this also introduces a potential side-channel vulnerability: The attacker might observe if, whenever the DRBG executes a conditional reseed, the DRBG actually reseeds or not. Thus, each time a conditional reseed is called, one bit of Shannon entropy may be lost.

8.3 Stronger Attack Model

A variation of the attack game could allow the attacker to compromise the DRBG state by *overwriting* it, or *resetting* it to a fixed inital state, rather than by just *reading* it. The attack game would otherwise be the same, including the role of the flag `corrupt`. But now, the attacker could benefit from a multitarget attack, similar to the attack from section 2.5: Fix a state V^* and repeat the following three-step sequence as often as possible: (1) set the DRBG to V^*, (2) reseed, and (3) generate some output bits. Eventually try $O(Q)$ times to guess any of the $O(R)$ seeds from step (2) and generate the output bits to detect a match.

As pointed out above, we doubt the plausibility of an attacker to the DRBG state by a chosen or known value V^*. But for readers who prefer to consider such attacks, we point out the following: Reseeding is never worse than instantiating, except that H_{rsd} may be smaller than H_{init}. In fact, the combination of setting the DRBG state to V^* and then to reseed with a seed S is equivalent to instantiating the DRBG state with a seed $S^* = V^* \parallel S$. Since V^* is known to (or even chosen by) the attacker, the distributions, which S and S^* are drawn from, have exactly the same min-entropy. Accordingly, we make two recommendations for the stronger attack model: (1) set $H_{\text{rsd}} \approx H_{\text{init}}$, and (2) if the XDRBG is personalized, then personalize both instantiate and reseed requests. In that context, observe that AIS 20/31 sets $H_{\text{init}} = H_{\text{rsd}}$.

Regarding the first recommendation, we revisited the proof for our main result. As it turned out, both proof and result still apply, with a slightly tweaked bound. In fact, for $H = H_{\text{init}} = H_{\text{rsd}}$ we can replace equation 2 by

$$\epsilon \le Q \left(\frac{\lambda_1}{2^H - Q} + \frac{1}{2^H - Q} + \frac{Q}{2 \times 2^{|V|}} \right) + \frac{\lambda_2}{2^H} + \frac{R^2}{2 \times 2^{|V|}},$$

which also applies to the attack model where the attacker can set the DRBG state.

9 Conclusions

Drawing on previous work in [BDPV10, CDKT19], we have proposed a new class of DRBG that can be based on any XOF, and analyzed its security in a model well adapted for DRBGs as defined in [BK15]. We have kept the specification of XDRBG general enough to adapt to the different requirements of SP 800-90 and AIS 20/31, and to work for any XOF. It is possible to make a more efficient DRBG by altering the internal workings of a sponge-based XOF, but this would result in a less generally-useful DRBG. We prefer to provide a more generic design.

XDRBG is quite efficient. Assuming reasonable choices for seed length, size of the additional input and output length, every XOF query made by XDRBG with SHAKE128 or SHAKE256 will result in at most two permutation calls. Even XDRBG-L-170, the lightweight variant based on a smallish 320-bit permutation, where we squeezed in as much quantum security as possible, is slow, but not prohibitively so: a typical request requires about 100 permutation calls. For the lightweight variant with standard security, namely for XDRBG-L-128, less than 20 permutation calls suffice.

Our hope is that XDRBG will be a useful addition to the set of DRBGs currently in use, especially in environments in which SHAKE or ASCON is the only cryptographic primitive available. Our focus in the design and analysis of XDRBG is to provide a practical DRBG that fits cleanly with the requirements of two widely-followed standards for cryptographic random bit generation, SP 800-90 and AIS 20/31. We have provided concrete parameter sets to meet these requirements, in the hopes of making XDRBG easy to incorporate into applications.

Open Questions.   We see at least three interesting directions for future research:

1. Our analysis for quantum security focused on the application of Grover's algorithm. Also, recall the cautionary note from Footnote 5 regarding the quantum security of sponge-based XOFs. Nevertheless, we conjecture that these bounds actually describe the security of the XDRBG against quantum adversaries very well. A *proof for those quantum bounds*, similar to our classical security analysis, would probably assume the compressed oracle model [Zha19].

2. Our security proof assumes oracle-independent entropy sources. While this seems entirely reasonable for the context of random number generation with trusted entropy sources (as in SP 800-90 or AIS 20/31), it would be interesting *to prove similar security bounds for oracle-dependent sources*, using techniques from [CDKT19].

3. We are not entirely satisfied with the performance of our high-security small-permutation proposal XDRBG-L-170. To match a given security level by applying our bounds, we seem to require a largish capacity (because $|V|$ is a lower bound for the capacity), which implies a low rate for XDRBG-L-170. *Is it possible to prove similar bounds with a significantly smaller $|V|$?*

Acknowledgments

References

[AMD23]   AMD – Advanced Micro Devices. AMD RNG ESV public use document, document version 0.4. Technical report, Advanced Micro Devices(AMS), 2023. https://csrc.nist.gov/CSRC/media/projects/cryptographic-module-validation-program/documents/entropy/E27_PublicUse.pdf.

[AMMS16]   Sergey Agievich, Vadim Marchuk, Alexander Maslau, and Vlad Semenov. Bash-f: another lrx sponge function. Cryptology ePrint Archive, Paper 2016/587, 2016. https://eprint.iacr.org/2016/587.

[BDPA07]   G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. Sponge functions. *Ecrypt Hash Workshop 2007*, 2007.

[BDPV10]   Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sponge-based pseudo-random number generators. In Stefan Mangard and François-Xavier Standaert, editors, *CHES 2010*, volume 6225 of *LNCS*, pages 33–47, Santa Barbara, CA, USA, August 17–20, 2010. Springer, Heidelberg, Germany.

[BDPVA08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indifferentiability of the sponge construction. In Nigel Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[BK15]   Elaine Barker and John Kelsey. Recommendation for random number generation using deterministic random bit generators. Technical report, National Institute of Standards & Technology, Gaithersburg, MD, United States, 2015.

[BKL+19]   D. J. Bernstein, S. Kölbl, Stefan Lucks, P. Maat Costa Massolino, F. Mendel, K. Nawaz, T. Schneider, P. Schwabe, F.-X. Standaert, Y. Todo, and B. Viguier. Gimli. Submission to the NIST Lightweight Cryptography Standardization Process, 2019. https://csrc.nist.gov/Projects/lightweight-cryptography/round-2-candidates.

[BKM+22]   Elaine Barker, John Kelsey, Kerry McKay, Allen Roginsky, and Meltem Sönmez Turan. Recommendation for random bit generator (rbg) constructions (3rd draft). Technical report, National Institute of Standards & Technology, Gaithersburg, MD, United States, 2022.

[CDK+22a] Sandro Coretti, Yevgeniy Dodis, Harish Karthikeyan, Noah Stephens-Davidowitz, and Stefano Tessaro. On seedless prngs and premature next. In Dana Dachman-Soled, editor, *3rd Conference on Information-Theoretic Cryptography, ITC 2022, July 5-7, 2022, Cambridge, MA, USA*, volume 230 of *LIPIcs*, pages 9:1–9:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[CDK+22b] Sandro Coretti, Yevgeniy Dodis, Harish Karthikeyan, Noah Stephens-Davidowitz, and Stefano Tessaro. On seedless prngs and premature next. *IACR Cryptol. ePrint Arch.*, page 558, 2022.

[CDKT19]   Sandro Coretti, Yevgeniy Dodis, Harish Karthikeyan, and Stefano Tessaro. Seedless fruit is the sweetest: Random number generation, revisited. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 205–234. Springer, 2019.

[Cza21]   Jan Czajkowski. Quantum indifferentiability of SHA-3. *IACR Cryptol. ePrint Arch.*, page 192, 2021.

[DEM+20]   Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, Bart Mennink, Robert Primas, and Thomas Unterluggauer. ISAP v2.0. *IACR Trans. Symm. Cryptol.*, 2020(S1):390–416, 2020.

[DEMS22]   Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläffer. Status Update on Ascon v1.2. Update to the NIST Lightweight Cryptography Standardization Process, 2022. https://csrc.nist.gov/csrc/media/Projects/lightweight-cryptography/documents/finalist-round/status-updates/ascon-update.pdf.

[DGH+04]   Yevgeniy Dodis, Rosario Gennaro, Johan Håstad, Hugo Krawczyk, and Tal Rabin. Randomness extraction and key derivation using the CBC, cascade and HMAC modes. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 494–510, Santa Barbara, CA, USA, August 15–19, 2004. Springer, Heidelberg, Germany.

[DVW20]   Yevgeniy Dodis, Vinod Vaikuntanathan, and Daniel Wichs. Extracting randomness from extractor-dependent sources. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 313–342. Springer, 2020.

[Fer19]   Niels Ferguson. The Windows 10 random number generation infrastructure. Technical report, Microsoft, 2019. url = https://aka.ms/win10rng.

[FS03]   Niels Ferguson and Bruce Schneier. *Practical cryptography*. Wiley, 2003.

[GT16a]   Peter Gaži and Stefano Tessaro. Provably robust sponge-based PRNGs and KDFs. Cryptology ePrint Archive, Report 2016/169, 2016. https://eprint.iacr.org/2016/169.

[GT16b]   Peter Gazi and Stefano Tessaro. Provably robust sponge-based PRNGs and KDFs. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 87–116, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.

[Hut16a]   Daniel Hutchinson. A robust and sponge-like prng with improved efficiency. Cryptology ePrint Archive, Paper 2016/886, 2016. https://eprint.iacr.org/2016/886.

[Hut16b]   Daniel Hutchinson. A robust and sponge-like PRNG with improved efficiency. In Roberto Avanzi and Howard M. Heys, editors, *Selected Areas in Cryptography - SAC 2016 - 23rd International Conference, St. John's, NL, Canada, August 10-12, 2016, Revised Selected Papers*, volume 10532 of *Lecture Notes in Computer Science*, pages 381–398. Springer, 2016.

[Kil11]   Killmann, Wolfgang and Schindler, Werner. A proposal for functionality classes for random number generators. Technical Report AIS20, Bundesamt für Sicherheit in der Informationstechnik (BSI), 2011.

[KSF99]   John Kelsey, Bruce Schneier, and Niels Ferguson. Yarrow-160: Notes on the design and analysis of the yarrow cryptographic pseudorandom number generator. In Howard M. Heys and Carlisle M. Adams, editors, *Selected Areas in Cryptography, 6th Annual International Workshop, SAC'99, Kingston, Ontario, Canada, August 9-10, 1999, Proceedings*, volume 1758 of *Lecture Notes in Computer Science*, pages 13–33. Springer, 1999.

[KSWH98]   John Kelsey, Bruce Schneier, David Wagner, and Chris Hall. Cryptanalytic attacks on pseudorandom number generators. In Serge Vaudenay, editor, *FSE'98*, volume 1372 of *LNCS*, pages 168–188, Paris, France, March 23–25, 1998. Springer, Heidelberg, Germany.

[Mec18]   John P Mechalas. Intel® Digital Random Number Generator (DRNG) Software Implementation Guide, 2018. https://www.intel.com/content/www/us/en/developer/articles/guide/intel-digital-random-number-generator-drng-software-implementation-guide.html.

[MMHH23] Stephan Müller, Sebastian Mayer, Caroline Holz auf der Heide, and Andreas Hohenegger. Documentation and analysis of the Linux random number generator. Technical report, Bundesamt für Sicherheit in der Informationstechnik (BSI), 2023.

[Nat01] National Institute of Standards and Technology. Advanced encryption standard (AES). Technical Report Federal Information Processing Standards (FIPS) Publication 197, U.S. Department of Commerce, Washington, D.C., 2001.

[Nat13] National Institute of Standards and Technology. Digital signature standard (dss). Technical Report Federal Information Processing Standards (FIPS) Publication 186-4, U.S. Department of Commerce, Washington, D.C., 2013.

[Nat15a] National Institute of Standards and Technology. Secure hash standard (SHS). (U.S. Department of Commerce, Washington, DC), Federal Information Processing Standards Publication (FIPS) 180-4, August 2015. https://doi.org/10.6028/NIST.FIPS.180-4.

[Nat15b] National Institute of Standards and Technology. Sha-3 standard: Permutation-based hash and extendable output functions. Technical Report Federal Information Processing Standards (FIPS) Publication 202, U.S. Department of Commerce, Washington, D.C., 2015.

[Nat16] National Institute of Standards and Technology. Submission requirements and evaluation criteria for the post-quantum cryptography standardization process, 2016. url = https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf.

[Nat23] National Institute of Standards and Technology. Lightweight cryptography standardization process: NIST selects Ascon, Feb 2023. url = https://csrc.nist.gov/News/2023/lightweight-cryptography-nist-selects-ascon.

[NL18] Yoav Nir and Adam Langley. ChaCha20 and Poly1305 for IETF Protocols. RFC 8439, June 2018.

[Pet23] Peter, Matthias and Schindler, Werner. A proposal for functionality classes for random number generators–version 2.35 draft. Technical Report AIS20, Bundesamt für Sicherheit in der Informationstechnik (BSI), 2023.

[RS16] Ronald L. Rivest and Jacob C. N. Schuldt. Spritz - a spongy rc4-like stream cipher and hash function. *IACR Cryptol. ePrint Arch.*, page 856, 2016.

[TBK+18] Meltem Sönmez Turan, Elaine Barker, John Kelsey, Kerry McKay, Mary Baish, and Mike Boyle. Recommendation for random number generation using deterministic random bit generators. Technical report, National Institute of Standards & Technology, Gaithersburg, MD, United States, 2018.

[WS19] Joanne Woodage and Dan Shumow. An analysis of NIST SP 800-90A. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 151–180, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.

[Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors,

Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II, volume 11693 of *Lecture Notes in Computer Science*, pages 239–268. Springer, 2019.

A   Unusual Use Cases

Our security analysis above discusses the normal use cases for a DRBG. In this appendix, we consider less common ways a DRBG may be used and how this might affect our security bounds.

A.1   Seeding from another DRBG

In some contexts, a DRBG's seed may come from another DRBG. Since our security proof assumes access to entropy for seed material, it is natural to consider the security impact of seeding from a DRBG. We can extend our security bounds to deal with the situation, by simply incorporating an additional term for violating the security of the DRBG providing the seed. Informally, if the attacker cannot distinguish the outputs of the DRBG providing the seed from ideal random outputs, it also cannot gain any advantage in distinguishing XDRBG outputs seeded from those DRBG outputs.

A.2   Adding Entropy via the Additional Input

XDRBG, like the DRBGs in SP 800–90A, allows for an optional additional input to each DRBG call. This input may be used in many different ways in practical applications. For example:

1. Some systems maintain a *seed file*, to save entropy across device restarts as a hedge against an entropy source failure. A natural way to incorporate the seed file into the DRBG state is to put it into the additional input of the instantiate call.

2. Secret information, such as the hash of a private key, can be incorporated into the DRBG during instantiation, again to provide a hedge against the failure of the entropy source.

3. Additional entropy can be drawn from some secondary entropy source, or even from the primary entropy source, and provided to the DRBG during instantiation or reseeding.

Intuitively, it is easy to see that entropy provided in the additional input is incorporated into the DRBG state in the same way as the seed, since the additional input is simply appended to the seed in instantiate and reseed calls. Thus, an INSTANTIATE or RESEED in which sufficient entropy is provided in the additional input will end up in a secure state, even if no entropy is provided in the seed.

Let h_1 the entropy in the seed, and h_2 be the entropy in the additional input. By virtue of our encoding function, and as long as the seed is independent from the additional input, the string input to the XOF must thus have $h_1 + h_2$ bits of min-entropy. And even if seed and additional input are statistically dependent, the min-entropy of the encoded input for the XOF has at least $\max(h_1, h_2)$ bits of min entropy.

B The Function encode

XDRBG requires a function ENCODE : $\{0,1\}^* \times \{0,1\}^* \times \{0,1,2\} \to \{0,1\}^*$ such that for all $(S, \alpha, n) \neq (S', \alpha', n')$ ENCODE$(S, \alpha, n) \neq (S', \alpha', n')$. (That is, the encoding function must not introduce any trivial collisions.) There are many suitable encodings possible, but for concreteness, we define a recommended encoding as follows:

Let $|\alpha|/8 \in \{0, \ldots, 84\}$. I.e., the additional input α is a sequence of bytes, and it is at most 84 bytes long. Then, the following encoding unambiguously encodes the inputs while adding only a single byte of stretch:

$$\text{ENCODE}(S, \alpha, n) = (S \parallel \alpha \parallel (n * 85 + |\alpha|/8)_8),$$

where $(\ldots)_8$ indicates an 8-bit (i.e., single-byte) encoding of a value in $\{0, 255\}$. Thus, $|\text{ENCODE}| = 8$, i.e., the stretch is constantly one byte.

This encoding is efficient and flexible, but does require that the additional input string is no longer than 84 bytes–a constraint that seems very easy to manage in practice. For example, IPV6 addresses and GUIDs are 16 bytes long, Ethernet addresses are 12 bytes long, and the most demanding requirement for unique randomly-generated device identifiers can be met with a 32-byte random value. Thus, we recommend this encoding for XDRBG.

C HashXOF

Although there are already multiple widely-used DRBGs based on a hash function, we can construct a XOF suitable for XDRBG from any standard hash function, such as SHA256. The design is as follows:

```
1: function HashXOF(x, ℓ)
2:     t ← Hash(x ‖ 0₆₄)
3:     Z ← ε
4:     i ← 1
5:     while |Z| < ℓ do
6:         Y ← Hash(t ‖ i₆₄)
7:         i ← i + 1
8:         Z ← Z ‖ Y
9:     return(Z truncated to ℓ bit)
```

Using XDRBG with HashXOF(SHA256) will provide comparable performance to either HMAC_DRBG(SHA256) or Hash_DRBG(SHA256).

D NIST Post-Quantum Security Categories

NIST defines five security categories for submissions to the Post-Quantum Cryptography process in [Nat16]. These categories are as follows:

| category | requirement | security | |
	any attack must require computational resources comparable to or greater than those required for	classical	quantum
1	key search on a block cipher with a 128-bit key	128	64
2	collision search on a 256-bit hash function	128	85
3	key search on a block cipher with a 192-bit key	192	96
4	collision search on a 384-bit hash function	192	128
5	key search on a block cipher with a 256-bit key	256	128

E Pool-Based DRNGs

In this paper, we assume the existence of trusted entropy sources that provide strings with a known amount of min-entropy on demand. In this, we follow the lead of NIST and BSI standards–SP 800-90 and AIS 20/31 specify techniques for evaluating entropy sources, and then assume the availability of entropy sources whose claims of entropy can be relied upon.

When entropy sources do not reliably provide a known amount of min-entropy, or when they may even be adversarially controlled, a very different approach is required. The Fortuna cryptographic PRNG, first described in [FS03], is designed to guarantee that its PRNG algorithm will eventually be seeded securely, as long as the entropy sources used are providing *some* entropy, even without any way to know how much is being provided. This kind of design is modeled in depth in [CDK+22a].

The Fortuna PRNG considers different entropy pools, which over time receive inputs from one or more sources of entropy. It is not exactly known how much entropy the sources provide. From time to time, the entropy gathered in some subset of these pools is used to update the state. If, after a state compromise, the pools have gathered a sufficient amount of entropy, then the PRNG recovers from the compromise. Otherwise, the entropy from those pools is essentially lost, since the attacker can use his knowledge of the previous state of the RNG and subsequent outputs to guess the entropy input. (This is referred to as "premature next" in [CDK+22b] and as the "iterative guessing attack" in [KSWH98].) The goal of Fortuna is to guarantee that the PRNG will eventually reach a secure state if there is any entropy being provided to it.

Potentially adversarial entropy sources, as well as "oracle-dependent" entropy sources (entropy sources whose distributions are not independent of the cryptographic functions used to extract entropy from them) are analyzed in [CDKT19, DVW20]. Since we assume trusted, already-analyzed entropy sources in this paper (in keeping with the SP 800-90 and AIS 20/31 standards), we do not consider adversarial or oracle-dependent sources here.