Physical Randomness Extractors: Generating Random Numbers with Minimal Assumptions

Kai-Min Chung[*]   Yaoyun Shi[†]   Xiaodi Wu[‡]

Abstract

How to generate provably true randomness with minimal assumptions? This question is important not only for the efficiency and the security of information processing, but also for understanding how extremely unpredictable events are possible in Nature. All current solutions require special structures in the initial source of randomness, or a certain independence relation among two or more sources. Both types of assumptions are impossible to test and difficult to guarantee in practice. Here we show how this fundamental limit can be circumvented by extractors that base security on the validity of physical laws and extract randomness from *untrusted* quantum devices. In conjunction with the recent work of Miller and Shi (arXiv:1402:0489), our *physical randomness extractor* uses just a single and general weak source, produces an arbitrarily long and near-uniform output, with a close-to-optimal error, secure against all-powerful quantum adversaries, and tolerating a constant level of implementation imprecision. The source necessarily needs to be unpredictable to the devices, but otherwise can even be known to the adversary.

Our central technical contribution, the Equivalence Lemma, provides a general principle for proving composition security of untrusted-device protocols. It implies that *unbounded* randomness expansion can be achieved simply by cross-feeding *any* two expansion protocols. In particular, such an unbounded expansion can be made robust, which is known for the first time. Another significant implication is, it enables the secure randomness generation and key distribution using *public* randomness, such as that broadcast by NIST's Randomness Beacon.

Our result has significant interpretations for fundamental physics. It implies that close-to-uniform randomness either does not exist in Nature or exist in abundance. It also provides the strongest known method for mitigating the Freedom-of-Choice loophole for refuting local hidden variable theories.

---

[*]Institute of Information Science, Academia Sinica, Taiwan.

[†]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48103, USA.

[‡]Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. Most of the research was conducted when the author was a student at the University of Michigan, Ann Arbor.

Version Differences. This draft differs substantially from both the first and the second versions of our arXiv posting (arXiv:1402.4797).

- V1 of arXiv:1402.4797 is our QIP 2014 submission (it was accepted and presented in a joint-plenary presentation.)

- V2 of arXiv:1402.4797 introduces and formally defines the notion of physical randomness extractors. The Master Protocol is also changed slightly so that for arbitrarily small min-entropy sources the protocol is robust. The new analysis requires some thought.

- This current version was a substantial re-writing of V2. The main technical new material is in the formal definition of physical extractors. The proofs are correspondingly changed. In particular, an abstract notion of error model is added, and the robustness claim in V2 is now rigorous.

1 Motivations

Randomness is a vital resource for modern day information processing. The wide range of its applications include cryptography, fast randomized algorithms, accurate physical simulations, and fair gambling. In practice, randomness is generated through a "random number generator" (RNG), such as Intel's on-chip hardware generator `RdRand` and Linux's software generator `/dev/random`. Since it is impossible to test if the output of a RNG is uniformly distributed or fixed, [1] one relies on the mathematical properties of the RNG to ensure the output quality under a set of assumptions that are hopefully true in reality. For example, Linux's RNG critically requires being seeded with a large amount of initial entropy and the unproven assumption that no adversary is computationally powerful enough to differentiate the output from uniform.

Those assumptions, however, have been repeatedly shown to cause failures of practical RNGs (see, e.g., [25, 39, 26, 33]). Such vulnerabilities of RNGs directly threaten the very foundation of digital security, and risk invalid conclusions drawn from computations assuming true randomness. Thus when security is of paramount importance, it is highly desirable to use RNGs that are secure under a minimal set of assumptions.

The classical theory for this objective is that of *randomness extractors* [42]. In this theory, an extractor is a deterministic algorithm that transforms several sources of weak randomness into near-perfect randomness. The amount of randomness in a weak source is quantified by *min-entropy*, or conditional quantum min-entropy when the adversary is quantum. More precisely, an (n, k) source is an n-bit binary string with (conditional quantum) min-entropy k, which means that the best chance for an adversary to guess the source correctly is $\leq 2^{-k}$ [44, 37, 38]. A fundamental limit known in this theory is that randomness extraction is possible only when *two or more independent* sources are available. In particular, deterministic extraction, i.e., using just one source, is known to be impossible to produce even 1 (near-prefect) random bit [40]. Since independence is impossible to check [2] and difficult to guarantee in practice, the classical theory of randomness extractors inevitably relies on *assuming* independence.

Quantum mechanics has perfect randomness in its postulate, thus appears to provide a simple solution to the problem [3]. Indeed, commercial products are already available (e.g., the Quantis generators of ID Quantique). However, users must trust the quantum devices in use for security. This is a strong assumption undesirable in certain circumstances for the following reasons. First, as classical beings, we can only directly verify classical information, thus cannot directly verify the inner-workings of quantum devices. Second, we may not want to trust the manufacturers or the certifying government agencies. Finally, even if the manufacturers are truthful, the devices themselves may not work properly due to technological limitations. No method is currently known for reliably implementing quantum devices in a large scale.

Recent works have shown that one can still leverage the quantum power for generating randomness even when the underlying quantum devices may be imperfect, or even malicious [19, 20, 34, 43, 30, 22, 21, 24, 16]. However, all those protocols also crucially rely on a certain form of independence. More specifically, *randomness expansion* protocols require that the source is globally uniform (thus independent from the devices) [19, 20, 34, 43, 30, 22]. *Randomness amplification* protocols [21, 24, 16] initiated by Colbeck and Renner [21] require that the source, when conditioned on the adversary's side information, is a highly random and highly structured SV-source [4]. In [21], the source in addition needs to satisfy certain causal relation among its blocks, while in [24, 16], conditioned on the adversary's information, the source and the devices are assumed to be independent. [1]

---

[1] The conditional independence condition is stated explicitly in the work Gallego et al. [24] (Supplementary Note

2 Our contributions

Physical Randomness Extractors: a model for extracting randomness without independence assumptions. To circumvent those fundamental limits and to minimize necessary assumptions, here we formulate a framework of extracting randomness from *untrusted* quantum devices in the quantum mechanical world, shown in Fig. 2. This framework of *Physical Randomness Extractors (PREs)* allows general and rigorous discussions of extracting randomness when the devices and the adversary are both bound by physical laws. This reliance on physical theories for security is a fundamental departure from the classical theory of randomness extraction. Since all cryptographic protocols will eventually be deployed in the physical world, no additional effort needs to be made to enforce the assumptions on the correctness of physical laws (as Nature automatically ensures that) [5].
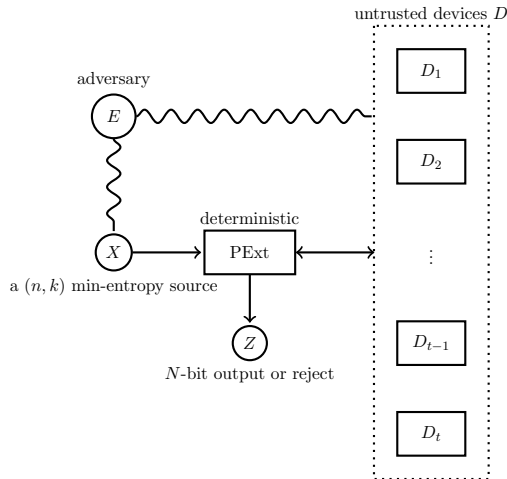

Figure 1: Physical Randomness Extractor (PRE). A PRE is a deterministic algorithm PExt that takes a classical source X as the input, interacts with a set of untrusted quantum devices D, and finally either aborts (aka rejects) or outputs a binary string Z. Each device is used through its classical input-output interface but its inner-working is unknown (and could be malicious). The Adversary E is quantum and all-powerful, may be in an unknown quantum correlation with the devices, and together with the devices may hold a certain amount of side information about X. After the protocol starts, no communication is allowed among the Adversary and the devices. The error of PExt upper-bounds both the probability of accepting an undesirable output (soundness error) or that of rejecting an honest implementation (completeness error). If X is globally uniformly random, PExt is said to be *seeded*; otherwise, it is *seedless*. PExt is *robust* if an honest implementation can deviate from an ideal implementation by a constant amount. See Section 4 for the formal definitions.

Our framework is built upon the above-mentioned two lines of research on randomness expansion and randomness amplification. It in particular includes the quantum restriction of those models [6]. As special cases. Randomness expansion is precisely *seeded* PRE-extraction, where the seed is uniform globally. Randomness amplification, when restricted to the quantum world, can be seen as *seedless* (i.e., the classical source is not uniform) PRE-extraction of a single bit with a restricted source. Our framework explicitly quantifies the various relevant resources. This allows richer analyses and comparisons of protocols, and raises new questions for optimizing the performance parameters and investigating their inherent tradeoffs. For example, the *extraction rate* introduced,

---

2) and Brandao et al. [16] (Section II.B). The causal relation assumed by Colbeck and Renner [21] is illustrated in their Fig. 2 and stated below it.

i.e., the ratio of the output length and the total length of the device output, is a natural measure for the efficiency of a PRE (See Section 4 for more details.) We discuss several fundamental open problem in Section 6.

We point out that it'd be more appropriate to consider the untrusted devices as the source of the output randomness, while the classical source is used to prevent cheating. This intuition is supported by the strong quantitative relations linking the min-entropy of the source to the error parameter, and the number of device usages to the output length. A useful comparison of a PRE with classical strong extractor is to consider the weak source of a PRE corresponds to the seed for the latter, while the devices correspond to the weak source. Under this correspondence there is a fundamental difference between those two models regarding the correlation vs independence of the corresponding two sources.
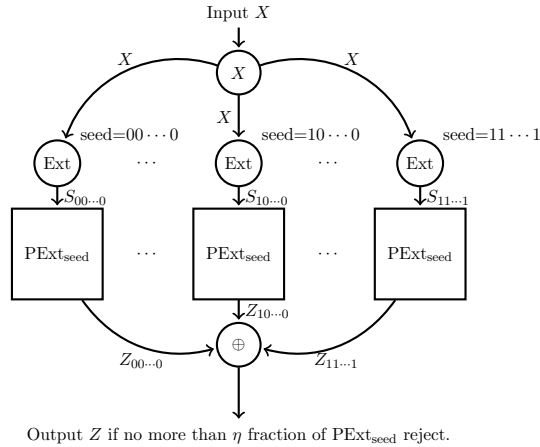

Output Z if no more than η fraction of $PExt_{seed}$ reject.

Figure 2: Our Physical Randomness Extractor PExt with parameters Ext, $PExt_{seed}$, and η. Ext is a quantum-proof strong extractor [7] and $PExt_{seed}$ a seeded-PRE whose input length equals the output length of Ext. For each distinct seed value i of Ext, run an instance of Ext with that seed value and X as the source. Use the output S_i as the input to a separate instance of $PExt_{seed}$. Output the XOR of the Z_i's, or abort if $\geq \eta$ fraction of $PExt_{seed}$ aborted.

An explicit construction. We further construct the first such PRE, as shown in Fig. 2, that needs only a single classical source and makes no independence assumptions. The source can be known completely to the adversary (and only has entropy to the devices). It can be arbitrarily correlated with the untrusted devices, with an almost optimal translation of the degree of correlation into the quality parameter for the output [8]. Our extractor, for the first time, circumvents any form of input-structural or *independence* assumptions underlying all existing solutions[9]. In conjunction with [30], our extractor is able to extract *arbitrarily long* randomness from untrusted devices using *any* weak source with *constant bits min-entropy with respect to the devices*. It is also robust against a constant level of device imprecision, a critical property for practical implementations. Given enough number of devices, the output error of our protocol can be made close to the minimum. Given a desirable output error ϵ, the number of devices can be made a polynomial in $1/\epsilon$.

The assumptions for our extractor to work form a minimal set in the following sense [10]. First, a single min-entropy source alone (i.e. without any additional resource) is insufficient due to the impossibility of deterministic extraction [40]. Untrusted devices alone (i.e., without any min-entropy source) are not sufficient either, because the devices can then pre-program their deterministic answers without generating any randomness. Without any communication restriction

between the adversary and the devices, our task would become impossible trivially. Such a restriction between the computational components of an extractor and the adversary is also implicitly assumed for classical extractors. If the devices can communicate freely, there would be effectively a single device. Then this single device's optimal strategy for minimizing the abort probability can be made deterministic [11]. Together with the extractor's deterministic algorithm, we would then have a deterministic extractor, which cannot even extract a single bit from a general min-entropy source. We note that it would be useful for practical considerations to relax the no-communication restriction. On the other hand, results assuming no-communication can be useful for those settings as well (e.g. the output min-entropy reduces by the amount among the devices and the adversary.) While in principle, there may be other incomparable minimal set of assumptions allowing for randomness extraction, we successfully remove assumptions required by all current methods: structural restrictions on the input, some forms of independence and the trust on the inner-working of the quantum device(s).

Our construction needs two existing ingredients: A *quantum-proof classical randomness extractor* Ext and a seeded PRE $\mathrm{PExt}_{\mathrm{seed}}$. Intuitively, Ext uses a globally uniform seed to transform a min-entropy source into an output that is close-to-uniform with respect to an all-powerful quantum adversary (see Section 5 for a formal definition.) Our main theorem provides a "Master Protocol" for constructing physical randomness extractors from *any* pair of Ext and $\mathrm{PExt}_{\mathrm{seed}}$ with matching input/output length. Our main technical contribution is a general principle for proving security when composing multiple untrusted-device protocols.

We introduce a few technical concepts in order to state our main theorem concretely. The *error of* Ext is the worst-case, over all (n, k) sources, standard distance (trace distance) of the input-output-adversary state to the ideal state, where the output is uniformly distributed with no correlation with the input-adversary subsystem. The *noise* in an untrusted device describe its deviation from the performance of an "ideal" device. We define noise fairly generally so that our result is applicable in a wide range of settings. One specific example is when performing a Bell-test, such as in the well-known CHSH game [18], the noise can be defined to be the gap between the device's success probability with that of the optimal quantum success probability. The *error of an untrusted-device protocol* is the maximum of two types of errors: the *completeness* error and the *soundness* error. The completeness error under a fixed level of noise is the probability of the protocol rejecting an implementation where the device(s) used are within the specified noise level to the ideal device(s). The soundness error quantifies the chance of accepting an undesirable output.

Theorem 2.1 (Main Theorem (Informal)) *Let* $(\mathrm{Ext}, \mathrm{PExt}_{\mathrm{seed}})$ *be a pair of quantum-proof classical randomness extractor and seeded PRE such that the output length of* Ext *is the same as the input length of* $\mathrm{PExt}_{\mathrm{seed}}$. *Suppose that* ϵ *upper-bounds both the errors of* Ext *on any* (n, k) *source and* $\mathrm{PExt}_{\mathrm{seed}}$ *for a certain noise level. Then the composition of multiple instances of* Ext *and* $\mathrm{PExt}_{\mathrm{seed}}$ *shown in Fig. 2 with* $\eta = \sqrt{\epsilon}$ *is a seedless PRE whose error for the same noise level and on an* (n, k) *source is* $O(\sqrt{\epsilon})$. *Furthermore, the source can be known to the adversary and the min-entropy required is with respect to the devices only.*

The last property of allowing the source to be public means that the randomness is extracted from the untrusted devices, and may have significant practical implication. This is because it allows one to use a reputable public service, such as NIST's Randomness Beacon [32], when the users are sufficiently confident that the device makers have little knowledge of the public randomness.

Different choices of Ext and $\mathrm{PExt}_{\mathrm{seed}}$ give different instantiations of the Master Protocol with different advantages. We highlight the following three instantiations. (1) *The weakest sources.* Using any (n, k) sources of (sufficiently large) constants n and k, we can achieve a constant extraction

rate with a constant error for an unbounded output length. (2) *Minimizing error.* Given sufficiently many devices, our method can reach an error $2^{-\Omega(k^\mu)}$, where $\mu \geq 1/2$ is a universal constant. (3) *High min-entropy sources.* For a polynomial entropy rate (i.e., $k \geq n^\alpha$ for $\alpha \in (0, 1)$), we can extract from poly(n) untrusted devices with an inverse polynomial error (i.e., $k^{-\beta}$ for $\beta \in (0, 1)$) in poly(n) time. The Miller-Shi expansion protocol [30] is the strongest known $\text{PExt}_{\text{seed}}$ in many aspects thus is used to achieve robustness and unbounded extraction. For Ext, (1) uses (repeatedly) a one-bit extractor [27], (2,3) use Trevisan's extractors [41, 23] (as in Corollary 5.3 and 5.6 of De *et al.* [23], respectively).

We sketch the proof for the Main Theorem here. A foundation for all known untrusted-device protocols is to test the super-classical behavior of the devices using the classical source. The main challenge for our seedless extraction is to perform such a test with only a given amount of min-entropy to the devices, without any structural or independence assumptions. Our solution is in essence a reduction of seedless extraction to the syntactically easier task of seeded extraction. We first improve the input randomness *locally*. By the property of the quantum-proof strong extractor Ext, the output $S_{00\ldots0} \cdots S_{10\ldots0} \cdots S_{11\ldots1}$ of the Ext instances forms a "quantum somewhere randomness (QSR)" source, in that most of the blocks S_i are almost uniform to the devices. Call such a block "good." Next, each good S_i is transformed by the corresponding PExt to be near uniform to the adversary. This transformation *decouples* the correlation between a good S_i with the rest of the blocks, ensuring the near-perfect randomness of the final output.

Equivalence Lemma: a principle for proving composition security. Note that for "decoupling" to be meaningful, the source is in general only (close to) uniform-to-device but may be arbitrarily correlated otherwise. Thus the decoupling feature of the seeded extractors does not follow directly from their definition or the original proof for their security [43, 30], which require a globally uniform input. Our main technical contribution is the following "Equivalence Lemma" that bridges the gap in the input requirements in full generality.

Lemma 2.2 (Equivalence Lemma (informal)) *The performance of a seeded physical randomness extractor remains the same when its uniform-to-all input is replaced by a uniform-to-device input.*

As a basic principle for securely composing untrusted-device protocols, Equivalence Lemma has found other applications. We describe two most striking applications (besides the main result).

The first is on *unbounded randomness expansion*, i.e., seeded extraction where then output length does not depend on the input length. Whether or not one could expand randomness securely beyond the exponential rate first shown by Vazrani and Vidick [43] was a natural question [12]. Intuitively, unbounded expansion is possible because the untrusted devices are randomness-generating. Indeed, for any N, repeating an expansion protocol $O(\log^* N)$ times using a different set of devices each time expands a seed of a fixed length N output bits. A folklore method for achieving unbounded expansion using a constant number of devices is to cross-feed two expansion protocols, i.e., using the output of one as the input to the other. Through an intricate analysis, Coudron and Yuen [22] showed that a specific cross-feeding protocol is indeed secure.

The Equivalence Lemma immediately implies that the cross-feeding protocol using *any* two expansion protocols is secure. This is simply because for each expansion protocol, the input is always (almost) uniformly random to its devices (thus the output is always (almost) uniformly random to the next set of devices.) This in particular implies that using the robust expansion protocol of Miller and Shi [30] gives a *robust* unbounded expansion protocol. The protocol analyzed in [22] requires that an honest implementation must tend to an ideal implementation as the output

length grows (thus if the output length is chosen after the device is given, either the device has to be perfect or the output length cannot be unbounded.) [2]

The second significant implication is that *public* randomness an be used to produce private randomness, as long as the public randomness is uniform to the untrusted devices. This implication holds for both random number generation and key distribution. A specific scenario that this implication can be of significant practical value is the following. The NIST Randomness Beacon project [32] aims to broadcast true randomness to the public. Since the bits become known after broadcast, one cannot use them directly for cryptographic applications. However, as long as one is willing to assume that the public randomness is uniform with respect to the untrusted-devices, it can be used securely to generate private randomness. A related yet subtly different application is that in adapting the Miller-Shi randomness expansion protocol [30] for key distribution, the Lemma allows the use of locally generated uniform randomness as the initial seed, despite the original expansion protocol requiring global randomness.

Physics Implications. Our result implies that unless the world is deterministic, we can in principle create arbitrarily many events and be confident that their joint distribution is close to uniform. This rules out a "weak randomness world," where randomness exists in Nature but not in a large and close-to-uniform scale [13]. The previous such dichotomy statements [21, 24] model weak randomness in Nature by the highly structured and highly random SV-sources [4], together with a structural or conditional independence assumption. We remove those assumptions. Our dichotomy statement is asymptotically optimal in the sense that it requires only a constant, as opposed to a linear, amount of uncertainty for certifying unbounded output randomness.

Our result also provides a practical and the strongest known approach for mitigating the "freedom-of-choice" loophole in Bell test experiments for refuting hidden local variable theories. Such experiments require the choice of the measurement settings to be nearly uniformly distributed. By using the output of our protocol, those experiments remain sound even when only extremely weak source of randomness is available. We can thus consider the composition of the protocol with the subsequent Bell tests as a combined test for refuting local hidden variable theory that, unlike the standard Bell tests, needs only a weak random source for choosing the experiment settings.

We stress that the above interpretations are made under the assumption that quantum mechanics is maximally informative for the adversary. That is, the adversary cannot obtain any information other through quantum operations.

3 Preliminaries

We assume familiarity with the standard concepts from quantum information and summarize our notation as follows.

Quantum States. We only consider finite dimensional Hilbert spaces as quantum states in infinite dimensions can be truncated to be within a finite dimensional space with an arbitrarily small error. The state space $\mathcal{A}$ of m-qubit is the complex Euclidean space $\mathbb{C}^{2^m}$. An m-qubit quantum state is represented by a density operator ρ, i.e., a positive semidefinite operator over $\mathcal{A}$ with trace 1. The set of all quantum states in $\mathcal{A}$ is denoted by $\mathrm{Dens}\,(\mathcal{A})$.

The Hilbert-Schmidt inner product on $\mathrm{L}\,(\mathcal{A})$, the operator space of $\mathcal{A}$, is defined by $\langle X, Y \rangle = \mathrm{tr}(X^*Y)$, for all $X, Y \in \mathrm{L}\,(\mathcal{A})$, where $*$ is the adjoint operator. Let $\mathsf{id}_{\mathcal{X}}$ denote the identity operator

---

[2]An earlier version of this work [17] containing the Equivalence Lemma and [22] were independent, though we did not state this application there.

6

over $\mathcal{X}$, which might be omitted from the subscript if it is clear in the context. An operator $U \in \mathrm{L}(\mathcal{X})$ is a unitary if $UU^* = U^*U = \mathrm{id}_\mathcal{X}$. The set unitary operations over $\mathcal{X}$ is denoted by $U(\mathcal{X})$.

For a multi-partite state, e.g. $\rho_{ABE} \in \mathrm{Dens}(\mathcal{A} \otimes \mathcal{B} \otimes \mathcal{E})$, its reduced state on some subsystem(s) is represented by the same state with the corresponding subscript(s). For example, the reduced state on $\mathcal{A}$ system of ρ_{ABE} is $\rho_A = \mathrm{tr}_{\mathcal{BE}}(\rho_{ABE})$, and $\rho_{AB} = \mathrm{tr}_{\mathcal{E}}(\rho_{ABE})$. When all subscript letters are omitted, the notation represents the original state (e.g., $\rho = \rho_{ABE}$).

A classical-quantum-, or cq-state $\rho \in \mathrm{Dens}(\mathcal{A} \otimes \mathcal{B})$ indicates that the $\mathcal{A}$ subsystem is classical and $\mathcal{B}$ is quantum. Likewise for ccq-, etc., states.

We use $|\psi\rangle$ to denote the density operator (i.e., $|\psi\rangle\langle\psi|$) for a pure state $|\psi\rangle$ when it is clear from the context. Use $\mathcal{U}_A$ to denote the completely mixed state on a space $\mathcal{A}$, i.e., $\mathcal{U}_A = \frac{1}{\dim(\mathcal{A})}\mathrm{id}_\mathcal{A}$.

Norms. For any $X \in \mathrm{L}(\mathcal{A})$ with singular values $\sigma_1, \cdots, \sigma_d$, where $d = \dim(\mathcal{A})$, the trace norm of $\mathcal{A}$ is $\|X\|_{\mathrm{tr}} = \sum_{i=1}^d \sigma_i$. The trace distance between two quantum states ρ_0 and ρ_1 is $\|\rho_0 - \rho_1\|_{\mathrm{tr}}$. Their *fidelity*, denoted by $\mathrm{F}(\rho_0, \rho_1))$, is

$$\mathrm{F}(\rho_0, \rho_1) = \|\sqrt{\rho_0}\sqrt{\rho_1}\|_{\mathrm{tr}}. \tag{3.1}$$

The trace distance and the fidelity satisfy the following relations.

Lemma 3.1 (Fuchs-van de Graaf) *For any $\rho_0, \rho_1 \in \mathrm{Dens}(\mathcal{A})$, we have*

$$1 - \frac{1}{2}\|\rho_0 - \rho_1\|_{\mathrm{tr}} \leq \mathrm{F}(\rho_0, \rho_1) \leq \sqrt{1 - \frac{1}{4}\|\rho_0 - \rho_1\|_{\mathrm{tr}}^2}. \tag{3.2}$$

The fidelity between subsystems of quantum states cancan be preserved in the following sense.

Lemma 3.2 (Folklore) *Let $\rho, \xi \in \mathrm{Dens}(\mathcal{A})$ and $\rho' \in \mathrm{Dens}(\mathcal{A} \otimes \mathcal{B})$ be density operators with $\mathrm{tr}_\mathcal{B}\,\rho' = \rho$. There exists a density operator $\xi' \in \mathrm{Dens}(\mathcal{A} \otimes \mathcal{B})$ with $\mathrm{tr}_\mathcal{B}\,\xi' = \xi$ and $\mathrm{F}(\rho', \xi') = \mathrm{F}(\rho, \xi)$.*

Quantum Operations. Let $\mathcal{X}$ and $\mathcal{Y}$ be state spaces. A *super-operator* from $\mathcal{X}$ to $\mathcal{Y}$ is a linear map

$$\Psi : \mathrm{L}(\mathcal{X}) \to \mathrm{L}(\mathcal{Y}). \tag{3.3}$$

Physically realizable *quantum operations* are represented by *admissible* super-operators, which are completely positive and trace-preserving. Thus any quantum protocol can be viewed as an admissible super-operator. We shall use this abstraction in our analysis and make use of the following observation.

Fact 3.3 (Monotonicity of trace distances) *For any admissible super-operator $\Psi : \mathrm{L}(\mathcal{X}) \to \mathrm{L}(\mathcal{Y})$ and $\rho_0, \rho_1 \in \mathrm{Dens}(\mathcal{X})$, we have*

$$\|\Psi(\rho_0) - \Psi(\rho_1)\|_{\mathrm{tr}} \leq \|\rho_0 - \rho_1\|_{\mathrm{tr}}. \tag{3.4}$$

A unitary operation $U \in U(\mathcal{X})$ is a special type of admissible quantum operations that are *invertible*. For any unitary U, its corresponding super-operator Ψ_U is defined as

$$\Psi_U(\cdot) = U \cdot U^\dagger. \tag{3.5}$$

Let $\{|i\rangle : 1 \leq i \leq \dim(\mathcal{X})\}$ be the computational basis for $\mathcal{X}$. An $\mathcal{X}$-controlled unitary on $\mathcal{Y}$ is a unitary $U \in U(\mathcal{X} \otimes \mathcal{Y})$ such that for some $U_i \in U(\mathcal{Y})$, $1 \leq i \leq \dim(\mathcal{X})$,

$$U = \sum_{1 \leq i \leq \dim(\mathcal{X})} |i\rangle\langle i| \otimes U_{x_i}. \tag{3.6}$$

Likewise define an $\mathcal{X}$-*controlled* admissible quantum operation T from $\mathcal{K}$ to $\mathbf{L}$ as an admissible quantum operation such that for some admission quantum operations $T_i : L(\mathcal{K}) \to L(\mathbf{L})$, $1 \le i \le \dim(\mathcal{X})$

$$T = \sum_{1 \le i \le \dim(\mathcal{X})} \langle i | \cdot | i \rangle | i \rangle \langle i | \otimes T_i(\cdot). \tag{3.7}$$

Min-entropy. For a cq state ρ_{XE}, the amount of *extractable* randomness (from X against E) is characterized by its *(smooth) conditional min-entropy*.

Definition 3.4 (conditional min-entropy) *Let* $\rho_{XE} \in \mathrm{Dens}(\mathcal{X} \otimes \mathcal{E})$. *The* min-entropy *of* X *conditioned on* E *is defined as*

$$\mathrm{H}_\infty(X|E)_\rho \overset{def}{=} \max\{\lambda \in \mathbb{R} : \exists \sigma_E \in \mathrm{Dens}(\mathcal{E}), \text{s.t. } 2^{-\lambda}\mathsf{id}_X \otimes \sigma_E \ge \rho_{XE}\}.$$

4 Formal definitions of Physical Randomness Extractors

We now proceed with formal definitions. We first formalize the notion of physical systems. By an input or output, we mean a finite length binary string.

A *quantum device D* is a Hilbert space, also denoted by D, together with an admissible quantum operation, called its *device operation*, which takes a classical input, conditioned on which applies a quantum operation on D, then produces a classical output. A *physical system* $\mathcal{S} = (X, D, E)$ consists of three disjoint subsystems: a source X, which is always classical, t quantum devices $D = (D_1, \cdots, D_t)$, for some $t \ge 0$, and a quantum adversary E. We write $\mathcal{S} = \mathcal{S}(\rho, \{A_{D_i}\})$ to denote that the device operations are $\{A_{D_i}\}$ and the system state is currently ρ. Likewise for writing $\mathcal{S} = \mathcal{S}(\rho)$. Note that the assumption of no-communication among the devices is formally captured by that each device algorithm A_{D_i} operates only on its corresponding space D_i.

As randomness is relative, we will say that in a multi-partite state, a certain classical component has a certain min-entropy *(with respect) to another component*. Similarly we add a scope of subsystems to which a certain classical component is an (n, k)-source or uniformly distributed. If the scope is the rest of the system, we refer to it as "global." With those conventions, we quantify the min-entropy of a physical system below.

A physical system $\mathcal{S}(\rho, \{A_{D_i}\})$ is an (n, k, t, m)-physical system with a random-to-devices source if (1) X is an (n, k)-source to the devices, and (2) Each device D_i can only output at most m bits in total or any additional bit of output will not be used.

By replacing "random-to-devices" with *globally random*, *uniform-to-devices*, and *globally uniform*, we define the corresponding physical systems similarly. For the latter two cases, we omit the min-entropy to call $\mathcal{S}$ an (n, t, m)-physical system. We now define the syntax of physical randomness extractors.

Definition 4.1 (Physical Randomness Extractor) *A physical randomness extractor* PExt *for a physical system* $\mathcal{S}(X, D, E)$ *is a classical deterministic algorithm that conditioned on* X, *classically interacts with the devices by invoking the device operations, and finally outputs a decision bit* $A \in \{0, 1\}$, *where* 0 *is for rejecting and* 1 *for accepting, and an output string* $Z \in \{0, 1\}^*$ *to the corresponding registers* $A \otimes Z$. *(See Fig. 2.)*

The extraction operation

$$\Phi_{\mathrm{PExt}} : \mathrm{L}(X \otimes D) \to \mathrm{L}(A \otimes Z \otimes X \otimes D) \tag{4.1}$$

is the X-*controlled admissible operation from* D *to* $A \otimes Z \otimes D$ *induced by the composition of* PExt *and the device operations.*

When discussing post-extraction states, it will be convenient to say that $\mathcal{S}$ is equipped with the registers $A \otimes Z$, and denote the extended physical system by $\mathcal{S}_{AZ}$. Denote by $A(\text{PExt}, \mathcal{S})$ the event that PExt accepts when applied to $\mathcal{S}$.

In order to discuss the quality of a PRE, we need the following relative notions of (approximate) uniform distribution. For an $\epsilon \in [0, 2]$, we say that X is ϵ-uniform-to-E in a cqq state $\rho_{XEE'}$ if there exists a $\rho'_{XEE'}$ where X is uniform-to-E and $\|\rho_{XEE} - \rho'_{XEE}\|_{\text{tr}} \leq \epsilon$. Let $\mathcal{S}_{AZ}$ be a physical system equipped with the decision-output registers $A \otimes Z$, and it is in a state $\gamma = \gamma_{AZXDE}$. Denote by $\gamma^A = \gamma^A_{AZXDE}$ the (subnormalized) projection of γ to the $A = 1$ subspace.

We will discuss noise model abstractly, i.e. independent of the technology implementing the devices.

Definition 4.2 (Implementation) *An* implementation *of devices* $D = (D_1, \cdots, D_t)$ *is a device state* ρ_D *together with the device operations* $\{A_{D_i}\}_{1 \leq i \leq t}$.

For a physical system $\mathcal{S}$ over a set D of devices, denote by $\Pi(\mathcal{S})$ its implementation. If D' is a subset of D, denote by the restriction of Π to D' by $\Pi_{D'}$.

Recall that a *premetric* on a set A is a function $\delta : A \times A \to \mathbb{R}$ such that $\delta(a, a') \geq 0$ and $\delta(a, a) = 0$, for all $a, a' \in A$. We require a noise model to be reasonable in that the noise of a larger system is no less than the noise in a smaller system.

Definition 4.3 (Noise Model) *A* noise model *is a premetric on implementations that takes values in* $[0, 1]$ *and is non-increasing under taking device restrictions.*

More precisely, let δ be a noise model, $t \geq 0$ be an integer, D be a set of devices, and Π and Π' be two implementations of D. Then (1) $\delta(\Pi, \Pi') \in [0, 1]$, (2) $\delta(\Pi, \Pi') = 0$ if $\Pi = \Pi'$, and (3) (*Reasonable Property*) [14]. If D' is a subset of D, then $\delta(\Pi_{D'}, \Pi'_{D'}) \leq \delta(\Pi, \Pi')$.

To define the soundness error of a PRE, we need to define that of a post-extraction state. Let γ be a post-extraction state described above. A subnormalized state $\alpha = \alpha_{AZXDE}$ is called *ideal* if $A = 1$ (i.e. $\alpha = |1\rangle\langle 1|_A \otimes \alpha_{ZXDE}$) and Z is uniform to XE. We say that γ has a soundness error ϵ if there exists an ideal post-extraction state α such that $\left\| \gamma^A - \alpha \right\|_{\text{tr}} \leq \epsilon$. We are now ready to define properties of physical randomness extractors.

Definition 4.4 (Soundness, Completeness, and Robustness of a PRE) *Let* $\mathbb{S}$ *be a non-empty set of physical systems,* δ *a noise model,* $\eta \in [0, 1]$, *and* PExt *a PRE. We say that* PExt *is an untrusted-device PRE for* $\mathbb{S}$ *and has a completeness error* ϵ_c *tolerating an* η *level of noise, and a soundness error* ϵ_s, *if the following completeness and soundness properties hold.*

- *(Completeness) There exists an implementation* Π^*, *referred to as the* ideal *implementation, in the implementations of* $\mathbb{S}$, *such that for all* $\mathcal{S} \in \mathbb{S}$ *whose implementation* Π *satisfies* $\delta(\Pi, \Pi^*) \leq \eta$, $\mathbf{Pr}[A(\text{PExt}, \mathcal{S})] \geq 1 - \epsilon_c$.

- *(Soundness) For any* $\mathcal{S}(\rho) \in \mathbb{S}$, $\Phi_{\text{PExt}}(\rho)$ *has a soundness error* $\leq \epsilon_s$.

We further call PExt *a* random-to-devices (n, k, t, m)-PRE, *for integers* $n, k, t, m \geq 0$, *if* $\mathbb{S}$ *is the set of all* (n, k, t, m)-*physical sources with a random-to-devices source. Likewise define the notions of a* random-to-all (n, k, t, m)-PRE, *a* seeded (n, t, m)-*PRE with a* uniform-to-devices *seed, and a* seeded (n, t, m)-*PRE with a* uniform-to-all *seed. If N is the (maximum) output length of* PExt, *the* (extraction) rate *of* PExt *is* $N/(mt)$.

Note that our soundness definition requires the output to be uniform with respect to both the source X and the adversary E, which implies that the randomness PExt extracts is from the devices D.

Previous works (e.g. [43]) define ϵ to be a soundness error if either the protocol accepts with $\leq \epsilon$ probability or the state conditioned on accepting has the desired amount of randomness. While our definition is essentially equivalent, syntactically it has several new and subtle features that greatly simplify the analysis of PRE compositions. First, a single inequality for the definition avoids the otherwise necessary argument about conditional property (conditioned on accepting). Second, use the whole state, as opposed to tracing out the device component, in calculating the distance to an ideal state (that has a uniform X). Third, the ideal state in comparison does not need to have the same Adversary subsystem. Those features allow the application of triangle inequality to the case of perturbed input state and consequently, composed protocols.

Previous randomness expansion protocols seen as seeded-PREs. By definitions, randomness expansion protocols [19, 20, 43, 30] are precisely seeded PREs with uniform-to-all seeds. There has been a large body of research on randomness expansion protocols since [19]. Our framework allows deeper quantitative analyses and comparisons of their performances.

Phrased in our framework, Vazirani-Vidick [43] showed that a quantum-secure 2-device PRE needs only a poly-logarithmic seed length (measured against the output length) and can achieve an inverse polynomial extraction rate and an inverse polynomial error (in the output length). The concurrent work of Miller-Shi [30] significantly improved the rate to be *linear*, and the error to be negligible (inverse quasi-polynomial), besides adding the constant-noise robustness feature. Another concurrent work of Coudron and Yuen [22] reduces the seed length to a *constant* at an inverse polynomial rate using 8 devices.Finally, combining our technique (Equivalence Lemma in the next section) with Miller-Shi [30] in a straightforward matter, we can achieve simultaneously a linear rate, a constant noise robustness and a constant seed length.

5   Results

In this section, we will present the precise statements of the Equivalence Lemma and the Main Theorem with the explicit construction of our main protocol and the necessary tools for completing our analysis. We leave all the proofs in the Appendix.

Our analysis uses the seeded PRE as a black-box. This differs from previous analyses that rely on the details of the untrusted-device protocols (e.g., those for randomness expansion such as [43, 30] or for other tasks, such as delegation of quantum computation [36] and certifying strong monogamy [29, 17]).

We fix a noise model, and refer to the triple $(\epsilon_c, \eta, \epsilon_s)$ of a completeness error ϵ_c, a noise level tolerated η, and a soundness error ϵ_s as the *performance parameters*. We shall prove the lemma below in Section **??**. We point out that while the proof is short, the result may appear surprising or even counter intuitive.

Lemma 5.1 (Equivalence Lemma) *Any seeded PRE for uniform-to-all seeds is also a seeded PRE for uniform-to-devices seeds with the same performance parameters under the same noise model.*

Our proof consists of two steps, the first of which is general and the second is specific for the setup of the Lemma. For a set of system states $\mathbb{S}$, denote by $\mathbb{S}'$ the set of system states that can be obtained from a $\mathcal{S}(\rho_{XDE}) \in \mathbb{S}$ by applying a X-controlled operation on E.

Proposition 5.2 *Let $\mathbb{S}$ be a set of physical systems and* PExt *a PRE. Then* PExt *has the same performance parameters on $\mathbb{S}$ and $\mathbb{S}'$.*

Proof [: Lemma **??**] Since $\mathbb{S} \subseteq \mathbb{S}'$, we need only to show that the performance parameters of $\mathbb{S}'$ are no worse than those of $\mathbb{S}$. Fix a $\mathcal{S}'(\rho'_{XDE'}) \in \mathbb{S}'$. Let M be an X-controlled E-to-E' operation and $\mathcal{S}(\rho_{XDE}) \in \mathbb{S}$ be such that $\rho' = M(\rho)$.

Completeness and robustness. Use the same ideal implementation Π for $\mathbb{S}$ as the ideal implement for $\mathbb{S}'$. Assume that $\delta(\Pi(\mathcal{S}'), \Pi) \leq \eta$. Since $\Pi(\mathcal{S}) = \Pi(\mathcal{S}')$, we have $\delta(\Pi(\mathcal{S}), \Pi) \leq \eta$. Thus $A(\mathcal{S}) \geq 1 - \epsilon_c$. Since the acceptance probability depends only on the reduced density operator on XD, and $\rho'_{XD} = \rho_{XD}$. We have $A(\mathcal{S}') \geq 1 - \epsilon_c$. This proves the claimed result on completeness and robustness.

Soundness. Note that Φ_{PExt} commute with M as both are X-controlled operations acting on two disjoint quantum subsystems. That is,

$$\Phi_{\mathrm{PExt}}(\rho') = M(\Phi_{\mathrm{PExt}}(\rho)). \tag{5.1}$$

Furthermore,

$$\Phi_{\mathrm{PExt}}(\rho')^A = M\left(\Phi_{\mathrm{PExt}}(\rho)^A\right). \tag{5.2}$$

By the soundness of PExt on $\mathbb{S}$, for an ideal post-extraction state δ,

$$\left\|\Phi_{\mathrm{PExt}}(\rho)^A - \delta\right\|_{\mathrm{tr}} \leq \epsilon_s. \tag{5.3}$$

Thus

$$\left\|\Phi_{\mathrm{PExt}}(\rho')^A - M(\delta)\right\|_{\mathrm{tr}} = \left\|M\left(\Phi_{\mathrm{PExt}}(\rho)^A\right) - M(\delta)\right\|_{\mathrm{tr}} \leq \epsilon_s. \tag{5.4}$$

Since $M(\delta)$ also has $A = 1$ and Y uniform to XE, $\Phi_{\mathrm{PExt}}(\rho')$ has also an ϵ_s soundness error. This proves the soundness claim. ∎

The Equivalence Lemma follows from the above together with the following proposition.

Proposition 5.3 *Let D be a set of devices. If $\mathbb{S}_D$ is the set of all global-uniform physical systems over D, $\mathbb{S}'_D$ is the set of all device-uniform physical systems over D.*

Proof. We shall omit the subscript D in this proof. Clearly all states in $\mathbb{S}'$ are device-uniform. We need only to show that for an arbitrary device uniform state $\rho' = \rho'_{XDE}$, there exists a global-uniform $\rho = \rho_{XDE'}$ and an X-controlled operation M on E' such that $\rho' = M(\rho)$.

We write $\rho' = \sum_{1 \leq x \leq \dim(X)} |x\rangle\langle x| \otimes \rho'^x_{DE}$. Let E' be a system of the same dimension as DE, and for each x, let $|\phi^x_{DEE'}\rangle$ be a purification of ρ'^x_{DE}. That is, with $\phi^x = |\phi^x\rangle\langle\phi^x|$, $\phi^x_{DE} = \rho'^x_{DE}$. By the assumption that $\rho'_{XD} = U_X \otimes \rho'_D$, ϕ^x_D are identical for all x. Set $|\phi\rangle = |\phi^{0^n}\rangle$. Thus by Uhlmann's Theorem (c.f. Chapter 9 of [31]), for each x, there exists a unitary operators U_x, such that

$$|\phi^x\rangle = U^x|\phi\rangle. \tag{5.5}$$

Define M to be the X-controlled E-operation that is the composition of applying U^x, as controlled by X, then tracing out E'. Then with $\rho = \sum_x |x\rangle\langle x| \otimes \phi \in \mathbb{S}$,

$$\rho' = M(\rho). \tag{5.6}$$

Therefore, $\rho' \in \mathbb{S}'$. ∎

Quantum-proof strong randomness extractors and Somewhere Random Source We will first review quantum-proof randomness extractors, which turn a min-entropy source to a quantum-secure output, with the help of a short seed. Then we will introduce the somewhere random sources in our protocol construction.

Definition 5.4 (Quantum-proof Strong Randomness Extractor) *A function* $\text{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* quantum-proof *(or simply* quantum*)* (k,ϵ)-*strong randomness extractor, if for all cq states* ρ_{XE} *with* $\text{H}_\infty(X|E) \geq k$, *and for a uniform seed* Y *independent of* ρ_{XE}, *we have*

$$\left\| \rho_{\text{Ext}(X,Y)YE} - \mathcal{U}_m \otimes \rho_Y \otimes \rho_E \right\|_{\text{tr}} \leq \epsilon. \tag{5.7}$$

It is known that Trevisan's extractors [41] are secure against quantum adversaries [23]. Those will be used for instantiating our main theorem. An apparent problem when one tries to apply those extractors in our setting is that we do not have the required uniform seed. Our solution is to enumerate all the possible seed values and run the extractors on the fixed seed values. The output property of the extractor now translates to a guarantee that the output of at least one instance (in fact, a large fraction of them) of the fixed-seeded extractors is close to uniform. The output together forms what we call *quantum somewhere randomness*. In classical setting, a somewhere random source $\mathbf{S}$ is simply a sequence of random variables $\mathbf{S} = (S_1, \ldots, S_r)$ such that the marginal distribution of some block S_i is uniform (but there can be arbitrary correlation among them). Somewhere random sources are useful intermediate objects for several constructions of randomness extractors (see, e.g., [35, 28]), but to the best of our knowledge, its quantum analogue has not been considered before.

Definition 5.5 (Quantum-SR Source) *A cq-state* $\rho \in \text{Dens}\,(S_1 \otimes \cdots \otimes S_r \otimes E)$ *with classical* $S_1,\,S_2,\,\cdots,\,S_r \in \{0,1\}^m$ *and quantum* E *is a* (r,m)-*quantum somewhere random (SR) source against* E *if there exists* $i \in [t]$ *such that*

$$\rho_{S_i E} = \mathcal{U}_m \otimes \rho_E. \tag{5.8}$$

We say that ρ *is a* (r,m,ϵ)-*quantum somewhere random source if there exists* $i \in [t]$ *such that*

$$\left\| \rho_{S_i E} - \mathcal{U}_m \otimes \rho_E \right\|_{\text{tr}} \leq \epsilon. \tag{5.9}$$

We remark that the fact that ρ is a (r,m,ϵ)-quantum somewhere random source does not necessarily imply that ρ is ϵ-close in trace distance to some (r,m)-quantum somewhere random source ρ'. In contrast, the analogous statement is true for classical somewhere random source. However, by Lemma 3.2, one can show that they are $2\sqrt{\epsilon}$ close. On the other hand, just like its classical counterpart, one can convert a weak source X to a somewhere random source by applying a (quantum-proof) strong randomness extractor to X with all possible seeds (Each seed yields one block).

Proposition 5.6 *Let* $\text{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a* quantum-proof (k,ϵ)-*strong extractor. Let* ρ_{XE} *be a cq-state with* $\text{H}_\infty(X|E) \geq k$. *For every* $i \in \{0,1\}^n$, *let* $S_i = \text{Ext}(X,i)$. *Then the cq-state*

$$\rho_{S_1 \ldots S_{2^d} E} \overset{def}{=} \sum_x p_x \, |S_1\rangle\langle S_1| \otimes \cdots \otimes |S_{2^d}\rangle\langle S_{2^d}| \otimes \rho_E^x, \tag{5.10}$$

is a $(2^d, m, \epsilon)$-*quantum SR source. Moreover, the expectation of* $\left\| \rho_{S_i E} - \mathcal{U}_m \otimes \rho_E \right\|_{\text{tr}}$ *over a uniform random index* $i \in \{0,1\}^n$ *is at most* ϵ.

Physical Randomness Extractor PExt

1. Let $\text{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ be a quantum-proof strong randomness extractor.
2. Let $\text{PExt}_{\text{seed}}$ be a seeded PRE with seed length m that uses t_{seed} devices. Let $0 < \eta < 1$.
3. PExt operates on an input source X over $\{0,1\}^n$ and $t_{\text{PRE}} = 2^d \cdot t_{\text{seed}}$ devices $D = (D_1, \ldots, D_{2^d})$, where each D_i denotes a set of t_{seed} devices, as follows.

---

1. For every $i \in \{0,1\}^d$, let $S_i = \text{Ext}(X, i)$ and invoke $(A_i, Z_i) \leftarrow \text{PExt}_{\text{seed}}(S_i, D_i)$.

2. If there exist η fraction of $A_i = 0$, then PExt outputs $A = 0$; otherwise, PExt outputs $(A, Z) = (1, \bigoplus_{i \in [2^d]} Z_i)$.

Figure 3: Our Main Construction of Physical Randomness Extractor PExt.

Proof [: Proposition 5.6] Since Ext is a quantum-proof (k, ϵ)-strong extractor and $H_\infty(X|E) \geq k$, we have that

$$\left\| \rho_{\text{Ext}(X,Y)YE} - \mathcal{U}_m \otimes \rho_Y \otimes \rho_E \right\|_{\text{tr}} \leq \epsilon, \tag{5.11}$$

which is equivalent to

$$\sum_{i=1}^{2^d} \frac{1}{2^d} \left\| \rho_{\text{Ext}(X,i)E} - \mathcal{U}_m \otimes \rho_E \right\|_{\text{tr}} \leq \epsilon. \tag{5.12}$$

Thus immediately we have that there exists an index $i \in [2^d]$ such that

$$\left\| \rho_{\text{Ext}(X,s_i)E} - \mathcal{U}_m \otimes \rho_E \right\|_{\text{tr}} \leq \epsilon, \tag{5.13}$$

or equivalently $\left\| \rho_{S_i E} - \mathcal{U}_m \otimes \rho_E \right\|_{\text{tr}} \leq \epsilon$. ∎

Construction of PREs for any min-entropy source We are now able to state precisely our main theorem. The Master Protocol is described in details in Fig. 3. The proof is presented in Appendix **??**.

Theorem 5.7 (Main Theorem) *Let* $\text{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a quantum* $(k, \epsilon_{\text{Ext}})$-*strong randomness extractor,* $\text{PExt}_{\text{seed}}$ *be a* (m, t_{seed}, l)-*seeded-PRE with a uniform-to-all seed of length* m, *a completeness error* ϵ_c *tolerating an* η *level of noise, and a soundness error* ϵ_s. *Let* $0 < \eta < 1$ *be the rejection threshold. Then* PExt *(as shown in Fig. 3) is a* $(n, k, 2^d t_{\text{seed}}, l)$-*PRE for random-to-device sources with a completeness error* $(\epsilon_c + \epsilon_{\text{Ext}})/\eta$ *and a soundness error* $\epsilon_s + 2\sqrt{\epsilon_{\text{Ext}}} + \eta$.

Proof. We first introduce some notations. We will use i, $i = 1, \cdots, 2^d$ as a subscript to index an instance of Ext or $\text{PExt}_{\text{seed}}$ corresponding to the seed i for Ext. Fix a physical system $\mathcal{S}(\rho_{XDE})$ with $H_\infty(X|D) \geq k$. We continue to write ρ_{XSDE} to denote the state after applying the Ext instances. We write $\bar{A}$ for $1 - A$ and similarly for $\bar{A}_i$ for each i. For each i, define

$$w_i = \left\| \rho_{S_i DE} - \mathcal{U}_{S_i} \otimes \rho_{DE} \right\|_{\text{tr}}. \tag{5.14}$$

By the strong extracting property of Ext, for a uniformly chosen i,

$$E_i[w_i] \leq \epsilon_{\text{Ext}}.$$

13

Also, by Propositions 3.2 and 3.1, there exists a state $\rho'^{(i)}_{XSDE}$ with $\rho'_{S_iDE} = \mathcal{U}_{S_i} \otimes \rho_{DE}$ and

$$\left\| \rho - \rho' \right\|_{\mathrm{tr}} \le 2\sqrt{w_i}.$$

Fix such a $\rho'^{(i)}$ for each i.

Completeness. For each i, fix an ideal implementation $\Pi^*_{D_i}$ for $\mathrm{PExt}_{\mathrm{seed}\,i}$. We will use their tensor product Π^* as the ideal implementation for PExt. Now suppose that $\delta(\Pi(\mathcal{S}), \Pi^*) \le \eta$. By the reasonable property of the noise model, for each i, $\delta(\Pi_{D_i}(\mathcal{S}), \Pi^*_i) \le \eta$. Let $\tilde{\rho} = \tilde{\rho}_{XSDE}$ be an arbitrary state with $\tilde{\rho}_{S_iDE} = U_{S_i} \otimes \rho_{DE}$. Then $\tilde{\rho}$ is uniform-to-D_i and

$$\delta(\Pi_{D_i}(\tilde{\rho}), \Pi^*_{D_i}) = \delta(\Pi_{D_i}(\rho), \Pi^*_{D_i}) \le \eta,$$

by the completeness of $\mathrm{PExt}_{\mathrm{seed}}$,

$$A_i(\tilde{\rho}) \ge 1 - \epsilon_c.$$

Since Φ_{PExt_i} acts on S_iD_i only,

$$\bar{A}_i(\rho) \le \bar{A}_i(\tilde{\rho}) + w_i \le \epsilon_c + w_i.$$

Thus

$$\mathbf{Pr}[\bar{A}] = \mathbf{Pr}[\sum_i \bar{A}_i \ge \eta 2^d] \le (\epsilon_c + E_i[w_i])/\eta \le (\epsilon_c + \epsilon_{\mathrm{Ext}})/\eta.$$

Soundness. The proof is based on the following two observations. The first is that if some Z_i is uniform (to XE), then so is Z. It follows that if for each i, Z_i is uniform in a (subnormalized) γ_i, then the Z obtained from $\sum_i \gamma_i$ is also uniform. Note that not all $\mathrm{PExt}_{\mathrm{seed}\,i}$ accepts when PExt accepts, thus the observation cannot be directly applied. This problem is resolved by an additional insight that for a randomly chosen i, the chance of PExt accepts but $\mathrm{PExt}_{\mathrm{seed}\,i}$ rejects is small. The details follow.

Denote by $\vec{A} = [A_i]_i$ and $\vec{Z} = [Z_i]_i$. Let $\gamma = \gamma_{X\vec{A}\vec{Z}DE}$ be the state after applying all instances of $\mathrm{PExt}_{\mathrm{seed}}$. Let $\gamma'^{(i)}$ be the same except that ρ is replaced by $\rho'^{(i)}$. By the soundness of PExt, there exists an ideal post-extraction (with respect to PExt_i) state γ_i such that

$$\left\| \gamma'^{(i)A_i} - \gamma_i \right\|_{\mathrm{tr}} \le \epsilon_c.$$

Thus

$$\left\| \gamma^{A_i} - \gamma_i \right\|_{\mathrm{tr}} \le \left\| \gamma^{A_i} - \gamma'^{(i)A_i} \right\|_{\mathrm{tr}} + \left\| \gamma'^{(i)A_i} - \gamma_i \right\|_{\mathrm{tr}} \le 2\sqrt{w_i} + \epsilon_c.$$

Applying the final acceptance projection (that is, accept when $< \eta 2^d$ of A_i's reject), we have

$$\left\| \gamma^{A \wedge A_i} - \gamma^A_i \right\|_{\mathrm{tr}} \le 2\sqrt{w_i} + \epsilon_c. \tag{5.15}$$

Note that γ^A_i still has Z_i uniform to XE. Thus with

$$\gamma' = E_i[\gamma^A_i], \tag{5.16}$$

and $Z(\cdot)$ the super-operator for outputting Z, we have that the subnormalized state $Z(\gamma')$ has its Z uniform to XE. We shall show that $Z(\gamma')$ approximates $\Phi_{\mathrm{PExt}}(\rho)^A$ well.

14

$$\left\| \Phi_{\mathrm{PExt}}(\rho)^A - Z(\gamma') \right\|_{\mathrm{tr}} \tag{5.17}$$

$$= \left\| \gamma^A - \gamma' \right\|_{\mathrm{tr}} \tag{5.18}$$

$$= \left\| E_i \left[\gamma^{A \wedge A_i} + \gamma^{A \wedge \bar{A}_i} - \gamma_i^A \right] \right\|_{\mathrm{tr}} \tag{5.19}$$

$$\leq E_i \left[\left\| \gamma^{A \wedge A_i} + \gamma^{A \wedge \bar{A}_i} - \gamma_i^A \right\|_{\mathrm{tr}} \right] \tag{5.20}$$

$$\leq E_i \left[\left\| \gamma^{A \wedge A_i} - \gamma_i^A \right\|_{\mathrm{tr}} \right] + E_i[A \wedge \bar{A}_i] \tag{5.21}$$

$$\leq 2\sqrt{\epsilon_{\mathrm{Ext}}} + \epsilon_c + \eta. \tag{5.22}$$

Eqn. (5.19) is because for any i,

$$\gamma^A = \gamma^{A \wedge A_i} + \gamma^{A \wedge \bar{A}_i}.$$

The last inequality is by Eqn. (5.15) and the acceptance criterion. Thus we conclude that the soundness error of PExt is $\leq 2\sqrt{\epsilon_{\mathrm{Ext}}} + \epsilon_c + \eta$. ∎

Instantiations. The Miller-Shi seeded PRE (and its unbounded expansion composition via "Equivalence Lemma") subsumes all other constructions, thus is preferred to use in our instantiations. Thus the main choice is the quantum-proof classical strong extractors. We use two known methods for constructing such extractors, both based on the work of König and Terhal [27] showing that any classically secure one-bit extractor is automatically secure against quantum adversaries (with slightly worse parameters.) The first method is to take a single-bit extractor and increase the output length by using independent copies of the seeds. The second is to apply Trevisan's compositions of the single-bit extractor, which was proved to be quantum-secure by De *et al.* [23].

The first instantiation uses the first method by setting the error parameter of the single-bit extractor (e.g. in Proposition C.5 of [23]) to be $\Theta(\epsilon / \log^c(1/\epsilon))$, where c is a universal constant from Miller-Shi [30], and the number of independent seeds to be $O(\log^c(1/\epsilon))$. This requires the min-entropy to be $O(\log^c 1/\epsilon)$. The number of devices is $(n/\epsilon)^{O(\log^c(1/\epsilon))}$, thus is efficient for constant ϵ.

The second instantiation uses the Trevisan's extractor in Corollary 5.1 of [23] for Ext. Fix an extractor "seed length index" ν, defined in Miller-Shi [30] as a real so that there exists a quantum-proof strong extractor on $(n, \Theta(n))$ sources, extracting $\Theta(n)$ bits with an error ϵ using a seed length $O(\log^{1/\nu}(n/\epsilon))$. Set the error parameter for Ext to be $\epsilon_{\mathrm{Ext}} = \exp(-k^\nu)$. It extracts $m = k - o(k)$ bits in the quantum somewhere randomness output. For $\mathrm{PExt}_{\mathrm{seed}}$, use Miller-Shi's expansion protocol with a constant q parameter (in their main theorem), an output length $N = k$, and an error $\epsilon_{MS} = 2^{-(ck)^\nu}$, for a constant $c > 0$ to be determined later. The extractor used inside Miller-Shi is the Trevisan's extractor in Corollary 5.4 of [23], which requires a seed length of $O(\log^{1/\nu}(N/\epsilon_{MS}))$. Thus the total randomness needed for Miller-Shi is $O(q \log(1/q) + c)k$, which can be made $\leq k$ by choosing sufficiently small q and c. Our Master Protocol now outputs k bits with an error $\exp(-\Omega(k^\nu))$. Applying the unbounded expansion protocol of Miller-Shi on this output, the final error remains $\exp(-\Omega(k^\nu))$. Note that the total number of devices is dominated by the number of Ext instances, which is $2^{O(\log^2 n + k^{2\nu}) \log k}$.

The third instantiation uses the Trevisan's extractor in Corollary 5.6 of [23], with an inverse polynomial error, and extracting a polynomial fraction of input min-entropy.

6 Future work

If the error is allowed to be a constant, our construction needs only a source of a sufficiently large constant min-entropy and length and the output can be arbitrarily long, using just a constant number of devices. However, for much smaller error, our construction does not achieve simultaneously close to optimal error parameter and efficiency in the number of devices and the running time. In particular, the construction cannot reach a cryptographic level of security as the number of devices is at least inverse polynomial of the error parameter. This raises a fundamental question: *is this high complexity necessary?* A preliminary result of the current authors together with Carl A. Miller shows that the number of devices has to be polynomially related to the input length for any untrusted-device protocol that works on all weak source of a sufficiently small *linear* min-entropy, and the devices are allowed to communicate in between playing two rounds of non-local games. While this does not answer our question when the devices are not allowed to communicate throughout the protocol, it indicates the difficulty of reducing the number of devices in seedless extraction. We do not yet have solid evidence to support a significant reduction in complexity. A strong lower bound (on the number of devices as function of the error parameter) would have a strong (negative) interpretation that we will have to resort to some stronger assumptions than those for our theorem in order to achieve cryptographic level of randomness generation.

Many other new questions arise from our framework of PRE. Is there an ideal PRE, where all parameters are simultaneously optimize? Or perhaps there are inherent tradeoffs. Other questions include, what quantities about the untrusted-device determine the maximum amount of output randomness? Can one quantify the restrictions on communication to shed light on its tradeoff with other parameters? Barrett, Colbeck and Kent [15] pointed out additional potential security pitfalls in composing untrusted-device protocols. An important direction is to develop a security model in which one can design PREs and prove composition security in a broad setting.

References

[1] This is because the uniform distribution is a convex combination of deterministic distributions. Thus any algorithm rejecting all deterministic distributions will also reject the uniform distribution.

[2] Because a convex combination of several independent distributions can be far from independent.

[3] According to quantum mechanics, certain simple measurement is guaranteed to produce perfect randomness.

[4] A Santha-Vazirani (SV) source has uncertainty in each individual bit even conditioned on the value of its previous bits. Namely, a source X over $\{0,1\}^n$ is ϵ-SV if for every $i \in [n]$ and $x_{\leq i} \in \{0,1\}^i$, $\epsilon \leq \Pr[X_i = x_i | X_{<i} = x_{<i}] \leq 1 - \epsilon$.

[5] The argument made here is on the reliance of physical laws in general. Specific theories may be proven not to be the true law of Nature, thus it remains possible that a PRE becomes insecure after a scientific revolution, e.g., one that disproves quantum mechanics or relativity.

[6] Works on those two threads may aim at proving security against an adversary not necessarily bounded by quantum mechanics.

[7] Intuitively a quantum-proof strong extractor uses a uniform seed and transform a min-entropy source to close-to-uniform output, against a quantum adversary. See the Supplementary Material for a formal definition.

[8] If the source is known completely to the devices then no secure extraction is possible and this is reflected in our translation from the correlation to the security guarantee. More precisely, if the probability for the devices to guess the source correctly is 2^{-k}, then the error parameter (to be defined precisely) has to be at least 2^{-k}. The error achieved by our protocol is 2^{-k^c} for some universal constant c.

[9] Note that trusted quantum devices can be used to generate independent random sources readily. Thus, it is a stronger assumption that directly implies independence.

[10] Since all classes of mathematical objects can be further divided into smaller classes according to certain features, a minimal set of requirements is always relevant to a collection of concepts. Consider for example the partition of protocols by whether or not communication among different devices are allowed. One can refine the partition by the amount of communication, or/and by the pairs of devices that communicate. Thus like all claims of minimal assumptions, ours is relative to a certain level of partition of protocols.

[11] First, the device can simulate any quantum process. Thus the minimum abort probability is the value of a classical zero-sum game, for which the device has an optimal and deterministic strategy.

[12] One should note that any super-exponential expansion would necessarily make the output not suitable for cryptographic applications. This is because when the seed length is k, the adversary has $\geq 2^{-k}$ chance of guessing the seed correctly. Thus when the output length is $\exp(\Omega(k))$, this probability would be only inverse polynomial of the output length.

[13] Our protocol does not apply for extremely small min-entropy k, e.g. $k < .1$. If the world is not deterministic, it appears reasonable to assume that the min-entropy should be extremely large. Thus such a small min-entropy scenario may not be physically relevant.

[14] We call this property *reasonable* as it is similar to the corresponding property of *reasonable* crossnorms.

[15] J. Barrett, R. Colbeck, and A. Kent. Memory attacks on device-independent quantum cryptography. *Phys. Rev. Lett.*, 110:010503, Jan 2013.

[16] F. G. Brandão, R. Ramanathan, K. H. Andrzej Grudka, M. Horodecki, and P. Horodecki. Robust device-independent randomness amplification with few devices. *QIP 2014*, arXiv:1310.4544.

[17] K.-M. Chung, X. Wu, and Y. Shi. Physical randomness extractors. *QIP 2014*, arXiv:1402.4797v1, 2014.

[18] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.*, 23:880–884, October 1969.

[19] R. Colbeck. *Quantum And Relativistic Protocols For Secure Multi-Party Computation*. PhD thesis, University of Cambridge, 2006.

[20] R. Colbeck and A. Kent. Private randomness expansion with untrusted devices. *Journal of Physics A: Mathematical and Theoretical*, 44(9):095305, 2011.

[21] R. Colbeck and R. Renner. Free randomness can be amplified. *Nature Physics*, 8:450–453, 2012.

[22] M. Coudron and H. Yuen. Infinite randomness expansion and amplification with a constant number of devices. *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC 2014)*, arXiv:1310.6755.

[23] A. De, C. Portmann, T. Vidick, and R. Renner. Trevisan's extractor in the presence of quantum side information. *SIAM Journal on Computing*, 067(258932), 2012.

[24] R. Gallego, L. Masanes, G. de la Torre, C. Dhara, L. Aolita, and A. Acín. Full randomness from arbitrarily deterministic events. *Nature Communications*, 4:2654, 2013.

[25] Z. Gutterman, B. Pinkas, and T. Reinman. Analysis of the linux random number generator. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, SP '06, pages 371–385, Washington, DC, USA, 2006. IEEE Computer Society.

[26] N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman. Mining your ps and qs: Detection of widespread weak keys in network devices. In *Proceedings of the 21st USENIX Security Symposium*, 2012.

[27] R. T. König and B. M. Terhal. The bounded-storage model in the presence of a quantum adversary. *IEEE Transactions on Information Theory*, 54(2):749–762, 2008.

[28] X. Li. New independent source extractors with exponential improvement. *In Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, 2013.

[29] L. Masanes. Universally composable privacy amplification from causality constraints. *Phys. Rev. Lett.*, 102:140501, Apr 2009.

[30] C. A. Miller and Y. Shi. Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices. Presented in part at the 17th Conference on Quantum Information Processing (QIP 2014) and 46th ACM Symposium on Theory of Computing (STOC 2014), arXiv:1402.0489., 2014.

[31] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, UK, 2000.

[32] NIST. Randomness beacon project. `http://www.nist.gov/itl/csd/ct/nist_beacon.cfm`.

[33] N. Perlroth, J. Larson, and S. Shane. N.S.A. able to foil basic safeguards of privacy on Web. *The New York Times*, September 5, 2013.

[34] S. Pironio, A. Acin, S. Massar, de la Giroday, A. Boyer, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe. Random numbers certified by bell's theorem. *Nature*, 464:1021, Oct. 19 2009. Comment: 10 pages, 3 figures, 16 page appendix. Version as close as possible to the published version following the terms of the journal.

[35] A. Rao. *Randomness Extractors for Independent Sources and Applications*. PhD thesis, The Univeristy of Texas at Austin, 2007.

[36] B. W. Reichardt, F. Unger, and U. Vazirani. Classical command of quantum systems. *Nature*, 496:456–460, April 2013.

[37] R. Renner. Security of quantum key distribution, Jan. 11 2005. Comment: PhD thesis; index added.

[38] R. Renner, S. Wolf, and J. Wullschleger. The single-serving channel capacity. In *Proceedings of the 2006 IEEE International Symposium on Information Theory (ISIT)*, Aug. 02 2006.

[39] T. Ristenpart and S. Yilek. When good randomness goes bad: Virtual machine reset vulnerabilities and hedging deployed cryptography. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2010, San Diego, California, USA, 28th February - 3rd March 2010*. The Internet Society, 2010.

[40] M. Santha and U. V. Vazirani. Generating quasi-random sequences from slightly-random sources. In *Proceedings of the 25th IEEE Symposium on Foundations of Computer Science (FOCS 1984)*, page 434, 1984.

[41] L. Trevisan. Extractors and pseudorandom generators. *Journal of the ACM*, 48(4):860–879, July 2001.

[42] S. Vadhan. The unified theory of pseudorandomness. *SIGACT News*, 38(3):39–54, September 2007.

[43] U. V. Vazirani and T. Vidick. Certifiable quantum dice: or, true random number generation secure against quantum adversaries. In H. J. Karloff and T. Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 61–76. ACM, 2012.

[44] D. Zuckerman. General weak random sources. In *Foundations of Computer Science, 1990. Proceedings., 31st Annual Symposium on*, pages 534–543 vol.2, 1990.