AICR · OPEN PROBLEM

One Haar-Random State Looks Like Many

Classical outputs extracted from t copies of a single Haar-random state should be statistically indistinguishable from outputs extracted from t independent ones

Status. Statement: AI-written from Amit Behera, Zvika Brakerski, Or Sattath and Omri Shmueli, *Pseudorandomness with Proof of Destruction and Applications*, IACR ePrint 2023/543, TCC 2023. Not yet checked by a human. The source states this as a conjecture it believes true, says it could neither prove it nor find it formalised anywhere, and proves a strictly weaker pointwise domination bound (its Lemma 2) that suffices for its own applications.

Category. *Quantum information; quantum cryptography.*

Conjecture (informal). *Hand an algorithm one copy of a random quantum state and let it print a classical string. Do this t times. There are two ways to supply the states: draw t independent Haar-random states, or draw one Haar-random state and hand out t copies of it. In the second, the t printed strings are correlated — they all come from the same underlying state. The conjecture is that no one can tell: for any algorithm and any polynomial t, the two joint distributions of classical strings are within negligible total variation distance. It is an entirely information-theoretic statement, with no computational assumption anywhere, and the source could neither prove it nor find it in the literature.*

1 The setting

Let $\mathcal{H}_n$ be the state space of n qubits, of dimension $N = 2^n$, and let $\mu_{\mathcal{H}_n}$ be the Haar measure on its pure states. Let $\mathcal{A}$ be any algorithm that takes a single n-qubit state and outputs a classical string; it may be an arbitrary quantum channel followed by a measurement, and it may use ancillae, so it need not be a projective measurement on the input register alone.

Two distributions over t-tuples of classical strings arise, and the conjecture compares them.

- The *product* distribution: draw $|\phi_1\rangle,\ldots,|\phi_t\rangle$ independently from $\mu_{\mathcal{H}_n}$ and output $\mathcal{A}(|\phi_1\rangle)\otimes\cdots\otimes\mathcal{A}(|\phi_t\rangle)$.
- The *correlated* distribution: draw a single $|\phi\rangle \sim \mu_{\mathcal{H}_n}$, and output $\mathcal{A}(|\phi\rangle)\otimes\cdots\otimes\mathcal{A}(|\phi\rangle)$, where t independent invocations of $\mathcal{A}$ each receive one copy of the same $|\phi\rangle$.

Why the two might differ is easy to see. If $\mathcal{A}$ measures in the computational basis, the correlated experiment draws t samples from one Porter-Thomas distribution while the product experiment draws each sample from a different one; second moments already separate the two families, and the question is whether the separation is visible in total variation at polynomial t.

2 Notation and parameters

- n is the number of qubits and $N = 2^n$ the dimension; *negligible* is in n.
- $t = t(n)$ is any polynomial; the conjecture is not claimed for superpolynomial t, and cannot hold for t comparable to N.
- $\mathcal{A}$ ranges over *all* algorithms with classical output, not only efficient ones — the statement is information-theoretic.
- TV denotes total variation distance between the two distributions over t-tuples of classical strings.

3 The conjecture

Conjecture 1 ([BBSS23, §1.3]). *For any algorithm $\mathcal{A}$ that outputs a classical string, and any polynomial $t = t(n)$,*

$$\mathrm{TV}\big(\mathcal{A}(|\phi_1\rangle)\otimes\cdots\otimes\mathcal{A}(|\phi_t\rangle),\ \mathcal{A}(|\phi\rangle)\otimes\cdots\otimes\mathcal{A}(|\phi\rangle)\big)$$

is negligible in n, where $|\phi_1\rangle,\ldots,|\phi_t\rangle \sim \mu_{\mathcal{H}_n}$ are independent and the same $|\phi\rangle \sim \mu_{\mathcal{H}_n}$ is used in all the invocations on the right.

Remark 1 (How the source states it). Conjecture 1 is the source's own sentence, page 10: "In some of the applications, we required the following quantum information-theoretic conjecture regarding Haar random states that we believe is true: For any algorithm $\mathcal{A}$ that outputs a classical string, and any polynomial $t = t(n)$, the total variation distance between $\mathcal{A}(|\phi_1\rangle)\otimes\cdots\otimes\mathcal{A}(|\phi_t\rangle)$, where

$|\phi_1\rangle,\ldots,|\phi_t\rangle \sim \mu_{\mathcal{H}_n}$ and $\mathcal{A}(|\phi\rangle)\otimes\cdots\otimes\mathcal{A}(|\phi\rangle)$, where the same state $|\phi\rangle \sim \mu_{\mathcal{H}_n}$ is used in all the algorithms, is negligible in n."

And on its status: "We could not prove nor find any previous work in the literature that proves or even formalizes this conjecture." The notation $\mathcal{A}$, TV and N above is this statement's; the mathematics is the source's, unchanged.

Theorem 1 (What is proved instead, [BBSS23, Lemma 2]). *Let $\mathcal{A}$ output c-bit strings. For every $t \in \mathrm{poly}(\lambda)$ and all $a_1,\ldots,a_t \in \{0,1\}^c$,*

$$\Pr_{\mathrm{corr}}[(f_1,\ldots,f_t)=(a_1,\ldots,a_t)] \;\le\; \frac{N^t}{\binom{N+t-1}{t}}\cdot \Pr_{\mathrm{prod}}[(f_1,\ldots,f_t)=(a_1,\ldots,a_t)],$$

where corr *and* prod *are the correlated and product distributions of §1.*

Remark 2 (Why Theorem 1 is much weaker than Conjecture 1). For $t \ll N$ the factor $N^t/\binom{N+t-1}{t}$ is about t factorial, so the lemma says only that the correlated distribution never exceeds the product distribution by more than a factorial factor, pointwise. That is enough to transfer a bound of the form "this bad event has probability η under the product distribution" into "at most t factorial times η under the correlated one", which is what the source's unforgeability and collision-freeness arguments need. It says nothing at all about total variation distance, which the factorial factor swamps.

The source is explicit that the lemma is a stand-in: "We believe that the distributions are in fact, statistically close due to the strong concentration of the Haar measure, but we have not been able to prove it. The lemma is a weaker version of this statement, but it suffices for our purposes."

The proof of Theorem 1 in the special case where $\mathcal{A}$ is a computational-basis measurement is two lines of symmetric-subspace counting: for each outcome tuple there is a unique symmetric type contributing, giving probability at most $1/\binom{N+t-1}{t}$ under the correlated distribution and exactly N^{-t} under the product one. That argument gives the ratio and, being pointwise, cannot give more.

Remark 3 (Where the difficulty is). Three features make this harder than it first looks.

The algorithm is arbitrary. $\mathcal{A}$ need not be a projective measurement on the input register: it may use ancillae, so the overall map

on the input is not unitary. The source flags exactly this elsewhere — "the destruct algorithms may use ancillae qubits, and therefore the overall process becomes non-unitary" — and non-unitary processing does not preserve Haar-randomness of a residual state, which is what a reduction to the projective case would want.

The output may be long. For a single-bit output, concentration of measure settles it quickly. For c-bit outputs the correlated distribution has a genuinely larger collision probability — of order $2/N$ against $1/N$ for the product distribution — and any proof must show that all such second- and higher-order discrepancies stay below negligible in total variation once summed over the $\binom{t}{2}$ pairs and beyond.

The bound must be uniform in $\mathcal{A}$. The claim quantifies over all algorithms with no efficiency restriction, so it is a statement about the t-copy Haar ensemble itself and not about what a bounded observer can do with it.

Remark 4 (Why it is worth settling on its own). The source's own reason: "We think this is an interesting open question on its own, and if proven, this result can be a useful tool for quantum cryptography." It is the kind of statement that gets reproved ad hoc, in a weaker form each time, wherever a security proof needs to move between one shared Haar-random state and many independent ones. A clean version with an explicit bound in t and n would replace those.

A refutation would be at least as interesting: it would exhibit an algorithm whose classical output distinguishes "one state, t copies" from "t states" at polynomial t, which is a statement about the t-copy Haar ensemble that no one currently expects.

4 Bibliography

[BBSS23] Amit Behera, Zvika Brakerski, Or Sattath and Omri Shmueli. *Pseudorandomness with proof of destruction and applications.* IACR ePrint 2023/543; TCC 2023. The source. The conjecture and its status are on page 10, in the Open Problems section; the correlated and product distributions are Definition 3; the pointwise bound is Lemma 2 with its proof sketch and footnote on page 18, and the general case in Appendix B.

[JLS18] Zhengfeng Ji, Yi-Kai Liu and Fang Song. *Pseudorandom quantum states.* CRYPTO 2018, Part III, LNCS 10993, pages 126–152, Springer, 2018. The paper that started the study of pseudorandom states, and the source of the earlier Haar-indistinguishability conjecture in this line.

[BS20b] Zvika Brakerski and Omri Shmueli. *Scalable pseudorandom quantum states.* CRYPTO 2020, Part II, LNCS 12171, pages 417–440, Springer, 2020. Cited by the source for scalable pseudorandom states; the same authors' work in this line is where statistical closeness to the Haar ensemble is proved for a particular efficiently generated family.

[AQY21] Prabhanjan Ananth, Luowen Qian and Henry Yuen. *Cryptography from pseudorandom quantum states.* arXiv:2112.10020, 2021. The post-selection construction of short-input function-like states from pseudorandom states that the source discusses as not obviously commuting with a destruct algorithm.

[AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian and Henry Yuen. *Pseudorandom (function-like) quantum state generators: new definitions and applications.* arXiv:2211.01444, 2022. Cited by the source for quantum adaptive pseudorandomness, one of the neighbouring open questions in the same list.