quantum pseudo-randomness

based on:
0709.1142
0802.1919 (with Richard Low)
0803.soon (with Matt Hastings)

Aram Harrow
Univ. of Bristol
26 March, 2008

Outline

1. Random unitaries are amazing.
2. We can't produce them.
3. But we can fake them.
4. Now what?

Random unitaries can...

- Create random states.
- Perform random measurements.
- Randomize quantum states (in L_1, L_2 or L_∞)
- Hide data in bipartite states (accessible to global operators but not local operations and classical communication (LOCC))
- Lock accessible information
- Encode (or decode) for pretty much any problem in quantum Shannon theory: [quant-ph/0606225]
  - Sending through [multiple access / broadcast] noisy quantum channels.
  - Entanglement-assisted channel coding.
  - State merging, fully quantum Slepian-Wolf, the quantum reverse Shannon theorem, entanglement distillation, etc....
- Perform remote state preparation / super-dense coding of quantum states
- Create thermal states (if we approximately conserve energy).

Random means

Haar uniform:

i.e. for any integrable function f on U(d) and any $V \in U(d)$,

$$E_{U\sim Haar}\, f(U) = E_{U\sim Haar}\, f(VU)$$

More on this later...

application: state randomization

Fix random elements $U_1, \ldots, U_n$ from $U(d)$.

$$n = \text{const} \cdot \frac{d \log 1/\epsilon}{\epsilon^2}$$ = a little more than d

State randomization map: $$\mathcal{E}(\rho) = \frac{1}{n} \sum_{i=1}^{n} U_i \rho U_i^\dagger$$

Result: $$\left\| \mathcal{E}(\rho) - \frac{I}{d} \right\|_\infty \leq \frac{\epsilon}{d}$$

$$\Longrightarrow \quad \left\| \mathcal{E}(\rho) - \frac{I}{d} \right\|_2 \leq \frac{\epsilon}{\sqrt{d}}$$

$$\Longrightarrow \quad \left\| \mathcal{E}(\rho) - \frac{I}{d} \right\|_1 \leq \epsilon$$

Compare:
d^2 Paulis suffice for exact state randomization.

Hayden, Shor, Leung, Winter. "Randomizing quantum states." quant-ph/0307104
Aubrun. "A remark on the [above] paper." 0802.4193

why this is remarkable

$$n = \text{const} \cdot \frac{d \log 1/\epsilon}{\epsilon^2} \qquad \mathcal{E}(\rho) = \frac{1}{n}\sum_{i=1}^{n} U_i \rho U_i^\dagger \qquad \left\| \mathcal{E}(\rho) - \frac{I}{d} \right\|_\infty \leq \frac{\epsilon}{d}$$

1. ($\mathcal{E} \otimes$ I) destroys LOCC-accessible correlations

Proof: Consider a measurement operator (A$\otimes$B) that is part of a separable measurement. Then ($\mathcal{E}^\dagger \otimes$I)(A$\otimes$B) ≈ (I$\otimes$B) (tr A/d).

2. But ($\mathcal{E} \otimes$ I)(Φ) is far from I/d $\otimes$ I/d.

Proof: ($\mathcal{E} \otimes$ I)(Φ) has rank n, which is ≪ d^2.

3. Data hiding: We can find ≈ d^2/n ≈ almost d orthogonal mixed states on $\mathbb{C}^d \otimes \mathbb{C}^d$ that are LOCC-indistinguishable.

Hayden, Shor, Leung, Winter. "Randomizing quantum states." quant-ph/0307104
Aubrun. "A remark on the [above] paper." 0802.4193

information locking

now take n = poly(log(d)). $\epsilon \gg \log(\log(d)) / \log(d)$

$$\rho^{XKQ} = \frac{1}{dn} \sum_{x=1}^{d} \sum_{k=1}^{n} |x\rangle\langle x|^X \otimes |k\rangle\langle k|^K \otimes (U_k|x\rangle\langle x|U_k^\dagger)^Q$$

English	Math
Q holds information about X that is "locked" by K.	accessible information $I_{acc}(X;Q) \approx \epsilon \log(d)$.
Revealing key K unlocks the information about X.	$I_{acc}(X;KQ) = \log(d)$

Interpretations

Optimistic: exponentially shorter quantum one-time pads!

Pessimistic: accessible information is an unstable security definition.

Non-normative: statement about entropic uncertainty relations.

Hayden, Shor, Leung, Winter. "Randomizing quantum states." quant-ph/0307104

unfortunately

We can't implement Haar-random unitaries on n qubits.

Approximating within ϵ requires $\exp(4^n \log(1/\epsilon))$ different unitaries and so an exponential amount of time and randomness.

(c.f. Shannon 1949 result about how most classical functions require exponential size circuits)

Knill. "Approximation by quantum circuits." quant-ph/9508006

pseudo-random unitaries

k-designs: A distribution μ on U(d) is a unitary k-design if it looks random whenever we take ≤k copies.

Three equivalent definitions:

1. $E_{U\sim\mu}\, U^{\otimes k} \otimes (U^*)^{\otimes k} = E_{U\sim \mathrm{Haar}}\, U^{\otimes k} \otimes (U^*)^{\otimes k}$

2. $E_{U\sim\mu}\, U^{\otimes k} \rho\, (U^\dagger)^{\otimes k} = E_{U\sim \mathrm{Haar}}\, U^{\otimes k} \rho\, (U^\dagger)^{\otimes k}$ for all states ρ

3. When k=2, $E_{U\sim\mu}\, U\, \Lambda(U^\dagger \rho U) U^\dagger = E_{U\sim \mathrm{Haar}}\, U\, \Lambda(U^\dagger \rho U) U^\dagger$ for all channels Λ and all states ρ. (twirling)

approximate k-designs:

$$\left\| \left(\mathbb{E}_{U\sim\mu} U^{\otimes k} \otimes (U^*)^{\otimes k}\right) - \left(\mathbb{E}_{U\sim \mathrm{Haar}} U^{\otimes k} \otimes (U^*)^{\otimes k}\right) \right\|_1 \leq \epsilon$$

Gross, Audenart, Eisert. "...On the structure of unitary designs" quant-ph/0611002

Variants of k-designs

Classical analogue: k-wise independent permutations

μ is a distribution on S_d such that for all distinct $i_1,\ldots,i_k \in \{1,\ldots,d\}$ $(\pi(i_1),\ldots,\pi(i_k))_{\pi\sim\mu}$ is uniform over k-element subsets of $\{1,\ldots,d\}$.

State analogue: state k-designs

μ is a distribution on unit vectors in $\mathbb{C}^d$ such that

$E_{\psi\sim\mu}\, \psi^{\otimes k} = E_{\psi\sim\text{Haar}}\, \psi^{\otimes k}$, where $\psi = |\psi\rangle\langle\psi|$.

Ambainis and Emerson. "Quantum t-designs..." quant-ph/0701126.
Aaronson. "Quantum copy protection." talk at QIP'08

Expanders

Like designs, but weaker and using fewer unitaries.

Gap: $\|(\mathbb{E}_{U\sim\mu} U \otimes U^*) - (\mathbb{E}_{U\sim\text{Haar}} U \otimes U^*)\|_\infty \quad =$

$$\|(\mathbb{E}_{U\sim\mu} U \otimes U^*) - |\Phi\rangle\langle\Phi|\|_\infty \quad \leq \quad \lambda < 1$$

This condition is analogous to the spectral gap property of random walks on classical expander graphs.

Degree: the degree of an expander is the size of the support of μ. Ideally this will be a constant.

Generalization: k-tensor product expanders (k-TPE)

$$\left\|\left(\mathbb{E}_{U\sim\mu} U^{\otimes k} \otimes (U^*)^{\otimes k}\right) - \left(\mathbb{E}_{U\sim\text{Haar}} U^{\otimes k} \otimes (U^*)^{\otimes k}\right)\right\|_\infty \leq \lambda < 1$$

Note: A k-TPE is also a k′-TPE for k′≤k.
An ∞-TPE is an expander on ℂ[U(d)], the group algebra of U(d).

Expanders vs. designs

number of copies	trace distance (L_1)	operator distance (L_∞)
1	approximate 1-design	expander
k	approximate k-design	k-tensor product expander
∞	Haar measure	U(d) expander (or S_n classically)

Also: repeatedly applying an expander yields a design.

k=∞ tensor product expanders

Define $\mathbb{C}[U(d)]$ to be the space of square-integrable functions on U(d). U(d) acts on $\mathbb{C}[U(d)]$ according to $g \cdot f(x) = f(gx)$. $\mathbb{C}[U(d)]$ is a (reducible) representation of U(d) which contains one copy of the trivial irrep (spanned by the uniform distribution) and at least one copy of every other irrep of U(d).

And every irrep of U(d) appears in some $U^{\otimes k} \otimes (U^*)^{\otimes k}$.

Therefore: rapidly mixing on U(d) ⇔ gapped on $\mathbb{C}[U(d)]$ ⇔ ∞-TPE

⇔ $\| E_{U\sim\mu} R(U) \|_\infty \leq \lambda < 1$ for all nontrivial irreps R(U).

Partial converse: If $\{U_1,\ldots,U_m\}$ are a k-TPE with $k \gg N^3/\epsilon$ then $\{U_1,\ldots,U_m\}$ can ϵ-approximate any $V \in U(d)$ with a string of length $O(\log(1/\epsilon))$. (c.f. $O(\log^3(1/\epsilon))$ from Solovay-Kitaev)

Uses of k-designs

- L_1 state randomization makes use of 1-designs, since we want to approximate $\mathbb{E}\, U\rho U^\dagger$.
- Coding / entanglement generation / decoupling / thermalization require a 2-design (details to follow).
- Twirling (used to efficiently estimate how noisy a channel is) requires a 2-design.
- Random measurements require 4-designs to achieve the state identification results of [Sen, quant-ph/0512085].
- Locking and L_∞-state randomization require ???
- Remote state preparation / super-dense coding of quantum states require 2-designs plus ???.

Entanglement generation from 2-designs

Draw bipartite ψ^{AB} from a state 2-design so

$$\mathbb{E}_{\psi\sim\mu}\psi^{A_1B_1}\otimes\psi^{A_2B_2}\approx\mathbb{E}_{\psi\sim\mathrm{Haar}}\psi^{A_1B_1}\otimes\psi^{A_2B_2}$$

Entanglement = $S(\psi^A) = -\mathrm{tr}\,\psi^A \log \psi^A$

$\geq -\log \mathrm{tr}\,(\psi^A)^2 = S_2(\psi^A)$

$$\begin{aligned}\mathbb{E}\,\mathrm{tr}(\psi^A)^2 &= \mathbb{E}\,\mathrm{tr}\,\mathrm{SWAP}^{\mathrm{A_1A_2}}(\psi^{\mathrm{A_1}}\otimes\psi^{\mathrm{A_2}}) \\ &= \mathbb{E}\,\mathrm{tr}(\mathrm{SWAP}^{\mathrm{A_1A_2}}\otimes\mathrm{I}^{\mathrm{B_1B_2}})(\psi^{\mathrm{A_1B_1}}\otimes\psi^{\mathrm{A_2B_2}})\approx\frac{1}{\mathrm{d_A}}+\frac{1}{\mathrm{d_B}}\end{aligned}$$

And by convexity $S(\psi^A) \geq -\log \mathrm{tr}\, E\, (\psi^A)^2 \approx \log(d_A) - O(d_A/d_B)$

Efficient designs

Efficient: On n qubits, run-time should be poly(n).

1-designs:

-Paulis are exact 1-designs. Require 2n random bits.

-Subsets of the Paulis yield approximate 1-designs using $n + O(\log n/\epsilon)$ bits. Use a δ-biased subset of $\{0,1\}^{2n}$ or an approximately 2-universal hash function to choose the Paulis.

Ambainis, Smith. "...derandomizing approximate quantum encryption." quant-ph/0404075
Desrosiers, Dupuis. "Quantum entropic security and approx. q. encryption" 0707.0691

2-designs:

-Cliffords are exact 1-designs. Require $O(n^2)$ random bits.

-Random quantum circuits yield approximate 2-designs using $O(n \log 1/\epsilon)$ bits.

DiVincenzo, Leung, Terhal. "Quantum data hiding" quant-ph/0103098
Dankert, Cleve, Emerson, Livine. "Exact and approximate 2-designs..." quant-ph/0606161
Dahlsten, Oliveira, Plenio. "The emergence of typical entanglement..." quant-ph/0701125
Harrow, Low. "Random circuits are 2-designs" 0802.1919

Efficient expanders

- Random unitaries [Hastings. 0706.0556]
  Optimal gap ($\lambda \approx (\#\text{unitaries})^{-1/2}$) but not efficient.

- Margulis expander. [Gross and Eisert. 0710.0651]
  Set of 8 affine transformations on $Z_N \times Z_N$. $\lambda \leq 2\sqrt{5}/8$.

- zig-zag product [Ben-Aroya, Schartz and Ta-Shma. 0709.0911]
  Iterative construction. Start with an O(1)-dim random expander.

- Cayley graph expanders [Harrow. 0709.1142]
  Apply R(g) for R an irrep and g a generator of a Cayley graph.
  Use the fact that $R \otimes R^*$ contains only one trivial irrep and that
  gapped on $\mathbb{C}[G] \Leftrightarrow \| \mathbb{E}_{g\sim\mu} R'(g) \|_\infty \leq \lambda < 1$ for R′ a nontrivial irrep.

- classical 2-tensor product expanders [Hastings, Harrow. 0803.soon]
  A 2-TPE mixes the $|i\rangle\langle j|$ terms over all $i \neq j$. Then apply a phase.

Open problems

- Efficient constructions of k-TPE's and k-designs.
- Efficient implementations of L_∞ state randomization, information locking and remote state preparation.
- Hamiltonian analogues of random circuits.
- Creating the Gibbs state on a quantum computer. (Finding a quantum Metropolis algorithm.)
- Constructing efficient Ramanujan expanders (meaning they have an optimal relationship between gap and degree). This would improve L_1 state randomization.

application: super-dense coding of quantum states

SDC: share n ebits, send n qubits --> send 2n cbits
SDCQS: --> prepare a 2n qubit state in Bob's lab
??!

caveat: To send $|\psi\rangle$ Alice holds not $|\psi\rangle$ but “ψ” (a classical description). This prevents iterating the protocol and sending an unlimited amount of information.

proof: Start with n ebits and let $|\psi\rangle$ be a 2n-qubit state. If $|\psi\rangle$ is maximally entangled then Alice can locally convert the n ebits to $|\psi\rangle$ and then she can send her half to Bob using n qubits of communication. Since most states are maximally entangled, we can use random unitaries in a clever way to make this work for all states.

Harrow, Hayden, Leung. “Super-dense coding of quantum states” quant-ph/0307221
Abeyesinghe, Hayden, Smith, Winter. “Optimal SDC of entangled states.” quant-ph/0407061