Quantum State Commitments (QSCs)

So far, we've only looked at commitments to classical bits.
Can we commit to quantum states?

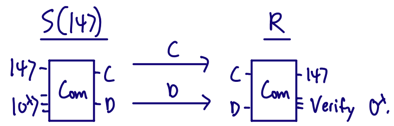

Notation: $|\psi\rangle$ register is M (for "message")
$|0^n\rangle$ register is V (for "verify")

Syntax looks the same as quantum bit commitments (QBCs), except we replaced m with $|\psi\rangle$. But this one change makes a big difference (e.g., if the sender has D and the receiver has C, who has $|\psi\rangle$?)

Not obvious: what does security mean for QSCs?

[GJMZ23]: Binding for QSCs should guarantee that once the sender gives away C, any efficient operation it applies to D either
- makes the opening invalid, or
- is identity on the message space.

Definition: If W is an operator on M (the message space), $\widetilde{W} := \mathrm{Com}(W_M \otimes I_V)\,\mathrm{Com}^\dagger$ denotes a "logical W."

Adv — C, D → **Chal**
1) Verify (C,D). — Measure $\mathrm{Com}(I \otimes |0^n\rangle\langle 0^n|)\,\mathrm{Com}^\dagger$
2) $b \leftarrow \{0,1\}$. Sample $W \leftarrow \{I, X, Z, XZ\}$
Apply $\widetilde{W}^b$

Adv ← D — **Chal**

b'

QSC binding: $\forall$ QPT Adv, $\Pr[b'=b] \le \frac{1}{2} + \mathrm{negl}$

Equivalently, Adv can't detect a random $\widetilde{Z}$ (binding in the standard basis) or a random $\widetilde{X}$ (binding in the Hadamard basis).

Why this definition?
Let A be the adversary's attack:

$$A = \widetilde{I}_M \otimes W^{(I)} + \widetilde{X}_M \otimes W^{(X)} + \widetilde{Z}_M \otimes W^{(Z)} + \widetilde{XZ}_M \otimes W^{(XZ)}$$

If A maps valid commitments to valid commitments, then
- Inability to detect $\widetilde{Z} \Leftrightarrow \widetilde{X}_M \otimes W^{(X)}, \widetilde{XZ}_M \otimes W^{(XZ)}$ are negl (operator norm)
- Inability to detect $\widetilde{X} \Leftrightarrow \widetilde{Z}_M \otimes W^{(Z)}, \widetilde{XZ}_M \otimes W^{(XZ)}$ are negl

So all A can do on M is $\widetilde{I}_M$!

Notice that $X^r Z^s$ completely scrambles any quantum state.
So QSC binding says that D alone hides the committed message.

Thus, hiding and binding are "dual" notions for QSCs.

Hiding for QSCs: C alone hides the committed message.

Adv — C, D → **Chal**
1) Verify (C,D) — Measure $\mathrm{Com}(I \otimes |0^n\rangle\langle 0^n|)\,\mathrm{Com}^\dagger$
2) $b \leftarrow \{0,1\}$. Sample $W \leftarrow \{I, X, Z, XZ\}$
Apply $\widetilde{W}^b$

Adv ← C — **Chal**

b'

Hiding: $\forall$ QPT Adv, $\Pr[b'=b] \le \frac{1}{2} + \mathrm{negl}$

Application: If $\mathrm{Com} \to C, D$ is stat. binding / comp. hiding, we can switch C and D (i.e., send D to commit, C to open) to get a comp. binding / stat. hiding scheme!

Constructing QSCs

① Hiding QSCs:

To commit to $|\psi\rangle$, sample $(a,b) \leftarrow \{0,1\}^2$, send
$$X^a Z^b |\psi\rangle, \underbrace{\mathrm{Com}(a,r), \mathrm{Com}(b,r')}_{\text{Can use QBCs.}}$$
Decommit by sending r, r'.

② Succinct QSCs:

(Succinct = Commitment is smaller than the message)
QSC binding def suggests: encrypt $|\psi\rangle$ with short keys.

Suppose $G: \{0,1\}^{n/2} \to \{0,1\}^{2n}$ is a PRG.
To commit to n-qubit $|\psi\rangle$, prepare
$$\sum_{k \in \{0,1\}^{n/2}} |k\rangle_C\, X^{G_0(k)} Z^{G_1(k)} |\psi\rangle_D$$
and send C ($n/2$ qubits).

Binding: the D part looks maximally mixed by PRG security.
(In fact, weaker-than-OWF assumptions suffice.)

Cryptography from QSCs

Commitments to classical messages enable zero-knowledge proofs and succinct arguments for NP.

Using QSCs, we can generalize these protocols to QMA.

QMA = Classical statements with efficiently verifiable quantum proofs

quantum witness/quantum PCP — $P(|\pi\rangle)$ | QMA statement x — V

1) P commits to a proof $|\pi\rangle$. $\xrightarrow{\text{commit}}$
2) V sends random challenge $\xleftarrow{\text{challenge}}$
3) P opens some of the commitments, revealing part of the proof $\xrightarrow{\text{response}}$

ZK for QMA [BG20, GJMZ23]: Instantiate template with hiding QSCs (and a particular QMA-complete problem).

Quantum Succinct Arguments [CM22, GJMZ23]: Instantiate template with tree of succinct QSCs + quantum PCP.
(Quantum analogue of Kilian's protocol)

Quantum PCP = quantum proof that can be checked by looking at a few qubits.

Quantum PCP Conjecture: These exist for all of QMA.

Proving Security

For classical protocols, we prove security by rewinding the prover to extract the proof π (witness/PCP).

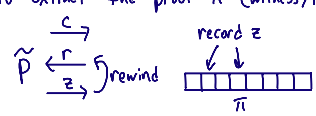

We proved post-quantum security in two steps:

1. If the commitments are collapse-binding, measuring z $\approx_c$ measuring the bit $V(c,r,z)$. (Lectures 2+4)
2. If we only measure 1 bit, we can repair the prover's state before the next query. (Lecture 4)

For QMA protocols, Step 1 no longer suffices.

It's not enough to measure the response if we want to extract a quantum proof $|\pi\rangle$.

Intuitively, we need to swap out the prover's answers onto an external register, which will eventually become the extracted proof $|\pi\rangle$.

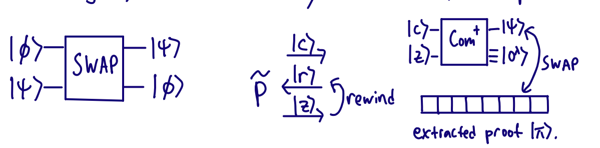

For this to work, this SWAP must be undetectable.

Fortunately, that's exactly what QSC binding guarantees!

QSC binding: Adv can't detect if their committed message is swapped with maximally mixed junk.

While this is the high-level idea, the full proof must handle several other issues that we won't discuss in detail today.

Roughly, the problem is that the [CMSZ21] "repair" from Lecture 4 requires that the reduction:

(a) can verify/open commitments (to estimate success prob.), but
(b) can't break binding of the commitments.

But for QBCs/QSCs, (b) only holds if the reduction doesn't have C, but the reduction needs C for (a).

Note: This isn't an issue when the commitment is classical.

[GJMZ23] resolves this by giving the reduction a verification oracle, but proving that this oracle doesn't compromise security is quite subtle.

Pseudorandom States (PRS)

PRS = Quantum states that appear more random than they really are.
(+ efficiently constructible)

[JLS18]: A t-copy PRS is a set of 2^λ efficiently preparable n-qubit pure states $\{|\psi_k\rangle\}_{k\in\{0,1\}^\lambda}$ s.t. the following are indistinguishable:

- Pseudorandom: Sample $k \leftarrow \{0,1\}^\lambda$ and output $|\psi_k\rangle^{\otimes t}$.
- Haar random: Sample n-qubit Haar random $|\psi\rangle \leftarrow$ Haar, output $|\psi\rangle^{\otimes t}$.

More compactly, $\underset{k\in\{0,1\}^\lambda}{\mathbb{E}} \psi_k^{\otimes t} \approx_c \underset{\psi\leftarrow \text{Haar}}{\mathbb{E}} \psi^{\otimes t}$ (For a pure state $|\psi\rangle$, we'll write $\psi := |\psi\rangle\langle\psi|$)

Recall: An n-qubit Haar-random state $|\psi\rangle$ is a random point on the unit sphere in $\mathbb{C}^{2^n}$. We can sample $|\psi\rangle$ by picking 2^n independent complex Gaussians $\{g_x\}$ and outputting $|\psi\rangle \propto \sum_{x\in\{0,1\}^n} g_x |x\rangle$.

- Single-copy PRS ($t=1$): In this case, $\underset{\psi\leftarrow \text{Haar}}{\mathbb{E}} \psi$ is the maximally mixed state, so Haar-randomness doesn't "matter".
  In fact, quantumness isn't necessary: Any PRG satisfies single-copy PRS security.
- Multi-copy PRS: For any $t = \text{poly}(\lambda)$,
  $$\underset{k\in\{0,1\}^\lambda}{\mathbb{E}} \psi_k^{\otimes t} \approx_c \underset{\psi\leftarrow \text{Haar}}{\mathbb{E}} \psi^{\otimes t}$$
  In particular, construction of ψ_k can't depend on t.

Construction: $|\psi_k\rangle := \sum_{x\in\{0,1\}^n} (-1)^{f_k(x)} |x\rangle$ where $f_k : \{0,1\}^n \to \{0,1\}$ is a post-quantum pseudorandom function (PRF)

In a moment, we'll show that this $|\psi_k\rangle$ is "computationally" Haar random.

What We Know About PRSs

1) When do these primitives require computational hardness?

- For single-copy PRS: $n > \lambda$ requires computational hardness (inefficient attacker can measure $\Pi = \sum_k |\psi_k\rangle\langle\psi_k|$ to distinguish).
  $n = \lambda$ can be unconditionally secure (i.e., trivial). Set $|\psi_k\rangle = |k\rangle$.
- For multi-copy PRS: $n \geq \log \lambda$ requires computational hardness. [AGQY22]
  This can be proved by analyzing the symmetric subspace, which we'll see later today.

2) What kinds of assumptions are needed?

- PRGs imply single-copy, but could even weaker assumptions suffice?
  Probably! [KQST23] (from last lecture) showed that single-copy PRS can exist relative to an oracle that solves any NP problem.
- There's also evidence that multi-copy PRS can exist w/o OWFs, but the evidence is weaker (see [Kretschmer21, KQST23]).

3) What are the applications?

- Single-copy PRS implies QBC/QSCs, and hence suffice for zero-knowledge proofs, oblivious transfer, and multiparty computation [BBCS92, BCKM21, GLSV21, AQY22].
- Multi-copy PRS implies single-copy PRS and much more:
  Ex: - private-key quantum money w/ verification queries [JLS18],
  - one-time secure encryption w/ short keys [AQY22],
  - succinct QSCs / succinct arguments [GJMZ23],

Constructing (Multi-Copy) PRS from OWFs

An n-qubit Haar-random state requires $\sim 2^n$ random bits to specify.
Idea [JLS18]: Use (post-quantum) pseudo-random function (PRF).

Detour: Post-Quantum PRFs

(Classical) PRFs: Efficiently computable family of functions $\{f_k\}$ s.t. for all efficient adversaries Adv,

For today, $f_k : \{0,1\}^n \to \{0,1\}$

oracle access to f_k / f

$$\left| \Pr_{k\in\{0,1\}^\lambda}\left[\text{Adv}^{f_k} \to 1\right] - \Pr_{f \text{ random}}\left[\text{Adv}^{f} \to 1\right] \right| = \text{negl}(\lambda).$$

(f is a uniformly random function)

Post-quantum PRF [Zhandry12]:
QPT Adv can query f in superposition:
$$\sum \alpha_x |x\rangle|b\rangle \xrightarrow{f} \sum \alpha_x |x\rangle|b \oplus f(x)\rangle$$

Theorem [Zhandry12]: The [GGM84] PRF is post-quantum if the underlying OWF is post-quantum.

Warning: Security against superposition queries does not follow generically from classical PRF security + post-quantum assumption.
(e.g. [Z12] shows that you can embed a random period into a PRF to make it insecure against quantum queries, but still secure against classical-query quantum attackers.)

Multi-Copy PRS from Post-Quantum PRFs [JLS18, BS20, AGQY23]:

Key ingredient: binary phase states

For $f:\{0,1\}^n \to \{0,1\}$, let $|\psi_f\rangle := \sum_{x\in\{0,1\}^n} (-1)^{f(x)}|x\rangle$

PRS Construction: Pseudorandom binary phase states, i.e.,

$\{|\psi_{f_k}\rangle\}_{k\in\{0,1\}^n}$ where $\{f_k\}$ is a post-quantum PRF.

Note: $|\psi_{f_k}\rangle$ is efficiently constructible given k.

① Prepare $\left(\sum_{x\in\{0,1\}^n}|x\rangle\right)\otimes|-\rangle$ (Hadamard on $|0^n\rangle|1\rangle$).

② Evaluate $f_k: |x\rangle|b\rangle \to |x\rangle|b\oplus f(x)\rangle$ in superposition:

$$\sum_{x\in\{0,1\}^n}|x\rangle\otimes(|0\rangle-|1\rangle) \xrightarrow{f_k} \sum_x |x\rangle\otimes(-1)^{f_k(x)}(|0\rangle-|1\rangle) = \underbrace{\sum_x(-1)^{f_k(x)}|x\rangle}_{|\psi_{f_k}\rangle}\otimes|-\rangle$$

Discard this.

For multi-copy security, we need to show for any $t=\mathrm{poly}(\lambda)$,

$$\mathop{\mathbb{E}}_{k\leftarrow\{0,1\}^\lambda}\psi_{f_k}^{\otimes t} \approx_c \mathop{\mathbb{E}}_{\psi\leftarrow\mathrm{Haar}}\psi^{\otimes t}$$

Proof overview:

$$\underset{(k\leftarrow\{0,1\}^\lambda)}{|\psi_{f_k}\rangle^{\otimes t}} \overset{\text{PRF security}}{\approx_c} \underset{(f\leftarrow\{0,1\}^{2^n})}{|\psi_f\rangle^{\otimes t}} \overset{\text{statistical}}{\approx} \underset{(|\psi\rangle\leftarrow\mathrm{Haar})}{|\psi\rangle^{\otimes t}}$$

More formally:

1) Switch to a uniformly random binary phase state:

$$\mathop{\mathbb{E}}_{k\leftarrow\{0,1\}^\lambda}\psi_{f_k}^{\otimes t} \approx_c \mathop{\mathbb{E}}_{f\leftarrow\{0,1\}^{2^n}}\psi_f^{\otimes t}$$

If Adv can distinguish these mixed states, we can break PRF security by querying the oracle (either f_k or random f) t times to construct $|\psi_{f_k}\rangle^{\otimes t}$ or $|\psi_f\rangle^{\otimes t}$ and running Adv.

2) Theorem: $\underset{\text{trace distance}}{\mathrm{TD}}\left(\mathop{\mathbb{E}}_{\psi\leftarrow\mathrm{Haar}}\psi^{\otimes t}, \mathop{\mathbb{E}}_{f\leftarrow\{0,1\}^{2^n}}\psi_f^{\otimes t}\right) = \Theta\left(\frac{t^2}{2^n}\right)$

This was conjectured by [JLS18] and later proven by [BS20]. [AGQY23] + [Ma] gave a simpler proof, which we'll see today.

First, we'll need to understand the symmetric subspace.

Detour: Symmetric Subspace

For d = local dimension ($d=2^n$ in our setting), t = # of systems, the symmetric subspace $\mathrm{Sym}^{d,t}$ is the span of all states invariant under permuting the t systems.

Def: $\mathrm{Sym}^{d,t} := \{|\psi\rangle\in(\mathbb{C}^d)^{\otimes t} : \pi|\psi\rangle = |\psi\rangle \ \forall \text{ permutations } \pi\in S_n\}$

(π permutes the t systems)

Claim: $\mathrm{Sym}^{d,t} = \mathrm{span}\{|\psi\rangle^{\otimes t} : |\psi\rangle\in\mathbb{C}^d\}$. (We won't prove this)

Example: $d=3, t=2$, $\mathrm{Sym}^{3,2} := \mathrm{span}\{(\alpha_1|1\rangle+\alpha_2|2\rangle+\alpha_3|3\rangle)^{\otimes 2} : (\alpha_1,\alpha_2,\alpha_3)\in\mathbb{C}^3\}$

$$\mathrm{Sym}^{3,2} = \mathrm{Span}\left\{\begin{matrix}\alpha_1^2\underline{|11\rangle}+\alpha_2^2\underline{|22\rangle}+\alpha_3^2\underline{|33\rangle}+\alpha_1\alpha_2\underline{(|12\rangle+|21\rangle)} \\ +\alpha_1\alpha_3\underline{(|13\rangle+|31\rangle)}+\alpha_2\alpha_3\underline{(|23\rangle+|32\rangle)}\end{matrix} : (\alpha_1,\alpha_2,\alpha_3)\in\mathbb{C}^3\right\}$$

Underlined vectors are a basis for $\mathrm{Sym}^{3,2}$.

For $\vec{v}=(v_1,v_2,\dots,v_t)\in[d]^t$, let $\mathrm{type}(\vec{v})$ be the d-dim vector whose i^{th} entry is the # of times i appears in $(v_1,\dots,v_t)$.

Let $\mathrm{Part}_{d,t} := \{(c_1,\dots,c_d) : \sum_{i\in[d]}c_i = t \text{ and each } c_i\geq 0\}$

for "partition"

Note: By a balls-and-bins argument, $|\mathrm{Part}_{d,t}| = \binom{t+d-1}{t}$.

For any $T\in\mathrm{Part}_{d,t}$, let $|\mathrm{Type}_T\rangle := \sum_{\mathrm{type}(\vec{v})=T}|\vec{v}\rangle$.

Type vectors give an extremely useful characterization of $\mathrm{Sym}_{d,t}$.

Claim: $\mathrm{Sym}_{d,t} = \mathrm{span}\{|\mathrm{Type}_T\rangle : T\in\mathrm{Part}_{d,t}\}$

Claim: $\mathop{\mathbb{E}}_{\psi\leftarrow\mathrm{Haar}}\psi^{\otimes t} = \mathop{\mathbb{E}}_{T\leftarrow\mathrm{Part}_{d,t}}|\mathrm{Type}_T\rangle\langle\mathrm{Type}_T|$

Roughly, $\mathop{\mathbb{E}}_{\psi\leftarrow\mathrm{Haar}}\psi^{\otimes t}$ is an operator that (i) has image $\subseteq\mathrm{Sym}_{d,t}$, and (ii) commutes w/ $U^{\otimes t}$ for any unitary U on $\mathbb{C}^d$.

Representation theory (Schur's lemma) says that the only operators satisfying (i) and (ii) is proportional to identity on $\mathrm{Sym}_{d,t}$.

(For proofs, see "The Church of the Symmetric Subspace" by Harrow.)

Thus, it suffices to show

$$\mathrm{TD}\left(\mathop{E}_{f\in\{0,1\}^{2^n}} \psi_f^{\otimes t},\ \mathop{E}_{T\leftarrow \mathrm{Part}_{2^n,t}} |\mathrm{Type}_T\rangle\langle \mathrm{Type}_T|\right) = \Theta\left(\frac{t^2}{2^n}\right).$$

Define $|\mathcal{X}\rangle := \sum_{f\in\{0,1\}^{2^n}} \left(\sum_{x_1} (-1)^{f(x_1)}|x_1\rangle\right)\otimes\cdots\otimes\left(\sum_{x_t}(-1)^{f(x_t)}|x_t\rangle\right)\otimes|f\rangle_F$

(interpretting f as a function $f:\{0,1\}^n\to\{0,1\}$ or vector $f\in\{0,1\}^{2^n}$)

Observe: $\mathrm{Tr}_F(\mathcal{X}) = \mathop{E}_{f\in\{0,1\}^{2^n}} \psi_f^{\otimes t}$

Write $|\mathcal{X}\rangle$ with the phase on the F part:

$$|\mathcal{X}\rangle = \sum_{x_1,\dots,x_t} |x_1,\dots,x_t\rangle \sum_{f\in\{0,1\}^{2^n}} (-1)^{f\cdot(e_{x_1}\oplus e_{x_2}\oplus\cdots\oplus e_{x_t})}|f\rangle_F$$

($f(x) = f\cdot e_x$ where $e_x\in\{0,1\}^{2^n}$ is 1 only at index x)

So $|\mathcal{X}\rangle = \sum_{x_1,\dots,x_t} |x_1,\dots,x_t\rangle\otimes H^{\otimes 2^n}|e_{x_1}\oplus\cdots\oplus e_{x_t}\rangle_F$

Since F gets traced out, we can ignore the $H^{\otimes 2^n}$.

Key Observation: $e_{x_1}\oplus\cdots\oplus e_{x_t} = \underbrace{\mathrm{type}(x_1,\dots,x_t) \bmod 2}_{:=\,\mathrm{bintype}(x_1,\dots,x_t)}$.

Aside: This observation is the basis of "compressed oracles" [Zhandry19], which says that querying a random oracle is equivalent to applying a unitary that records the type of all queries made so far (mod size of oracle range).

Moreover, if all entries of $\vec{x} = (x_1,\dots,x_t)$ are distinct, $\mathrm{type}(\vec{x}) = \mathrm{bintype}(\vec{x})$.

Let $|\mathrm{bintype}_T\rangle := \sum_{\mathrm{bintype}(\vec{x})=T} |\vec{x}\rangle$.

Then $\mathop{E}_{f\in\{0,1\}^{2^n}} \psi^{\otimes t} = \mathrm{Tr}_F(\mathcal{X}) = \sum p_T\, |\mathrm{bintype}_T\rangle\langle \mathrm{bintype}_T|$, where $p_T = \frac{\#\text{ of } \vec{x} : \mathrm{bintype}(\vec{x})=T}{2^{nt}}$

Let $\mathrm{Distinct} := \{T : T\in\{0,1\}^{2^n} \text{ and HammingWeight}(T)=t\}$.

For all $T\in \mathrm{Distinct}$, $|\mathrm{type}_T\rangle = |\mathrm{bintype}_T\rangle$.

To conclude:

(1) A random type $T\leftarrow \mathrm{Part}_{2^n,t}$ is in Distinct w/ prob $\Theta(t^2/2^n)$,

so $\mathrm{TD}\left(\mathop{E}_{\psi\leftarrow \mathrm{Haar}} \psi^{\otimes t},\ \mathop{E}_{T\leftarrow \mathrm{Distinct}} |\mathrm{Type}_T\rangle\langle\mathrm{Type}_T|\right) = \Theta(t^2/2^n)$.

(2) A random $(x_1,\dots,x_t)$ has $\mathrm{Type}(x_1,\dots,x_t)\in \mathrm{Distinct}$ w/ prob $\Theta(t^2/2^n)$.

so $\mathrm{TD}\left(\mathop{E}_{f\in\{0,1\}^{2^n}} \psi_f^{\otimes t},\ \mathop{E}_{T\leftarrow \mathrm{Distinct}} |\mathrm{Type}_T\rangle\langle\mathrm{Type}_T|\right) = \Theta(t^2/2^n)$.

Since TD is transitive, $\mathrm{TD}\left(\mathop{E}_{\psi\leftarrow \mathrm{Haar}} \psi^{\otimes t},\ \mathop{E}_{f\in\{0,1\}^{2^n}} \psi_f^{\otimes t}\right) = \Theta(t^2/2^n)$. Done!

Exercise: Extend this to the case where (± 1) is replaced with random k^{th} roots of unity for any $k\geq 2$.

Pseudoentangled States

Recently, [ABFGVZZ23] proposed pseudo-entangled states, which are states that appear more entangled than they really are.

Construction: $\sum_{x\in S_k}(-1)^{f_k(x)}|x\rangle$ where $S_k\subseteq\{0,1\}^n$ is a pseudorandom subset of size 2^ℓ, where ℓ can be any $\omega(\lg n)$

Observation: The "entanglement entropy" across any cut is $\leq \ell$

(Same construction appears in [BBSS23], but w/ different applications)

Such states can be built from OWFs (but requires some care).

Key technical step: Show that for random $S\subseteq\{0,1\}^n$ of size 2^ℓ ($\ell=\omega(\lg n)$) and random $f: S\to\{0,1\}$, $t=\mathrm{poly}(n)$ copies of $|\psi_{S,f}\rangle := \sum_{x\in S}(-1)^{f(x)}|x\rangle$ is statistically close to $\mathop{E}_{\psi\leftarrow \mathrm{Haar}} \psi^{\otimes t}$.

To prove this, we just need to show the following are stat. close:

- Output t independent random $x_1,\dots,x_t\leftarrow\{0,1\}^n$
- Sample random size $2^{\omega(\lg n)}$-size subset $S\subseteq\{0,1\}^n$, then output t independent random $x_1,\dots,x_t\leftarrow S$.

These are statistically close because the two distributions are identical if we condition on the event that all x_i are distinct, which occurs w/ $1-\mathrm{negl}(n)$ prob in both cases.

Using this fact, the type vector proof above goes through.

One last thing: Why do multi-copy PRSs require computational hardness at output length $\log\lambda$?

Answer: At output length $\log\lambda$, there exists $t=\mathrm{poly}\,\lambda$ s.t. $\mathrm{Sym}_{d=\lambda,t}$ has dimension $\binom{t+\lambda-1}{t} > 2^\lambda$.

Then this gives 2^λ vectors $\{|\psi_k\rangle^{\otimes t}\}_{k\in\{0,1\}^n}$ that look indistinguishable from the maximally mixed state on $\mathrm{Sym}_{\lambda,t}$, which has dimension $>2^\lambda$. This can be distinguished inefficiently (if $\dim > 2\cdot 2^\lambda$, say).

This is also why any non-trivial multi-copy PRS implies single-copy PRS with any desired $\mathrm{poly}(\lambda)$ output length.