3. Bell Sampling and Bell Difference Sampling

We now introduce the key primitive behind the algorithm that we will present for learning stabilizer states. We begin by defining an important measurement basis:

Definition 178 (Bell basis)**.** *Define the **Bell states***

$$\sigma_{00} = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \qquad \sigma_{01} = \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$$
$$\sigma_{10} = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle) \qquad \sigma_{11} = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)\,.$$

*The **Bell basis** consists of all 2^n states on $2n$ qubits which take the form*

$$\sigma_{\vec{s}} = \sigma_{s_1} \otimes \cdots \otimes \sigma_{s_n}\,,$$

where $s_1, \ldots, s_n \in \{00, 01, 10, 11\}$.

*Given an n-qubit state ρ, **Bell sampling** is the operation of performing a measurement of $\rho^{\otimes 2}$ in the Bell basis.*

Definition 179 (Characteristic distribution)**.** *Given a pure state $\rho = |\psi\rangle\langle\psi|$, define the **characteristic distribution** q_ρ to be the distribution over $\vec{s} \in \{0,1\}^{2n}$ with probability mass function*

$$q_\rho(\vec{s}) \triangleq 2^{-n} \langle\psi| P_{\vec{s}} |\psi\rangle^2\,.$$

(Exercise: verify that this is indeed a probability distribution)

Proposition 180. *If $|\psi\rangle$ is stabilizer, then the characteristic distribution is uniform over* $\mathrm{Weyl}(|\psi\rangle)$.

Proof. For $P \in \mathrm{Weyl}(|\psi\rangle)$, we have $2^{-n}\langle\psi| P |\psi\rangle^2 = 2^{-n}$. There are exactly 2^n Paulis in the unsigned stabilizer group, so the claim follows. □

Lemma 181. *When one performs Bell sampling on pure state $\rho = |\psi\rangle\langle\psi|$, one observes measurement outcome $s \in \{0,1\}^{2n}$ with probability $2^{-n}|\langle\psi| P_{\vec{s}} |\psi^*\rangle|^2$, where $|\psi^*\rangle$ denotes the entrywise conjugation of $|\psi\rangle$ in the computational basis.*

Proof. Note that the Bell states satisfy $\sigma_s = \frac{1}{\sqrt{2}}\mathrm{vec}(P_s)$ for all $s \in \{00, 01, 10\}$ and $\sigma_{11} = i \cdot \frac{1}{\sqrt{2}}\mathrm{vec}(P_{11})$, so

$$\sigma_{\vec{s}} = C_{\vec{s}} \cdot \mathrm{vec}(P_{\vec{s}}) \qquad \text{for } C_{\vec{s}} \triangleq 2^{-n/2} i^{\#\{j:(s_{2j}, s_{2j+1}) = (1,1)\}}\,.$$

As $|\psi\rangle|\psi\rangle = \mathrm{vec}(|\psi\rangle\langle\psi^*|)$ – where vec denotes the operation that flattens a matrix into a vector, and $|\psi^*\rangle$ denotes the entrywise conjugation of $|\psi\rangle$ in the computational basis – we have

$$|\langle\sigma_{\vec{s}}|\psi\rangle|\psi\rangle|^2 = 2^{-n}|\mathrm{tr}(P_{\vec{s}}|\psi\rangle\langle\psi^*|)|^2 = 2^{-n}|\langle\psi| P_{\vec{s}} |\psi^*\rangle|^2\,.$$

as claimed. □

Corollary 182. *When one performs Bell sampling on stabilizer state $\rho = |\psi\rangle\langle\psi|$, there is some $\vec{t} \in \{0,1\}^{2n}$ such that the distribution over measurement outcomes places mass $q_\rho(\vec{s} \oplus \vec{t})$ on any string $s \in \{0,1\}^{2n}$.*

PROOF. By Lemma 170, $|\psi\rangle$ takes the form of Eq. (62). As $i^{\ell(x)} = \prod_{j\in S} i^{x_j}$ for some $S \subseteq [n]$, and $\overline{i^a} \cdot |a\rangle = i^a \cdot Z|a\rangle$ for $a \in \{0,1\}$, we find that up to phase, $|\psi^*\rangle$ is given by $Z^{\otimes S}|\psi\rangle$. By Lemma 181, we conclude that under Bell sampling, we observe outcome s with probability

$$2^{-n}|\langle\psi| P_{\vec{s}} Z^{\otimes S}|\psi\rangle|^2 = q_\rho(\vec{s}\oplus\vec{t})$$

for some fixed string t satisfying $P_{\vec{s}}Z^{\otimes S} = P_{\vec{s}\oplus\vec{t}}$, as claimed. □

The upshot is that if one performs Bell sampling on a stabilizer state, one obtains a sample from the characteristic distribution, but with an unknown *shift* $\vec{t}$. By Proposition 180, this is the uniform distribution over strings of the form $\vec{s}\oplus\vec{t}$ where $\vec{s}$ ranges over strings corresponding to elements of Weyl($|\psi\rangle$).

If we could get rid of this shift, we would be able to access uniformly random elements of Weyl($|\psi\rangle$). The key idea is to run Bell sampling *twice*, and then XOR the two samples, which exactly cancels out the shift! This trick, due to [Mon07], has a name:

Definition 183 (Bell difference sampling)**.** ***Bell difference sampling*** *is the procedure of measuring a given state ρ in the Bell basis twice to get measurement outcomes $\vec{s}\oplus\vec{t}$ and $\vec{s'}\oplus\vec{t}$, and then outputting their XOR, namely $\vec{s}\oplus\vec{s'}$. We will denote the distribution over strings (or equivalently, their associated Pauli operators) obtained in this fashion by $\mathcal{B}_\rho$.*

Lemma 184. *If $|\psi\rangle$ is a stabilizer state, then Bell difference sampling results in a random sample from* Weyl($|\psi\rangle$).

PROOF. By design, Bell difference sampling results in a sample of the form $\vec{s}\oplus\vec{s'}$, where both $\vec{s},\vec{s'}$ are uniform over Weyl($|\psi\rangle$). As the uniform measure over a symplectic linear subspace is invariant under translation by any fixed element of that subspace, $\vec{s}\oplus\vec{s'}$ is itself distributed as a uniform sample from Weyl($|\psi\rangle$). □

4. Learning Algorithm

This yields a simple algorithm, due to [Mon07], for learning the stabilizer group of an unknown stabilizer state:

Algorithm 7: LEARNSTABILIZERGROUP($|\psi\rangle$)

Input: Copies of stabilizer state $|\psi\rangle$
Output: Classical description of Clifford circuit preparing $|\psi\rangle$

Perform Bell difference sampling $2n$ times to get strings $\vec{s}^1,\dots,\vec{s}^{2n}$
Compute a basis $\{\vec{s}^{j_1},\dots,\vec{s}^{j_n}\}$ for $\vec{s}^1,\dots,\vec{s}^{2n}$
Let $U \in \mathcal{C}_n$ be the Clifford circuit for which $UP_{\vec{s}^{j_i}}U^{-1} = Z_i$ for all $i \in [n]$
for $i \in [n]$ **do**
    Measure $|\psi\rangle$ in the eigenbasis of $P_{\vec{s}^{j_i}}$ to determine whether $P_{\vec{s}^{j_i}}|\psi\rangle = -|\psi\rangle$ or $P_{\vec{s}^{j_i}}|\psi\rangle = |\psi\rangle$
end
Let Q be the operator which acts as X in all qubits $i \in [n]$ for which $P_{\vec{s}^{j_i}}|\psi\rangle = -|\psi\rangle$
Output $U^\dagger Q$

Theorem 185. *Given access to copies of an unknown stabilizer state $|\psi\rangle$, the algorithm* LearnStabilizerGroup$(|\psi\rangle)$ *performs $O(n)$ two-copy measurements and with probability at least $1-2^{-n}$ outputs the classical description of a Clifford circuit preparing $|\psi\rangle$.*

Proof. The algorithm fails if and only if the strings $\vec{s}^1, \ldots, \vec{s}^{2n}$ are all contained in a subspace of dimension at most $n-1$. There are 2^n such subspaces, so by a union bound, this happens with probability at most $2^{-2n} \cdot 2^n = 2^{-n}$.

The operator Q is meant to correct for the fact that the group generated by $\{P_{\vec{s}^{j_i}}\}_{i\in[n]}$ is only the *unsigned* stabilizer group for $|\psi\rangle$. In particular, while $U^\dagger |0^n\rangle$ is the stabilizer state associated to those Paulis, the true stabilizer state is given by $U^\dagger |x\rangle$ for a string $x \in \{0,1\}^n$ whose 1 entries are precisely those i for which $P_{\vec{s}^{j_i}} |\psi\rangle = -|\psi\rangle$. □

While the algebraic structure in this learning gives rise to a very elegant protocol, it is clear that the guarantee we have proven is incredibly brittle. In particular, even if $|\psi\rangle$ were corrupted by a small amount of noise, it is no longer clear which parts of this lecture can be salvaged, if any. In the next lecture, we will show that remarkably, there is a way to redeem Bell difference sampling even in the presence of such "model misspecification."