A Revisit of Optimistic Online Learning

In this section, we briefly review some important properties in the classical optimistic online learning algorithms. Some of the propositions in this section will be frequently used in the proof of the regret bound.

For convenience, we will use $\psi(\cdot)$ to denote the negative entropy function, i.e., $\psi \colon \Delta_n \to \mathbb{R}$, $\psi(p) = \sum_{i=1}^{n} p_i \log p_i$. Note that log stands for the natural logarithm function with base e.

For a vector norm $\|\cdot\|$, its dual norm is defined as:

$$\|y\|_* = \max_x \{\langle x, y\rangle : \|x\| \leq 1\}.$$

Proposition A.1. *Let L be a vector in n-dimensional space. If $p^* = \arg\min_{p\in\Delta_n} \{\langle p, L\rangle + \psi(p)\}$, then p^* can be written as:*

$$p^* = \frac{\exp(-L)}{\|\exp(-L)\|_1}$$

and vice versa.

Proof. Write $L = (L_1, L_2, \ldots, L_n)$. By definition, we know that p^* is the solution to the following convex optimization problem:

$$\begin{aligned} \underset{p_1, p_2, \ldots, p_n}{\text{minimize}} \quad & \sum_{i=1}^{n} p_i L_i + \sum_{i=1}^{n} p_i \log p_i \\ \text{subject to} \quad & \sum_{i=1}^{n} p_i = 1, \\ & \forall i, p_i \geq 0 \end{aligned}$$

The Lagrangian is

$$\mathcal{L}(p, u, v) = \sum_{i=1}^{n} p_i L_i + \sum_{i=1}^{n} p_i \log p_i - \sum_{i=1}^{n} u_i p_i + v\left(\sum_{i=1}^{n} p_i - 1\right)$$

From KKT conditions, we know that the stationarity is:

$$L_i + 1 + \log p_i - u_i + v = 0. \tag{1}$$

The complementary slackness is:

$$u_i p_i = 0.$$

The primal feasibility is

$$\forall i, p_i \geq 0; \sum_{i=1}^{n} p_i = 1.$$

The dual feasibility is

$$u_i \geq 0.$$

If $u_i \neq 0$ then $p_i = 0$, from stationarity we know $u_i = -\infty$, but that violates the dual feasibility. So we can conclude that $u_i = 0$ for all $i \in [n]$, thus $p_i \propto \exp(-L_i)$ and the result follows. □

Now we present a generalized version of the optimistic multiplicative weight algorithm called optimisitically follow the regularized leader (Opt-FTRL) in Algorithm 3. In the algorithm, m_t has the same meaning as $m(t)$ for notation consistency.

Algorithm 3 Optimistic follow-the-regularized-leader

```
Input: The closed convex domain $\mathcal{X}$.
Output: Step size $\lambda$, loss gradient prediction $m$.
  Initialize $L_0 \leftarrow 0$, choose appropriate $m_1$.
  for $t = 1, \ldots, T$ do
    Choose $x_t = \arg\min_{x\in\mathcal{X}} \{\lambda \langle L_{t-1} + m_t, x\rangle + \psi(x)\}$.
    Observe loss $l_t$, update $L_t = L_{t-1} + l_t$.
    Compute $m_{t+1}$ using observations till now.
  end for
```

Now we study a crucial property that leads to the fast convergence of the algorithm, called the Regret bounded by Variation in Utilities (RVU in short). For simplicity, we only consider the linear loss function $l_t(x) = \langle l_t, x\rangle$. (There is a little abuse of notation here.)

Definition A.2 (Regret bounded by Variations in Utilities (RVU), Definition 3 in Syrgkanis et al. [46])**.** Consider an online learning algorithm $\mathcal{A}$ with regret $\mathcal{R}(T) = o(T)$, we say that it has the property of regret bound by variation in utilities if for any linear loss sequence $l_1, l_2, \dots, l_T$, there exists parameters $\alpha > 0, 0 < \beta \leq \gamma$ such that the algorithm output decisions $x_1, x_2, \dots, x_T, x_{T+1}$ that satisfy:

$$\sum_{i=1}^{T} \langle l_i, x_i\rangle - \min_{x\in\mathcal{X}} \sum_{i=1}^{T} \langle l_i, x\rangle \leq \alpha + \beta \sum_{i=1}^{T-1} \|l_{i+1} - l_i\|_*^2 - \gamma \sum_{i=1}^{T-1} \|x_{i+1} - x_i\|^2 ,$$

where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$.

We do not choose the norm to be any specific one here. In fact, Syrgkanis et al. [46] have already shown that the above optimistic follow-the-regularized-leader algorithm has the RVU property with respect to any norm $\|\cdot\|$ in which the negative entropy function ψ is 1-strongly convex. So, from Pinsker's inequality, for l_2 norms the following result holds:

Proposition A.3 (Proposition 7 in Syrgkanis et al. [46])**.** *If we choose $m_t = l_{t-1}$ in the optimistic follow-the-regularized-leader algorithm with step size $\lambda \leq 1/2$, then it has the regret bound by variation in utilities property with the parameters $\alpha = \log n/\lambda$, $\beta = \lambda$ and $\gamma = 1/(4\lambda)$, where n is the dimension of $\mathcal{X}$.*

B Regret Bound and Time Complexity of Our Algorithm

B.1 Ideal Samplers

We assume that after the execution of our algorithm, the sequences we get are $\{(x_t, y_t)\}_{t=1}^{T+1}$ and $\{(g_t, h_t)\}_{t=1}^{T+1}$, respectively. We denote $u_t := \frac{\exp(-\mathbf{A}h_t)}{\|\exp(-\mathbf{A}h_t)\|_1}$ and $v_t := \frac{\exp(\mathbf{A}^\intercal g_t)}{\|\exp(\mathbf{A}^\intercal g_t)\|_1}$ to be the corresponding Gibbs distribution, we will first assume that the Gibbs oracle in our algorithm has no error (i.e. $\epsilon_G = 0$) until Theorem B.5 is proved.

Observation B.1. *The sequence $\{u_t\}_{t=1}^{T}$ can be seen as the decision result of applying optimistic FTRL algorithm to the linear loss function $\mathbf{A}\eta_t$ with linear prediction function $\mathbf{A}\eta_{t-1}$, and similarly for $\{v_t\}_{t=1}^{T+1}$ with the loss function $-\mathbf{A}^\intercal\zeta_t$, the prediction function $-\mathbf{A}^\intercal\zeta_{t-1}$.*

Proof. By symmetry, we only consider u_t. Since $u_t = \frac{\exp(-\mathbf{A}h_t)}{\|\exp(-\mathbf{A}h_t)\|_1}$, from Proposition A.1 we can write

$$u_t = \arg\min_{u\in\Delta_m} \{\langle \mathbf{A}h_t, u\rangle + \psi(u)\} .$$

Then we notice the iteration of Algorithm 1 gives

$$h_t = \lambda\left(\sum_{i=1}^{t-1} \eta_i\right) + \lambda\eta_{t-1}.$$

So from the definition of the Algorithm 3, we know that our observation holds. □

This observation, together with Proposition A.3, gives the following inequalities. For any $u \in \Delta_m$, $v \in \Delta_n$, we have:

$$\sum_{t=1}^{T} \langle u_t - u, \mathbf{A}\eta_t\rangle \leq \frac{\log m}{\lambda} + \lambda \sum_{t=1}^{T-1} \|\mathbf{A}(\eta_{t+1} - \eta_t)\|^2 - \frac{1}{4\lambda} \sum_{t=1}^{T-1} \|u_{t+1} - u_t\|^2 , \tag{2}$$

$$\sum_{t=1}^{T} \langle v_t - v, -\mathbf{A}^\intercal\zeta_t\rangle \leq \frac{\log n}{\lambda} + \lambda \sum_{t=1}^{T-1} \|\mathbf{A}^\intercal(\zeta_{t+1} - \zeta_t)\|^2 - \frac{1}{4\lambda} \sum_{t=1}^{T-1} \|v_{t+1} - v_t\|^2 . \tag{3}$$

However, we find that the loss function is slightly different from what we expect.

Let us consider the difference $q_t := \mathbf{A}(v_t - \eta_t)$ and $p_t := -\mathbf{A}^\mathsf{T}(u_t - \zeta_t)$, we have the decomposition of the regret:

$$\sum_{t=1}^{T} \langle u_t - u, \mathbf{A}v_t \rangle = \sum_{t=1}^{T} \langle u_t - u, \mathbf{A}\eta_t \rangle + \sum_{t=1}^{T} \langle u_t - u, q_t \rangle .$$

Notice that $\mathbb{E}[q_t] = \mathbb{E}[p_t] = 0$, we have:

Lemma B.2.

$$\mathbb{E}\left[\sum_{t=1}^{T} \langle u_t - u, q_t \rangle\right] = 0, \mathbb{E}\left[\sum_{t=1}^{T} \langle v_t - v, p_t \rangle\right] = 0$$

Proof. By symmetry, we only prove the case for u. It suffices to prove that for every t, $\mathbb{E}\left[\langle u_t - u, q_t \rangle\right] = 0$. Since u is fixed, $\mathbb{E}\left[\langle u, q_t \rangle\right] = \langle u, \mathbb{E}\left[q_t\right]\rangle = 0$.

Now consider $\mathbb{E}\left[\langle u_t, q_t \rangle\right]$, notice that given $\eta_1, \ldots, \eta_{t-1}$ then u_t is a constant. We have:

$$\begin{aligned}
\mathbb{E}\left[\langle u_t, q_t \rangle\right] &= \mathbb{E}\left[\mathbb{E}\left[\langle u_t, q_t \rangle \,|\eta_1, \eta_2, \ldots, \eta_{t-1}\right]\right] \\
&= \mathbb{E}\left[\langle u_t, \mathbb{E}\left[q_t | \eta_1, \eta_2, \ldots, \eta_{t-1}\right]\rangle\right] \\
&= \mathbb{E}\left[\langle u_t, 0 \rangle\right] = 0.
\end{aligned}$$

□

Now we are going to bound the term $\sum_{t=1}^{T-1} \|\mathbf{A}(\eta_{t+1} - \eta_t)\|^2$.

Lemma B.3.

$$\sum_{t=1}^{T-1} \|\mathbf{A}(\eta_{t+1} - \eta_t)\|^2 \leq 6 + 3\sum_{t=1}^{T-1} \|v_{t+1} - v_t\|^2 , \tag{4}$$

$$\sum_{t=1}^{T-1} \|\mathbf{A}^\mathsf{T}(\zeta_{t+1} - \zeta_t)\|^2 \leq 6 + 3\sum_{t=1}^{T-1} \|u_{t+1} - u_t\|^2 . \tag{5}$$

Proof. Recall that by rescaling we have $\|\mathbf{A}\| \leq 1$. Hence,

$$\sum_{t=1}^{T-1} \|\mathbf{A}(\eta_{t+1} - \eta_t)\|^2 \leq \sum_{t=1}^{T-1} \|\eta_{t+1} - \eta_t\|^2 .$$

Write $\eta_{t+1} - \eta_t = (\eta_{t+1} - v_{t+1}) + (v_{t+1} - v_t) + (v_t - \eta_t)$. Using the triangle inequality of the l_1 norm and the Cauchy inequality $(a+b+c)^2 \leq 3(a^2+b^2+c^2)$, we get

$$\sum_{t=1}^{T-1} \|\eta_{t+1} - \eta_t\|^2 \leq 6\sum_{t=1}^{T} \|\eta_t - v_t\|^2 + 3\sum_{t=1}^{T-1} \|v_{t+1} - v_t\|^2 . \tag{6}$$

Similarly, we have:

$$\sum_{t=1}^{T-1} \|\zeta_{t+1} - \zeta_t\|^2 \leq 6\sum_{t=1}^{T} \|\zeta_t - u_t\|^2 + 3\sum_{t=1}^{T-1} \|u_{t+1} - u_t\|^2 . \tag{7}$$

Observing that in our algorithm we collect T independent and identically distributed samples and take their average, we have:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T} \|\zeta_t - u_t\|^2\right] &\leq 1, \\
\mathbb{E}\left[\sum_{t=1}^{T} \|\eta_t - v_t\|^2\right] &\leq 1.
\end{aligned}$$

Combining the result above, we just get the desired equation. □

We also need the following lemma to guarantee that the sum of the regret is always non-negative.

Lemma B.4. *The sum of the regrets of two players in Algorithm 1 is always non-negative. In other words:*

$$\max_{u \in \Delta_m} \max_{v \in \Delta_n} \left(\sum_{t=1}^{T} \langle u_t - u, \mathbf{A} v_t \rangle + \sum_{t=1}^{T} \langle v_t - v, -\mathbf{A}^\intercal u_t \rangle \right) \geq 0.$$

Proof.

$$\begin{aligned}
&\max_{u \in \Delta_m} \max_{v \in \Delta_n} \left(\sum_{t=1}^{T} \langle u_t - u, \mathbf{A} v_t \rangle + \sum_{t=1}^{T} \langle v_t - v, -\mathbf{A}^\intercal u_t \rangle \right) \\
&= \max_{u \in \Delta_m} \max_{v \in \Delta_n} \left(\sum_{t=1}^{T} \langle -u, \mathbf{A} v_t \rangle + \sum_{t=1}^{T} \langle v, \mathbf{A}^\intercal u_t \rangle \right) \\
&= \max_{v \in \Delta_n} \sum_{t=1}^{T} \langle v, \mathbf{A}^\intercal u_t \rangle - \min_{u \in \Delta_m} \sum_{t=1}^{T} \langle u, \mathbf{A} v_t \rangle \geq 0
\end{aligned}$$

The last step is because

$$\max_{v \in \Delta_n} \sum_{t=1}^{T} \langle v, \mathbf{A}^\intercal u_t \rangle \geq \left\langle \mathbf{A} \sum_{t=1}^{T} v_t / T, \sum_{t=1}^{T} u_t \right\rangle,$$

and

$$\min_{u \in \Delta_m} \sum_{t=1}^{T} \langle u, \mathbf{A} v_t \rangle \leq \left\langle \mathbf{A} \sum_{t=1}^{T} v_t, \sum_{t=1}^{T} u_t / T \right\rangle.$$

□

Combining the result above, we finally have the following theorem.

Theorem B.5. *Suppose that in our Algorithm 1, we choose the episode $T = \widetilde{\Theta}(1/\varepsilon)$, and choose a constant learning rate λ that satisfies $\lambda < \sqrt{3}/6$. Then with probability at least $2/3$ the total regret of the algorithm is $\widetilde{O}(1)$. To be more clear, we have:*

$$T \left(\max_{v \in \Delta_n} \langle v, \mathbf{A}^\intercal \hat{u} \rangle - \min_{u \in \Delta_m} \langle u, \mathbf{A} \hat{v} \rangle \right) \leq 36\lambda + \frac{3 \log(mn)}{\lambda},$$

and so our algorithm returns an ε-approximate Nash equilibrium.

Proof. Adding the inequalities (2) and (3) together, we get

$$\begin{aligned}
\sum_{t=1}^{T} \langle u_t - u, \mathbf{A} \eta_t \rangle + \sum_{t=1}^{T} \langle v_t - v, -\mathbf{A}^\intercal \zeta_t \rangle &\leq \frac{\log m}{\lambda} + \frac{\log n}{\lambda} \\
&+ \lambda \sum_{t=1}^{T-1} \|\mathbf{A}(\eta_{t+1} - \eta_t)\|^2 - \frac{1}{4\lambda} \sum_{t=1}^{T-1} \|v_{t+1} - v_t\|^2 \\
&+ \lambda \sum_{t=1}^{T-1} \|\mathbf{A}^\intercal(\zeta_{t+1} - \zeta_t)\|^2 - \frac{1}{4\lambda} \sum_{t=1}^{T-1} \|u_{t+1} - u_t\|^2.
\end{aligned} \tag{8}$$

Taking expectation, and using the inequalities (6) we have

$$\begin{aligned}
&\mathbb{E} \left[\lambda \sum_{t=1}^{T-1} \|\mathbf{A}(\eta_{t+1} - \eta_t)\|^2 - \frac{1}{4\lambda} \sum_{t=1}^{T-1} \|v_{t+1} - v_t\|^2 \right] \\
&\leq \left(3\lambda - \frac{1}{4\lambda} \right) \mathbb{E} \left[\sum_{t=1}^{T-1} \|v_{t+1} - v_t\|^2 \right] + 6\lambda \cdot \mathbb{E} \left[\sum_{t=1}^{T} \|\eta_t - v_t\|^2 \right] \\
&\leq 6\lambda.
\end{aligned}$$

Similarly we can prove

$$\mathbb{E}\left[\lambda\sum_{t=1}^{T-1}\|\mathbf{A}^{\mathsf{T}}(\zeta_{t+1}-\zeta_t)\|^2-\frac{1}{4\lambda}\sum_{t=1}^{T-1}\|u_{t+1}-u_t\|^2\right]\leq 6\lambda.$$

So, taking expectations of Equation (8), and using the above inequalities and the Lemma B.2, we get

$$\mathbb{E}\left[\max_{u\in\Delta_m}\sum_{t=1}^{T}\langle u_t-u,\mathbf{A}v_t\rangle+\max_{v\in\Delta_n}\sum_{t=1}^{T}\langle v_t-v,-\mathbf{A}^{\mathsf{T}}u_t\rangle\right]\leq 12\lambda+\frac{\log(mn)}{\lambda}. \tag{9}$$

Using the fact that

$$\mathbb{E}[\hat{u}]\cdot T=\sum_{t=1}^{T}\mathbb{E}[u_t],$$
$$\mathbb{E}[\hat{v}]\cdot T=\sum_{t=1}^{T}\mathbb{E}[v_t],$$

we have

$$\mathbb{E}\left[\max_{v\in\Delta_n}\langle v,\mathbf{A}^{\mathsf{T}}\hat{u}\rangle-\min_{u\in\Delta_m}\langle u,\mathbf{A}\hat{v}\rangle\right]\cdot T\leq 12\lambda+\frac{\log(mn)}{\lambda}. \tag{10}$$

By Lemma B.4, we know that the regret is always non-negative. So applying Markov's inequality, we know with probability at least $2/3$, the following inequality holds:

$$\max_{v\in\Delta_n}\langle v,\mathbf{A}^{\mathsf{T}}\hat{u}\rangle-\min_{u\in\Delta_m}\langle u,\mathbf{A}\hat{v}\rangle\leq\frac{1}{T}\left(36\lambda+\frac{3\log(mn)}{\lambda}\right).$$

□

B.2 Samplers with Errors

Theorem B.6 (Restatement of Theorem 3.2)**.** *Suppose that in our Algorithm 1, we choose the episode $T=\widetilde{O}(1/\varepsilon)$, and choose a constant learning rate λ that satisfies $0<\lambda<\sqrt{3}/6$. The quantum implementation of the oracle in the algorithm will return T independent and identically distributed samples from a distribution that is ϵ_{G}-close to the desired distribution in total variational distance in quantum time $T_{\mathrm{G}}^{\mathrm{Q}}$.*

Then with probability at least $2/3$ the total regret of the algorithm is $\widetilde{O}(1+\epsilon_{\mathrm{G}}/\varepsilon)$ and the algorithm returns an $\widetilde{O}(\varepsilon+\epsilon_{\mathrm{G}})$-approximate Nash equilibrium in quantum time $\widetilde{O}(T_{\mathrm{G}}^{\mathrm{Q}}/\varepsilon)$.

Proof. We will follow similar steps of proof for Theorem B.5. Since the sampling is not from the ideal distribution, we must bound the terms where η_t and ζ_t take place.

Notice that in this case, we have

$$\|A(v_t-\mathbb{E}[\eta_t])\|\leq\|v_t-\mathbb{E}[\eta_t]\|\leq\epsilon_{\mathrm{G}}.$$

So for the term q_t in Lemma B.2 we now have the bound:

$$\begin{aligned}&\mathbb{E}\left[\sum_{t=1}^{T}\langle u_t-u,A(v_t-\eta_t)\rangle\right]\\&=\mathbb{E}\left[\sum_{t=1}^{T}\langle u_t-u,A(v_t-\mathbb{E}[\eta_t])\rangle\right]+\mathbb{E}\left[\sum_{t=1}^{T}\langle u_t-u,A(\mathbb{E}[\eta_t]-\eta_t)\rangle\right]\\&=\mathbb{E}\left[\sum_{t=1}^{T}\langle u_t-u,A(v_t-\mathbb{E}[\eta_t])\rangle\right]\leq 2T\epsilon_{\mathrm{G}},\end{aligned}$$

where the last step is by Hölder's inequality.

Then for the other term, we have

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T}\|\eta_t - v_t\|^2\right] &\le 2\cdot\mathbb{E}\left[\sum_{t=1}^{T}\|\eta_t - \mathbb{E}[\eta_t]\|^2\right] + 2\cdot\mathbb{E}\left[\sum_{t=1}^{T}\|v_t - \mathbb{E}[\eta_t]\|^2\right]\\ &\le 2 + 2T\epsilon_{\mathrm{G}}^2.\end{aligned}$$

So following the similar steps of proof for Theorem B.5, and using the above bounds, we can get

$$\begin{aligned}&\mathbb{E}\left[\max_{u\in\Delta_m}\sum_{t=1}^{T}\langle u_t - u, \mathbf{A}v_t\rangle + \max_{v\in\Delta_n}\sum_{t=1}^{T}\langle v_t - v, -\mathbf{A}^\intercal u_t\rangle\right]\\ &\le 24\lambda + 24\lambda T\epsilon_{\mathrm{G}}^2 + \frac{\log(mn)}{\lambda} + 4T\epsilon_{\mathrm{G}}.\end{aligned}$$

Again using linearity of expectation and Markov's inequality, we conclude that with probability at least $2/3$

$$T\left(\max_{v\in\Delta_n}\langle v, \mathbf{A}^\intercal\hat{u}\rangle - \min_{u\in\Delta_m}\langle u, \mathbf{A}\hat{v}\rangle\right) \le 72\lambda + \frac{3\log(mn)}{\lambda} + 72T\lambda\epsilon_{\mathrm{G}}^2 + 12T\epsilon_{\mathrm{G}}.$$

□

C Consistent Quantum Amplitude Estimation

Theorem C.1 (Consistent phase estimation, [2, 47])**.** *Suppose U is a unitary operator. For every positive reals ϵ, δ, there is a quantum algorithm (a unitary quantum circuit) $\mathcal{A}$ such that, on input $O(\log(\epsilon^{-1}))$-bit random string s, it holds that*

- *For every eigenvector $|\psi_\theta\rangle$ of U (where $U|\psi_\theta\rangle = \exp(\mathrm{i}\theta)|\psi_\theta\rangle$), with probability $\ge 1-\epsilon$:*
$$\langle\psi_\theta|\langle f(s,\theta)|\mathcal{A}|\psi_\theta\rangle|0\rangle \ge 1-\epsilon;$$
- *$f(s,\theta)$ is a function of s and θ such that $|f(s,\theta)-\theta| < \delta$,*

with time complexity $\widetilde{O}(\delta^{-1})\cdot\mathrm{poly}(\epsilon^{-1})$.

Theorem C.2 (Consistent quantum amplitude estimation)**.** *Suppose U is a unitary operator such that*

$$U|0\rangle_A|0\rangle_B = \sqrt{p}|0\rangle_A|\phi_0\rangle_B + \sqrt{1-p}|1\rangle_A|\phi_1\rangle_B.$$

where $p \in [0,1]$ and $|\phi_0\rangle$ and $|\phi_1\rangle$ are normalized pure quantum states. Then for every positive reals ϵ, δ, there is a quantum algorithm that, on input $O(\log(\epsilon^{-1}))$-bit random string s, outputs $f(s,p) \in [0,1]$ such that

$$\mathbf{Pr}\left[|f(s,p)-p| \le \delta\right] \ge 1-\epsilon,$$

with time complexity $\widetilde{O}(\delta^{-1})\cdot\mathrm{poly}(\epsilon^{-1})$.

Proof. Suppose U is a unitary operator such that

$$U|0\rangle_A|0\rangle_B = \sqrt{p}|0\rangle_A|\phi_0\rangle_B + \sqrt{1-p}|1\rangle_A|\phi_1\rangle_B.$$

Let

$$Q = -U(I - 2|0\rangle_A\langle 0|\otimes|0\rangle_B\langle 0|)U^\dagger(I - 2|0\rangle_A\langle 0|\otimes I_B).$$

Similar to the analysis in Brassard et al. [9], we have

$$U|0\rangle_A|0\rangle_B = \frac{-\mathrm{i}}{\sqrt{2}}\left(\exp(\mathrm{i}\theta_p)|\psi_+\rangle_{AB} - \exp(-\mathrm{i}\theta_p)|\psi_-\rangle_{AB}\right),$$

where $\sin^2(\theta_p) = p$ ($0 \le \theta_p < \pi/2$), and

$$|\psi_\pm\rangle_{AB} = \frac{1}{\sqrt{2}}(|0\rangle_A|\phi_0\rangle_B \pm \mathrm{i}|1\rangle_A|\phi_1\rangle_B).$$

Note that $|\psi_\pm\rangle_{AB}$ are eigenvectors of Q, i.e., $Q|\psi_\pm\rangle_{AB} = \exp\left(\pm \mathrm{i}2\theta_p\right)|\psi_\pm\rangle_{AB}$.

Now applying the algorithm $\mathcal{A}$ of consistent phase estimation of Q by Theorem C.1 on state $U|0\rangle_A|0\rangle_B \otimes |0\rangle_C$ (with an $O\left(\log\left(\epsilon^{-1}\right)\right)$-bit random string s), we obtain

$$\mathcal{A}(U|0\rangle_A|0\rangle_B \otimes |0\rangle_C) = \frac{-\mathrm{i}}{\sqrt{2}}\Big(\exp(\mathrm{i}\theta_p)\mathcal{A}(|\psi_+\rangle_{AB}|0\rangle_C) - \exp(-\mathrm{i}\theta_p)\mathcal{A}(|\psi_-\rangle_{AB}|0\rangle_C)\Big).$$

Since each of $|\psi_\pm\rangle_{AB}$ is an eigenvector of Q, it holds that, with probability $\geq 1-\epsilon$,

$$\langle\psi_\pm|_{AB}\langle f(s, \pm 2\theta_p)|_C\mathcal{A}(|\psi_\pm\rangle_{AB}|0\rangle_C) \geq 1-\epsilon.$$

which implies that $\mathcal{A}(U|0\rangle_A|0\rangle_B \otimes |0\rangle_C)$ is $O\left(\sqrt{\epsilon}\right)$-close to

$$\frac{-\mathrm{i}}{\sqrt{2}}\Big(\exp(\mathrm{i}\theta_p)|\psi_+\rangle_{AB}|f(s, 2\theta_p)\rangle_C - \exp(-\mathrm{i}\theta_p)|\psi_-\rangle_{AB}|f(s, -2\theta_p)\rangle_C\Big)$$

in trace distance, where $\left|f\left(s, \pm 2\theta_p\right) \mp 2\theta_p\right| < \delta$. Measuring register C, we denote the outcome as γ, which will be either $f\left(s, 2\theta_p\right)$ or $f\left(s, -2\theta_p\right)$. Finally, output $\sin^2(\gamma/2)$ as the estimate of p (which is consistent). Since $\sin^2(\cdot)$ is even and 2-Lipschitz, the additive error is bounded by

$$\left|\sin^2\left(\frac{\gamma}{2}\right) - p\right| \leq 2\left|\left|\frac{\gamma}{2}\right| - |\theta_p|\right| < \delta.$$

Note that $\mathcal{A}$ makes $\widetilde{O}\left(\delta^{-1}\right) \cdot \mathrm{poly}\left(\epsilon^{-1}\right)$ queries to Q, thus our consistent amplitude estimation has quantum time complexity $\widetilde{O}\left(\delta^{-1}\right) \cdot \mathrm{poly}\left(\epsilon^{-1}\right)$. □

Theorem C.3 (Error-Reduced Consistent quantum amplitude estimation)**.** *Suppose U is a unitary operator such that*

$$U|0\rangle_A|0\rangle_B = \sqrt{p}|0\rangle_A|\phi_0\rangle_B + \sqrt{1-p}|1\rangle_A|\phi_1\rangle_B.$$

where $p \in [0,1]$ and $|\phi_0\rangle$ and $|\phi_1\rangle$ are normalized pure quantum states. Then for every positive integers r and positive real δ, there is a quantum algorithm that, on input $O(r)$-bit random string s, outputs $f^*(s,p) \in [0,1]$ such that

$$\mathbf{Pr}\left[|f^*(s,p) - p| \leq \delta\right] \geq 1 - O(\exp(-r)),$$

with time complexity $\widetilde{O}\left(\delta^{-1}\right) \cdot \mathrm{poly}(r)$.

Proof. Consider that we divide the input random string s into r strings $s_1, s_2, \ldots, s_r$ of length $O(1)$. For each $i \in [r]$, we use Theorem C.2 with input string s_i and parameter $\epsilon = 1/10$. So we get, for each $i \in [r]$,

$$\mathbf{Pr}\left[|f(s_i,p) - p| \leq \delta\right] \geq \frac{9}{10}.$$

Now we set $f^*(s,p)$ to be the median of the estimations $f(s_i,p)$ for $i \in [r]$. We claim it satisfies the desired property. To show that, we define random variables X_i for $i \in [r]$ as follows:

$$X_i = \begin{cases} 1, & \text{if } |f(s_i,p) - p| \leq \delta, \\ 0, & \text{otherwise.} \end{cases}$$

Noticing $\mathbb{E}\left[\sum_{i=1}^r X_i\right] \geq 9r/10$, and by Chernoff bound, we have:

$$\mathbf{Pr}\left[\sum_{i=1}^r X_i < \frac{r}{2}\right] \leq \exp\left(-\frac{8r}{45}\right).$$

Thus with probability at least $1 - \exp(-8r/45)$, we know that at least half of the estimations fall into the interval $[p-\delta, p+\delta]$, and then $f^*(s,p)$ returns a correct answer. □

D Details and Proofs of Fast Quantum Multi-Gibbs Sampling

We present the detailed version of the fast quantum multi-Gibss sampling. Here, we use the shorthand $\mathcal{O}_p^{\text{Gibbs}} = \mathcal{O}_p^{\text{Gibbs}}(1, 0)$, and it also means the distribution of the sample.

We first define the notion of amplitude-encoding (a unitary operator that encodes a vector in its amplitudes).

Definition D.1 (Amplitude-encoding)**.** A unitary operator V is said to be a β-amplitude-encoding of a vector $u \in \mathbb{R}^n$ with non-negative entries, if

$$\langle 0|_C V |0\rangle_C |i\rangle_A |0\rangle_B = \sqrt{\frac{u_i}{\beta}} |i\rangle_A |\psi_i\rangle_B$$

for all $i \in [n]$.

Then, as shown in Algorithm 4, we can construct a quantum multi-Gibbs sampler for a vector u if an amplitude-encoding of the vector u is given. To complete the proof of Theorem 4.2, we only have to construct an amplitude-encoding of $\mathbf{A}z$ (see Appendix D.2 for details).

```
Algorithm 4 Quantum Multi-Gibbs Sampling implementing $\mathcal{O}_u^{\text{Gibbs}}(k, \epsilon_{\text{G}})$
Input: Sample count $k$, a $\beta$-amplitude-encoding $V$ of vector $u \in \mathbb{R}^n$, polynomial $P_{2\beta} \in \mathbb{R}[x]$ that
    satisfies Lemma 4.1 with parameter $\epsilon_P = k\epsilon_{\text{G}}^2/300n$.
Output: $k$ independent samples $i_1, i_2, \ldots, i_k$.
Obtain $\mathcal{O}_{\tilde{u}} \colon |i\rangle|0\rangle \mapsto |i\rangle|\tilde{u}_i\rangle$ using $\widetilde{O}(\beta)$ queries to $V$, where $u_i \leq \tilde{u}_i \leq u_i + 1$, by consistent
    quantum amplitude estimation (Theorem C.3).
Find the $k$ largest $\tilde{u}_i$'s by quantum $k$-maximum finding (Theorem D.3) and let $S$ be the set of
    their indexes. This can be done with $\widetilde{O}(\sqrt{nk})$ queries to $\mathcal{O}_{\tilde{u}}$.
Compute $\tilde{u}^* = \min_{i \in S} \tilde{u}_i$, and $W = (n-k)\exp(\tilde{u}^*) + \sum_{i \in S} \exp(\tilde{u}_i)$.
for $\ell = 1, \ldots, k$ do
    Prepare the quantum state
        $|u_{\text{guess}}\rangle = \sum_{i \in S} \sqrt{\frac{\exp(\tilde{u}_i)}{W}} |i\rangle + \sum_{i \notin S} \sqrt{\frac{\exp(\tilde{u}^*)}{W}} |i\rangle.$
    Obtain $U_u = (V_{CAB}^\dagger \otimes I_D)(V_{DAB} \otimes I_C)$ being a block-encoding of $\operatorname{diag}(u)/\beta$. Similarly,
        obtain $U_{\tilde{u}}^{\max}$ being a block-encoding of $\operatorname{diag}(\max\{\tilde{u}, \tilde{u}^*\})/2\beta$.
    Obtain $U^-$ being a block-encoding of $\operatorname{diag}(u - \max\{\tilde{u}, \tilde{u}^*\})/4\beta$ by the LCU (Linear-
        Combination-of-Unitaries) technique (Theorem D.6), using $O(1)$ queries to $U_u$ and $U_{\tilde{u}}^{\max}$.
    Obtain $U^{\exp}$ being a block-encoding of $P_{2\beta}(\operatorname{diag}(u - \max\{\tilde{u}, \tilde{u}^*\})/4\beta)$ by the QSVT tech-
        nique (Theorem D.7), using $O(\beta \log(\epsilon_P^{-1}))$ queries to $U^-$.
    Post-select $|\tilde{u}_{\text{post}}\rangle = \langle 0|^{\otimes a} U^{\exp} |u_{\text{guess}}\rangle |0\rangle^{\otimes a}$ by quantum amplitude amplification (Theo-
        rem D.8), and obtain $|\tilde{u}_{\text{Gibbs}}\rangle = |\tilde{u}_{\text{post}}\rangle / \| |\tilde{u}_{\text{post}}\rangle \|$. (Suppose $U^{\exp}$ has $a$ ancilla qubits.)
    Measure $|\tilde{u}_{\text{Gibbs}}\rangle$ in the computational basis and let $i_\ell \in [n]$ be the outcome.
end for
Return $i_1, i_2, \ldots, i_k$.
```

D.1 Useful Theorems

Theorem D.2 (Quantum state preparation, [23, 32])**.** *There is a data structure implemented on QRAM maintaining an array $a_1, a_2, \ldots, a_\ell$ of positive numbers that supports the following operations.*

- *Initialization: For any value c, set $a_i \leftarrow c$ for all $i \in [\ell]$.*
- *Assignment: For any index i and value c, set $a_i \leftarrow c$.*

- *State Preparation: Prepare a quantum state*

$$|a\rangle = \sum_{i\in[\ell]} \sqrt{\frac{a_i}{\|a\|_1}}|i\rangle.$$

Each operation costs $\mathrm{polylog}(\ell)$ *time.*

Theorem D.3 (Quantum k-maximum finding, Theorem 6 of Dürr et al. [19])**.** *Given $k \in [n]$ and quantum oracle $\mathcal{O}_u$ for an array $u_1, u_2, \ldots, u_n$, i.e., for every $i \in [n]$,*

$$\mathcal{O}_u|i\rangle|0\rangle = |i\rangle|u_i\rangle,$$

there is a quantum algorithm that, with probability ≥ 0.99, finds a set $S \subseteq [n]$ of cardinality $|S| = k$ such that $u_i \geq u_j$ for every $i \in S$ and $j \notin S$, using $O\left(\sqrt{nk}\right)$ queries to $\mathcal{O}_u$.

We now recall the definition of block-encoding, a crucial concept in quantum singular value transformation [20], which is used in line 9 to 12 in Algorithm 4.

Definition D.4 (Block-encoding)**.** Suppose A is a linear operator on b qubits, $\alpha, \epsilon \geq 0$ and a is a positive integer. A $(b+a)$-qubit unitary operator U is said to be an (α, ϵ)-block-encoding of A, if

$$\left\|\alpha\langle 0|^{\otimes a} U|0\rangle^{\otimes a} - A\right\|_{\mathrm{op}} \leq \epsilon.$$

Definition D.5 (State Preparation Pair, Definition 28 of Gilyén et al. [20])**.** Let $y \in \mathbb{R}^n$ be a vector, specially in this context the number of coordinates starts from 0. Suppose $\|y\|_1 \leq \beta$. Let ϵ be a positive real. We call a pair of unitaries (P_L, P_R) acting on b qubits a (β, ϵ)-state-preparation pair for y if

$$P_L|0\rangle^{\otimes b} = \sum_{j=0}^{2^b-1} c_j|j\rangle,$$

$$P_R|0\rangle^{\otimes b} = \sum_{j=0}^{2^b-1} d_j|j\rangle,$$

such that:

$$\sum_{j=0}^{m-1} \left|\beta c_j^* d_j - y_j\right| \leq \epsilon$$

and for $j \in [2^b]$, $j \geq m$, we require $c_j^* d_j = 0$.

We now state a theorem about linear combination of unitary operators, introduced by Berry et al. [5] and Childs and Wiebe [15]. The following form is from Gilyén et al. [20]. Again we restrict ourselves to the case of real linear combinations.

Theorem D.6 (Linear Combination of Unitaries, Lemma 29 of Gilyén et al. [20])**.** *Let ϵ be a positive real number and $y \in \mathbb{R}^n$ be a vector as in Definition D.5 with (β, ϵ_1) state preparation pair (P_L, P_R). Let $\{A_j\}_{j=0}^{m-1}$ be a set of linear operators on s qubits, and forall j, we have U_j as an (α, ϵ_2)-block-encoding of A_j acting on $a+s$ qubits. Let*

$$W = \left(\sum_{j=0}^{m-1} |j\rangle\langle j| \otimes U_j\right) + \left(I - \sum_{j=0}^{m-1} |j\rangle\langle j|\right) \otimes I_{a+s},$$

Then we can implement a $(\alpha\beta, \alpha\epsilon_1 + \alpha\beta\epsilon_2)$-block-encoding of $A = \sum_{j=0}^{m-1} y_j A_j$, with one query from $P_L^\dagger$, P_R, and W.

Theorem D.7 (Eigenvalue transformation, Theorem 31 of Gilyén et al. [20])**.** *Suppose U is an (α, ϵ)-block-encoding of an Hermitian operator A. For every $\delta > 0$ and real polynomial $P \in \mathbb{R}[x]$ of degree d such that $|P(x)| \leq 1/2$ for all $x \in [-1, 1]$, there is an efficiently computable quantum circuit $\tilde{U}$, which is a $\left(1, 4d\sqrt{\epsilon/\alpha} + \delta\right)$-block-encoding of $P(A/\alpha)$, using $O(d)$ queries to U.*

Finally, for quantum amplitude amplification without knowing the exact value of the amplitude, we need the following theorem:

Theorem D.8 (Quantum amplitude amplification, Theorem 3 of Brassard et al. [9])**.** *Suppose U is a unitary operator such that*

$$U|0\rangle_A|0\rangle_B = \sqrt{p}|0\rangle_A|\phi_0\rangle_B + \sqrt{1-p}|1\rangle_A|\phi_1\rangle_B.$$

where $p \in [0,1]$ is unknown and $|\phi_0\rangle$ and $|\phi_1\rangle$ are normalized pure quantum states. There is a quantum algorithm that outputs $|0\rangle_A|\phi_0\rangle_B$ with probability ≥ 0.99, using $O\big(1/\sqrt{p}\big)$ queries to U.

D.2 Main Proof

We generalize Theorem 4.2 as follows.

Theorem D.9. *Algorithm 4 will produce k independent and identical distributed samples from a distribution that is ϵ_{G}-close to $\mathcal{O}_u^{\mathrm{Gibbs}}$ in total variation distance, in quantum time $\widetilde{O}\Big(\beta\sqrt{nk}\Big)$.*

It is immediate to show Theorem 4.2 from Theorem D.9 by constructing a β-amplitude-encoding V of $\mathbf{A}z$. To see this, let $u = \mathbf{A}z$, then $u_i = (\mathbf{A}z)_i \in [0,\beta]$. By Theorem D.2, we can implement a unitary operator U_z^{QRAM} such that

$$U_z^{\mathrm{QRAM}} \colon |0\rangle_C|0\rangle_B \mapsto |0\rangle_C \sum_{j\in[n]} \sqrt{\frac{z_j}{\beta}}|j\rangle_B + |1\rangle_C|\phi\rangle_B.$$

Using two queries to $\mathcal{O}_{\mathbf{A}}$, we can construct a unitary operator $\mathcal{O}'_{\mathbf{A}}$ such that

$$\mathcal{O}'_{\mathbf{A}} \colon |0\rangle_E|i\rangle_A|j\rangle_B \mapsto \Big(\sqrt{A_{i,j}}|0\rangle_E + \sqrt{1-A_{i,j}}|1\rangle_E\Big)|i\rangle_A|j\rangle_B.$$

Let

$$V = \big(|0\rangle_C\langle 0| \otimes \mathcal{O}'_{\mathbf{A}} + |1\rangle_C\langle 1| \otimes I_{EAB}\big)\Big(U_z^{\mathrm{QRAM}} \otimes I_{EA}\Big). \tag{11}$$

It can be verified (see Proposition D.10) that

$$\langle 0|_C\langle 0|_E V|0\rangle_C|0\rangle_E|i\rangle_A|0\rangle_B = \sum_{j\in[n]} \sqrt{\frac{A_{i,j}z_j}{\beta}}|i\rangle_A|j\rangle_B,$$

and thus $\langle 0|_C\langle 0|_E V|0\rangle_C|0\rangle_E|i\rangle_A|0\rangle_B = \sqrt{u_i/\beta}|i\rangle_A|\psi_i\rangle_B$ for some $|\psi_i\rangle$. Therefore, V is a β-amplitude-encoding of $\mathbf{A}z$.

Now, we will show Theorem D.9 in the following.

Proof of Theorem D.9. Now we start to describe our algorithm. By our consistent quantum amplitude estimation (Theorem C.3), we choose an $O(r)$-bit random string s, then we can obtain a quantum algorithm $\mathcal{O}_{\hat{u}}$ such that, with probability $1 - O(\exp(-r))$, for every $i \in [n]$, it computes $f^*(s, u_i/\beta)$ with $\widetilde{O}\big(\delta^{-1}\big) \cdot \mathrm{poly}(r)$ queries to V, where $f^*(s,p)$ is a function that only depends on s and p, and it holds that

$$|f^*(s,p) - p| \leq \delta$$

for every $p \in [-1,1]$. Here, r, δ are parameters to be determined. Note that

$$\frac{u_i}{\beta} = \|\langle 0|_C V|0\rangle_C|i\rangle_A|0\rangle_B\|^2,$$

so when applying consistent quantum amplitude estimation, we just use a controlled-XOR gate conditioned on the index and with A the target system, before every query to V.

By quantum k-maximum finding algorithm (Theorem D.3), we can find a set $S \subseteq [n]$ with $|S| = k$ such that $f^*(s, u_i/\beta) \geq f^*(s, u_j/\beta)$ for every $i \in S$ and $j \notin S$ with probability $0.99 - O\Big(\sqrt{nk}\exp(-r)\Big)$, using $O\Big(\sqrt{nk}\Big)$ queries to $\mathcal{O}_{\hat{u}}$. To obtain a constant probability, it is sufficient to choose $r = \Theta(\log(n))$.

For each $i \in S$, again applying our consistent quantum amplitude estimation (Theorem C.3), we can obtain the value of $f^*(s, u_i/\beta)$ with probability $1 - O(\exp(-r))$, using $\widetilde{O}(\delta^{-1}) \cdot \mathrm{poly}(r)$ queries to V; then we set

$$\hat{u}_i = \beta f^* \left(s, \frac{u_i}{\beta} \right)$$

for all $i \in S$, with success probability $1 - O(k \exp(-r))$ and using $\widetilde{O}(k\delta^{-1}) \cdot \mathrm{poly}(r)$ queries to V in total. It can be seen that $|\hat{u}_i - u_i| \leq \beta\delta$ for every $i \in S$.

Let $\tilde{u}_i = \hat{u}_i + \beta\delta$, and then we store $\tilde{u}_i$ for all $i \in S$ in the data structure as in Theorem D.2 (which costs $O(k)$ QRAM operations). Then, we calculate

$$W = (n-k)\exp(\tilde{u}^*) + \sum_{i \in S} \exp(\tilde{u}_i)$$

by classical computation in $\widetilde{O}(k)$ time, where

$$\tilde{u}^* = \min_{i \in S} \tilde{u}_i.$$

By Theorem D.2, we can prepare the quantum state

$$|u_{\text{guess}}\rangle = \sum_{i \in S} \sqrt{\frac{\exp(\tilde{u}_i)}{W}} |i\rangle + \sum_{i \notin S} \sqrt{\frac{\exp(\tilde{u}^*)}{W}} |i\rangle$$

in $\widetilde{O}(1)$ time.

Now we introduce another system D, and then let

$$U_u = (V_{CAB}^\dagger \otimes I_D)(V_{DAB} \otimes I_C).$$

It can be shown (see Proposition D.11) that U_u is a $(1,0)$-block-encoding of $\mathrm{diag}(u)/\beta$. By QRAM access to $\tilde{u}_i$, we can implement a unitary operator

$$V_{\tilde{u}} \colon |i\rangle_A |0\rangle_B \mapsto |i\rangle_A \left(\sqrt{\frac{\max\{\tilde{u}_i, \tilde{u}^*\}}{2\beta}} |0\rangle_B + \sqrt{1 - \frac{\max\{\tilde{u}_i, \tilde{u}^*\}}{2\beta}} |1\rangle_B \right)$$

in $\widetilde{O}(1)$ time by noting that $\max\{\tilde{u}_i, \tilde{u}^*\} = \tilde{u}_i$ if $i \in S$ and $\tilde{u}^*$ otherwise. We introduce one-qubit system C, and let

$$U_{\tilde{u}}^{\max} = \left(V_{\tilde{u}}^\dagger \otimes I_C \right) (\mathrm{SWAP}_{BC} \otimes I_A)(V_{\tilde{u}} \otimes I_C).$$

It can be shown that $U_{\tilde{u}}^{\max}$ is a $(1,0)$-block-encoding of $\mathrm{diag}(\max\{\tilde{u}, \tilde{u}^*\})/2\beta$. Applying the LCU technique (Theorem D.6), we can obtain a unitary operator U^- that is a $(1,0)$-block-encoding of $\mathrm{diag}(u - \max\{\tilde{u}, \tilde{u}^*\})/4\beta$, using $O(1)$ queries to U_u and $U_{\tilde{u}}^{\max}$. By the QSVT technique (Theorem D.7 and Lemma 4.1), we can construct a unitary operator $U^{\exp}$ that is a $(1,0)$-block-encoding of $P_{2\beta}(\mathrm{diag}(u - \max\{\tilde{u}, \tilde{u}^*\})/4\beta)$, using $O(\beta \log(\epsilon_P^{-1}))$ queries to U^-, where

$$\left| P_{2\beta}(x) - \frac{1}{4}\exp(2\beta x) \right| \leq \epsilon_P$$

for every $x \in [-1, 0]$ and $\epsilon_P \in (0, 1/2)$ is to be determined. Suppose $U^{\exp}$ has an a-qubit ancilla system, and let $|\tilde{u}_{\text{post}}\rangle = \langle 0|^{\otimes a} U^{\exp} |u_{\text{guess}}\rangle |0\rangle^{\otimes a}$. Note that

$$|\tilde{u}_{\text{post}}\rangle = \sum_{i \in S} P_{2\beta}\left(\frac{u_i - \tilde{u}_i}{4\beta} \right) \sqrt{\frac{\exp(\tilde{u}_i)}{W}} |i\rangle + \sum_{i \notin S} P_{2\beta}\left(\frac{u_i - \tilde{u}^*}{4\beta} \right) \sqrt{\frac{\exp(\tilde{u}^*)}{W}} |i\rangle.$$

It can be shown (Proposition D.12) that $\| |\tilde{u}_{\text{post}}\rangle \|^2 \geq \Theta(k/n)$; thus by quantum amplitude amplification (Theorem D.8), we can obtain

$$|\tilde{u}_{\text{Gibbs}}\rangle = \frac{|\tilde{u}_{\text{post}}\rangle}{\| |\tilde{u}_{\text{post}}\rangle \|}$$

using $O(\sqrt{n/k})$ queries to $U^{\exp}$. By measuring $|\tilde{u}_{\text{Gibbs}}\rangle$ on the computational basis, we return the outcome as a sample from the distribution $\tilde{u}_{\text{Gibbs}}$; it can be shown (Proposition D.13) that the total variation distance between $\tilde{u}_{\text{Gibbs}}$ and $\mathcal{O}_u^{\text{Gibbs}}$ is bounded by

$$d_{\text{TV}}\left(\tilde{u}_{\text{Gibbs}}, \mathcal{O}_u^{\text{Gibbs}}\right) \leq \sqrt{\frac{88n\epsilon_P}{k\exp(-2\beta\delta)}}.$$

By taking $\delta = 1/2\beta$ and $\epsilon_P = k\epsilon_{\text{G}}^2/300n$, we can produce one sample from $\tilde{u}_{\text{Gibbs}}$, using $\widetilde{O}(\beta\sqrt{n/k})$ queries to U_u and $U_{\tilde{u}}^{\max}$, with $\widetilde{O}(\beta\sqrt{nk})$-time precomputation.

Finally, by applying k times the above procedure (with the precomputation processed only once), we can produce k independent and identically distributed samples from $\tilde{u}_{\text{Gibbs}}$ that is ϵ_{Gibbs}-close to the Gibbs distribution $\mathcal{O}_u^{\text{Gibbs}}$, with total time complexity

$$\widetilde{O}\left(\beta\sqrt{nk}\right) + k \cdot \widetilde{O}\left(\beta\sqrt{\frac{n}{k}}\right) = \widetilde{O}\left(\beta\sqrt{nk}\right).$$

□

D.3 Technical Lemmas

Proposition D.10. *Let V defined by Equation* (11)*, we have*

$$\langle 0|_C\langle 0|_D V|0\rangle_C|0\rangle_D|i\rangle_A|0\rangle_B = \sum_{j\in[n]}\sqrt{\frac{A_{i,j}z_j}{\beta}}|i\rangle_A|j\rangle_B.$$

Proof.

$$\begin{aligned}
&V|0\rangle_C|0\rangle_D|i\rangle_A|0\rangle_B\\
=&\left(|0\rangle_C\langle 0|\otimes\mathcal{O}'_{\mathbf{A}} + |1\rangle_C\langle 1|\otimes I_{AB}\right)\left(|0\rangle_C|0\rangle_D|i\rangle_A\sum_{j\in[n]}\sqrt{\frac{z_j}{\beta}}|j\rangle_B + |1\rangle_C|0\rangle_D|i\rangle_A|\phi\rangle_B\right)\\
=&|0\rangle_C\sum_{j\in[n]}\left(\sqrt{A_{i,j}}|0\rangle_D + \sqrt{1-A_{i,j}}|1\rangle_D\right)\sqrt{\frac{z_j}{\beta}}|i\rangle_A|j\rangle_B + |1\rangle_C|0\rangle_D|i\rangle_A|\phi\rangle_B\\
=&|0\rangle_C|0\rangle_D\sum_{j\in[n]}\sqrt{\frac{A_{i,j}z_j}{\beta}}|i\rangle_A|j\rangle_B + |0\rangle_C|1\rangle_D\sum_{j\in[n]}\sqrt{\frac{(1-A_{i,j})z_j}{\beta}}|i\rangle_A|j\rangle_B + |1\rangle_C|0\rangle_D|i\rangle_A|\phi\rangle_B.
\end{aligned}$$

□

Proposition D.11. *In the proof of Theorem D.9, U_u is a $(1,0)$-block-encoding of* $\operatorname{diag}(u)/\beta$.

Proof. To see this, for every $i, j \in [n]$,

$$\begin{aligned}
&\langle 0|_C\langle 0|_D\langle j|_A\langle 0|_B U_u|0\rangle_C|0\rangle_D|i\rangle_A|0\rangle_B\\
=&\langle 0|_C\langle 0|_D\langle j|_A\langle 0|_B(V_{CAB}^\dagger\otimes I_D)(V_{DAB}\otimes I_C)|0\rangle_C|0\rangle_D|i\rangle_A|0\rangle_B\\
=&\left(\sqrt{u_j/\beta}\langle 0|_C\langle 0|_D\langle j|_A\langle\psi_j|_B + \langle 1|_C\langle 0|_D\langle g_j|_{AB}\right)\left(\sqrt{u_i/\beta}|0\rangle_C|0\rangle_D|i\rangle_A|\psi_i\rangle_B + |0\rangle_C|1\rangle_D|g_i\rangle_{AB}\right)\\
=&\ \langle j|i\rangle_A\,\frac{u_i}{\beta}.
\end{aligned}$$

□

Proposition D.12. *In the proof of Theorem D.9, if $\delta = 1/2\beta$, $E = \sum_{j\in[n]}\exp(u_j)$, and $\epsilon_P = k\epsilon_{\text{G}}^2/300n$, then*

$$\Theta\left(\frac{k}{n}\right) \leq \frac{E}{16W} - 2\epsilon_P \leq \||u_{\text{post}}\rangle\|^2 \leq \frac{E}{16W} + 3\epsilon_P.$$

Proof. We first give an upper bound for W in terms of u_i and $\tilde{u}^*$. Notice that $\tilde{u}_i \le u_i + 2\beta\delta$ for all $i \in S$, we have:

$$W = (n-k)\exp(\tilde{u}^*) + \sum_{i\in S}\exp(\tilde{u}_i) \le \exp(2\beta\delta)\Bigg((n-k)\exp(u^*) + \sum_{i\in S}\exp(u_i)\Bigg).$$

Note that

$$\frac{(n-k)\exp(u^*) + \sum\limits_{i\in S}\exp(u_i)}{\sum\limits_{i\in[n]}\exp(u_i)} \le \frac{n-k}{k} + 1 = \frac{n}{k},$$

then we have

$$\frac{E}{W} \ge \sum_{i\in[n]} \frac{\exp(u_i)}{\exp(2\beta\delta)((n-k)\exp(u^*) + \sum_{i\in S}\exp(u_i))} \ge \frac{k}{n}\exp(-2\beta\delta). \tag{12}$$

With this, noting that $(a-b)^2 \ge a^2 - 2ab$ for any real a and b, we have

$$\begin{aligned}
\||u_{\text{post}}\rangle\|^2 &= \sum_{i\in S}\left(P_{2\beta}\left(\frac{u_i - \tilde{u}_i}{4\beta}\right)\right)^2 \frac{\exp(\tilde{u}_i)}{W} + \sum_{i\notin S}\left(P_{2\beta}\left(\frac{u_i - \tilde{u}^*}{4\beta}\right)\right)^2 \frac{\exp(\tilde{u}^*)}{W} \\
&\ge \sum_{i\in S}\left(\left(\frac{1}{4}\exp\left(\frac{u_i - \tilde{u}_i}{2}\right)\right)^2 - 2\epsilon_P\right)\frac{\exp(\tilde{u}_i)}{W} \\
&\quad + \sum_{i\notin S}\left(\left(\frac{1}{4}\exp\left(\frac{u_i - \tilde{u}^*}{2}\right)\right)^2 - 2\epsilon_P\right)\frac{\exp(\tilde{u}^*)}{W} \\
&= \frac{1}{16}\left(\sum_{i\in S}\exp(u_i - \tilde{u}_i)\frac{\exp(\tilde{u}_i)}{W} + \sum_{i\notin S}\exp(u_i - \tilde{u}^*)\frac{\exp(\tilde{u}^*)}{W}\right) \\
&\quad - 2\epsilon_P\left(\sum_{i\in S}\frac{\exp(\tilde{u}_i)}{W} + \sum_{i\notin S}\frac{\exp(\tilde{u}^*)}{W}\right) \\
&\ge \frac{E}{16W} - 2\epsilon_P \\
&\ge \Theta\left(\frac{k}{n}\right).
\end{aligned}$$

On the other hand, a similar argument using the inequality $(a+b)^2 \le a^2 + 3ab$ for positive real $a \ge b$ gives

$$\||u_{\text{post}}\rangle\|^2 \le \frac{E}{16W} + 3\epsilon_P.$$

These yield the proof. □

Proposition D.13. *In the proof of Theorem D.9, the total variation distance between the two distributions $\tilde{u}_{Gibbs}$ and $\mathcal{O}_u^{\text{Gibbs}}$ is bounded by*

$$d_{TV}\left(\tilde{u}_{\text{Gibbs}}, \mathcal{O}_u^{\text{Gibbs}}\right) \le \sqrt{\frac{88n\epsilon_P}{k\exp(-2\beta\delta)}}.$$

Proof. Define $E = \sum\limits_{j\in[n]}\exp\big(u_j\big)$. Let

$$|u_{\text{Gibbs}}\rangle = \sum_{i\in[n]}\sqrt{\frac{\exp(u_i)}{E}}\,|i\rangle$$

be the intended quantum state with amplitudes the same as the Gibbs distribution $\mathcal{O}_u^{\text{Gibbs}}$. The inner product between $|\tilde{u}_{\text{post}}\rangle$ and $|u_{\text{Gibbs}}\rangle$ can be bounded by:

$$\begin{aligned}
\langle \tilde{u}_{\text{post}} | u_{\text{Gibbs}} \rangle &= \sum_{i \in S} P_{2\beta}\left(\frac{u_i - \tilde{u}_i}{4\beta}\right) \sqrt{\frac{\exp(\tilde{u}_i)}{W}} \sqrt{\frac{\exp(u_i)}{E}} \\
&\quad + \sum_{i \notin S} P_{2\beta}\left(\frac{u_i - \tilde{u}^*}{4\beta}\right) \sqrt{\frac{\exp(\tilde{u}^*)}{W}} \sqrt{\frac{\exp(u_i)}{E}} \\
&\geq \sum_{i \in S} \left(\frac{1}{4} \exp\left(\frac{u_i - \tilde{u}_i}{2}\right) - \epsilon_P\right) \sqrt{\frac{\exp(\tilde{u}_i)}{W}} \sqrt{\frac{\exp(u_i)}{E}} \\
&\quad + \sum_{i \notin S} \left(\frac{1}{4} \exp\left(\frac{u_i - \tilde{u}_i}{2}\right) - \epsilon_P\right) \sqrt{\frac{\exp(\tilde{u}^*)}{W}} \sqrt{\frac{\exp(u_i)}{E}} \\
&\geq \frac{1}{4\sqrt{WE}} \left(\sum_{i \in [n]} \exp(u_i)\right) - \epsilon_P.
\end{aligned}$$

The last step is by Cauchy's inequality. By Proposition D.12 and Equation (12), we have

$$\begin{aligned}
|\langle \tilde{u}_{\text{Gibbs}} | u_{\text{Gibbs}} \rangle|^2 = \frac{|\langle \tilde{u}_{\text{post}} | u_{\text{Gibbs}} \rangle|^2}{\| |\tilde{u}_{\text{post}}\rangle \|^2} &\geq \frac{E}{16W \| |\tilde{u}_{\text{post}}\rangle \|^2} - \frac{\epsilon_P}{2 \| |\tilde{u}_{\text{post}}\rangle \|^2} \\
&\geq \frac{E}{16W\left(\frac{E}{16W} + 3\epsilon_P\right)} - \frac{\epsilon_P}{2 \| |\tilde{u}_{\text{post}}\rangle \|^2} \\
&\geq 1 - \frac{48\epsilon_P}{E/W} - \frac{8\epsilon_P}{E/W - 32\epsilon_P} \\
&\geq 1 - \frac{48n\epsilon_P}{k\exp(-2\beta\delta)} - \frac{8n\epsilon_P}{k\exp(-2\beta\delta) - 32n\epsilon_P} \\
&\geq 1 - \frac{88n\epsilon_P}{k\exp(-2\beta\delta)}.
\end{aligned}$$

Finally, we have

$$\begin{aligned}
d_{\text{TV}}\left(\tilde{u}_{\text{Gibbs}}, \mathcal{O}_u^{\text{Gibbs}}\right) &\leq \frac{1}{2} \operatorname{tr}\left(\Big| |\tilde{u}_{\text{Gibbs}}\rangle\langle \tilde{u}_{\text{Gibbs}}| - |u_{\text{Gibbs}}\rangle\langle u_{\text{Gibbs}}| \Big|\right) \\
&= \sqrt{1 - |\langle \tilde{u}_{\text{Gibbs}} | u_{\text{Gibbs}} \rangle|^2} \\
&\leq \sqrt{\frac{88n\epsilon_P}{k\exp(-2\beta\delta)}},
\end{aligned}$$

which is bounded by ϵ_G by the choice of ϵ_P. □