Quantum rejection sampling

Maris Ozols[1,2], Martin Roetteler[1], and Jérémie Roland[1]
[1]*NEC Laboratories America, Inc.*
[2]*University of Waterloo & Institute for Quantum Computing*

NEC Laboratories America
Relentless passion for innovation

IQC Institute *for* Quantum Computing

The resampling problem

The goal of *resampling* is to produce samples from a target distribution, given the ability to sample from another distribution. To study the query complexity of this problem, we consider an oracle that attaches unknown data to each element of the base set $[n] := \{1, \dots, n\}$. We formalize the classical and quantum version of the problem as follows:

Classical
Given a black-box oracle producing pairs $(\xi(k), k) \in [d] \times [n]$ where k is distributed according to P and $\xi : [n] \to [d]$ is an unknown function, $\textsc{Sampling}_{P \to Q}$ is the problem to produce a sample $(\xi(k), k)$ such that k is distributed according to Q.

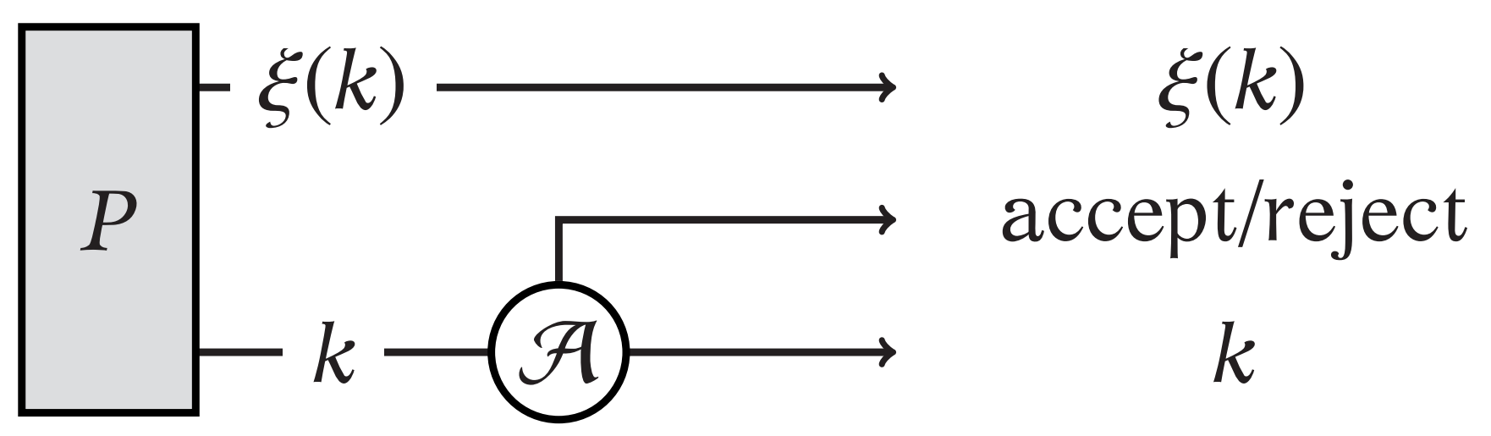

Quantum
Let $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{R}^n$ ($\|\boldsymbol{\alpha}\|_2 = \|\boldsymbol{\beta}\|_2 = 1, \forall k : \alpha_k, \beta_k \geq 0$) and O be a unitary that acts as $O : |\bar{0}\rangle \mapsto |a\rangle := \sum_{k=1}^n \alpha_k |\xi_k\rangle|k\rangle$, where $|\xi_k\rangle \in \mathbb{C}^d$ are some fixed unknown quantum states. Given oracle access to unitary black boxes $O, O^\dagger$, the $\text{Q}\textsc{Sampling}_{\alpha \to \beta}$ problem is to prepare the state $|b\rangle := \sum_{k=1}^n \beta_k |\xi_k\rangle|k\rangle$.

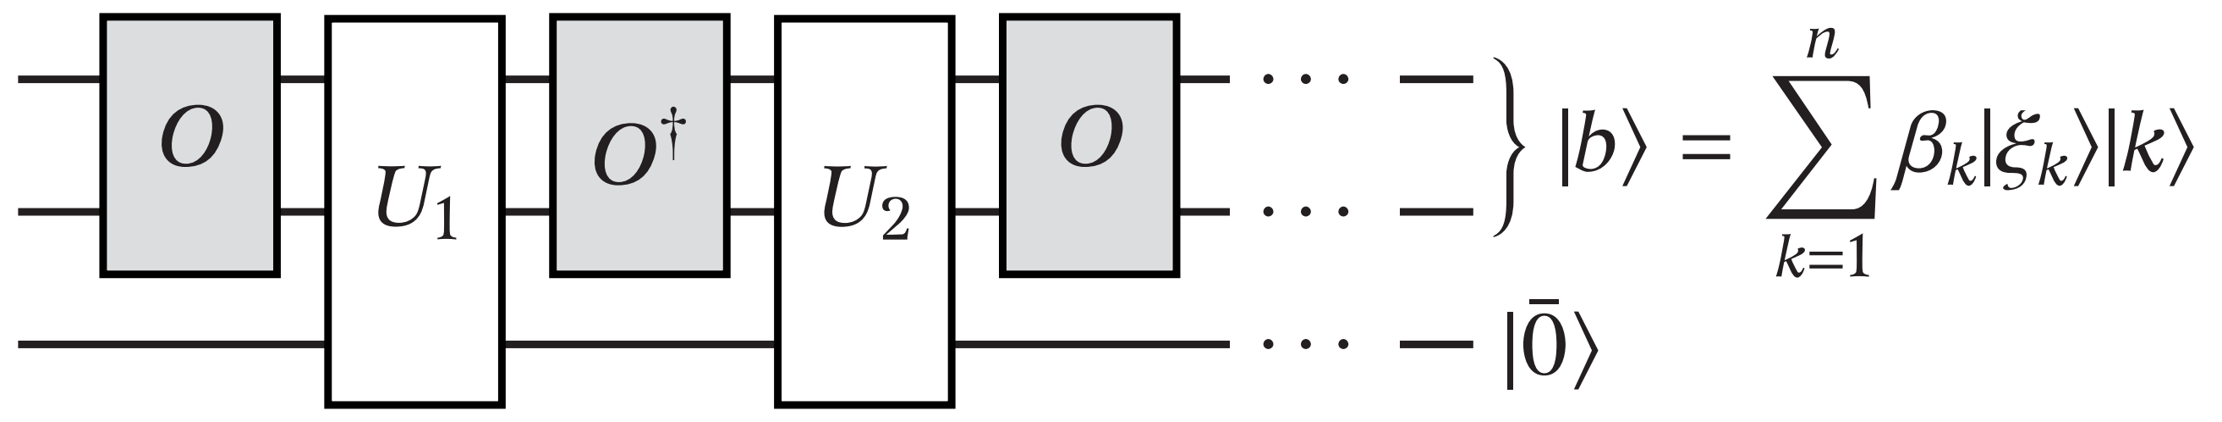

For any $p > 0$, $\text{Q}\textsc{Sampling}^p_{\alpha \to \beta}$ is the problem where it suffices to prepare $\sqrt{p}|b\rangle|\bar{0}\rangle + |\text{error}\rangle$, where $|\bar{0}\rangle$ is a default state for the workspace and $|\text{error}\rangle$ is an arbitrary (unnormalized) error state.

Classical algorithm

The classical *rejection sampling* works as follows [1]: let $\gamma \leq 1$ be the largest scaling factor such that γQ lies under P, formally, $\gamma = \min_k(p_k/q_k)$. We accept a sample k from P with probability $\gamma q_k/p_k$, otherwise we reject it and repeat. The expected number T of samples from Q to produce one sample from P is then given by $T = 1/\gamma = \max_k(q_k/p_k)$.

Main result

Theorem. Let $p_{\min} := (\boldsymbol{\beta}^\mathsf{T} \cdot \boldsymbol{\alpha})^2$ and $p_{\max} := \sum_{k:\alpha_k > 0} |\beta_k|^2$. The quantum query complexity of $\text{Q}\textsc{Sampling}^p_{\alpha \to \beta}$ for $p \in [p_{\min}, p_{\max}]$ is $\Theta(1/\|\boldsymbol{\varepsilon}\|_2)$, where $\boldsymbol{\varepsilon}$ is the largest vector such that $\boldsymbol{\beta}^\mathsf{T} \cdot \boldsymbol{\varepsilon} = \sqrt{p}\,\|\boldsymbol{\varepsilon}\|_2$ and $\varepsilon_k = \min\{\alpha_k, \gamma\beta_k\}$ for some $\gamma > 0$. For $p \leq p_{\min}$, the query complexity is 1, and for $p > p_{\max}$, it is infinite.

SDP characterization

Quantum query complexity of $\text{Q}\textsc{Sampling}^p_{\alpha \to \beta}$ for $p_{\min} \leq p \leq p_{\max}$ is $O(1/\sqrt{q})$, where q is the optimal value of the following SDP:

$$\max_{M \geq 0} \operatorname{Tr} M \quad \text{s.t. } \forall k : \alpha_k^2 \geq M_{kk}, \\ \operatorname{Tr}[(\boldsymbol{\beta} \cdot \boldsymbol{\beta}^\mathsf{T} - pI)M] \geq 0.$$

Quantum algorithm

Our algorithm is based on *amplitude amplification* [2]. Fix some $\boldsymbol{\varepsilon} \in \mathbb{R}^n$ ($\|\boldsymbol{\varepsilon}\|_2 \leq 1, \forall k : \varepsilon_k \geq 0$) and let $U_{\boldsymbol{\varepsilon}} := (I_d \otimes R_{\boldsymbol{\varepsilon}}) \cdot (O \otimes I_2)$ be the following operation:

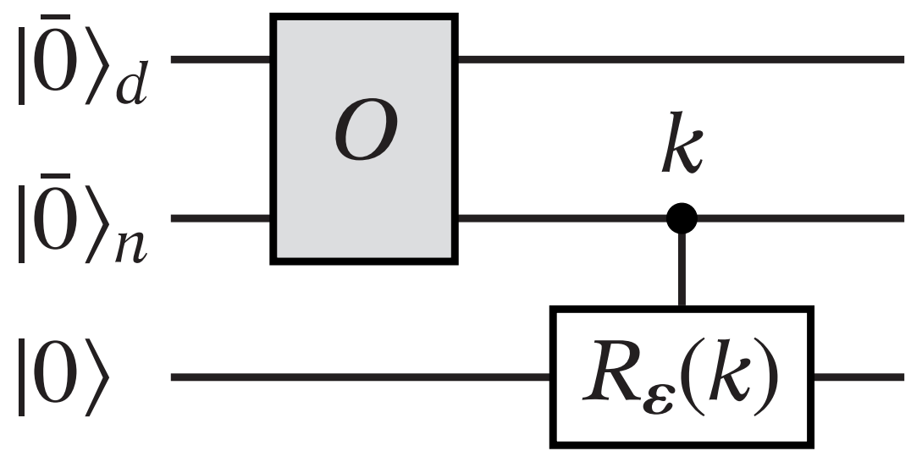

where each $R_{\boldsymbol{\varepsilon}}(k)$ is a 2-dimensional rotation. It prepares state

$$|\Psi_{\boldsymbol{\varepsilon}}\rangle := (I_d \otimes R_{\boldsymbol{\varepsilon}}) \cdot |a\rangle|0\rangle = \sum_{k=1}^{n} |\xi_k\rangle|k\rangle\Big(\sqrt{|\alpha_k|^2 - \varepsilon_k^2}\,|0\rangle + \varepsilon_k|1\rangle\Big).$$

Quantum rejection sampling algorithm

```
1. Start in initial state $|\bar{0}\rangle_d|\bar{0}\rangle_n|0\rangle$
2. Apply $U_{\varepsilon}$
3. Perform the following steps $t$ times:
     ▸ perform $\mathrm{ref}_{I_d \otimes I_n \otimes |1\rangle\langle 1|}$ by applying Pauli $Z$ on the coin register
     ▸ perform $\mathrm{ref}_{|\Psi_{\varepsilon}\rangle}$ by applying $U_{\varepsilon}^\dagger$, changing the relative phase of
       $|\bar{0}, \bar{0}, 0\rangle$ by a factor of $-1$, and then applying $U_{\varepsilon}$
4. Discard the last register
```

We need one query to implement the operation $U_{\boldsymbol{\varepsilon}}$ and two queries to implement the second reflection, thus in total we need $2t + 1 = O(1/\|\boldsymbol{\varepsilon}\|_2)$ calls to oracles O and $O^\dagger$.

Lower bound

We choose a set of oracles $\mathcal{O}_{\alpha,\beta} := \{O_{\vec{x},u} : \vec{x} \in \mathbb{F}_2^n, u \in S\}$, where $S := \{u \in \mathrm{U}(n) : u|\bar{0}\rangle = |\bar{0}\rangle\} \cong \mathrm{U}(n-1)$. The first oracle acts as $O_{\vec{0},I}|\bar{0}\rangle = \sum_{k=1}^n \alpha_k|k\rangle$, and arbitrarily on the orthogonal complement. The remaining oracles are

$$O_{\vec{x},u} := V_{\vec{x}}\, O_{\vec{0},I}\, u, \qquad \text{where} \qquad V_{\vec{x}} := \textstyle\sum_{k=1}^n (-1)^{x_k}|k\rangle\langle k|.$$

Next, we symmetrize the circuit over the automorphism group $\mathbb{F}_2^n \times \mathrm{U}(n-1)$ of $\mathcal{O}_{\alpha,\beta}$ as follows:

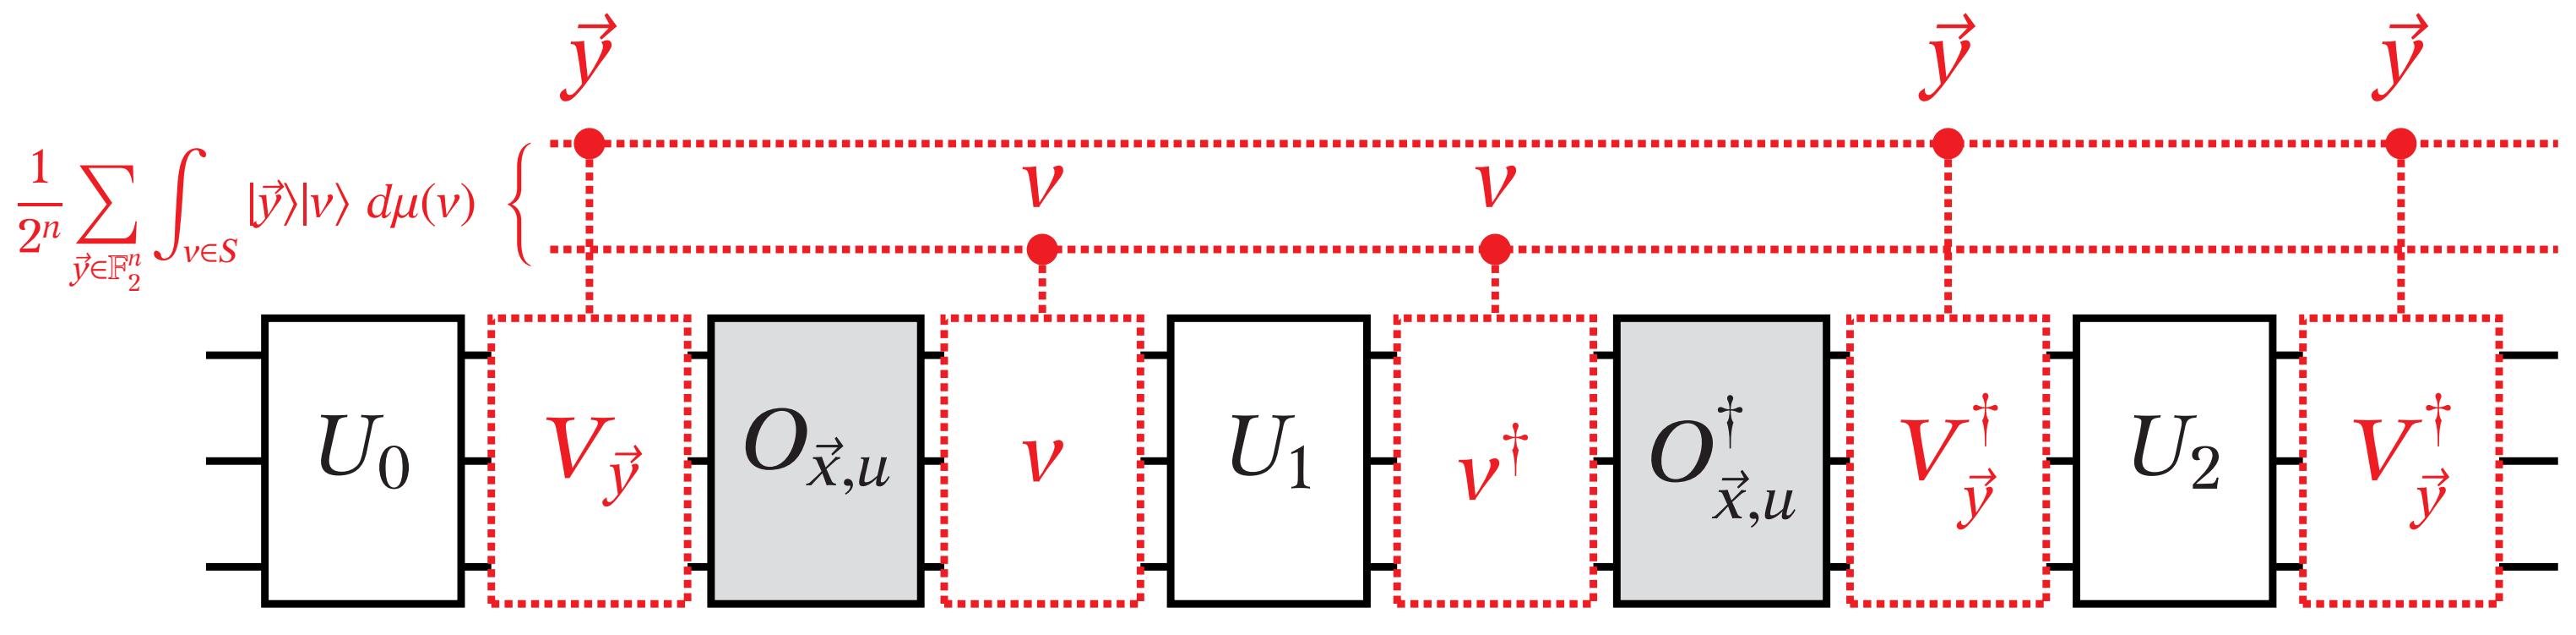

The rest of the proof is based on the *hybrid argument*, which yields the same SDP as analysis of the algorithm.

Application to Boolean hidden shift problem

The *hidden shift problem* for Boolean function $f : \mathbb{F}_2^n \to \mathbb{F}_2$ is the following problem [3]: given an oracle access to shifted function

$$f_{\vec{s}}(\vec{x}) := f(\vec{x} + \vec{s})$$

for some unknown value of $\vec{s} \in \mathbb{F}_2^n$, determine the value of $\vec{s}$ by querying the oracle on different inputs. Let BHSP_f^p denote the number of queries needed to determine $\vec{s}$ with probability p.

Theorem. $\mathrm{BHSP}_f^p = O(1/\|\boldsymbol{\varepsilon}\|_2)$ where $\boldsymbol{\varepsilon}$ is the largest vector such that $\|\boldsymbol{\varepsilon}\|_1 = \sqrt{2^n p}\,\|\boldsymbol{\varepsilon}\|_2$ and $\varepsilon_{\vec{w}} = \min\{|\hat{f}_{\vec{w}}|, \gamma/\sqrt{2^n}\}$ for some $\gamma > 0$, where $\hat{f}$ is the *Fourier transform* of f.

References

[1] J. von Neumann. Various techniques used in connection with random digits. *National Bureau of Standards, Applied Math Series*, 12:36–38, 1951.

[2] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation, 2000.

[3] M. Roetteler. Quantum algorithms for highly non-linear Boolean functions. *Proc. SODA'10*, 448–457, 2010.