Demonstration of Algorithmic Quantum Speedup for an Abelian Hidden Subgroup Problem

Phattharaporn Singkanipa,[1,2] Victor Kasatkin,[3,2] Zeyuan Zhou,[4] Gregory Quiroz,[4,5] and Daniel A. Lidar[6,7]

[1]*Department of Physics, University of Southern California, Los Angeles, California 90089, USA*
[2]*Center for Quantum Information Science and Technology, University of Southern California, Los Angeles, California 90089, USA*
[3]*Viterbi School of Engineering, University of Southern California, Los Angeles, California 90089, USA*
[4]*William H. Miller III Department of Physics and Astronomy, Johns Hopkins University, Baltimore, Maryland 21218, USA*
[5]*Johns Hopkins University Applied Physics Laboratory, Laurel, Maryland 20723, USA*
[6]*Departments of Electrical & Computer Engineering, Chemistry, Physics & Astronomy, and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, California 90089, USA*
[7]*Quantum Elements, Inc., Thousand Oaks, California, USA*

(Received 15 January 2024; revised 12 January 2025; accepted 28 April 2025; published 5 June 2025)

Simon's problem is to find a hidden period (a bitstring) encoded into an unknown 2-to-1 function. It is one of the earliest problems for which an exponential quantum speedup was proven for ideal, noiseless quantum computers, albeit in the oracle model. Here, using two different 127-qubit IBM Quantum superconducting processors, we demonstrate an algorithmic quantum speedup for a variant of Simon's problem where the hidden period has a restricted Hamming weight w. For sufficiently small values of w and for circuits involving up to 58 qubits, we demonstrate an exponential speedup, albeit of a lower quality than the speedup predicted for the noiseless algorithm. The speedup exponent and the range of w values for which an exponential speedup exists are significantly enhanced when the computation is protected by dynamical decoupling. Further enhancement is achieved with measurement error mitigation. This case constitutes a demonstration of a bona fide quantum advantage for an Abelian hidden subgroup problem.

Intro Simon's problem

In computational complexity theory and quantum computing, Simon's problem is a computational problem that is proven to be solved exponentially faster on a quantum computer than on a classical (that is, traditional) computer.

This problem yields an oracle separation between the complexity classes BPP (bounded-error classical query complexity) and BQP (bounded-error quantum query complexity).

Simon's problem considers access to a function $f : \{0,1\}^n \to \{0,1\}^m,\ m \geq n$ as implemented by a black box or an oracle. This function is promised to be either a one-to-one function, or a two-to-one function; if f is two-to-one, it is furthermore promised that two inputs x and x' evaluate to the same value if and only if x and x' differ in a fixed set of bits. I.e.,

> If f is not one-to-one, it is promised that there exists a non-zero s such that, for all $x \neq x'$, $f(x) = f(x')$ if and only if $x' = x \oplus s$

where $\oplus$ denotes bitwise exclusive-or. Simon's problem asks, in its decision version, whether f is one-to-one or two-to-one. In its non-decision version, Simon's problem asks whether f is one-to-one or what is the value of s (as defined above). The goal is to solve this task with the least number of queries (evaluations) of f.

Note that if $x' = x$, then $f(x') = f(x)$ and $x' = x \oplus s$ with $s = 0$. On the other hand (because $a \oplus b \oplus b = a$ for all a and b), $x' = x \oplus s \iff x' \oplus x = s$. Thus, Simon's problem may be restated in the following form:

> Given black-box or oracle access to f, promised to satisfy, for some s and all x, x', $f(x) = f(x')$ if and only if $x' \oplus x \in \{0, s\}$, determine whether $s \neq 0$ (decision version), or output s (non-decision version).

Note also that the promise on f implies that if f is two-to-one then it is a periodic function: $f(x) = f(x \oplus s)$.

That's their Hamming Weight w

Example [edit]

The following function is an example of a function that satisfies the required property for $n = 3$:

x	$f(x)$
000	101
001	010
010	000
011	110
100	000
101	110
110	101
111	010

In this case, $s = 110$ (i.e. the solution). Every output of f occurs twice, and the two input strings corresponding to any one given output have bitwise XOR equal to $s = 110$.

Simon's algorithm [edit]

The algorithm as a whole uses a subroutine to execute the following two steps:

1. Run the quantum subroutine an expected $O(n)$ times to get a list of linearly independent bitstrings $y_1, \ldots, y_{n-1}$.
2. Each y_k satisfies $y_k \cdot s = 0$, so we can solve the system of equations this produces to get s.

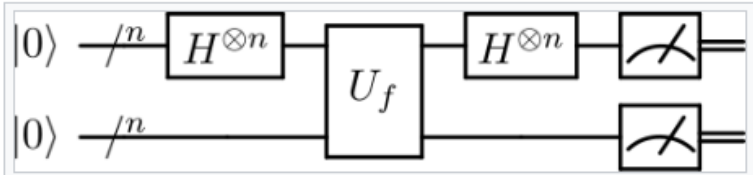

Quantum circuit representing/implementing Simon's algorithm

$$\frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} |k\rangle |0\rangle^{\otimes n}.$$

Query the oracle U_f to get the state

$$\frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} |k\rangle |f(k)\rangle.$$

Apply another Hadamard transform to the first register. This will produce the state

$$\frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} \left[\frac{1}{\sqrt{2^n}} \sum_{j=0}^{2^n-1} (-1)^{j \cdot k} |j\rangle \right] |f(k)\rangle = \sum_{j=0}^{2^n-1} |j\rangle \left[\frac{1}{2^n} \sum_{k=0}^{2^n-1} (-1)^{j \cdot k} |f(k)\rangle \right].$$

Then: measure first register

If $s = 0^n$ (one-to-one), amplitude is $1/2^n$.
If two-to-one: if j.s = 1, prob = 0, if j.s = 0, prob = $2/2^n$

2. Modified Problem in the Paper: Restricted Hamming Weight

The paper restricts the set of allowed hidden periods:

$$s \in \{0,1\}^n \quad \text{with} \quad |s| = w,$$

where $|s|$ is the **Hamming weight** (number of ones). arxiv +1

This modification is crucial because:

- It reduces circuit complexity.
- It makes experimental implementation feasible on **58-qubit circuits** and **127-qubit IBM processors**.
- Theoretical analysis in the paper proves exponential quantum speedup *still holds even under noise* for this restricted family. arxiv

In the restricted-Hamming-weight variant, the paper shows the same core scaling persists for sufficient small w. The experiment measures *time-to-solution* as a function of problem size and observes an exponential divergence consistent with:

$$T_{\text{quantum}}(n) \sim a \cdot n^{\alpha},$$

$$T_{\text{classical}}(n) \sim b \cdot c^{n},$$

with $c > 1$. This experimentally extracted exponent is the "speedup exponent." The paper reports exponential speedup for circuits involving up to **58 qubits**. arxiv

Typically in Simon's problem, linear vs exponential query separation

Results

The quantum circuit for each hidden bitstring b was constructed according to Fig. 1 and then compiled to fit the architecture of the devices using the as-late-as-possible (ALAP) schedule. This schedule initializes each qubit just before its first operation. We performed our demonstrations in two main modes: (i) “with DD,” i.e., error-suppressed demonstrations with different dynamical decoupling sequences, and (ii) “without DD,” i.e., the circuit as specified in Fig. 1 without any additional error suppression.

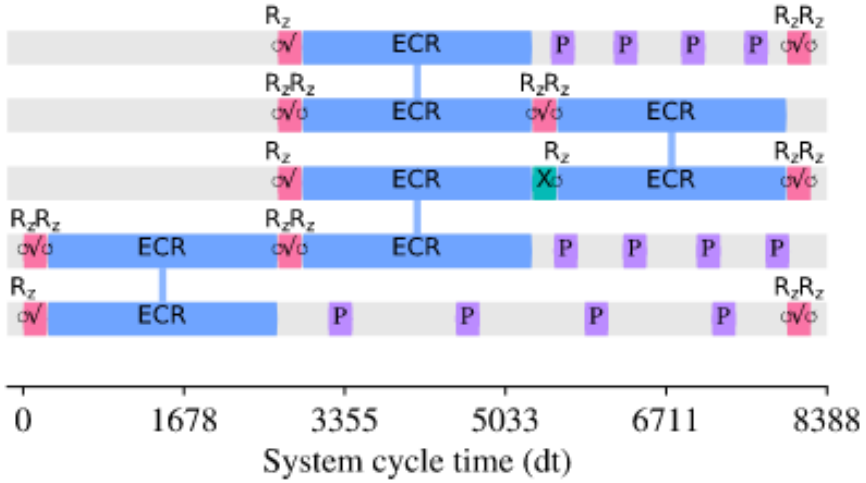

FIG. 2. Circuit for Simon-3 ($b = 111$) compiled into the Sherbrooke architecture using the ALAP schedule. The XY4 DD pulse sequences shown here as an example (the “P” inside the purple boxes) are placed such that they fill all available idle spaces after each qubit is initialized. Pulse intervals vary depending on the length of the idle period. Circular arrows denote R_z gates, and $\surd$ denotes the $\sqrt{X}$ gate. ECR denotes an echoed cross-resonance gate, which is equivalent to a CNOT up to single-qubit rotations.

They apply dynamical decoupling (DD) and measurement error mitigation (MEM): $p_{obs} = M^{-1} p_{meas}$

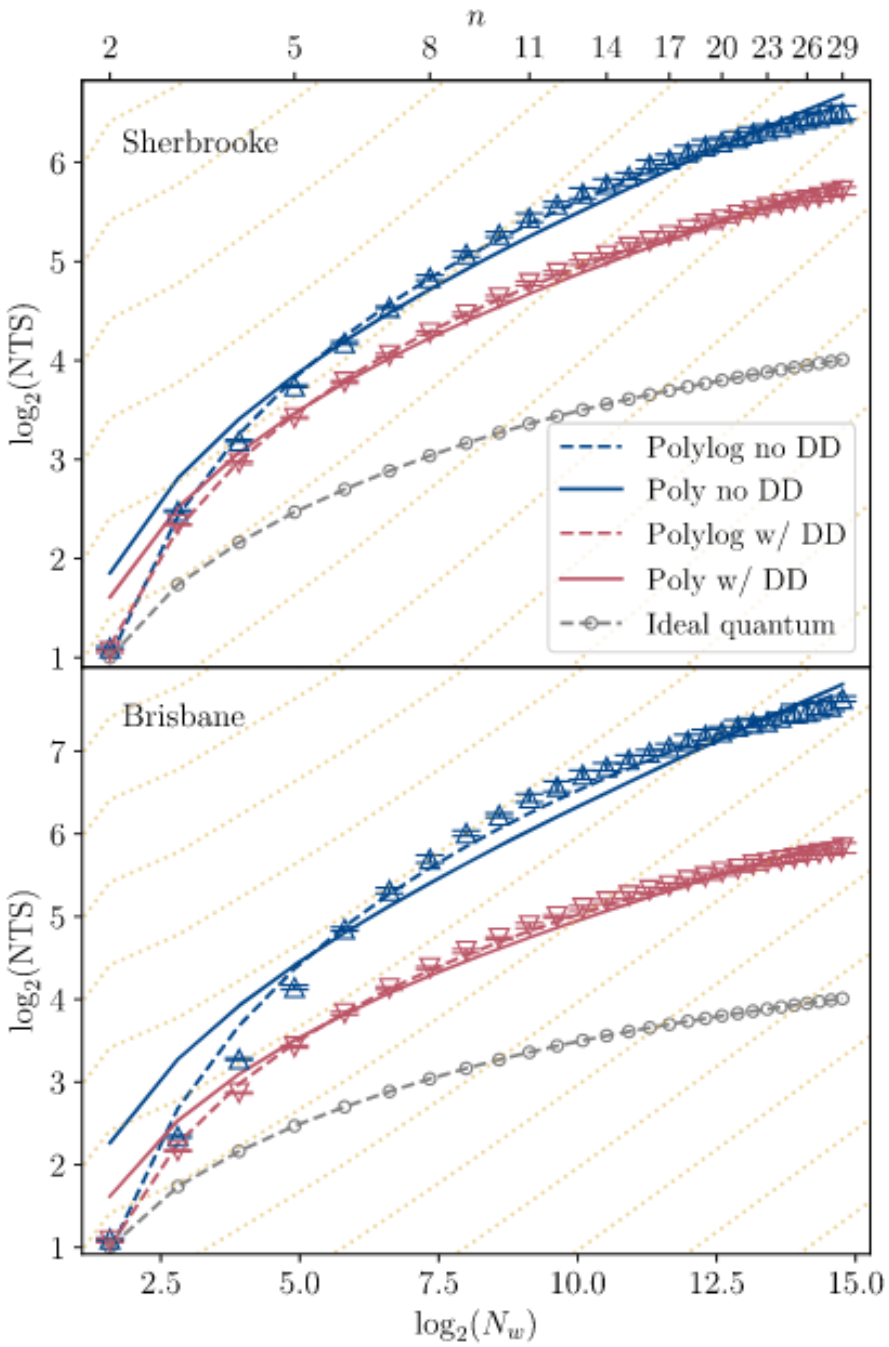

FIG. 3. NTS as a function of $\log_2(N_w)$ and problem size n on Sherbrooke (top) and Brisbane (bottom) for $\mathtt{w}_4$Simon-29, both with MEM. The blue lines represent the unprotected, no-DD form of the algorithm. The red lines represent the DD-protected form, i.e., a circuit where DD fills the idle gaps, with the optimal DD sequence from the set $\mathcal{D}$ used at each problem size n. The dashed lines denote fitting using the polylog model [Eq. (11)]; the solid lines denote fitting using the poly model [Eq. (12)]. The error bars, representing confidence intervals derived through bootstrapping, extend 1σ in both directions from each data point. The yellow dotted lines represent the scaling of the lower bound on NTS_C and serve as a visual guide for the scaling of the poly model. The gray dashed line is the theoretical performance of the quantum algorithm running on a noiseless device, given by Eq. (9).

For small Hamming weights, quantum runtimes scale exponentially better than classical ones. The observed speedup is weaker than the ideal noiseless prediction, but still exponential. With DD and MEM, the speedup becomes substantially clearer, confirming the crucial role of noise suppression.

NTS denotes the number-of-oracle-queries-to-solution

FIG. 5. Fitted scaling parameters α (left axis) and β (right axis) of the polylog and poly models, respectively, as a function of w for $\mathtt{W}_w$Simon-n on Sherbrooke and Brisbane, both with MEM. The vertical lines indicate the HW where the transition between the two models occurs, denoted w_t in the text. The polylog (poly) model provides a better fit to the left (right) of the vertical line corresponding to Brisbane with DD (blue solid line) and Sherbrooke with DD (red solid line). For both Brisbane and Sherbrooke without DD, the polylog model is always better but only for $w \in [1, 4]$, beyond which we do not have enough data (see text). When NTS_Q follows the polylog model, an exponential quantum speedup holds with a scaling exponent given by α [Eq. (11)]. For the poly model [Eq. (12)], we find that the quantum slope is always above the classical slope indicated by the dashed-dotted yellow line, corresponding to a polynomial quantum slowdown. The gray curve corresponds to the α value of an ideal (noise-free) quantum implementation. The error bars, representing the standard deviation of the fitted parameters on the bootstrapped data, extend 1σ in each direction from each data point. The Brisbane results are generally better than Sherbrooke's. The exponential speedup "quality," as quantified by the value of α, generally deteriorates as w increases.