Lecture 9: Quantum Search Algorithm

Grover's Algorithm and Amplitude Amplification

Moody T. Chu

North Carolina State University

Spring 2026

Lecture Overview

1. The Unstructured Search Problem
2. The Quantum Oracle and Phase Kickback
3. The Grover Iteration Operator (G)
4. Mathematical Analysis: Inversion About the Mean
5. Geometric Interpretation: 2D Subspace Rotations
6. Performance Analysis and Iteration Bounds
7. Generalization: Multiple Target Solutions

The Unstructured Search Problem

- **The Setup:** Given a black-box function $f : \{0,1\}^n \to \{0,1\}$,
  - we are promised there is a unique $\omega \in \{0,1\}^n$ such that $f(\omega) = 1$.
  - Find ω.
- **The Challenge:** The database is completely unstructured.
  - There is no sorting, no alphabetical order, and no heuristic to guide the search.
  - There are $N = 2^n$ elements in the database.
- **Why is this important?** Many computational problems reduce to exhaustive search over an unstructured space: satisfiability, graph coloring, optimization over combinatorial spaces. A faster search subroutine gives a quantum speedup across all of them (Section 7).

Complexity: Classical vs. Quantum

- **Classical Approach:**
  - The only classical strategy is linear enumeration.
    - Comparison-based sorts (Quicksort) or bit-based sorts (Radix Sort) are inapplicable — there is no key structure to exploit.
  - Since each of the N items is equally likely to be ω, the number of queries X before finding it is uniformly distributed: $P(X = k) = \frac{1}{N}$ for $k = 1, \ldots, N$.
  - Expected number of queries: $E[X] = \sum_{k=1}^{N} k \cdot \frac{1}{N} = \frac{N+1}{2} \approx \frac{N}{2}$.
  - Complexity: $\mathcal{O}(N)$ evaluations of f.
- **Quantum Approach (Grover's Algorithm):**
  - Uses quantum superposition and interference to locate ω with high probability.
  - Complexity: $\mathcal{O}(\sqrt{N})$ evaluations of the quantum oracle.
  - This *quadratic speedup* over the best possible classical algorithm is provably optimal (BBBV theorem, Section 7).

The Quantum Oracle

- We evaluate $f(x)$ via a reversible **oracle** U_f acting on a data register $|x\rangle$ and an ancilla qubit $|y\rangle$:

$$U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle \, . \tag{1}$$

  - $x = \omega$: ancilla flips ($f(x) = 1$). $x \neq \omega$: ancilla unchanged ($f(x) = 0$).

- **Quantum parallelism:** initialize the data register to $|\phi\rangle = \frac{1}{\sqrt{N}} \sum_x |x\rangle$. Linearity gives:

$$U_f(|\phi\rangle |y\rangle) = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle |y \oplus f(x)\rangle \, . \tag{2}$$

A single oracle call entangles every $|x\rangle$ with $|f(x)\rangle$ simultaneously. The challenge is extracting useful information via measurement — which is where phase kickback enters.

The Grover Oracle (Phase Kickback)

- Prepare the auxiliary qubit in $|-\rangle = H|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$.
- Applying U_f yields the **Phase Kickback** effect:

$$U_f(|x\rangle|-\rangle) = |x\rangle\left(\frac{|0 \oplus f(x)\rangle - |1 \oplus f(x)\rangle}{\sqrt{2}}\right) = (-1)^{f(x)}|x\rangle|-\rangle\,.$$

- The auxiliary qubit is unchanged; the phase $(-1)^{f(x)}$ is *kicked back* onto the data register. Abbreviate the action on the n-qubit data register alone:

$$U_\omega|x\rangle = (-1)^{f(x)}|x\rangle\,. \tag{3}$$

  - U_ω flips the sign of the amplitude of $|\omega\rangle$ only, leaving all other states unchanged.

The Oracle in Matrix Representation

- The phase oracle U_ω can be written explicitly as:

$$U_\omega = I - 2\,|\omega\rangle\,\langle\omega|\,. \tag{4}$$

- Verify its action:
  - ◇ For $x \neq \omega$: $|x\rangle \perp |\omega\rangle$, so $U_\omega\,|x\rangle = |x\rangle - 2\,|\omega\rangle\,\underbrace{\langle\omega|x\rangle}_{=0} = |x\rangle\,.$
  - ◇ For $x = \omega$: $\langle\omega|\omega\rangle = 1$, so $U_\omega\,|\omega\rangle = |\omega\rangle - 2\,|\omega\rangle = -\,|\omega\rangle\,.$
- In linear algebra, $U_\omega = I - 2\,|\omega\rangle\,\langle\omega|$ is a **Householder reflector**: it reflects any vector across the hyperplane orthogonal to $|\omega\rangle$, reversing the component along $|\omega\rangle$ and leaving all orthogonal components unchanged. Since all $|x\rangle$ with $x \neq \omega$ are orthogonal to $|\omega\rangle$ (orthonormal basis), U_ω acts as $-I$ on $|\omega\rangle$ and as I on everything else — exactly what we need.

State Preparation and the Diffusion Operator

- Initialize the n-qubit data register to the uniform superposition:

$$|\phi\rangle := H^{\otimes n} |0\rangle_n = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle . \tag{5}$$

- Construct the n-qubit **diffusion operator**:

$$W := 2 |\phi\rangle \langle\phi| - I. \tag{6}$$

- W is also a Householder reflector, this time reflecting *about* $|\phi\rangle$:

$$WH^{\otimes n} |x\rangle = \begin{cases} |\phi\rangle , & \text{if } |x\rangle = |0\rangle , \\ -H^{\otimes n} |x\rangle , & \text{if } |x\rangle \neq |0\rangle . \end{cases} \tag{7}$$

  - ◇ If $|x\rangle \neq |0\rangle$, then $|x\rangle \perp |0\rangle$, so $H^{\otimes n} |x\rangle \perp H^{\otimes n} |0\rangle = |\phi\rangle$. Hence W reflects $H^{\otimes n} |x\rangle$ through the hyperplane orthogonal to $|\phi\rangle$, giving $-H^{\otimes n} |x\rangle$.

The Grover Iterate (G)

- Define the *Grover iterate* as the $(n+1)$-qubit unitary:

$$G := (W \otimes I_2)\, U_f\,, \tag{8}$$

whose effective action on the n-qubit data register is:

$$G_{\text{eff}} = WU_\omega. \tag{9}$$

- The Grover algorithm applies G approximately $\frac{\pi}{4}\sqrt{N}$ times, then measures:

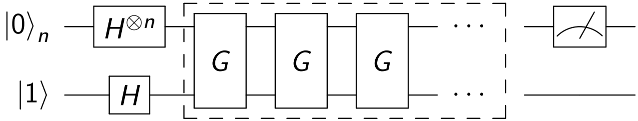

- The goal: use the two reflections (U_ω and W) to *amplify* the amplitude of $|\omega\rangle$ step by step, without ever knowing ω explicitly.

The Two Reflectors: Basic Mechanics

- The two operators act as follows on $|\omega\rangle$ and $|\phi\rangle$:

$$U_\omega |\omega\rangle = -|\omega\rangle\,, \qquad U_\omega |\phi\rangle = |\phi\rangle - \frac{2}{\sqrt{N}}|\omega\rangle\,,$$

$$W|\phi\rangle = |\phi\rangle\,, \qquad W|\omega\rangle = \frac{2}{\sqrt{N}}|\phi\rangle - |\omega\rangle\,.$$

- **Interpretation:**
  - ◇ U_ω negates $|\omega\rangle$ (marks the target) and leaves $|\phi\rangle$ almost unchanged (subtracting only a tiny $\frac{2}{\sqrt{N}}|\omega\rangle$ component).
  - ◇ W fixes $|\phi\rangle$ (it is the reflection axis) and reflects $|\omega\rangle$ slightly toward $|\phi\rangle$.
  - ◇ Together, each application of $G_{\text{eff}} = WU_\omega$ rotates the state a small angle $\theta \approx 2/\sqrt{N}$ closer to $|\omega\rangle$, as the next section makes precise.

Analysis: One Step

- Before the first application of G:

$$|0\rangle_n |1\rangle \;\longrightarrow\; |\phi\rangle |-\rangle = \frac{1}{\sqrt{N}} \sum_x |x\rangle |-\rangle .$$

- After U_f (phase kickback on the $|-\rangle$ ancilla):

$$\longrightarrow\; \left(|\phi\rangle - \frac{2}{\sqrt{N}} |\omega\rangle \right) |-\rangle .$$

- After W (apply the diffusion operator):

$$\begin{aligned} &\longrightarrow\; \left(\left(1 - \frac{4}{N}\right) |\phi\rangle + \frac{2}{\sqrt{N}} |\omega\rangle \right) |-\rangle \\ &= \left(\left(\frac{1}{\sqrt{N}} - \frac{4}{\sqrt{N^3}} \right) \sum_{x \neq \omega} |x\rangle + \left(\frac{3}{\sqrt{N}} - \frac{4}{\sqrt{N^3}} \right) |\omega\rangle \right) |-\rangle . \end{aligned}$$

- After one step, the amplitude of $|\omega\rangle$ increases from $\frac{1}{\sqrt{N}}$ to $\frac{3}{\sqrt{N}} - \frac{4}{\sqrt{N^3}} \approx \frac{3}{\sqrt{N}}$, while total probability is conserved.

Analysis: Repeated Steps

- Write the general state as $|\Psi\rangle = s\,|\phi\rangle + t\,|\omega\rangle$, where normalization requires:

$$\left(\frac{s}{\sqrt{N}}\right)^2 (N-1) + \left(\frac{s}{\sqrt{N}} + t\right)^2 = 1.$$

- Applying $G_{\text{eff}} = WU_\omega$:

$$\begin{aligned} WU_\omega\,|\Psi\rangle &= W\Big(s\Big(|\phi\rangle - \tfrac{2}{\sqrt{N}}\,|\omega\rangle\Big) - t\,|\omega\rangle\Big) \\ &= s\,|\phi\rangle - \Big(t + \tfrac{2s}{\sqrt{N}}\Big)\Big(\tfrac{2}{\sqrt{N}}\,|\phi\rangle - |\omega\rangle\Big) \\ &= \Big(s - \tfrac{4s}{N} - \tfrac{2t}{\sqrt{N}}\Big)\,|\phi\rangle + \Big(t + \tfrac{2s}{\sqrt{N}}\Big)\,|\omega\rangle\,. \end{aligned}$$

- The amplitude of $|\omega\rangle$ increases by $\frac{2s}{\sqrt{N}}$ at each step, driving the state steadily toward the target.

Probability Improvement per Step

- The probability of measuring $|\omega\rangle$ evolves as:

$$\begin{cases} P_{\text{before}} = \left(\frac{s}{\sqrt{N}} + t\right)^2, \\ P_{\text{after}} = \left(\frac{3s}{\sqrt{N}} + t - \frac{t}{2^{n-1}} - \frac{4s}{\sqrt{N^3}}\right)^2. \end{cases} \tag{10}$$

- After algebraic simplification:

$$P_{\text{after}} \approx P_{\text{before}} + \frac{4}{\sqrt{N}}\sqrt{P_{\text{before}}(1 - P_{\text{before}})}.$$

  - ◇ The per-step gain ΔP is maximized when $P_{\text{before}} = 1/2$, giving $\Delta P_{\max} = \frac{1}{\sqrt{N}}$.
  - ◇ Since the total probability range is $[0, 1]$ and each step adds at most $1/\sqrt{N}$, the algorithm needs at least $\mathcal{O}(\sqrt{N})$ iterations to drive P from ≈ 0 to ≈ 1.

Dynamics of Probabilities

- Starting from $s_0 = 1$, $t_0 = 0$, the recurrence is:

$$\begin{pmatrix} s_{k+1} \\ t_{k+1} \end{pmatrix} = \underbrace{\begin{pmatrix} 1 - \frac{4}{N} & -\frac{2}{\sqrt{N}} \\ \frac{2}{\sqrt{N}} & 1 \end{pmatrix}}_{=: M} \begin{pmatrix} s_k \\ t_k \end{pmatrix}. \tag{11}$$

  - ◇ One can show the invariant $s^2 + \frac{2st}{\sqrt{N}} + t^2 = 1$ holds for all k (an ellipse).
- The eigenvalues of M are $e^{\pm i\theta}$ where $\theta = 2 \arcsin\left(\frac{1}{\sqrt{N}}\right) \approx \frac{2}{\sqrt{N}}$.
  - ◇ Both have modulus 1, so the iterates (s_k, t_k) cycle around the ellipse without converging.
  - ◇ The state is periodic with period $\approx \pi/\theta \approx \frac{\pi}{2}\sqrt{N}$: at $k \approx \frac{\pi}{4}\sqrt{N}$ the probability of $|\omega\rangle$ reaches its maximum; at $k \approx \frac{\pi}{2}\sqrt{N}$ it returns to near zero.
  - ◇ **This is why** iterating more than $\mathcal{O}\left(\frac{\pi}{4}\sqrt{N}\right)$ times is counterproductive: the state over-rotates past $|\omega\rangle$ and the success probability falls.

Geometric Interpretation: The 2D Subspace

- Although the Hilbert space has dimension $N = 2^n$, the entire search dynamics unfold within a **real 2-dimensional subspace** spanned by:
  - ◇ $|\beta\rangle = |\omega\rangle$ (the target state),
  - ◇ $|\alpha\rangle = \frac{1}{\sqrt{N-1}} \sum_{x \neq \omega} |x\rangle$ (the uniform superposition over all non-targets).
- The initial state $|\phi\rangle$ lies in this plane:

$$|\phi\rangle = \cos\left(\frac{\theta}{2}\right)|\alpha\rangle + \sin\left(\frac{\theta}{2}\right)|\beta\rangle, \quad \sin\left(\frac{\theta}{2}\right) = \frac{1}{\sqrt{N}}. \qquad (12)$$

  - ◇ Since $\sin(\theta/2) = 1/\sqrt{N} \ll 1$ for large N, we have $\theta/2 \approx 1/\sqrt{N}$, hence $\theta \approx 2/\sqrt{N}$. The state $|\phi\rangle$ starts very close to the non-target axis $|\alpha\rangle$ — the target has negligible initial amplitude.

The Mechanics of Rotation

- Within the 2D subspace $\{|\alpha\rangle, |\beta\rangle\}$, each operator acts geometrically:
  - ◇ U_ω **reflects** the state across the $|\alpha\rangle$ axis (it negates the $|\beta\rangle = |\omega\rangle$ component).
  - ◇ W **reflects** the state across the $|\phi\rangle$ axis (it fixes $|\phi\rangle$, which lies at angle $\theta/2$ from $|\alpha\rangle$).
- **Theorem:** the composition of two reflections whose axes meet at angle φ is a rotation by 2φ.
  - ◇ *Proof sketch:* reflecting $\mathbf{v}$ across the x-axis then across a line at angle φ is equivalent to rotating $\mathbf{v}$ by 2φ counterclockwise (each reflection reverses the sign of the angle to the respective axis; composing the two sign-reversals gives a net rotation).
  - ◇ The axes of U_ω ($|\alpha\rangle$) and W ($|\phi\rangle$) meet at angle $\theta/2$.
- Therefore $G_{\text{eff}} = WU_\omega$ rotates the state by exactly θ toward $|\beta\rangle$ at every step.

Iterative Evolution and Optimality

- Since each application of G_{eff} is a rotation by θ in the $\{|\alpha\rangle, |\beta\rangle\}$ plane, after k steps:

$$G_{\text{eff}}^k |\phi\rangle = \cos\left(\frac{2k+1}{2}\theta\right)|\alpha\rangle + \sin\left(\frac{2k+1}{2}\theta\right)|\beta\rangle . \tag{13}$$

  - *Derivation:* $|\phi\rangle$ starts at angle $\theta/2$ from $|\alpha\rangle$; each step adds θ, giving total angle $(2k+1)\theta/2$ from $|\alpha\rangle$ after k steps. This follows by induction using the addition formula $\sin(A+\theta) = \sin A \cos\theta + \cos A \sin\theta$ at each step.

- Measuring $|\omega\rangle$ with near certainty requires $\sin^2\left(\frac{2k+1}{2}\theta\right) \approx 1$, i.e.:

$$\frac{2k+1}{2}\theta \approx \frac{\pi}{2} \implies k_{\text{opt}} \approx \frac{\pi}{2\theta} - \frac{1}{2} \approx \frac{\pi}{4}\sqrt{N}. \tag{14}$$

- The explicit construction of G maps an $\mathcal{O}(N)$ classical search directly into an $\mathcal{O}(\sqrt{N})$ quantum rotation — a genuine quadratic speedup.

The Danger of Over-Rotation

- Unlike classical search, where checking more items always increases confidence, quantum search is a *continuous rotation* in the 2D subspace.
- If $k > k_{\text{opt}}$, the state vector rotates *past* $|\beta\rangle$ and begins heading back toward $|\alpha\rangle$.
  - ◇ Success probability is periodic: $P(\text{success}) = \sin^2\left(\frac{2k+1}{2}\theta\right)$.
  - ◇ It drops back to near zero at $k \approx \frac{\pi}{2}\sqrt{N}$ (half a full cycle past $|\alpha\rangle$), and oscillates indefinitely.
- The algorithm must stop at *exactly* $k_{\text{opt}} \approx \frac{\pi}{4}\sqrt{N}$ iterations.
  - ◇ In practice, if N is not known precisely (e.g., when $M > 1$), quantum counting techniques can be used to estimate k_{opt} before running the search.
- Are there ways to accelerate convergence? Fixed-point variants of Grover's algorithm (e.g., using phase rotations other than π) exist and avoid over-rotation at the cost of a small constant factor.

Generalization: M Target Solutions

- Suppose the database has $M \geq 1$ valid solutions.
- The structure of the algorithm is identical; only the definitions of $|\alpha\rangle$ and $|\beta\rangle$ change:
  - ◇ $|\beta\rangle = \frac{1}{\sqrt{M}} \sum_{\omega} |\omega\rangle$ (uniform superposition over all targets),
  - ◇ $|\alpha\rangle = \frac{1}{\sqrt{N-M}} \sum_{x \notin \{\omega\}} |x\rangle$ (uniform superposition over non-targets).
- The initial angle satisfies:

$$\sin\left(\frac{\theta}{2}\right) = \sqrt{\frac{M}{N}}. \tag{15}$$

- The rotation per step is still θ, giving optimal iterations:

$$k_{\text{opt}} \approx \frac{\pi}{4}\sqrt{\frac{N}{M}}. \tag{16}$$

- Measuring the final state yields one of the M solutions uniformly at random.

Optimality of Grover's Algorithm

- **Can we do better than $\mathcal{O}(\sqrt{N})$?** No.
- **BBBV Theorem** (Bennett, Bernstein, Brassard, Vazirani, 1997): Any quantum algorithm that identifies a target in an unstructured search space of size N must query the oracle at least $\Omega(\sqrt{N})$ times. Grover's algorithm is therefore *provably optimal*.
- **Broader Implications:**
  - ◇ **Amplitude amplification** generalizes Grover's algorithm: if a classical subroutine $\mathcal{A}$ finds a "good" solution with probability p, then quantum amplitude amplification finds one in $\mathcal{O}(1/\sqrt{p})$ evaluations of $\mathcal{A}$ — a quadratic speedup over the $\mathcal{O}(1/p)$ classical repetitions needed. Grover's algorithm is the special case $\mathcal{A} = H^{\otimes n}$, $p = M/N$.
  - ◇ This provides a generic quadratic speedup for any NP problem solvable by exhaustive search, including SAT, graph coloring, and combinatorial optimization.
  - ◇ **Important caveat:** $\mathcal{O}(\sqrt{N})$ is still exponential in $n = \log_2 N$ bits. Grover does *not* place NP in BQP; it merely reduces the constant factor.

Summary of Lecture 9

- **Unstructured search** is classically $\mathcal{O}(N)$; Grover's algorithm achieves $\mathcal{O}(\sqrt{N})$ — provably optimal by the BBBV theorem.
- **Phase kickback** via $U_\omega = I - 2\left|\omega\right\rangle\left\langle\omega\right|$ marks the target by flipping its amplitude sign.
- **Diffusion operator** $W = 2\left|\phi\right\rangle\left\langle\phi\right| - I$ reflects the state about the uniform superposition, amplifying the marked amplitude.
- **Geometrically,** $G_{\text{eff}} = WU_\omega$ is a rotation by $\theta \approx 2/\sqrt{N}$ in the 2D plane $\{\left|\alpha\right\rangle, \left|\beta\right\rangle\}$; after $k_{\text{opt}} \approx \frac{\pi}{4}\sqrt{N}$ steps the state reaches $\left|\omega\right\rangle$ with near certainty.
- **Over-rotation is fatal:** success probability is periodic; the algorithm must stop at exactly k_{opt}.
- **Amplitude amplification:** any subroutine with success probability p is boosted to $\mathcal{O}(1/\sqrt{p})$ queries. Grover does *not* place NP in BQP — $\mathcal{O}(\sqrt{N})$ is still exponential in n bits.

Looking Ahead: Lecture 10

- Grover's algorithm shows that quantum mechanics can speed up *search* quadratically. The next question: can it speed up *linear algebra*?
- **Quantum Linear Systems (HHL Algorithm):** given a sparse $N \times N$ matrix A and vector $\mathbf{b}$, prepare the quantum state $|x\rangle \propto A^{-1}|b\rangle$ in time $\mathcal{O}(\log N)$ — an exponential speedup over classical Gaussian elimination.
- Key prerequisites from this course that Lecture 10 will build on: phase estimation (Lecture 7), Hamiltonian simulation (Lecture 6), and the block-encoding framework.