




**Manish Kesarwani**

**Senior Research Scientist  
IBM Research, India  
Ph.D. Student, IISc**




# Relational Database Engines on Quantum Platforms



**Jayant Haritsa**

**Database Systems Lab  
Indian Institute of Science**

# Tutorial Overview



**Stage I: Quantum Computing  
Fundamentals**



**Stage II: Quantum Computing  
Challenges & Benefits**



**Stage III: Quantum Optimization  
for DBMS**



**Stage IV: Quantum Query  
Execution**



**Stage V: Demo + Hands-On  
Exercise (Afternoon session)**

# Thanks



**Sameep Mehta**  
(IBM Research, India)



**Sieglinde Pfaendler**  
(IBM Quantum, Zürich)



**Immanuel Trummer**  
(Cornell University, USA)



**Wolfgang Mauerer**  
(OTH Regensburg, Germany)




**Rihan Hai**  
(TU Delft, Netherlands)



**Manuel Schönberger**  
(OTH Regensburg, Germany)



**Sven Groppe**  
(University of Lübeck, Germany)




**SheshaShayee Raghunathan**  
(IBM Quantum, India)

# 2025 ACM Turing Award

For their *essential role in establishing the foundations of quantum information science* and transforming secure communication and computing



**Charles H. Bennett**



**Gilles Brassard**

# Nobel Prize in Physics 2025

For the **discovery of macroscopic quantum mechanical tunnelling** and energy quantisation in an electric circuit



**John Clarke**

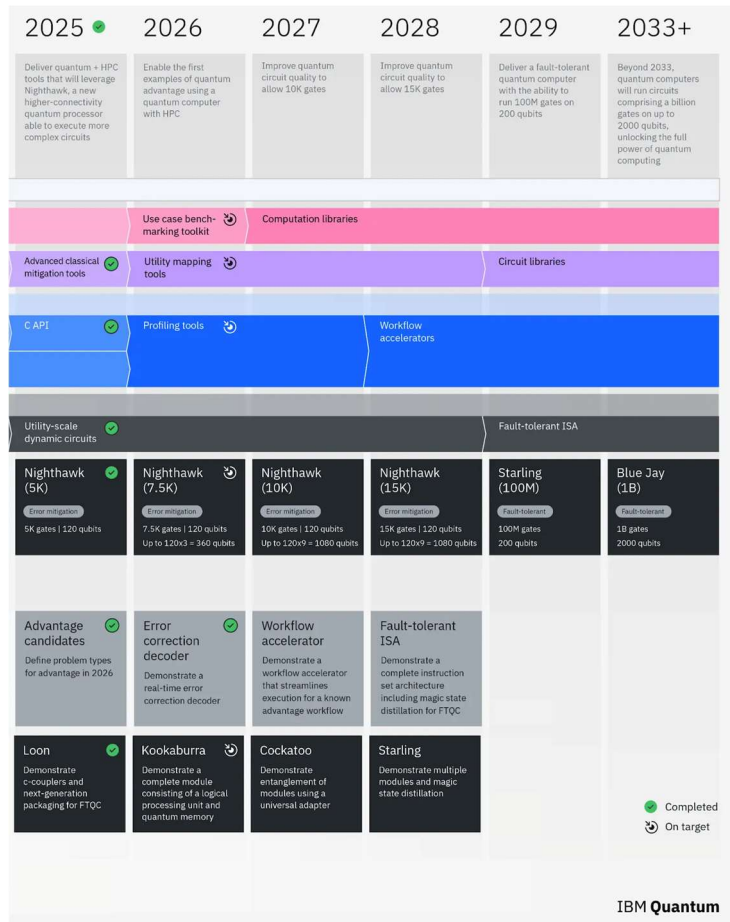


**Michel H. Devoret**



**John M. Martinis**

# IBM Quantum Roadmap



**Fault-Tolerant By 2029 !**

**IBM Starling [2029]  
200 Qubits, 100M Gates**

**IBM Blue Jay [2033]  
2000 Qubits, 1B Gates**

# Should we invest in Quantum Now?

## Timelines for business advantage from quantum computing

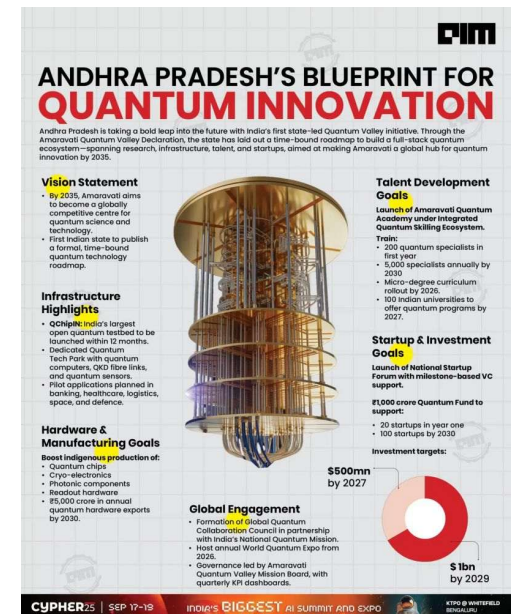


Source: Gartner  
© 2024 Gartner, Inc. and/or its affiliates. All rights reserved. 3188863

**Gartner**

# Government Investments

- **USA:** Recently announced equity stakes tied to **\$2B in CHIPS Act grants**
  - Half the grant money to IBM, for creating a new quantum chip foundry called **Anderon**
  - Other **eight** recipients include GlobalFoundries, D-Wave, Rigetti, Infleqtion, Quantinuum, PsiQuantum, Atom Computing, and Diraq
- **FRANCE:** Announced investment of **~\$1.2B** in quantum technologies
- **INDIA:**
  - **National Quantum Mission: \$600M** for scientific and industrial R&D
  - **Karnataka Quantum Mission: \$100M** to transform state into the “quantum capital of Asia” by 2035
  - **Amaravati (Andhra Pradesh): \$1B** quantum technology park and hub



*Stage I: Quantum  
Computing Fundamentals*

---

# *1.1 Concepts*

# 1. Qubit

- Binary Digit (**Bit**) is the fundamental concept in Classical Computing
  - It is **deterministic**: always either 0 or 1
  - Realized using tiny electronic switches called **transistors**.




Bit  $\in \{0,1\}$

- Quantum computation is built upon the *quantum bit* or *qubit*
  - A qubit can exist in a **continuum of states**  $\psi$  between  $|0\rangle$  and  $|1\rangle$

$| \rangle$  → Dirac notation  
Pronounced as **KET**


Qubit Technology	Example Companies
Superconducting	IBM, Google, Rigetti
Trapped Ion	IonQ, Quantinuum
Photonic	Xanadu, PsiQuantum
Neutral Atom	QuEra, Pasqal
Silicon Spin	Intel



# 2. Quantum Superposition

- **Quantum superposition**

- A qubit can exist in a continuum of states  $\psi$  between  $|0\rangle$  and  $|1\rangle$
- $\alpha$  and  $\beta$  are complex numbers such that  $|\alpha|^2 + |\beta|^2 = 1$ , and
- $\phi \in [0, 2\pi]$  is the quantum phase (angle of rotation around the Z-axis).

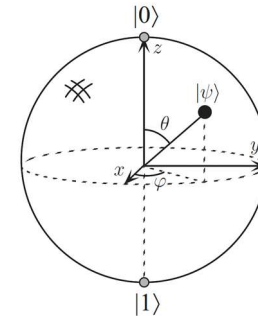


$$|\psi\rangle = \alpha|0\rangle + e^{i\phi}\beta|1\rangle$$

(Superposition)

- **The Bloch Sphere**

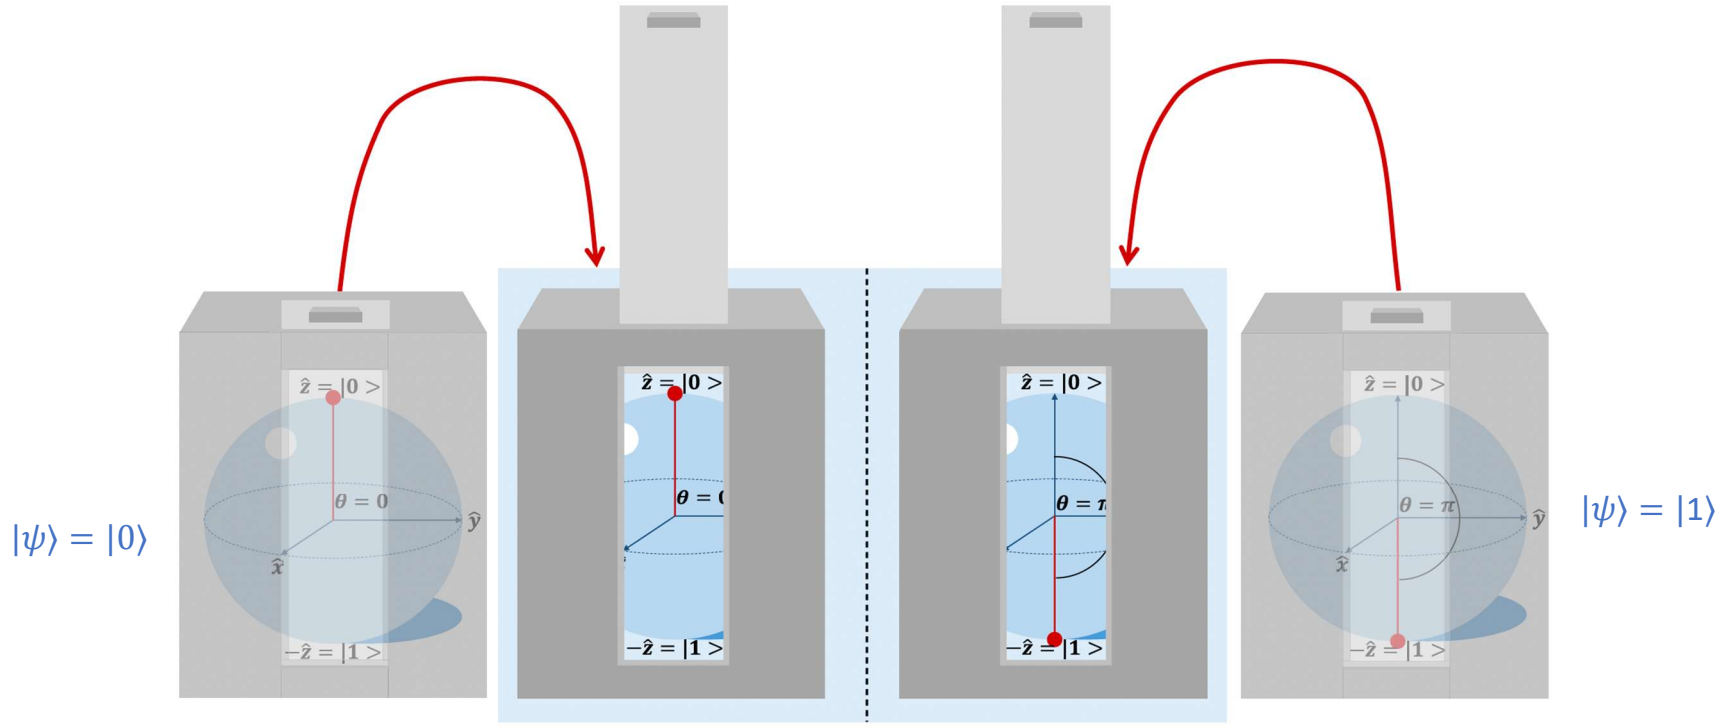
- Qubits are represented as vectors pointing to any location on a 3D sphere.
- **Poles:** Represent the classical basis states  $|0\rangle$  and  $|1\rangle$
- **Surface:** Points on the surface represent pure superposition states.
- Probability amplitude  $\theta$  = vertical tilt
- Relative phase  $\phi$  = rotation around the equator



$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle$$

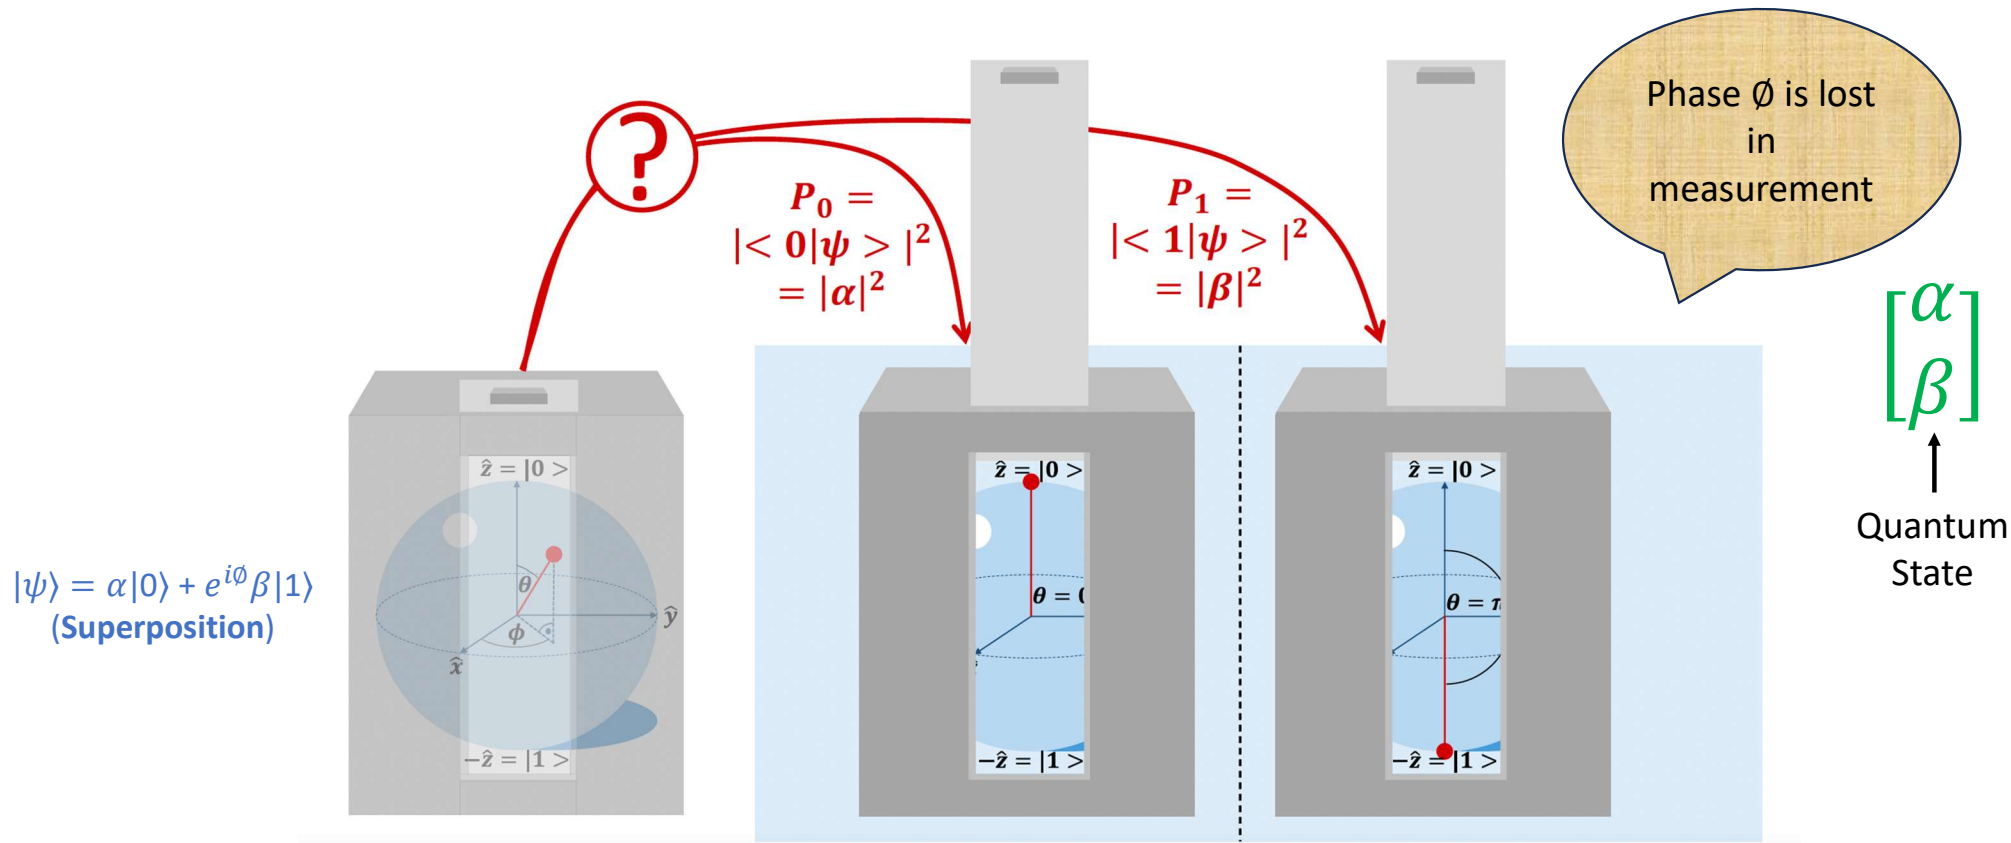
# 3. Quantum Measurement/Observation

➤ The state is **NOT** destroyed by a measurement for state  $|0\rangle$  and  $|1\rangle$



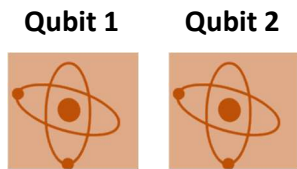
# 3. Quantum Measurement/Observation

- A superposition state **collapses** to  $|0\rangle$  or  $|1\rangle$  according to corresponding probabilities



# 4. Quantum Entanglement

- It is a non-local property of two or more qubits
- Higher correlation than is possible in classical system
- Consider two qubits in the following quantum state:


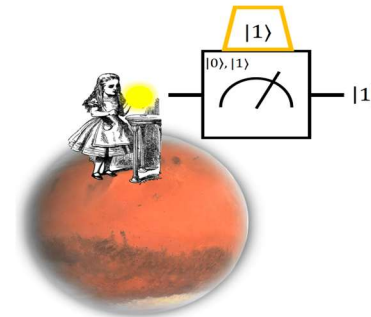


$$|\psi\rangle = \frac{1}{\sqrt{2}} |00\rangle + \frac{1}{\sqrt{2}} |11\rangle,$$
$$P(|00\rangle) = \frac{1}{2} \text{ and } P(|11\rangle) = \frac{1}{2}$$

Both qubits are correlated

- if you measure one as 0, the other becomes 0
  - if you measure one as 1, the other becomes 1
- BUT before measurement:** neither qubit individually has a definite classical value

- The two coins are not independently spinning anymore — they become one connected quantum object.
- If entangled qubits are separated given to Alice and Bob
- If after measurement Alice gets  $|0\rangle$
- Then Bob will get  $|0\rangle$  no matter when he measures

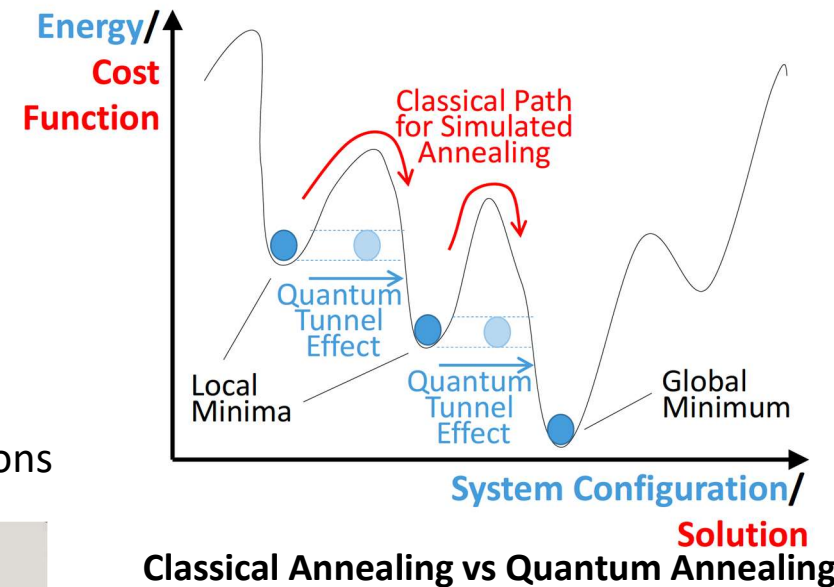



# 5. Quantum Tunneling

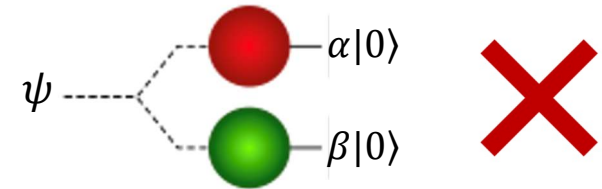
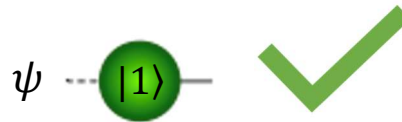
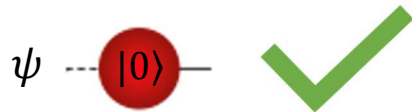
- In quantum mechanics,
  - particles are described as wavefunctions
  - the particle decays inside the barrier and
  - there is a probability of finding the particle on the other side

*The process where a particle "tunnels" through the barrier to appear on the other side is quantum tunneling*

- Without gaining additional energy, contrary to classical expectations



# 6. Qubit Cloning



## No Cloning Theorem [1982]

Given: Two quantum states  $|\psi\rangle$  and  $|\emptyset\rangle$

**Assumption:** if there exists a universal cloning machine



**Steps:**

1.  $U(|\psi\rangle|s\rangle) = |\psi\rangle|\psi\rangle$
2.  $U(|\emptyset\rangle|s\rangle) = |\emptyset\rangle|\emptyset\rangle$
3. Inner product of above:  $\langle\psi|\emptyset\rangle = (\langle\psi|\emptyset\rangle)^2$

**Conclusion:** Cloning is impossible for arbitrary states.

No quantum operation can perfectly copy an arbitrary unknown quantum state !

# *1.2 Quantum Gates*

# Quantum Gate

- Suppose the initial quantum state is:  $|\psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}$  ← State Vector
- A quantum gate ( $U$ ) transforms one quantum state into another  $|\psi'\rangle = U |\psi\rangle$ 
  - Initial Probabilities:  $|\alpha|^2 + |\beta|^2 = 1$
  - After applying a quantum gate, **the final probabilities must still sum to 1**
  - Quantum gates must preserve the length of the state vector  $\|U |\psi\rangle\| = \| |\psi\rangle \|$

- Can every classical gate become a quantum gate?

The evolution of an isolated quantum system is reversible

- Mathematically this means:  $U^\dagger U = I$
- Therefore  $U$  should be a **unitary matrix**.

Any unitary matrix is a valid quantum gate!

# Quantum NOT Gate

As two-by-two unitary matrix

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

$$\psi = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}, \text{ then } X \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \beta \\ \alpha \end{bmatrix}$$

$$\alpha|0\rangle + \beta|1\rangle \rightarrow \beta|0\rangle + \alpha|1\rangle$$

# Standard Single Qubit Quantum Gates

- **Pauli-Z Gate:** Leaves  $|0\rangle$  unchanged, but flips the sign of  $|1\rangle$

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

- **Hadamard (H) Gate:** It turns  $|0\rangle \rightarrow \frac{(|0\rangle+|1\rangle)}{\sqrt{2}}$  (halfway between  $|0\rangle$  and  $|1\rangle$ ), and  $|1\rangle \rightarrow \frac{(|0\rangle-|1\rangle)}{\sqrt{2}}$  (also halfway between  $|0\rangle$  and  $|1\rangle$ )

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

- **Phase (P) Gate:** Takes a real number  $\phi$  and performs a rotation of qubit by  $\phi$  radians around the Z-axis

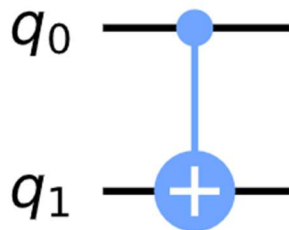
$$P(\phi) = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{bmatrix}$$

# Multi Qubit Gate: CNOT Gate

- Recall: a single qubit state is represented as:  $|\psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \alpha|0\rangle + \beta|1\rangle$
- For two qubits  $|\psi_0\rangle = \begin{bmatrix} \alpha_0 \\ \beta_0 \end{bmatrix}$  and  $|\psi_1\rangle = \begin{bmatrix} \alpha_1 \\ \beta_1 \end{bmatrix}$

$$|\psi_0\psi_1\rangle = |\psi_0\rangle \otimes |\psi_1\rangle = \begin{bmatrix} \alpha_0\alpha_1 \\ \alpha_0\beta_1 \\ \beta_0\alpha_1 \\ \beta_0\beta_1 \end{bmatrix} = \alpha_0\alpha_1|00\rangle + \alpha_0\beta_1|01\rangle + \beta_0\alpha_1|10\rangle + \beta_0\beta_1|11\rangle$$














- CNOT Gate:** Performs an X-gate on the second qubit (target), if the state of the first qubit (control) is  $|1\rangle$



Input (t,c)	Output (t,c)
00	00
01	11
10	10
11	01

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$


# Summary of Quantum Gates

Operator	Gate(s)	Matrix
Pauli-X (X)	 	$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$
Pauli-Y (Y)		$\begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$
Pauli-Z (Z)		$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$
Hadamard (H)		$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$
Phase (S, P)		$\begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}$
$\pi/8$ (T)		$\begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}$
Controlled Not (CNOT, CX)		$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$
Controlled Z (CZ)	 	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$
SWAP	 	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
Toffoli (CCNOT, CCX, TOFF)		$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$

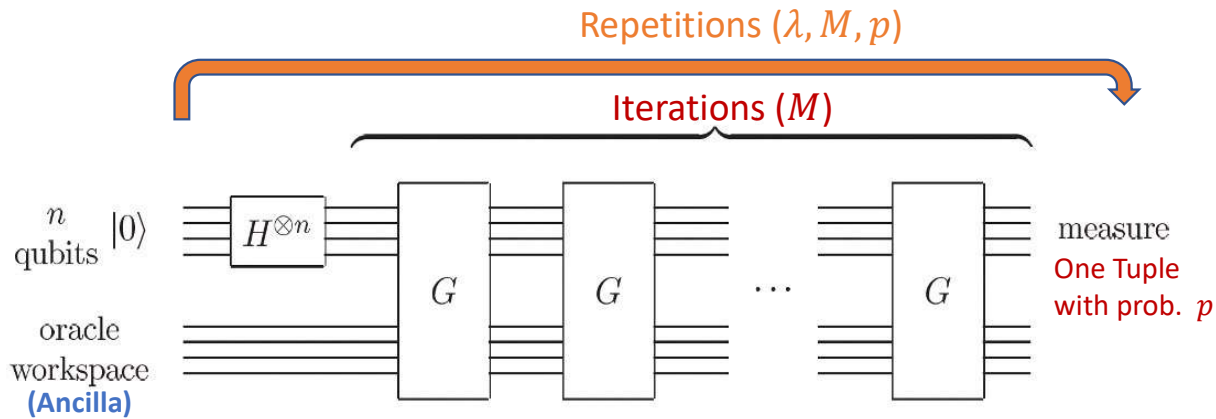
# ***1.3 Quantum Algorithms***

# Grover Search[1996]

Given a large unordered list of  $N$  items, identifies the items of interest **quadratically faster** than the best classical algorithm



# Overview of Grover's Search



$$\left. \begin{aligned} k_j &= \frac{1}{\sqrt{t}} \sin((2j+1)\theta) \\ \ell_j &= \frac{1}{\sqrt{N-t}} \cos((2j+1)\theta) \end{aligned} \right\}$$

Say we need  $R$  iterations to detect the answer, then at solution,  $k_m$  will be close to  $\frac{1}{\sqrt{t}}$  while  $l_m$  will be close to 0

Figure 6.1. Schematic circuit for the quantum search algorithm. The oracle may employ work qubits for its implementation, but the analysis of the quantum search algorithm involves only the  $n$  qubit register.

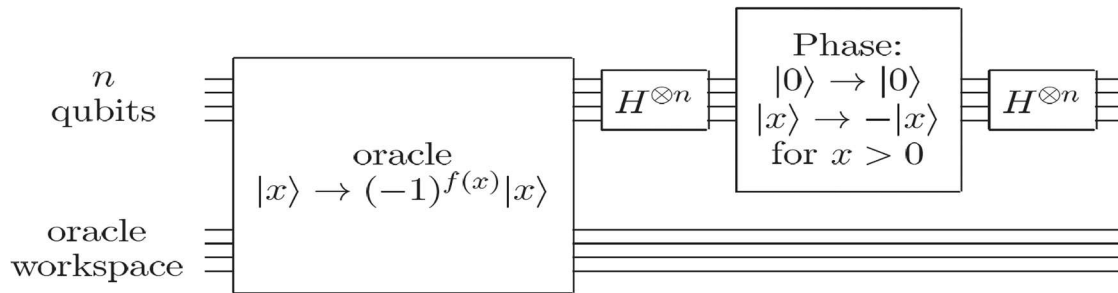
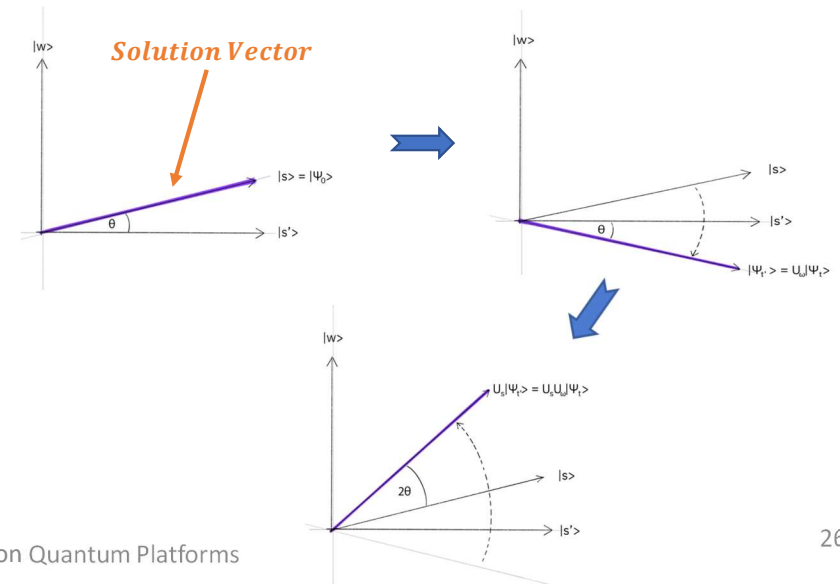


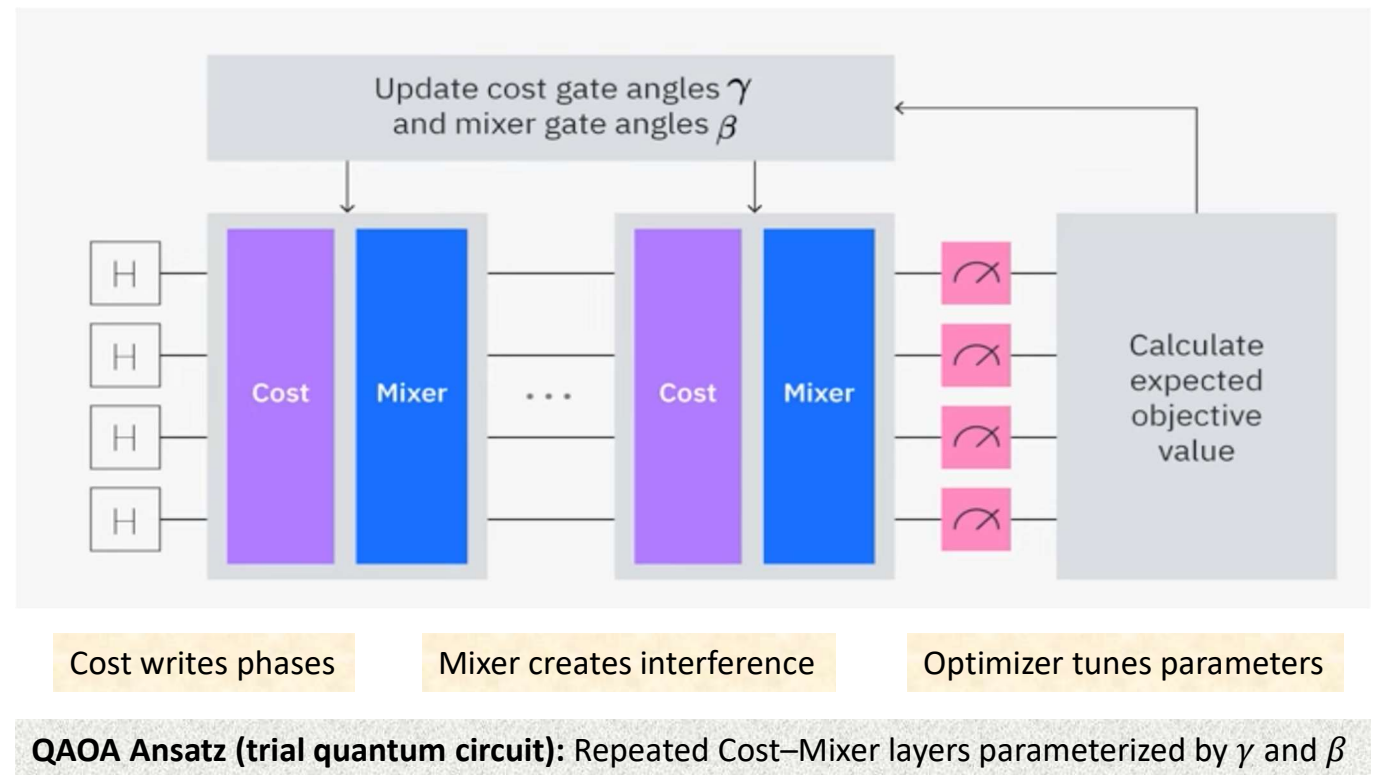
Figure 6.2. Circuit for the Grover iteration,  $G$ .



# Quantum Approximate Optimization Algorithm (QAOA)

- **Proposed By:** Edward Farhi, Jeffrey Goldstone, and Sam Gutmann in 2014
- **Engineers interference patterns** so that low-cost solutions become more likely after measurement
- **Input:** Quadratic Unconstrained Binary Optimization (QUBO)
- **Output:** Approximate solution
- **Approach:** Quantum-classical hybrid
- **Quantum Part:**
  - Encodes optimization cost into quantum phases, and
  - Redistributes amplitudes to create quantum interference
  - Repeats Cost–Mixer layers  $p$  times
- **Classical Part:**
  - Updates  $\gamma, \beta$  to improve the interference pattern

$\gamma_k$ : Cost Hamiltonian evolution at layer  $k$   
 $\beta_k$ : Amplitude redistribution strength at layer  $k$   
 $p$ : Number of Cost–Mixer layers



# ***1.4 Quantum Platforms***

---

**QUANTUM ANNEALING**

**D:wave**

**Topological**

**Adiabatic**

**Measurement Based**

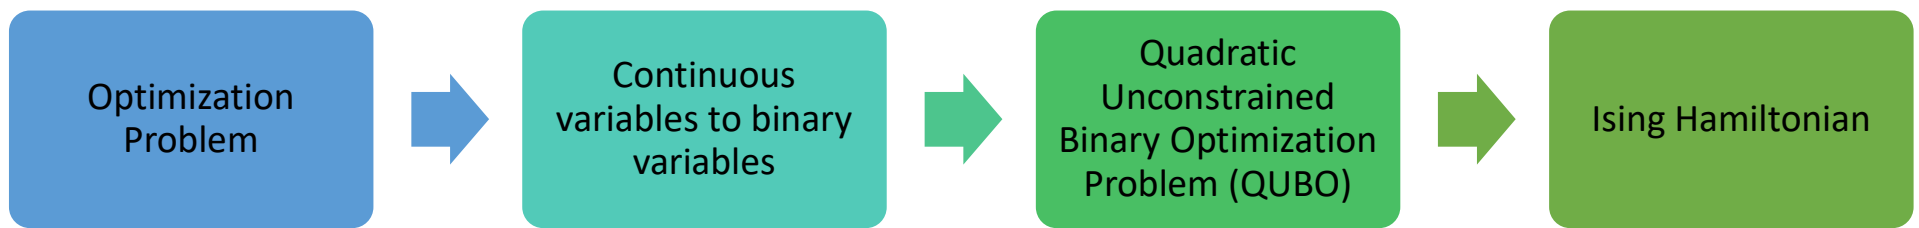
**Gate Model**

**UNIVERSAL QUANTUM COMPUTERS**

# Annealing Based: DWave

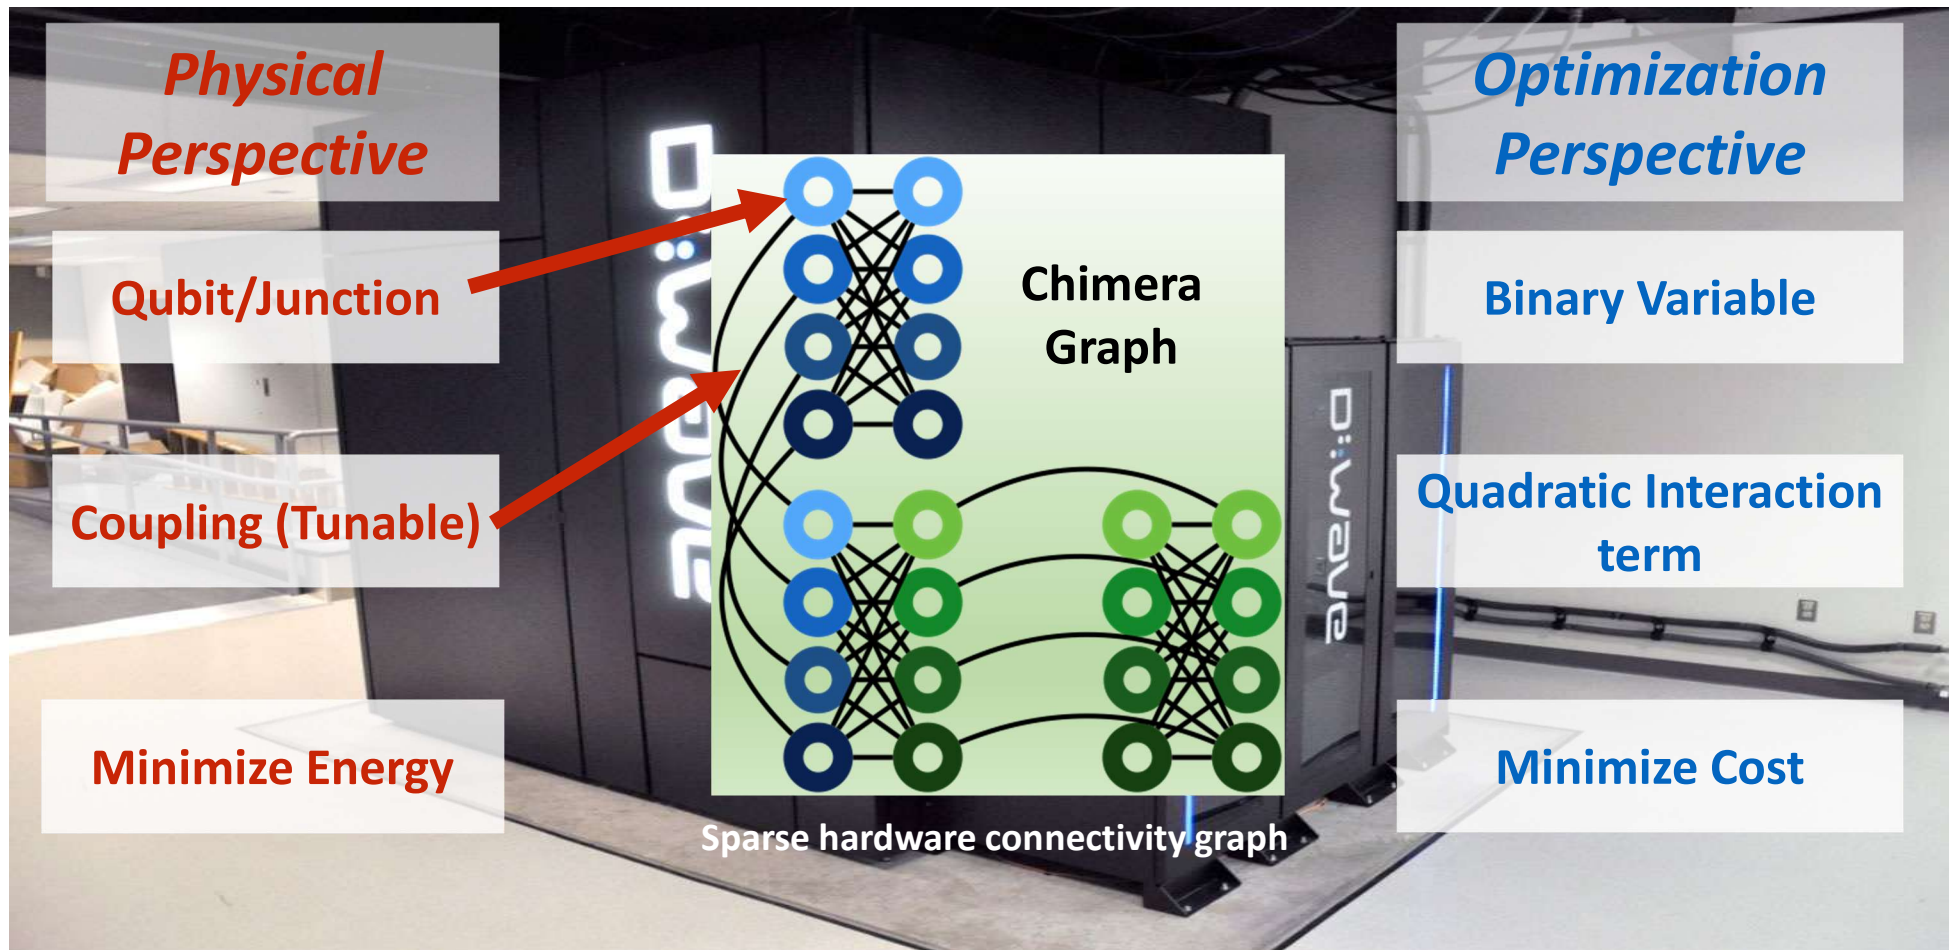
- First company to commercially release a quantum annealing system → D-Wave One (2011)
- Solves **optimization problems** formulated as **Ising Hamiltonian**
- An Ising Hamiltonian, is a specific type of Hamiltonian that involves spin variables
  - Spin variables can take values of +1 or -1
- Quantum Annealing starts from the known ground state of a simple Hamiltonian and slowly evolves the Hamiltonian toward the optimization Hamiltonian.
- According to the adiabatic theorem, the system tends to remain near the ground state during the evolution.
- **Quantum tunneling may help the system avoid poor local minima**

**Optimization Problem** → **Ising Hamiltonian**



$$x_i = \frac{1 - z_i}{2}$$

# Annealing Based: DWave



# Annealing Limitations

- **Not a universal** quantum computer
- **Restricted** mainly to Ising/QUBO **optimization problems**
- **Small energy gaps** may require extremely slow annealing
- **Sparse hardware connectivity** complicates problem embedding
- **Noise and decoherence** can reduce ground-state fidelity

# Gate-Based: Universal Quantum Computing

rigetti



QUANTINUUM



IONQ

Google

IBM



Microsoft

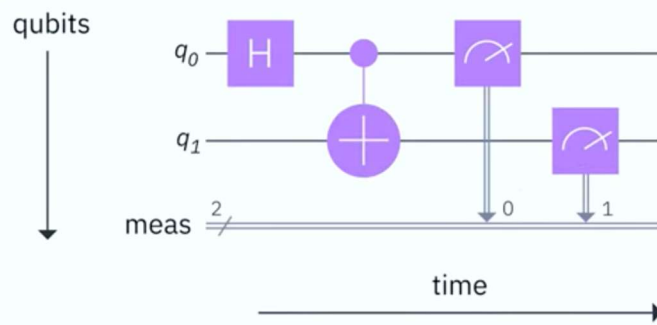
TU Delft

FUJITSU

D:wave

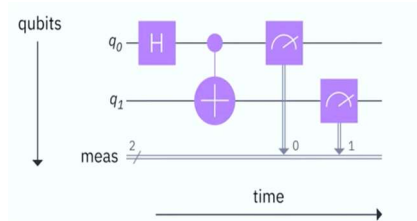
# Exemplar: IBM Quantum Computer

Quantum logic gates arranged in a computational circuit



# Qiskit Overview

- A collection of software for executing programs on quantum computers



**Q<sup>+</sup>**  
Map problem to quantum circuits & operators

- Circuit library
- Quantum info library
- And more...

**X**  
Optimize for target hardware

- Transpiler
- AI-enhanced transpiler
- And more...

**⌚**  
Execute on target hardware

- Runtime primitives
- Execution modes
- And more...

**📈**  
Post-process results

- Quantum info library
- Visualization module
- And more...

**Output**

Probability distribution

Expectation value

■ Qiskit SDK   
 ■ Qiskit Runtime Service   
 ■ Qiskit Transpiler Service

- **Qiskit SDK** (open-source SDK) → quantum circuits, operators, and primitives.
- **Qiskit Runtime** → execute quantum circuits on IBM Quantum™ hardware.
- **Qiskit Aer** → execute quantum circuits on simulators with realistic noise models

# 127 Qubit Eagle Processor

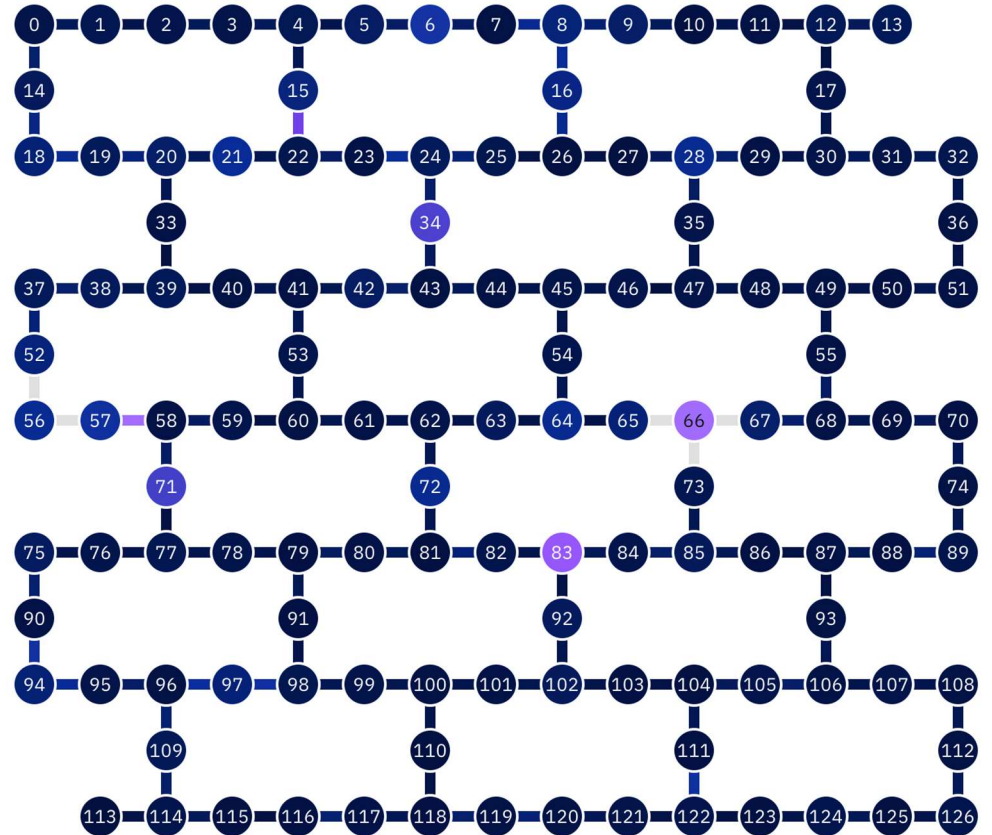
ibm\_sherbrooke



QPU status ● Online

Processor type Eagle r3

Qubits	2Q error (best/layered)	CLOPS
127	3.07e-3/2.12e-2	30K



# Gate-Based Quantum Platform Limitations

- **Limited Qubit Connectivity**

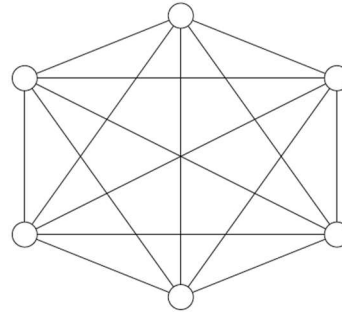
- Sparse connections between qubits
- Additional gates for missing connections
- Qubit interactions  $\uparrow \Leftrightarrow$  Circuit depth  $\uparrow$

- **Limited Coherence Times**

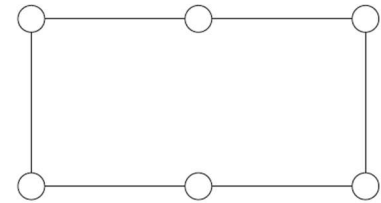
- Decaying quantum states
- Circuit depth  $\uparrow \Leftrightarrow$  Execution time  $\uparrow$
- Execution time  $\uparrow \Leftrightarrow$  Decoherence errors  $\uparrow$

- Imperfect Qubits: Faulty or Noisy

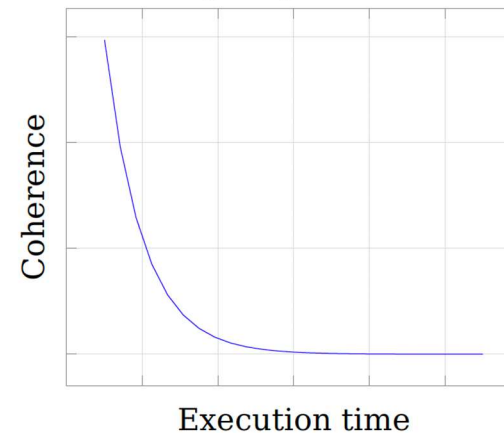
## Optimal topology



## Real topology



Circuit Transpilation  
Adding Swap Gates  $\rightarrow$



*Stage II: Quantum  
Computing Challenges &  
Potential Benefits for DBMS*

---

# Impact of Quantum Computing on ACID Semantics

- RDBMS contract is underpinned by the ACID properties
  - Atomicity, Consistency, Isolation, and Durability
- RDBMS ensures that SQL queries are processed reliably and consistently, producing deterministic results

## Atomicity:

- Quantum computing has the **capability to execute complex operations rapidly**
- Reduces the likelihood of partial transactions and **enhancing atomicity**
- However, quantum operations are inherently error-prone and require specialized error correction
  - Can introduce new computational overheads
  - Which in turn increases the number of operations to complete the transaction
- **Good News:** The field is continually evolving
  - More robust and efficient error correction methods are being developed
  - Improvements in the stability of qubits and the fidelity of measurements

# Impact of Quantum Computing on ACID Semantics

## Consistency:

- Leveraging superposition and entanglement, can efficiently perform intricate consistency checks
- However, **quantum state fragility due to decoherence** can introduce inconsistencies as information degrades
- Additionally, the probabilistic nature of quantum computing can amplify these issues
- **Good News:** Mechanisms like Powering Lemma can help improve the success probability

## Isolation:

- Quantum parallelism allows multiple transactions to be processed simultaneously, enhancing isolation
- However, quantum entanglement can inadvertently link transactions, complicating true isolation
- Ancilla qubits are employed to entangle data qubits of different columns
- **Good News:** All quantum operations are reversible

## Durability:

- The inability to perfectly clone quantum states makes it difficult to create independent copies
- **Good News:** Potential future advances in quantum memory systems [2] could provide robust data storage

# Quantum Challenges

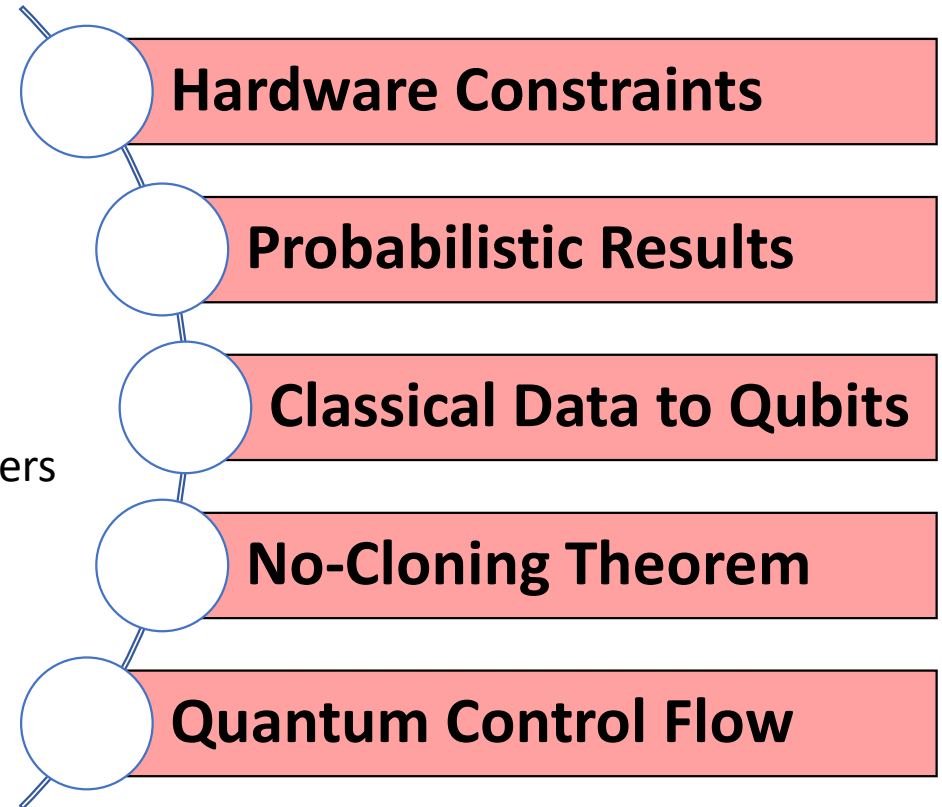
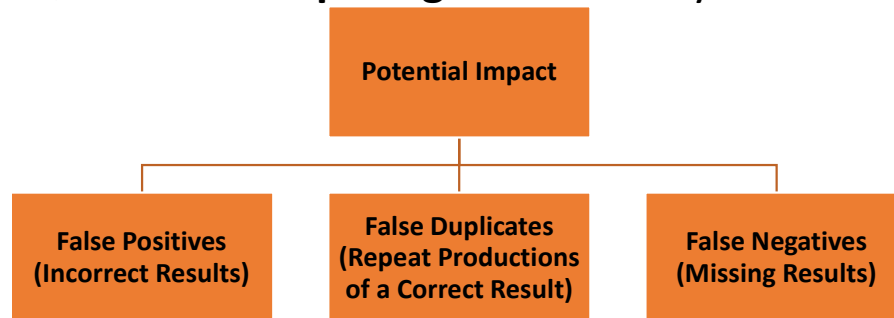
## ➤ Hardware Constraints

- Limited (156 Qubits)
- Noisy

*~1300 Qubits Required for Index Selection on TPC-H*

## ➤ Probabilistic Results

- **DBMS Expectation** → Deterministic, Exact Answers
- **Quantum Computing** → Inherently Probabilistic



# Quantum Challenges

## 2. Classical Data to Qubits

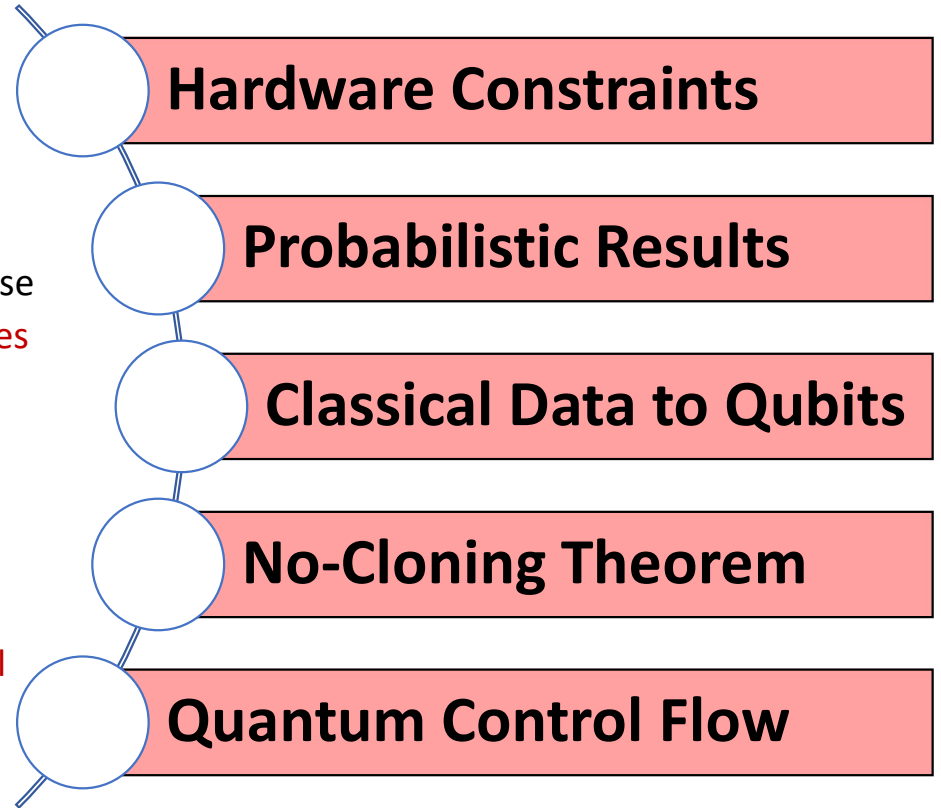
- Each existing proposal uniquely transforms input data
- Join ordering undergoes multiple transformations
- Re-evaluate data encoding techniques, like basis encoding

## 3. No-Cloning Theorem

- Restricts duplication of intermediate quantum state for reuse
- Restricts dynamic generation of multiple independent copies of a database for parallel processing

## 4. Quantum Control Flow

- Does not support iterative control structures
- Results in large quantum circuits with increased depth
- Qubits to maintain quantum states for extended durations
- Maintaining coherence long enough to perform meaningful calculations remains a significant technical challenge



# Quantum Benefits

## 1. Enhanced Execution Performance

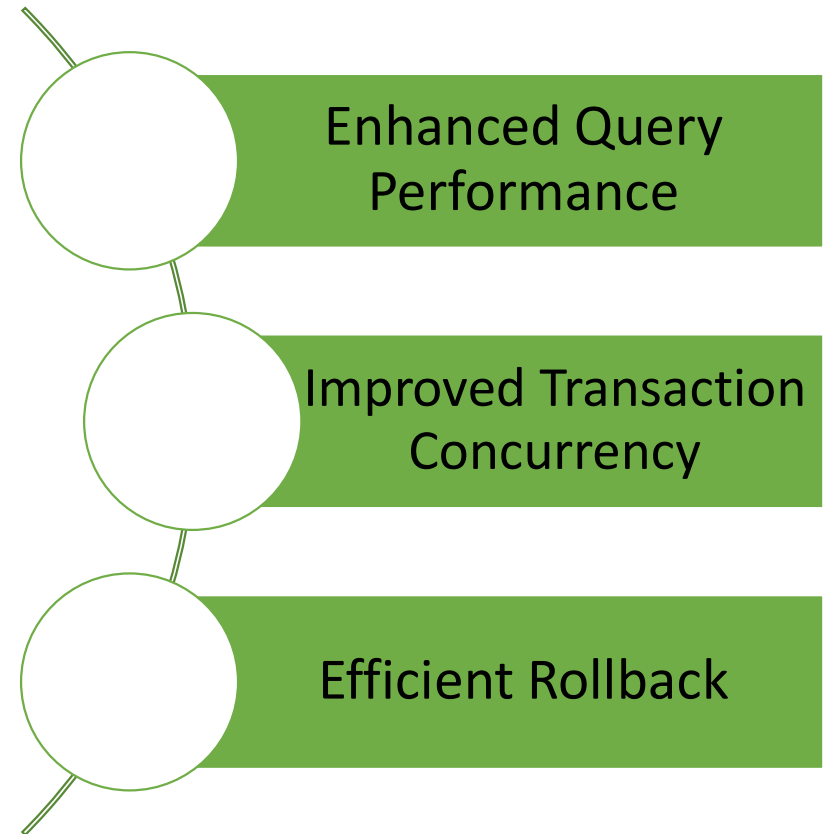
- GS algorithm could potentially offer quadratic speedups
- Simultaneous evaluation of multiple query paths
- Parallel execution of subqueries
- **Prime Candidates:** SQL queries involving complex join predicates and multiple complex filter conditions

## 2. Improved Transaction Concurrency

- Fast execution of SQL queries could lead to almost-instantaneous commits of transactions
- Reducing the number of parallel conflicting transactions
- Could significantly reduce the overheads to ensure serializability
- **Simpler & more efficient Transaction management systems**

## 3. Efficient Rollback

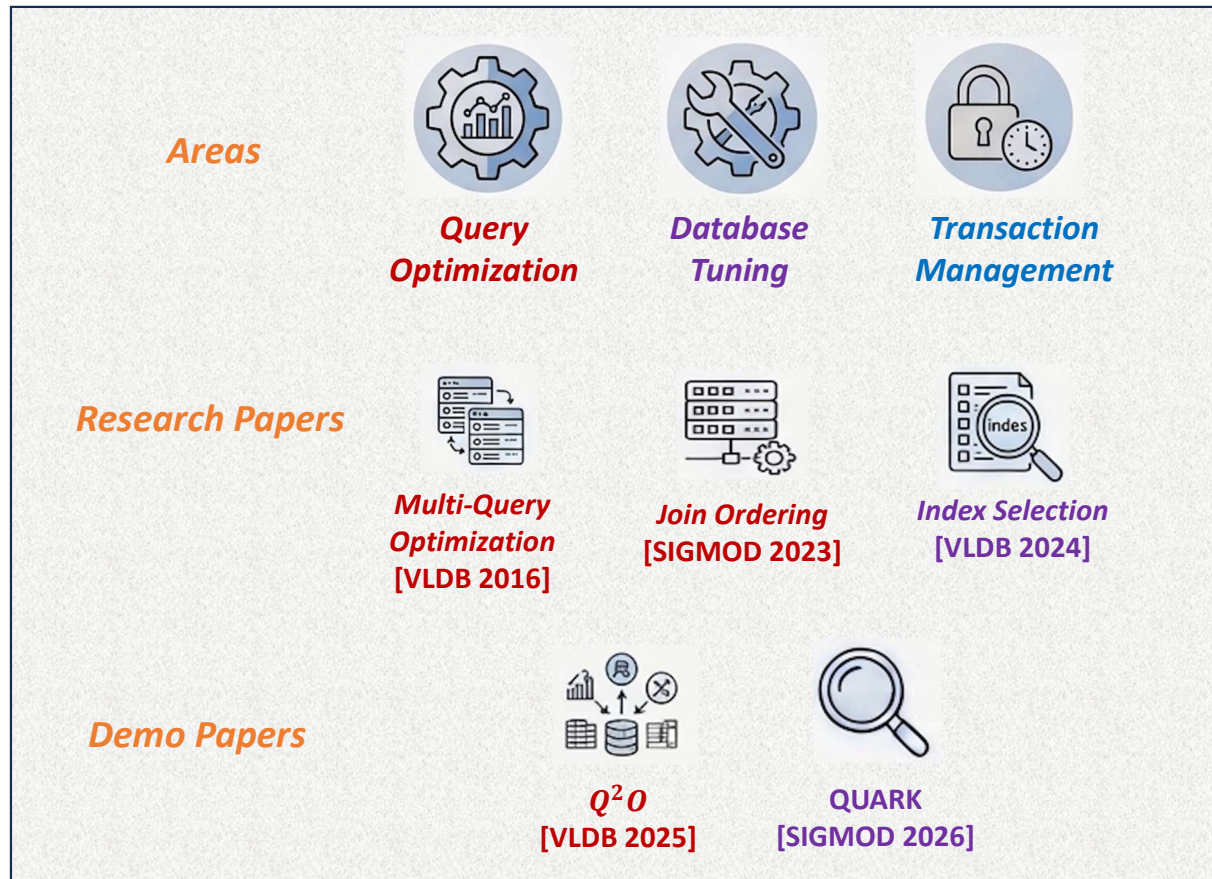
- The inherently reversible nature of quantum operators could facilitate efficient rollback of transactions
- **Allow transactions to efficiently retrace its steps**



*Stage III: Quantum  
Optimization for DBMS*

---

# Research At A Glance



*3.1 Multiple Query  
Optimization on the DWave  
2X Adiabatic Quantum  
Computer*

*Immanuel Trummer, and Christoph Koch, VLDB 2016*

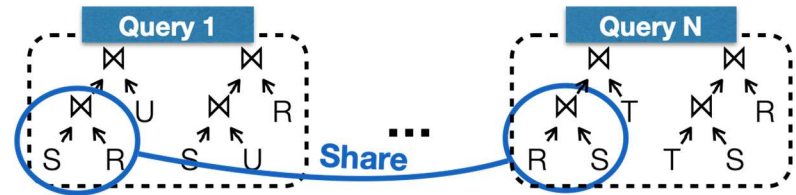
---

# Multiple Query Optimization (MQO)

Select the optimal combination of query plans to minimize execution cost for a batch of SQL queries

- ✓ From MQO to quantum-compliant format
- ✓ Embedding to qubit weights for optimization via quantum annealing
- ✓ Experiment on REAL Quantum Computer
- ✓ Showed three orders of magnitude improvement

**Find Optimal Plan Combination**

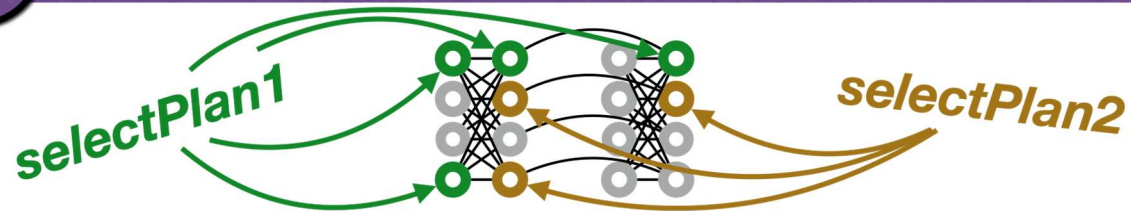


**1 Transform → Quadratic Formula**

**Binary Variables:**  
 $selectPlan_2 \dots selectPlan_1$

**Cost Formula:**  
 $costTerms \ savingsTerms$   
 $constraintTerms$

**2 Embedding → Weights on Qubits**



# MQO Problem Instance

- Given a set  $Q$  of SQL queries
- $q \in Q$  has a set of alternative plans  $P_q$  and  $P = \cup_q P_q$
- $p \in P$  is associated with an execution cost  $c_p$
- $s_{p_1, p_2} > 0$  is the cost reduction by plans  $p_1$  and  $p_2$  that can share intermediate results
- Solution is a subset of plans  $P_e \subseteq P$ 
  - Validity Constraint:  $\forall q \in Q: |P_q \cap P_e| = 1$
  - Cost:  $C(P_e) = \sum_{p \in P_e} c_p - \sum_{\{p_1, p_2\} \subseteq P_e} s_{p_1, p_2}$
- **Objective:** Find optimal set of plans to minimize cost
- **Assumptions:**
  - Alternative plans has been pre-identified
  - Execution costs can be reliably estimated

---

**Algorithm 1** How to solve multiple query optimization on an adiabatic quantum annealer.

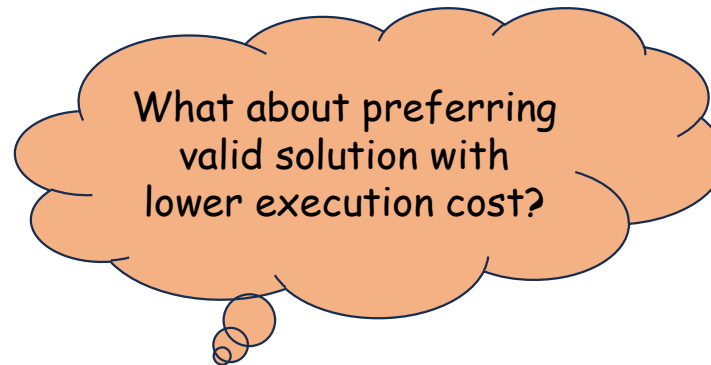
---

```
1: // Solves multiple query optimization problem  $M$ 
2: function QUANTUMMQO( $M$ )
3:   // Map MQO problem to logical energy formula
4:    $lef \leftarrow$  LOGICALMAPPING( $M$ )
5:   // Map logical into physical energy formula
6:    $pef \leftarrow$  PHYSICALMAPPING( $lef$ )
7:   // Minimize formula on quantum computer
8:    $b_i \leftarrow$  QUANTUMANNEALING( $pef$ )
9:   // Transform physical into logical solution
10:   $X_p \leftarrow$  PHYSICALMAPPING $^{-1}$ ( $b_i$ )
11:  // Transform logical solution to MQO solution
12:   $P_e \leftarrow$  LOGICALMAPPING $^{-1}$ ( $X_p$ )
13:  // Return best set of query plans to execute
14:  return  $P_e$ 
15: end function
```

---

# Logical Mapping: MQO $\rightarrow$ QUBO

- $X_p$ : binary variables for each query plan  $p$
- Validity Constraint:  $\sum_{p \in P_q} X_p = 1, \forall q \in Q \rightarrow$  QUBO does not allow to express constraints
- Decompose the constraint into two parts
  - At least one plan is selected  $\rightarrow E_L = -\sum_{p \in P} X_p$ 
    - This term motivates to set all variables  $X_p$  to one
  - At most one plan is selected  $\rightarrow E_M = \sum_{q \in Q} \sum_{\{p_1, p_2\} \subseteq P_q} X_{p_1} X_{p_2}$ 
    - This term takes value zero if at most one plan is selected per query
  - Both terms will have to be scaled by an appropriate factor



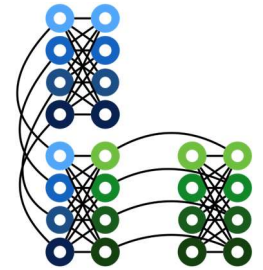
# Logical Mapping: MQO $\rightarrow$ QUBO

- At least One:  $E_L = -\sum_{p \in P} X_p$
- At most One:  $E_M = \sum_{q \in Q} \sum_{\{p_1, p_2\} \subseteq P_q} X_{p_1} X_{p_2}$
- Cost Term:  $E_C = \sum_{p \in P} c_p X_p$
- Sharing intermediate results between plans:  $E_S = -\sum_{\{p_1, p_2\} \subseteq P} s_{p_1, p_2} X_{p_1} X_{p_2}$
- QUBO:  $w_L E_L + w_M E_M + E_C + E_S$
- Priority of selecting at least one plan is higher than saving cost:  $w_L > \max_{p \in P} c_p$
- $w_L E_L + E_C + E_S$  would be minimized by executing each plan for each query
- So,  $w_M > w_L$  to restrict the number of plan selections per query to one  $\rightarrow$  How MUCH??
- What if generation cost of a query is lower than the cost reduction by sharing it ?

$$w_M > w_L + \max_{p_1 \in P} \sum_{p_2 \in P} s_{p_1, p_2}$$



# Physical Mapping: QUBO $\rightarrow$ QUBITS



## QUBIT Mapping:

- 1 logical binary variable to 1 Qubit? **NO**
- Each qubit is connected to at most six other qubits [Chimera Graph Slide #23]
- If for a plan  $p_1$  there are more than six plans  $p_2$  such that  $w_{p_1,p_2} \neq 0$
- Each logical variable  $X_p \rightarrow$  a group of physical qubits

## Constraints Mapping:

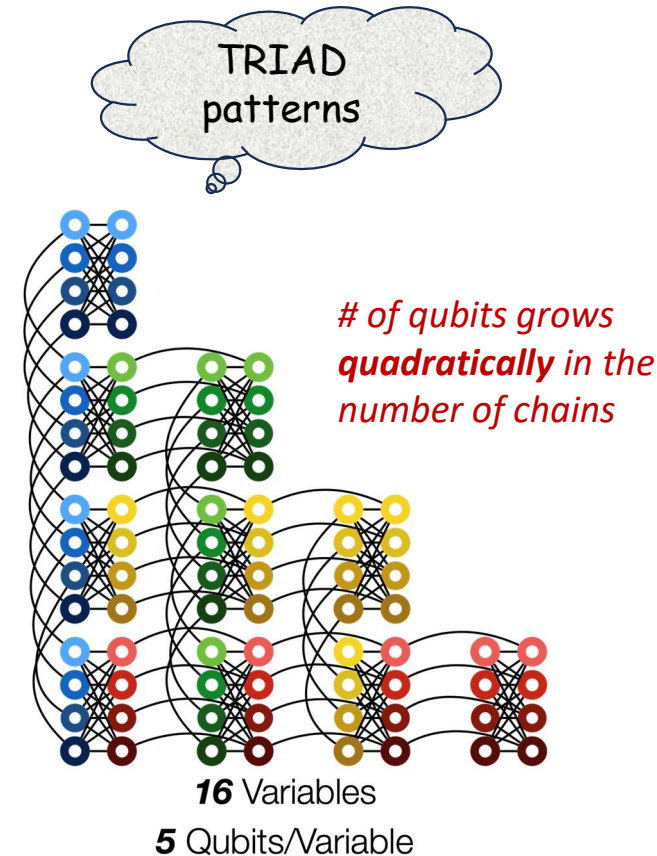
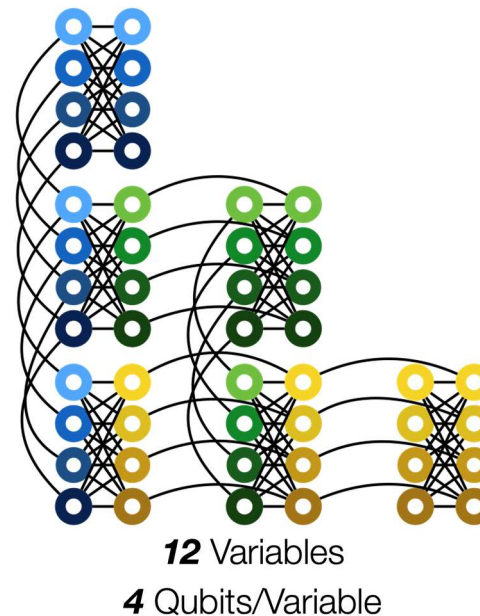
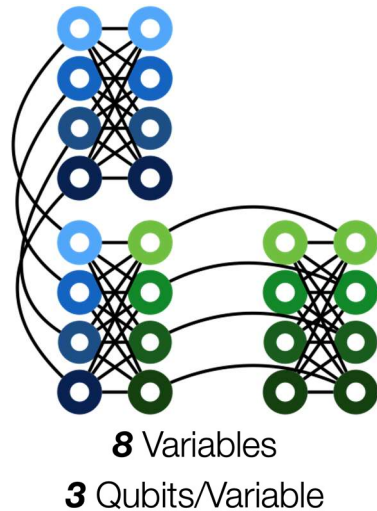

- For a term of the form  $w_i X_i \rightarrow$  distribute weight over all qubits representing  $X_i$
- $B_i$  denotes the set of qubits representing  $X_i \rightarrow w_i/|B_i|$  is added on each qubit in  $B_i$
- For  $w_{ij} X_i X_j \rightarrow$  select  $b_1 \in B_i$  and  $b_2 \in B_j$  such that  $b_1$  and  $b_2$  are connected by a coupling
  - Increase the strength of that coupling by  $w_{ij}$



Ensure all qubits representing the same variable **"behave as one"** and are assigned consistently to the same value after an annealing run

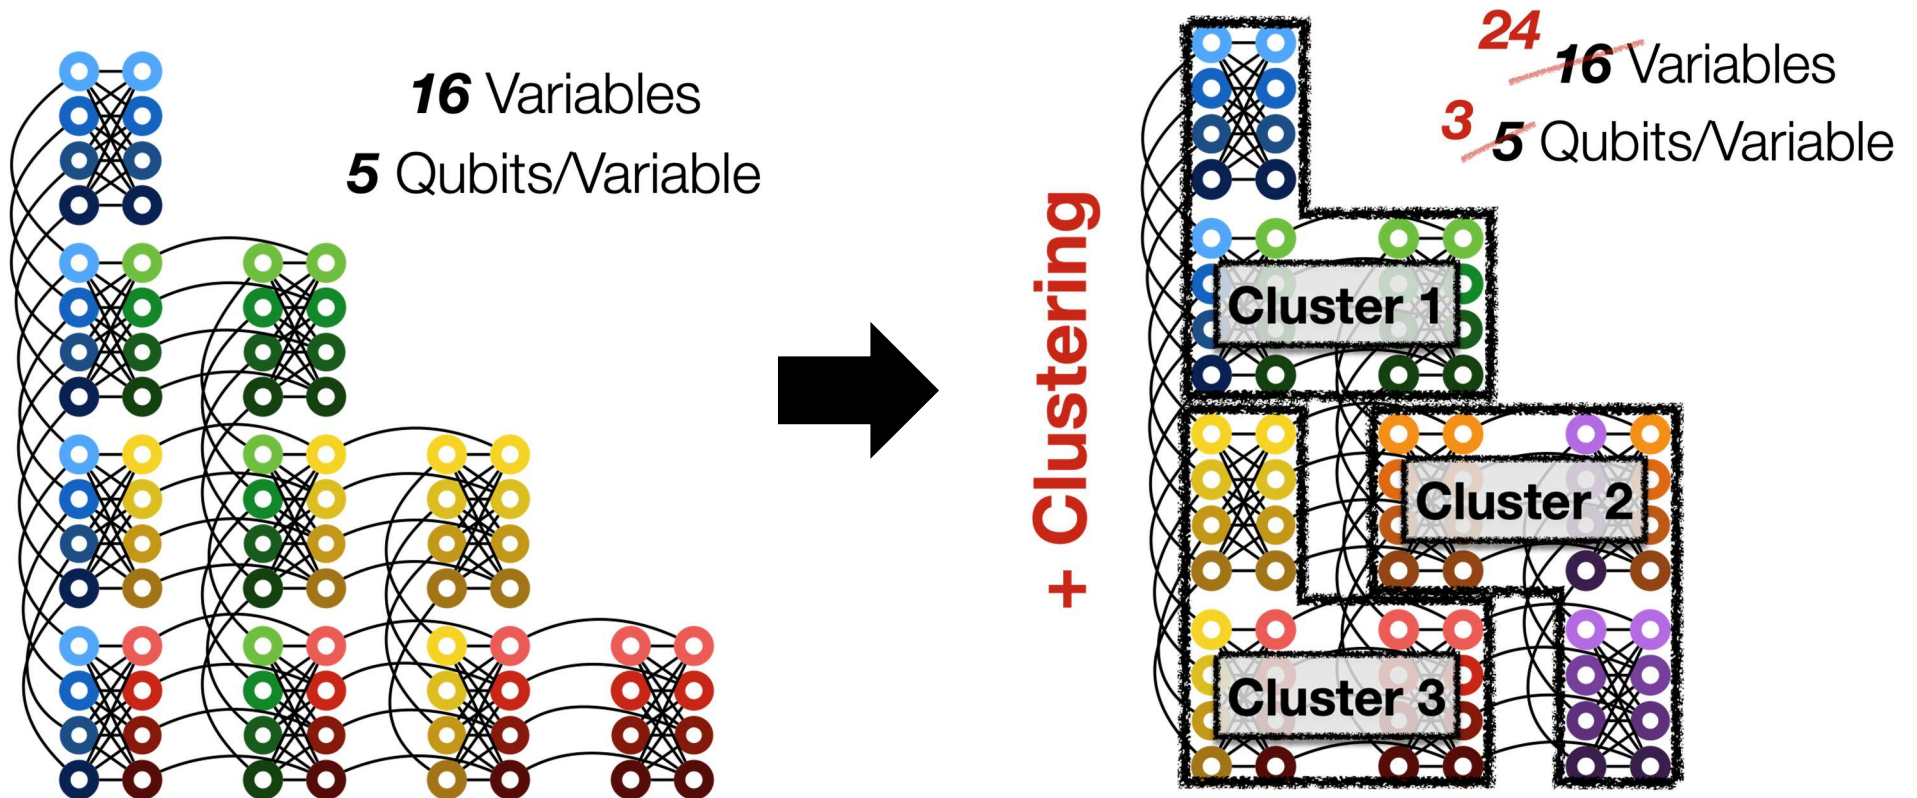
# Physical Mapping: QUBO $\rightarrow$ QUBITS

- Say qubits  $b_1$  and  $b_2$  represent the same variable
- Add the term  $b_1 + b_2 - 2b_1b_2$  to the QUBO
- It is 1 if qubits are assigned different values
- This requires that all qubits representing the same variable form a chain
- Finding optimal embedding is a NP-hard problem !



# Physical Mapping: QUBO $\rightarrow$ QUBITS

- Cluster queries based on structural properties  $\rightarrow$  One TRIAD pattern per cluster
- Queries in different clusters are less likely to share intermediate results

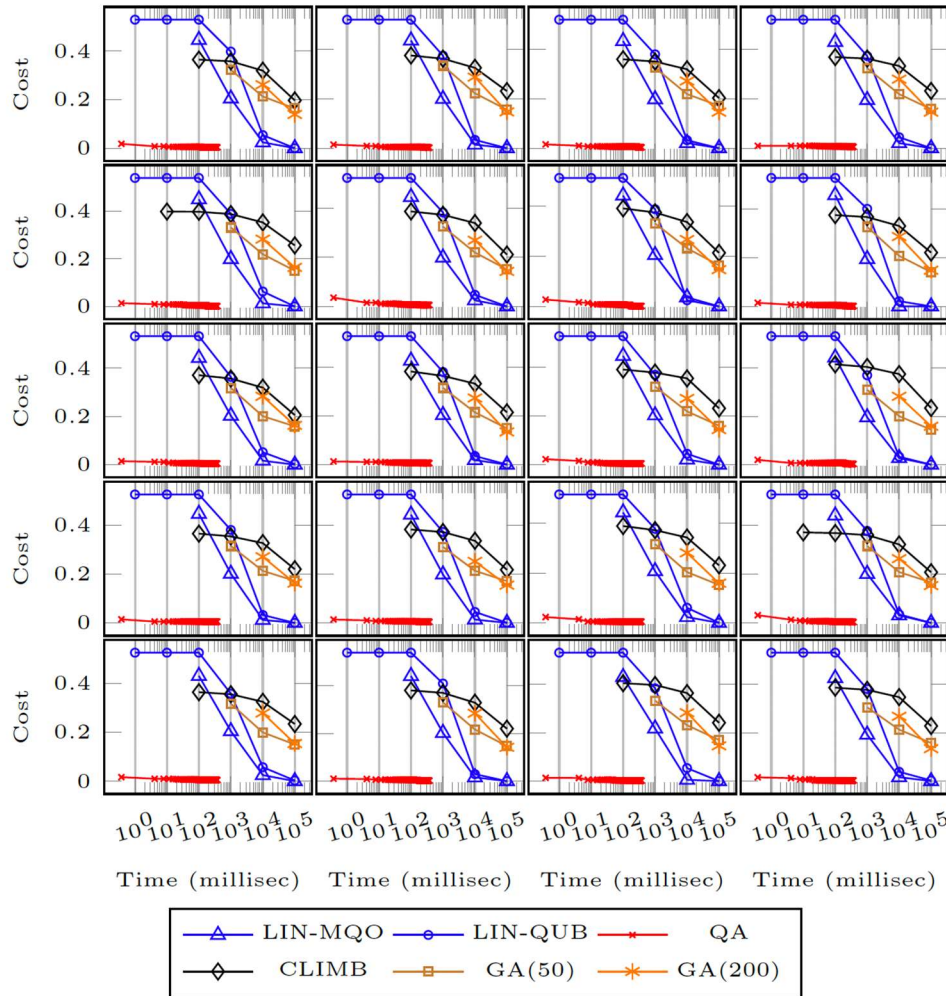


# Experiment Setup

- Platform: D-Wave 2X quantum annealer with 1097 qubits
- The size of the problems is inherently limited by the number of available qubits
- 20 MQO instances per test case
- 1000 annealing runs per experiment
- Metric: Solution quality as a function of optimization time
- Compared with classical MQO solutions:
  - LIN-MQO: Linear solver applied to MQO
  - LIN-QUB: Linear solver applied to QUBO
  - CLIMB: Iterated hill climbing
  - GEN(50): Genetic algorithm with population size 50
  - GEN(200): Genetic algorithm with population size 200

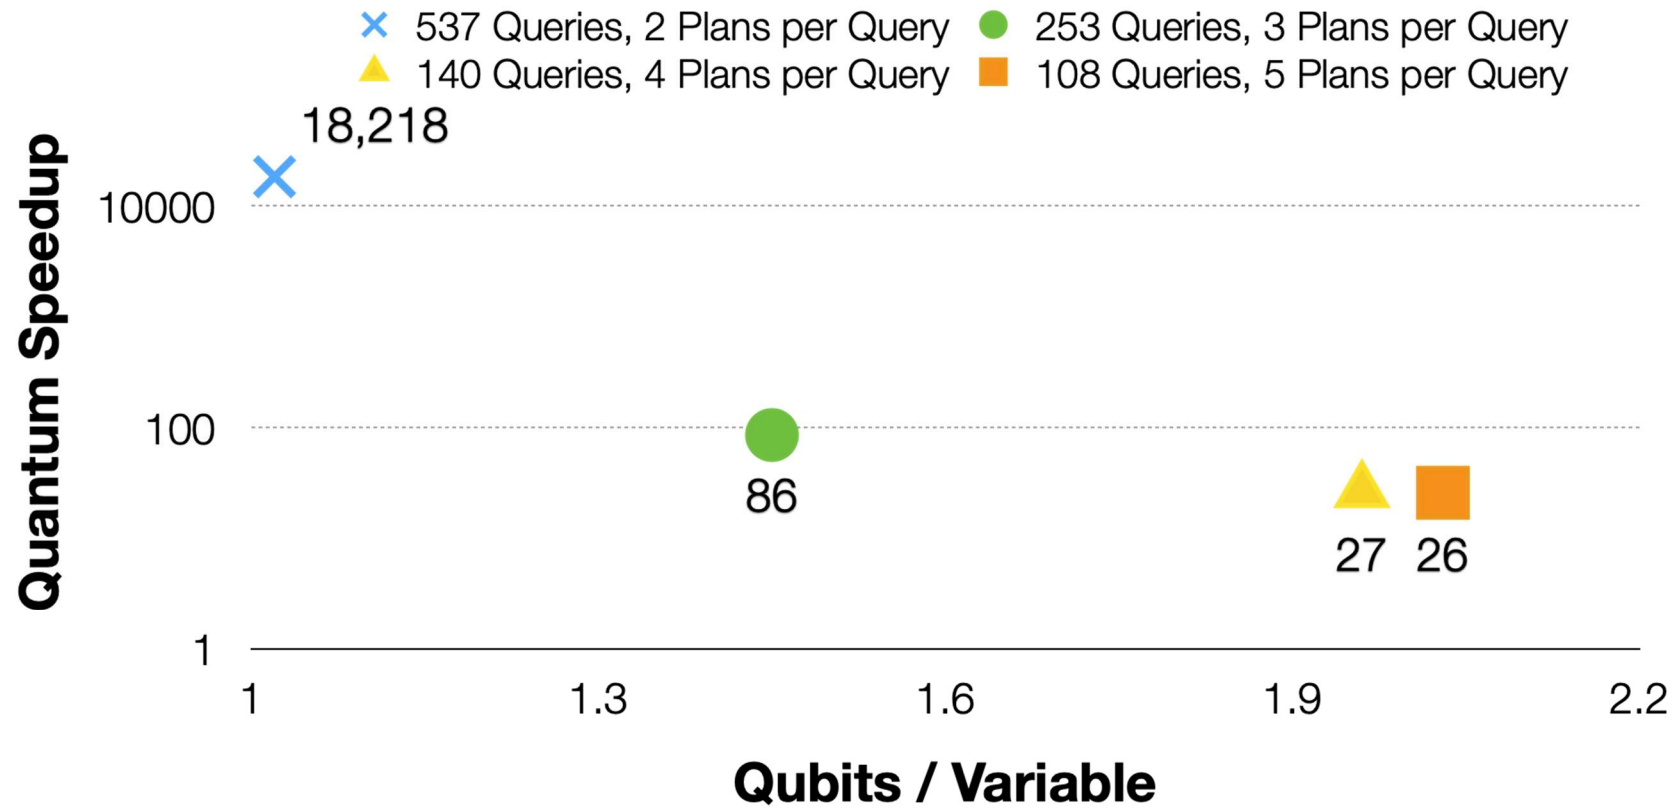
Test Set	Number of Queries	Plans per Query
1	537	2
2	253	3
3	140	4
4	108	5

# Results: Test Set 1



*Quantum solution after first annealing run (less than half millisecond) is equivalent to the classical solution after one second !*

# Results



*3.2 Ready to Leap  
(by Co-Design)?  
Join Order Optimisation  
on Quantum Hardware*

*Manuel Schonberger, Stefanie Scherzinger, Wolfgang Mauerer  
SIGMOD 2023*

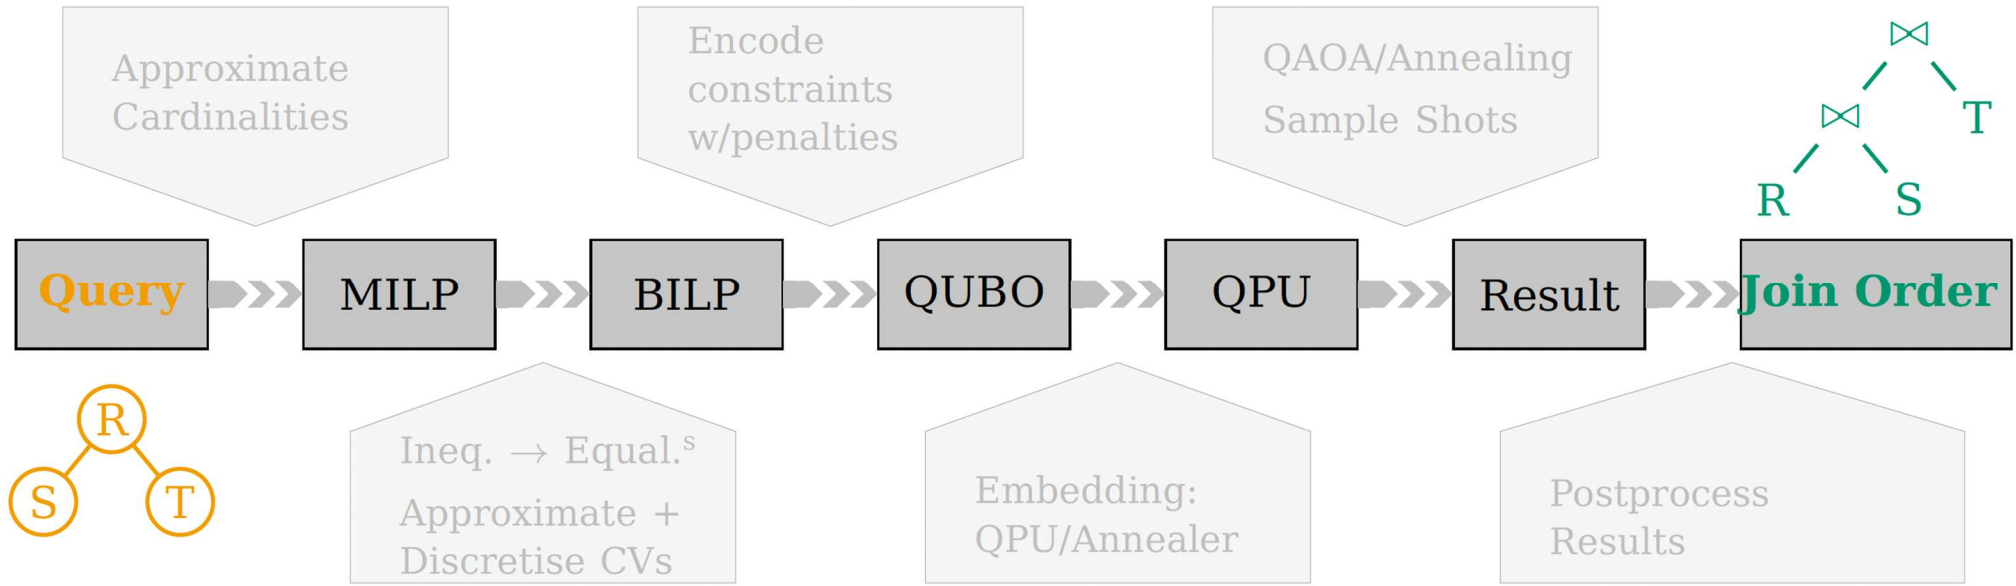
---

# Problem Statement

- **Mission:** Solve Join Ordering (JO) problem on Quantum Platform
- JO problem is given by a set of relations  $R_0, \dots, R_n$
- A query graph, where each node represents a relation
- An edge between relations  $R_i$  and  $R_j$  is labelled by a predicate  $p_{ij}$  with selectivity  $f_{ij}$ , where  $0 < f_{ij} \leq 1$
- **Objective:** Find optimal left-deep join trees (NP-complete)
  - Without restrictions on the query graph
- Cost Function:  $C_{out}(n_i, n_j) = n_i n_j f_{ij}$ 
  - $n_i$  denotes the cardinality of relation  $R_i$
- Cost Function for the optimal join order for a sequence  $s$  of  $n$  relations,  $s_1, \dots, s_n$  is

$$C(s) = \sum_2^n C_{out}(|s_1, \dots, s_{i-1}|, |s_i|)$$

# Join Ordering to QUBO

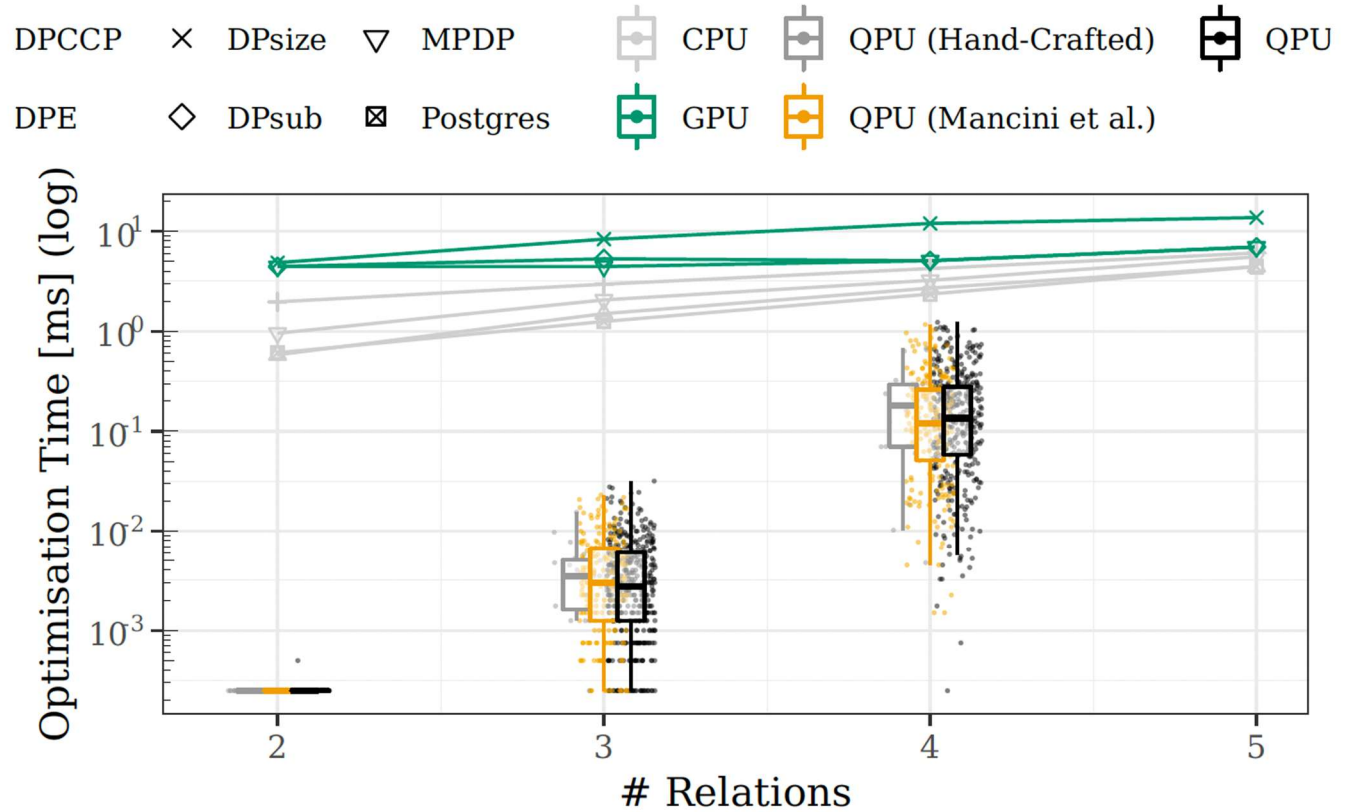


# Experiment

- **Problem Representation:**
  - A query graph (annotated with cardinalities),
  - applicable join predicates with (estimates of) their selectivities, and
  - the threshold values for approximating intermediate cardinalities
- **Problem Set:**
  - 20 queries for different problem size scenarios → between two and five relations
  - Chain, star and cycle queries using Steinbrunn et al. [VLDB'97]
  - Clique queries using Mancini et al. [SIGMOD'22]
  - Hand crafted to minimise imperfections of current QPUs
- **Evaluation:**
  - Run-time,
  - Result quality, and
  - Maximal problem sizes (qubit count, circuit depth)

# D-Wave Results

➤ Speedups for small queries  
➤ But limits reached quickly



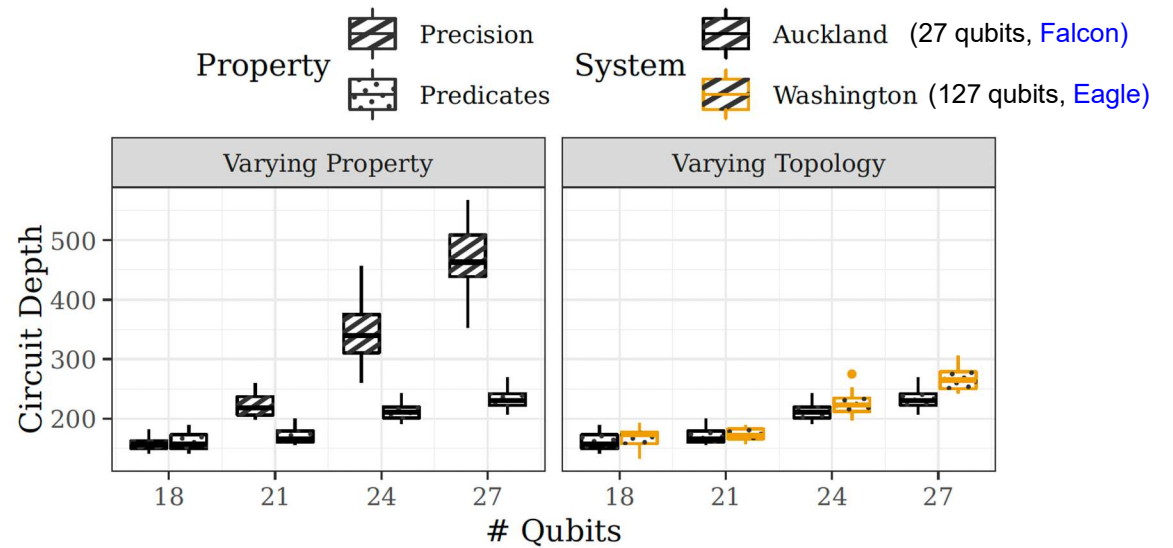
# D-Wave Results

Query Graph	$\Delta t$ [ $\mu s$ ]	3 Relations		4 Relations		5 Relations	
		Valid	Opt.	Valid	Opt.	Valid	Opt.
Chain	10	31.00%	8.69%	1.23%	0.07%	0.12%	0%
	100	37.01%	8.22%	1.53%	0.19%	0.07%	0%
	1000	41.25%	10.30%	1.78%	0.14%	0.15%	0%
Star	10	-	-	2.16%	0.33%	0.03%	0%
	100	-	-	2.01%	0.29%	0.02%	0%
	1000	-	-	2.48%	0.43%	0.04%	0%
Cycle	10	27.52%	11.46%	3.32%	0.41%	0.02%	0%
	100	31.10%	14.36%	3.77%	0.46%	0.02%	0%
	1000	34.58%	16.89%	4.21%	0.36%	0.04%	0%

Average fraction of valid and optimal solutions in 1,000 annealing runs (DWave) over 20 JO instances

# IBM Q Results

- Identical qubits, varying parameters:
  - Predicates
  - Discretization precision
- QPU topologies:
  - Auckland (27 qubits)
  - Washington (127 qubits)
- Approx. precision  $\uparrow \Leftrightarrow$  Circuit dept



# QAOA Results on IBM Auckland

		Predicates/Qubits				
		$n$	0/18	1/21	2/24	3/27
Valid $\text{sol}^n$	20		13%	11%	7%	13%
	50		12%	8%	10%	13%
Opt. $\text{sol}^n$	20		4%	3%	2%	5%
	50		3%	3%	5%	3%

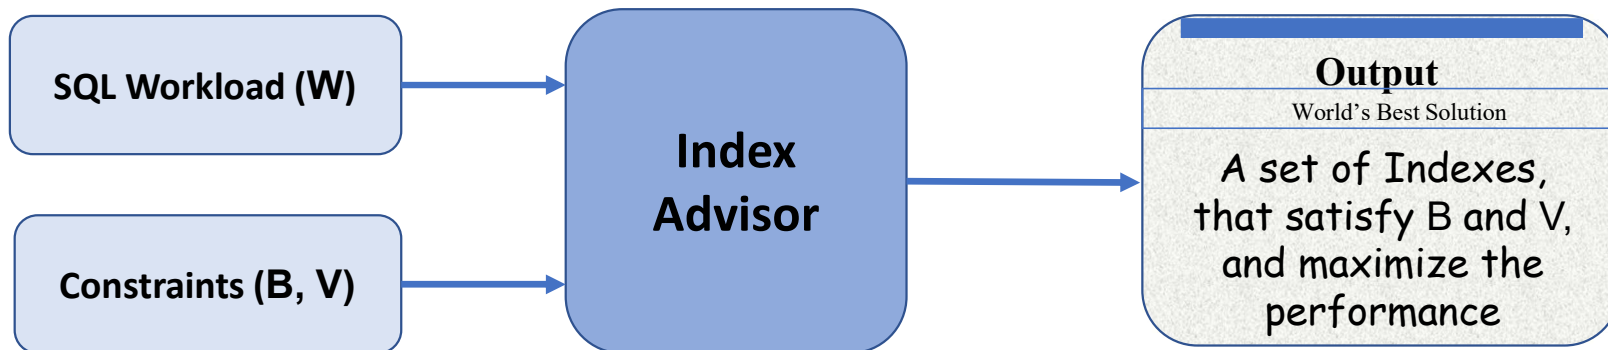
# 3.3 Index Advisors on Quantum Platforms

Manish Kesarwani, Jayant R. Haritsa  
[VLDB, 2024]

---

# Problem Statement

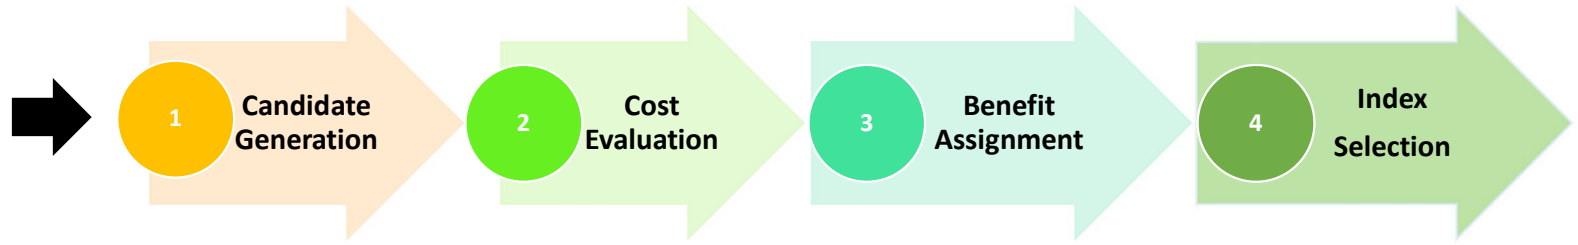
- Given an SQL query workload, recommend a configuration of (B-tree) indexes that maximize the performance benefit (measured with respect to execution time), subject to
  - ✓ **Space Constraint (B)**: The index choices must fit in a user-specified storage budget;
  - ✓ **Validity Constraint (V)**: It must satisfy validity requirements, which could be
    - a) *Intrinsic* – at most one clustered index per relation,
    - b) *Extrinsic* – mandatorily add all indexes listed in a pre-specified base configuration



- Due to *computational hardness of index selection*, arising from combinatorial candidate space
  - Heuristic solutions in current DBMS
  - Result in highly sub-optimal configurations

# Sub-optimality due to Heuristics

- SQL Workload
- Storage Constraint



For SQL workload comprising TPC-H queries Q6, Q14, Q22, and two instances of Q17 over a 1GB TPC-H database constraint to 140MB of storage for index creation

Index (I)	Columns in the Index	Time Benefit (V)	Storage Cost (W)
$i_0$	[L_DISCOUNT, L_SHIPDATE, L_QUANTITY, L_EXTENDEDPRICE]	165811	266
$i_1$	[L_SHIPDATE, L_DISCOUNT, L_EXTENDEDPRICE, L_PARTKEY]	178871	232
$i_2$	[P_CONTAINER, P_BRAND, P_PARTKEY]	1213770	8
$i_3$	[L_PARTKEY, L_QUANTITY]	1213770	8
$i_4$	[L_PARTKEY, L_QUANTITY, L_EXTENDEDPRICE]	1213770	8
$i_5$	[C_ACCTBAL, C_CUSTKEY, C_PHONE]	44370	8
$i_6$	[O_CUSTKEY]	44370	8

**Storage Capacity: 140**

**Commercial Database Recommendation**

Heuristic Index Set:  $\{i_2, i_5, i_6\}$   
 Heuristic Benefit:  $1.3 * 10^6$

**Optimal Configuration**

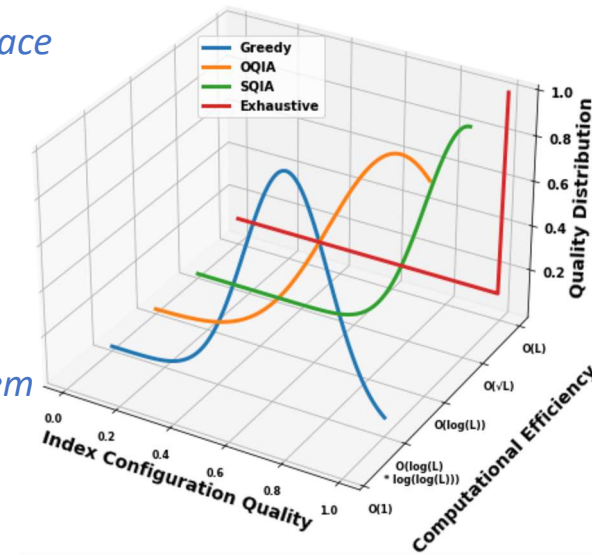
Optimal Index Set:  $\{i_2, i_3\}$   
 Optimal Benefit:  $2.4 * 10^6$

Quality = 0.54  
**Highly Sub-optimal**

Workload execution time could be **reduced** by 46%

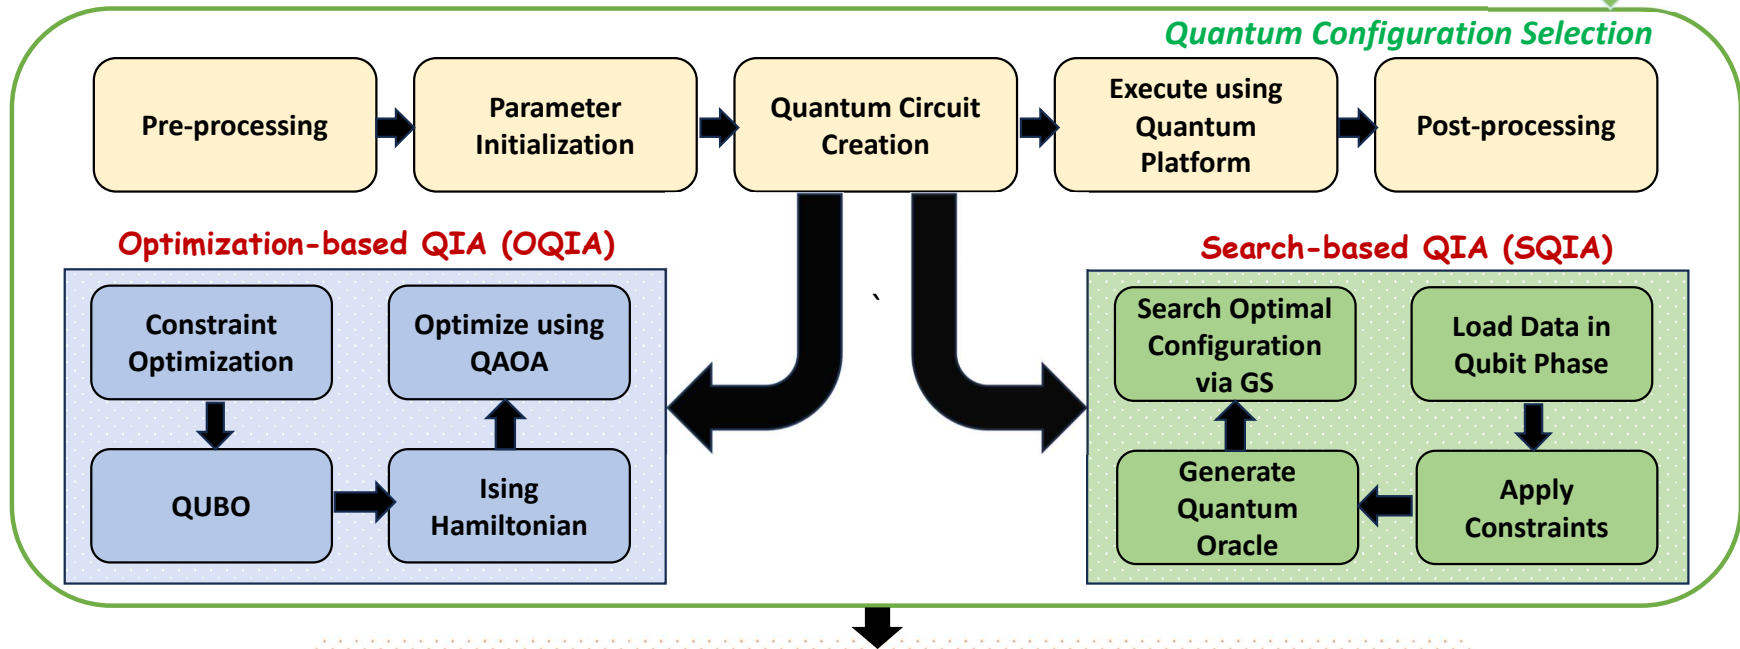
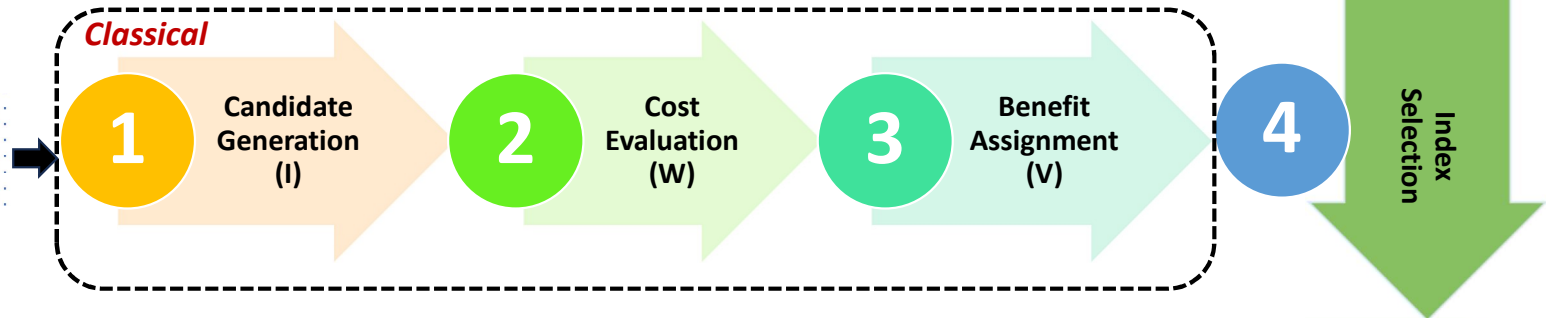
# Contributions

- Proposed the *first Hybrid Quantum-Classical Index Advisor (QIA)* architecture
- Logical design and circuit representations for *Search-based QIA (SQIA) approach*
  - Index Selection is modeled as *enumerative search over exponential configuration space*
    - Solution Space:  $L = 2^{|I|}$
  - Load the problem data in the qubit phases
  - Solved using Grover Search (GS)
  - Identify, *with high probability, optimal index configuration* in  $O(\sqrt{L})$  computations
- Transformation to Ising Hamiltonian for *Optimization-based QIA (OQIA) approach*
  - Index selection is modeled as a *Quadratic Unconstrained Binary Optimization problem*
  - Solved using the popular Quantum Approximate Optimization Algorithm (QAOA)
  - Identify, *better index configuration* in  $O(\log L)$  computations
- A pilot implementation and evaluation of the proposed approaches on both
  - *(noiseless) quantum simulators* and
  - *(noisy) quantum circuit processors*
- Obtained substantively improved index configurations, by a multiplicative factor of *1.5 to 2 and approaching optimality*
- Showed that the complexity of *our circuit design scales linearly* to large databases



# Quantum Index Advisor

SQL Workload  
Storage Constraint



Recommended Index Configuration

# *Search-based QIA (SQIA)*

---

# Search-based QIA (SQIA) Overview

Express the Index Selection problem *as an enumerative search over the exponential configuration space*

1. Leveraging the **quantum superposition** property generate the exponential space of candidate index configurations
  - Incurring only logarithmic qubit overhead
2. Next, we Load the problem instance data in the qubit **phases**
  - Harnesses the power of **quantum interference and entanglement** to attach data to the index configurations
3. Encode constraints into qubit phases and prepare a custom quantum oracle for the Index Selection problem
4. Configure the Grover Search algorithm to overcome instance-specific challenges
5. **Identify, with high probability, the optimal index configuration** with  $O(\sqrt{L})$  computations



# Generate All Index Configurations (Step 1)

**Problem Instance**

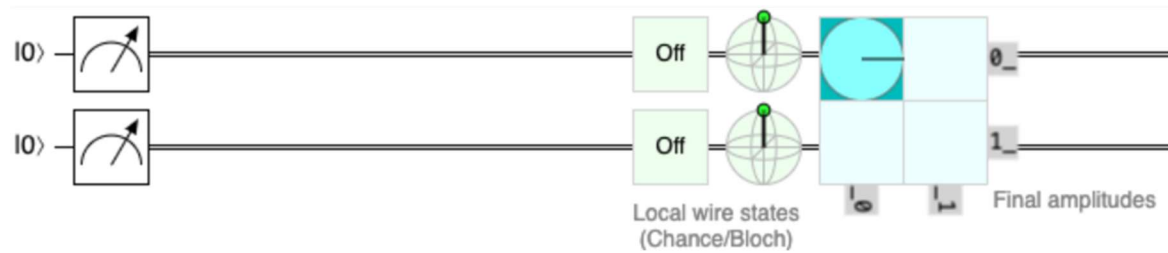
Index (I)	Benefit (V)	Weight (W)
$i_0$	2	1
$i_1$	4	4

Storage Capacity ( $W_{max}$ ): 4

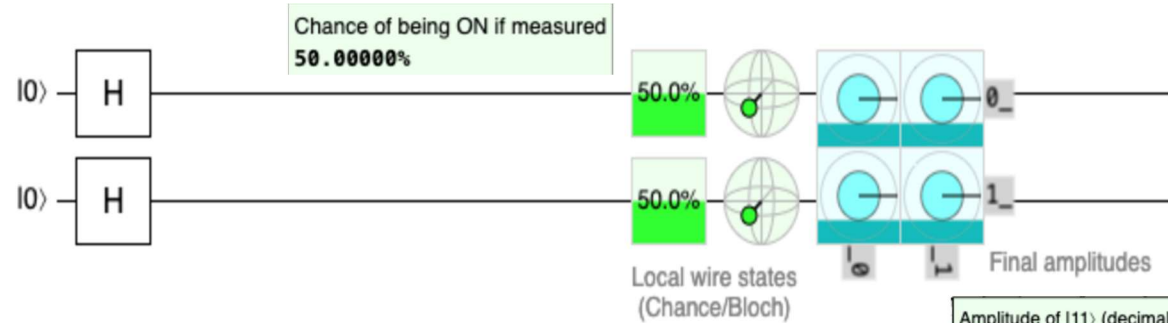
Possible Index Configurations

Index Configurations ( $i_1 i_0$ )	Aggregate Benefit	Aggregate Weight (W)
0 0	0	0
0 1	2	1
1 0	4	4
1 1	6	5

Initial state of Index Qubits



Implicit Generation of All Index Configurations in Superposition

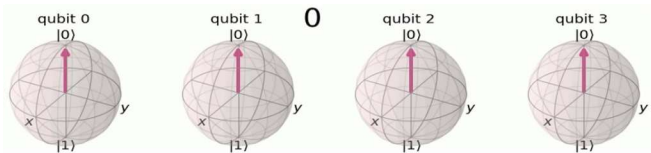


Amplitude of  $|11\rangle$  (decimal 3)  
 val:  $+0.50000+0.00000i$   
 mag<sup>2</sup>: 25.0000%, phase:  $+0.00^\circ$

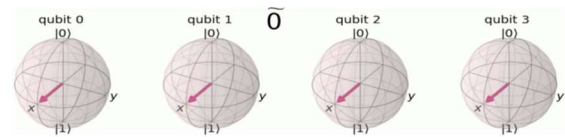
# Data Load in Qubit Phases (Step 2)

## Motivation

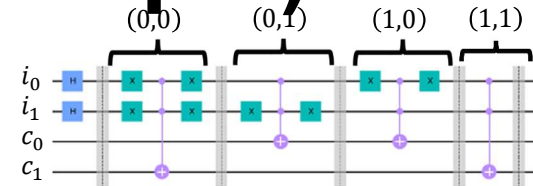
- Basis Encoding requires exponential pre-computation cost and quantum gates



Encoding in Qubit Basis State



Encoding in Fourier Basis using Qubit Phases

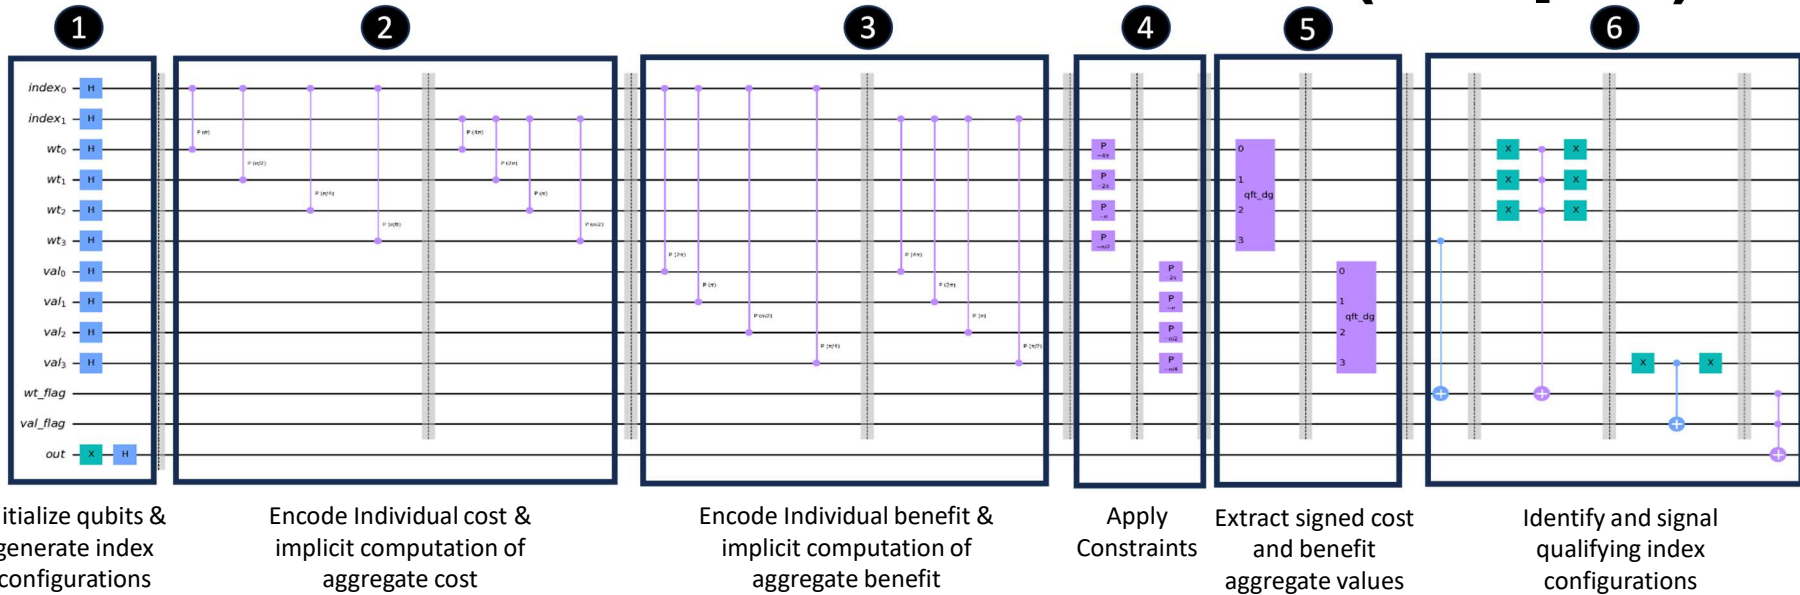


- Qubit phase manipulation, enables intrinsic computation of aggregate weights and benefits
- Additionally, simplifies identification of qualifying configurations

## Encoding data in Qubit Phases

- Consider an index  $i \in I$  with cost  $w_i$
- To load this cost into the  $m$  cost qubits, compute phase angles  $\gamma_0, \gamma_1, \dots, \gamma_{m-1}$  for each cost qubit
- Decomposing this angle computation into two parts:
  - **Rotation Angle ( $\theta_i$ ):** Rotation relative to the full rotation  $2\pi$ . Calculated as  $\theta_i = \frac{2\pi}{2^m} * w_i$
  - **Rotation Frequency:** Depends on the qubit position. For  $j^{th}$  qubit, the frequency is  $2^{m-j-1}$
  - Therefore,  $\gamma_i = 2^{m-j-1} * \theta_i$
- Use *quantum conditional phase gate* to *associate cost to index configurations present in quantum superposition state*

# Overall Quantum Oracle for IS (Step 3)



## Problem Instance

Index (I)	Cost (W)	Benefit (V)
$i_0$	1	2
$i_1$	4	4

Storage Capacity (B): 4  
Target Benefit: 2

Index Configurations ( $i_1 i_0$ )	After Data Load <b>3</b>			After Constraints <b>5</b>			After Qualification Test <b>6</b>		
	Cost	Benefit	Qualify	Cost	Benefit	Qualify	Cost	Benefit	Qualify
0 0	0	0	0	-4	-2	0	-4	-2	0
0 1	1	2	0	-3	0	0	-3	0	1
1 0	4	4	0	0	2	0	0	2	1
1 1	5	6	0	1	4	0	1	4	0

# *Experimental Evaluation*

---

# Problem Suite

Index (I)	Columns in the Index	Time Benefit (V)	Storage Cost (W)
$i_0$	[L_DISCOUNT, L_SHIPDATE, L_QUANTITY, L_EXTENDEDPRICE]	165811	266
$i_1$	[L_SHIPDATE, L_DISCOUNT, L_EXTENDEDPRICE, L_PARTKEY]	178871	232
$i_2$	[P_CONTAINER, P_BRAND, P_PARTKEY]	1213770	8
$i_3$	[L_PARTKEY, L_QUANTITY]	1213770	132
$i_4$	[L_PARTKEY, L_QUANTITY, L_EXTENDEDPRICE]	1213770	199
$i_5$	[C_ACCTBAL, C_CUSTKEY, C_PHONE]	44370	2
$i_6$	[O_CUSTKEY]	44370	9
<b>Storage Capacity: 140</b>			

**CDB\_I7**

Index (I)	Benefit (V)	Cost (W)
$i_0$	14	7
$i_1$	12	6
$i_2$	10	4
$i_3$	8	2
$i_4$	3	1
<b>Storage Capacity: 19</b>		

**I5**

Index (I)	Benefit (V)	Cost (W)
$i_0$	14	7
$i_1$	12	6
$i_2$	10	4
$i_3$	8	2
$i_4$	3	1
$i_5$	13	7
<b>Storage Capacity: 19</b>		


**I6**

Index (I)	Benefit (V)	Cost (W)
$i_0$	14	7
$i_1$	12	6
$i_2$	10	4
$i_3$	8	2
$i_4$	3	1
$i_5$	13	7
$i_6$	12	6
<b>Storage Capacity: 19</b>		

**I7**

# Evaluation on a noiseless Quantum Simulator

Problem	# of Candidate Configurations	IA Scheme	Configuration Quality			Quantum Resources		Normalized Overhead
			Weighted Average	Optimal Fraction	Worst Case	Qubits	Depth	
CDB_I7	128	Exhaustive	1.0	1.0	1.0	-		100%
		SQIA ( $\delta = 0.9, \alpha = 0.26, S = 1$ )	0.95	0.9	0.54	28	80	73%
		OQIA ( $p = 1, S = 100$ )	0.76	0.5	0.5	15	30	23%
		Greedy	0.54	0	0.54	-		-
I7	128	Exhaustive	1.0	1.0	1.0	-		100%
		SQIA ( $\delta = 0.9, \alpha = 0.26, S = 1$ )	1.0	1.0	1.0	25	68	75%
		OQIA ( $p = 1, S = 100$ )	0.97	0.4	0.89	12	24	23%
		Greedy	0.8	0	0.8	-		-
I6	64	Exhaustive	1.0	1.0	1.0	-		100%
		SQIA ( $\delta = 0.9, \alpha = 0.22, S = 1$ )	1.0	1.0	1.0	22	60	94%
		OQIA ( $p = 1, S = 100$ )	0.97	0.6	0.86	11	22	45%
		Greedy	0.8	0	0.8	-		-
I5	32	Exhaustive	1.0	1.0	1.0	-		100%
		SQIA ( $\delta = 0.9, \alpha = 0.18, S = 1$ )	0.9	0.8	0.34	21	58	111%
		OQIA ( $p = 1, S = 100$ )	0.99	0.9	0.89	10	20	91%
		Greedy	0.8	0	0.8	-		-



# Evaluation on IBM 127-qubit Eagle Circuit Processor

Problem	IA Scheme	Configuration Quality
CDB_I7	SQIA ( $\delta = 0.9, \alpha = 0.26, S = 1$ )	0.85
	OQIA ( $p = 1, S = 100$ )	0.68
I5	SQIA ( $\delta = 0.9, \alpha = 0.18, S = 1$ )	0.89
	OQIA ( $p = 1, S = 100$ )	1.0
I6	SQIA ( $\delta = 0.9, \alpha = 0.22, S = 1$ )	0.86
	OQIA ( $p = 1, S = 100$ )	0.98
I7	SQIA ( $\delta = 0.9, \alpha = 0.26, S = 1$ )	0.95
	OQIA ( $p = 1, S = 100$ )	1.0

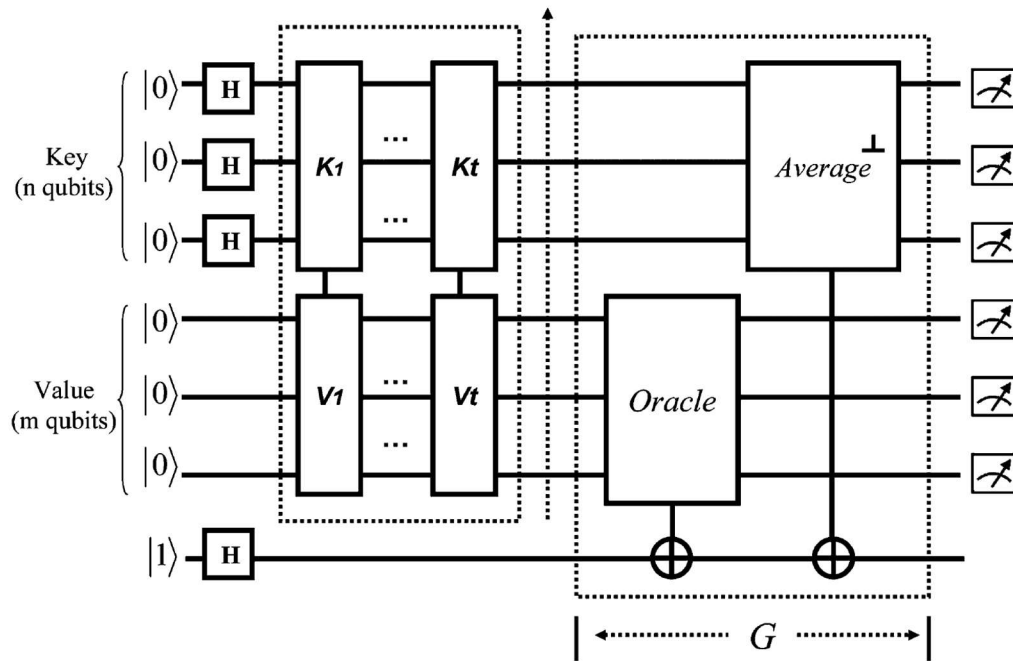
*Stage IV: Quantum  
Query Execution*

---

# Existing Literature

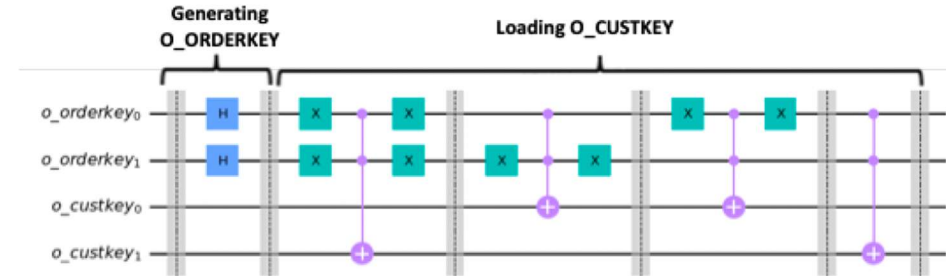
- Cockshott, Paul. ***Quantum relational databases***. arXiv:quant-ph/9712025, **1997**
- Ahmed Younes. ***Database Manipulation on Quantum Computers***. arXiv:0705.4303, **2007**
- Yi-Lin Ju, I-Ming Tsai, and Sy-Yen Kuo. 2007. ***Quantum circuit design and analysis for database search applications***. IEEE Transactions on Circuits and Systems, **2007**
- Amor Gueddana, Rihab Chatta, and Nouredine Boudriga. ***Optimized methods for inserting and deleting records and data retrieving in quantum database*** IEEE International Conference on Transparent Optical Networks, **2010**
- Szabolcs Jóczik and Attila Kiss. ***Quantum Computation and Its Effects in Database Systems***. In *New Trends in Databases and Information Systems – ADBIS*, **2020**
- Rihan Hai, Shih-Han Hung, Tim Coopmans, and Floris Geerts. ***Data Management in the Noisy Intermediate-Scale Quantum Era***. arXiv:2409.14111, **2024**
- Manish Kesarwani, and Jayant Haritsa. ***Is Quantum-Based SQL Query Execution Viable?*** In VLDB Workshop: The Second International Workshop on Quantum Data Science and Management (QDSM) **2024**.

# Classical-Quantum Data Loading



O_ORDERKEY	O_CUSTKEY
0	2
1	1
2	1
3	2

Orders Table



# Potential Benefits of SQL Query Execution on QP

- Grover Search could speed up *complex selection operations on non-indexed columns*
- Reduce time complexity for *arbitrary join predicates*
- *Union, Intersection, and Set Difference operators* could be benefited
- Could improve the *efficiency of the Sort operation* in space-bounded settings
- Quantum Fourier Transform (QFT) could offer *effective handling of temporal and spatial data queries*
- Quantum entanglement could *enable distributed SQL processing*, improving global data consistency

***Complex Selection &  
Join Predicates***

***Set & Sort Operators***

***Temporal &  
Distributed SQL  
query processing***

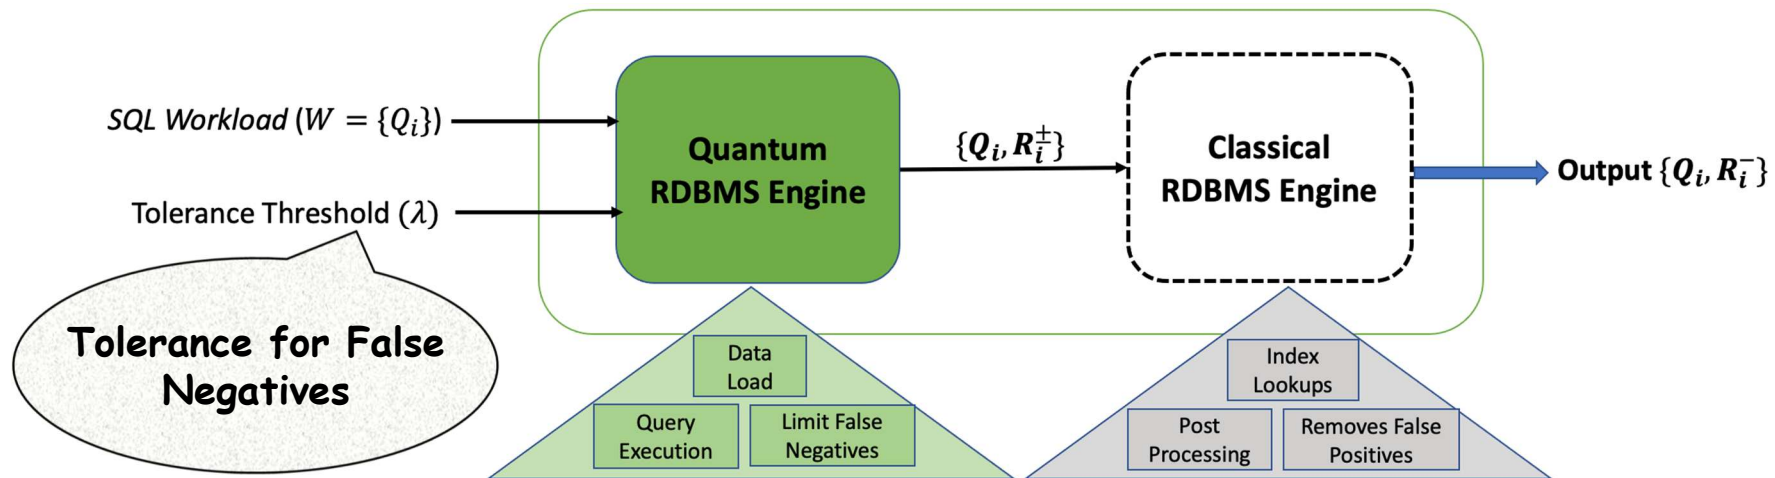
# *Is Quantum-Based SQL Query Execution Viable?*

Manish Kesarwani, Jayant R. Haritsa  
[QDSM (VLDB Workshop), 2024]

---

# Alternative SQL Query Contract

- For **quantum-enhanced RDBMS for SQL query processing** proposed updating the existing contract
- First, follow **a hybrid quantum-classical system**
  - Eliminate false negatives in the quantum domain and false positives in the classical space
- Second, specify **a nonzero tolerance threshold ( $\lambda \in (0, 1]$ )** along with the SQL query
  - Expected percentage of false negatives that the user is willing to accept



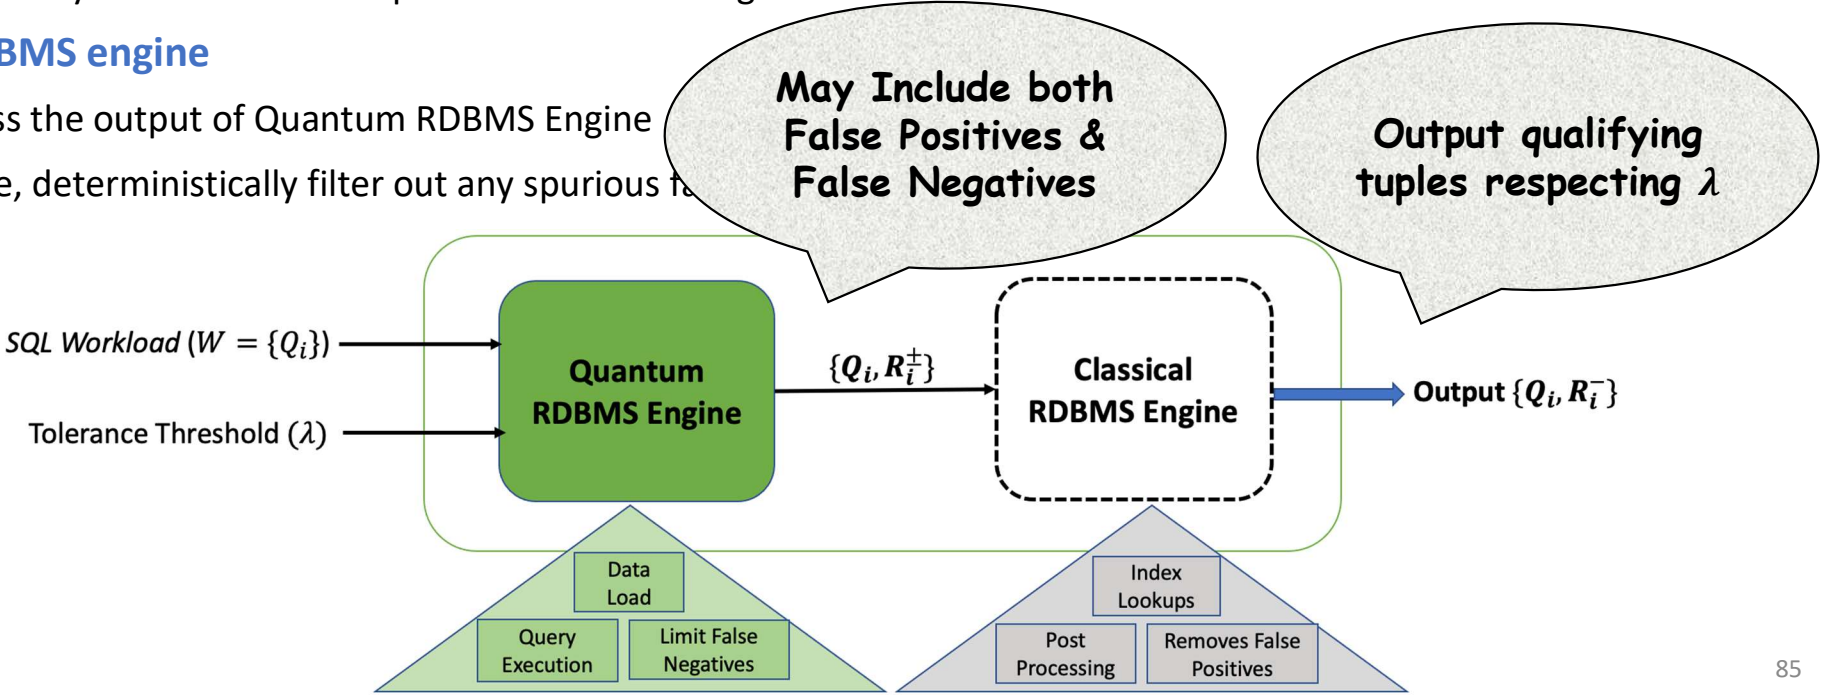
# Alternative SQL Query Contract

## Quantum RDBMS Engine

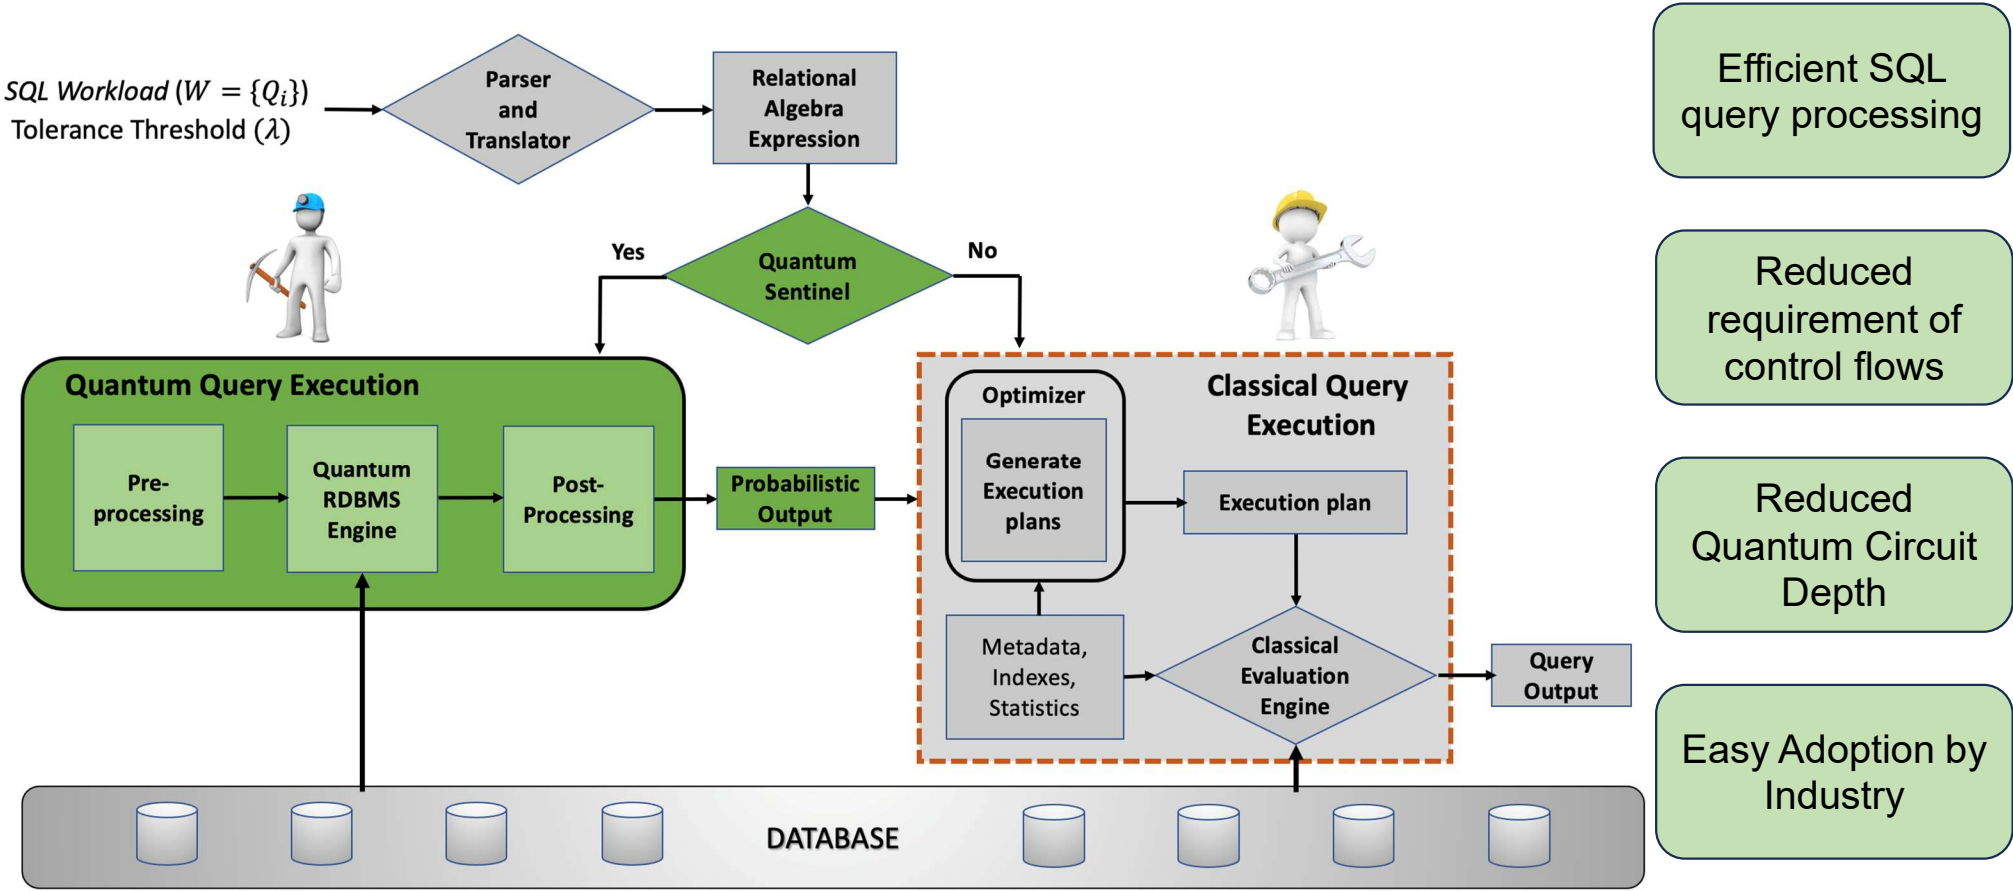
- Load the classical data into qubits and
- Construct quantum circuit encoding the SQL query
- Adapt a quantum algorithm, such as Grover Search (GS), to identify and produce qualifying tuples
  - Results may include both false positives and false negatives

## Classical RDBMS engine

- Post-process the output of Quantum RDBMS Engine
- For instance, deterministically filter out any spurious tuples



# Vision: Quantum-Classical Hybrid SQL Query Execution



Efficient SQL query processing

Reduced requirement of control flows

Reduced Quantum Circuit Depth

Easy Adoption by Industry

# Potential Application Areas

- **Approximate query processing**
  - Provide **quick, approximate answers** instead of exact results for queries
- **Sampling based query**
  - Quantum computers can generate **uniform and unbiased samples** efficiently
  - Quantum Amplitude Estimation (QAE) **requires quadratically fewer samples** than classical algorithms
  - For an error margin  $\epsilon$ , QAE needs only  $O(1/\epsilon)$  samples, while classically we require  $O(1/\epsilon^2)$  samples
- **Trend Analytics**
  - Classical Fourier Transform requires  $O(N \log N)$  operations for a dataset of size  $N$ .
  - Quantum Fourier Transform (QFT) operates in  $O(\log N)^2$  time providing exponential speedup
  - QFT can efficiently detect periodic patterns
- **NoSQL Databases**
  - Probabilistic quantum walks can optimize graph traversal tasks
  - Central to graph databases like Neo4j

*Stage V: Open problems  
and challenges for  
database research*

---

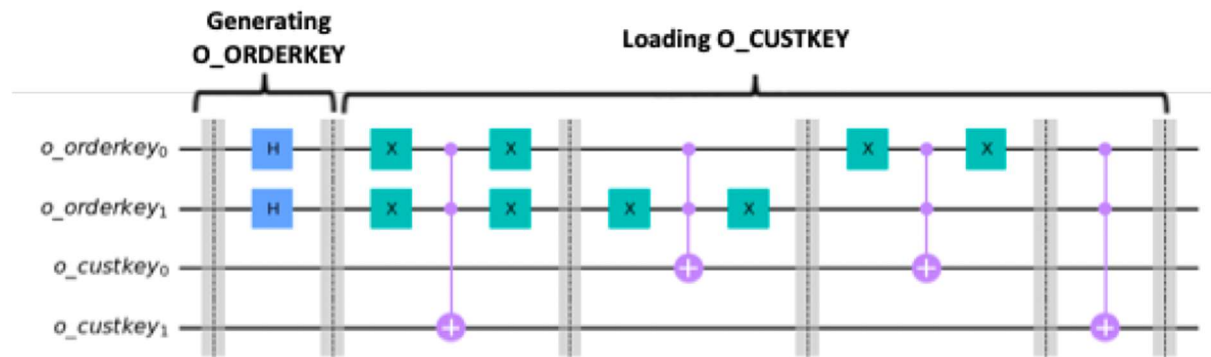
# Open problems and challenges

## 1. Data Encoding

- Efficiently encoding classical data into qubits
- Scalability to large datasets

O_ORDERKEY	O_CUSTKEY
0	2
1	1
2	1
3	2

Orders Table



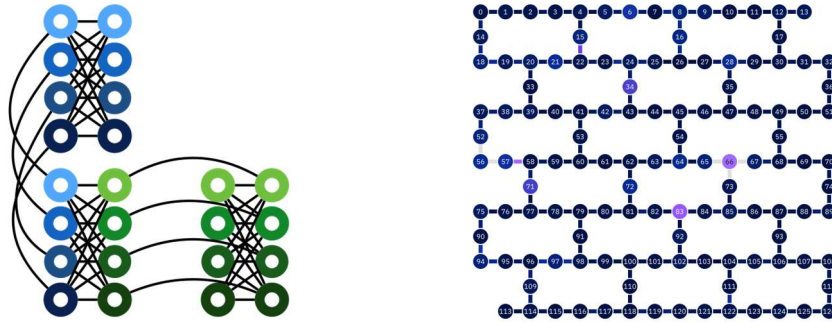
## 2. Custom Gates

- Arbitrary quantum gates are theoretically feasible
- Concerns feasibility of implementation on real quantum computers

## 3. Bounding the impact of Probabilistic Output

# Open problems and challenges

## 4. Efficient embedding to the Qubit Layout



## 4. Resource Efficiency

- Reducing qubit and gate overhead while maintaining accuracy
- Sparse qubit connectivity restricts quantum algorithms

## 5. Usability and Interfaces

- Bridging the gap between quantum computing and database professionals and clients

Knob Tuning

Cardinality Estimation

Index/View Benefit Estimation

Index/View Selection

Query Rewrite

Plan Enumeration

SQL Query Execution

# *Stage VI: Practical*

---

# *6.1 Demo*

---

## *6.2 Hands-On Exercise*

---