TAMING THE INSTRUCTION BANDWIDTH OF QUANTUM COMPUTERS VIA HARDWARE MANAGED ERROR CORRECTION

MICRO-50

Swamit Tannu

Zachary Myers
Douglas Carmean

Prashant Nair
Moinuddin Qureshi

Georgia Tech

Stanford University

Microsoft

Why Quantum Computers?

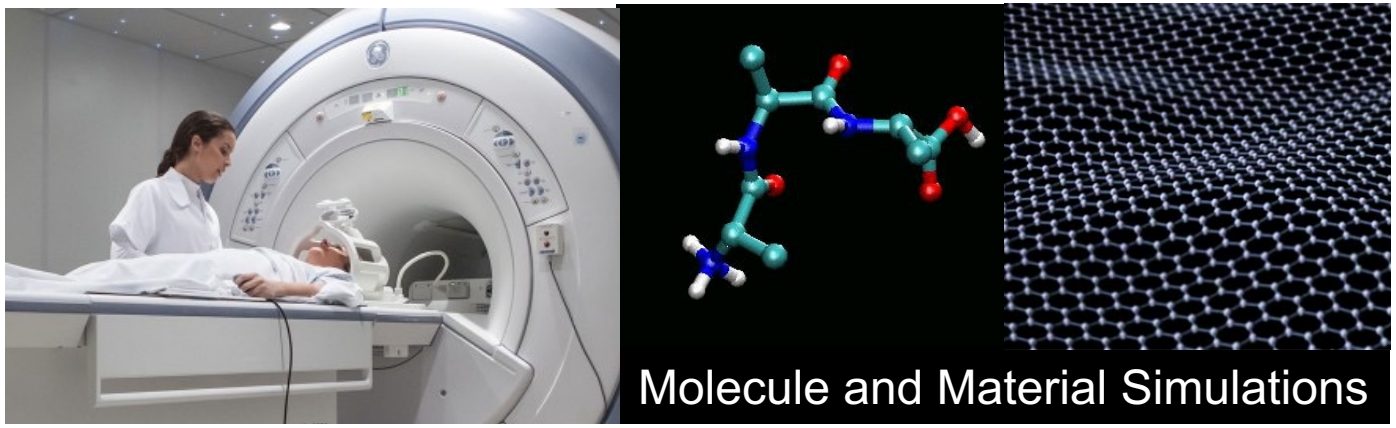

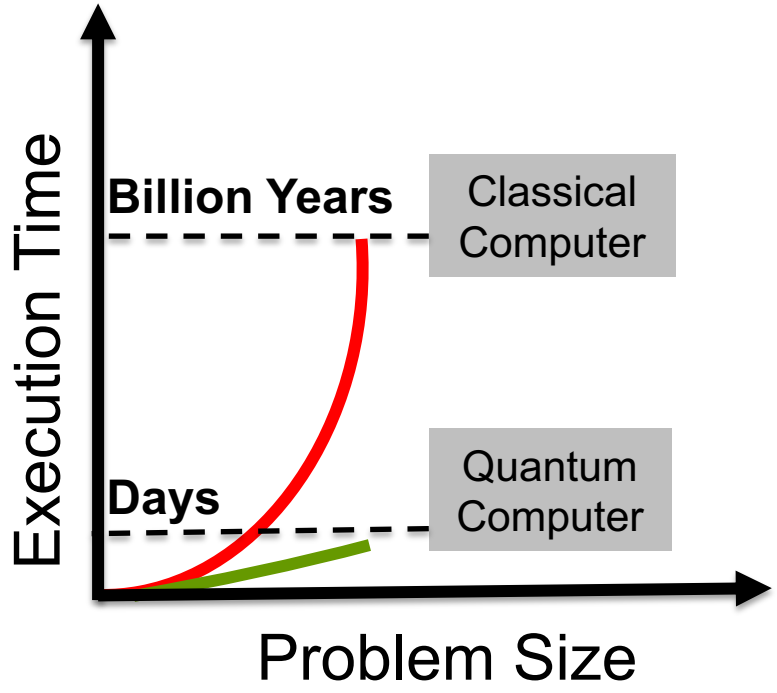

❖ Quantum computers provide large speedup for problems in material science, machine learning, and medicine

Quantum Computers enable solutions to important problems

Quantum Bits: Background

- State of a *Classical Bit*

→ 1 or 0 two points on sphere

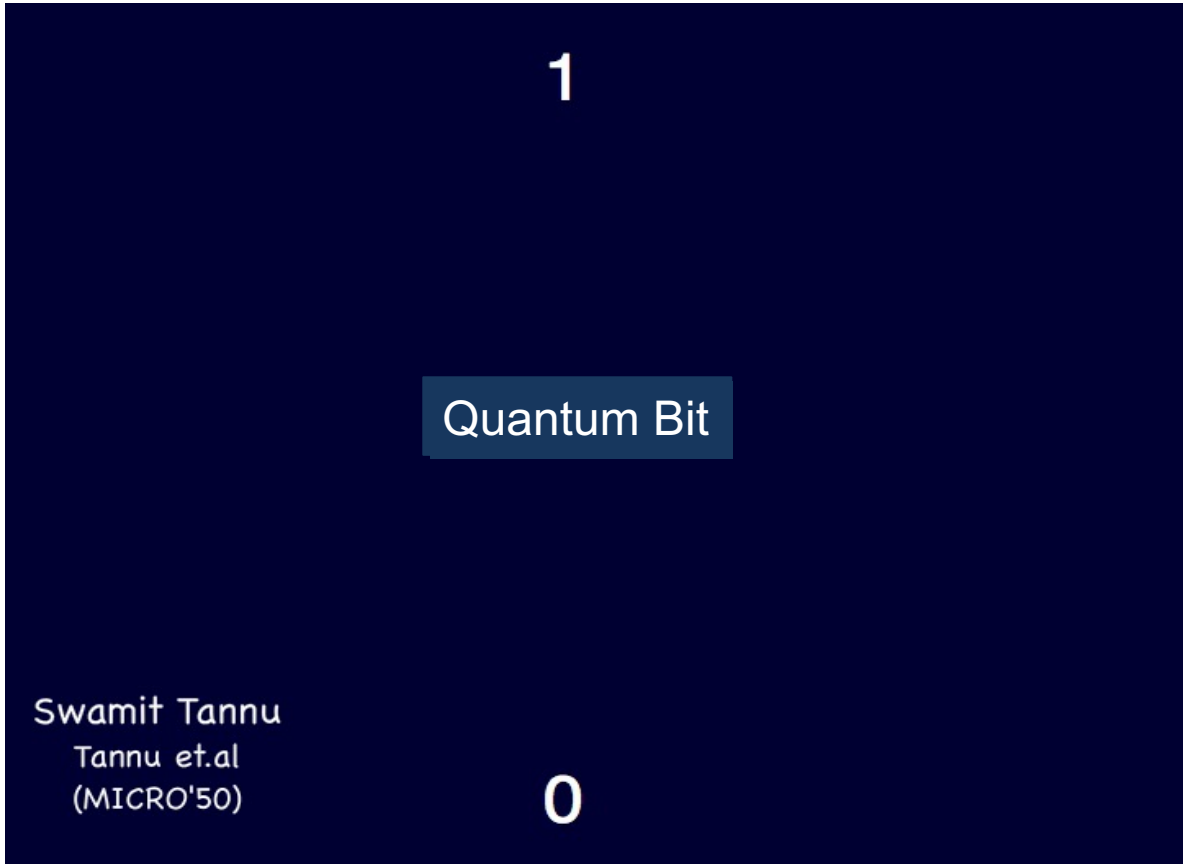

Quantum computer use quantum bits (qubits) to encode the information.

Quantum Bits: Background

- State of a *Classical Bit*

→ 1 or 0 two points on sphere

- State of a *Quantum Bit*

→ Any point on the sphere

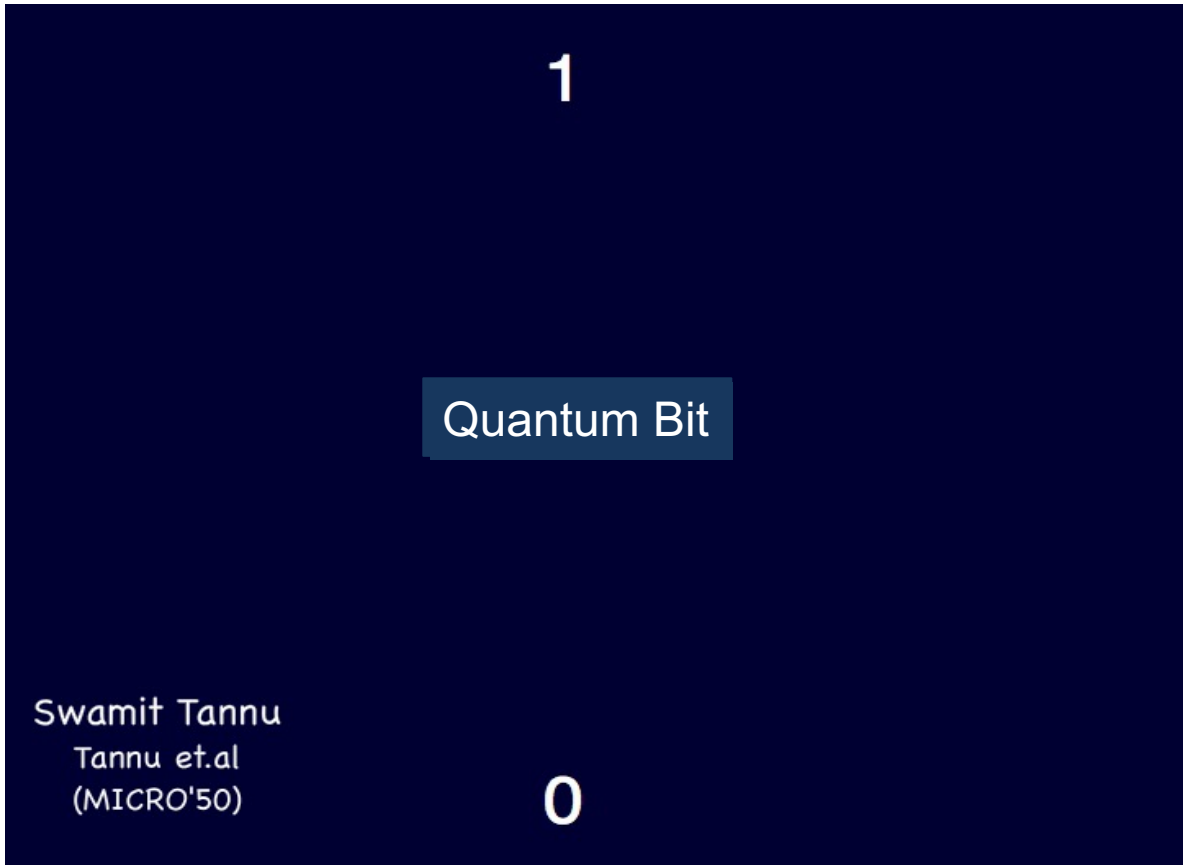

Quantum computer use quantum bits (qubits) to encode the information.

Quantum Instruction: Background

- Quantum Instructions manipulate the state of qubit

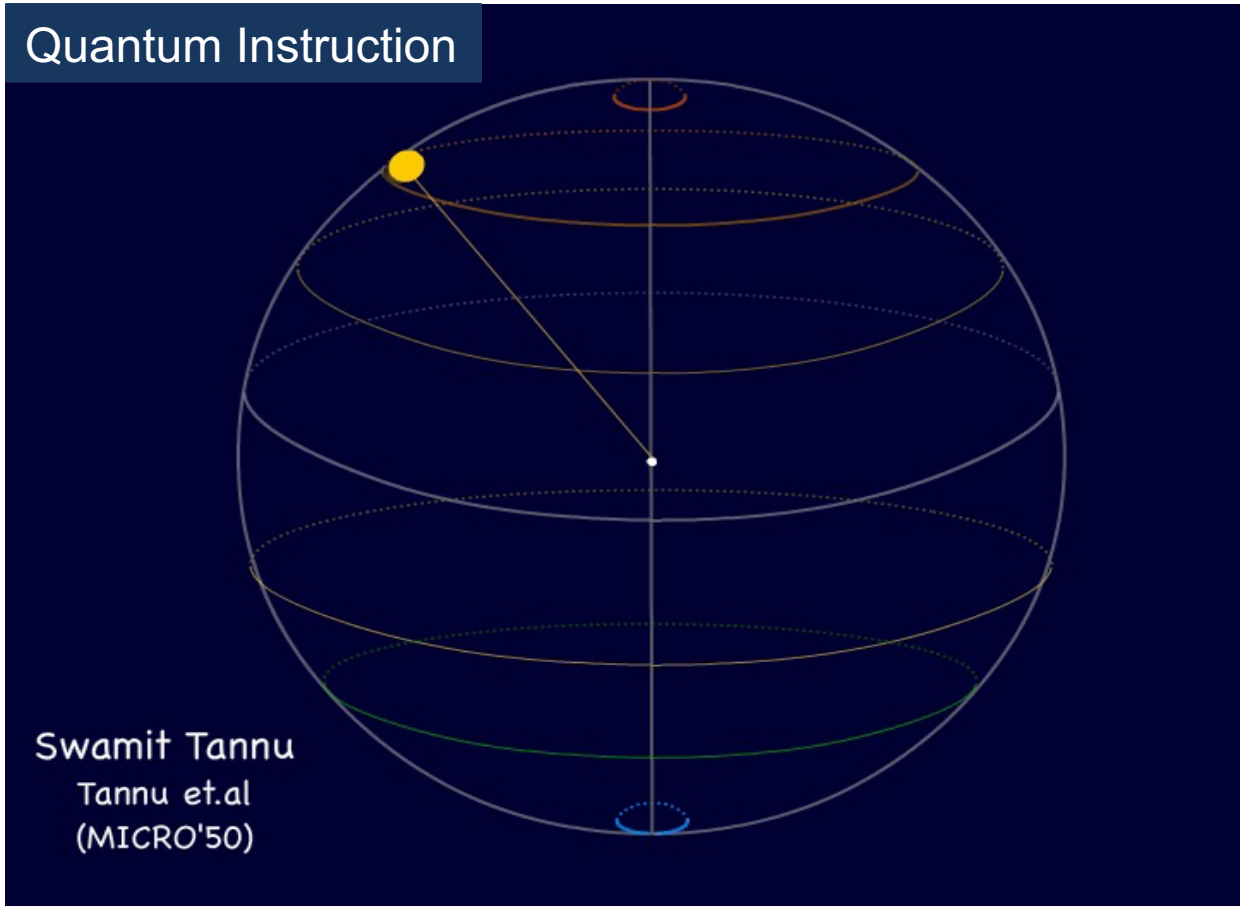

Quantum Instruction: Background

- ❖ Quantum Instructions manipulate the state of qubit
- → By rotating state vector in the Hilbert space

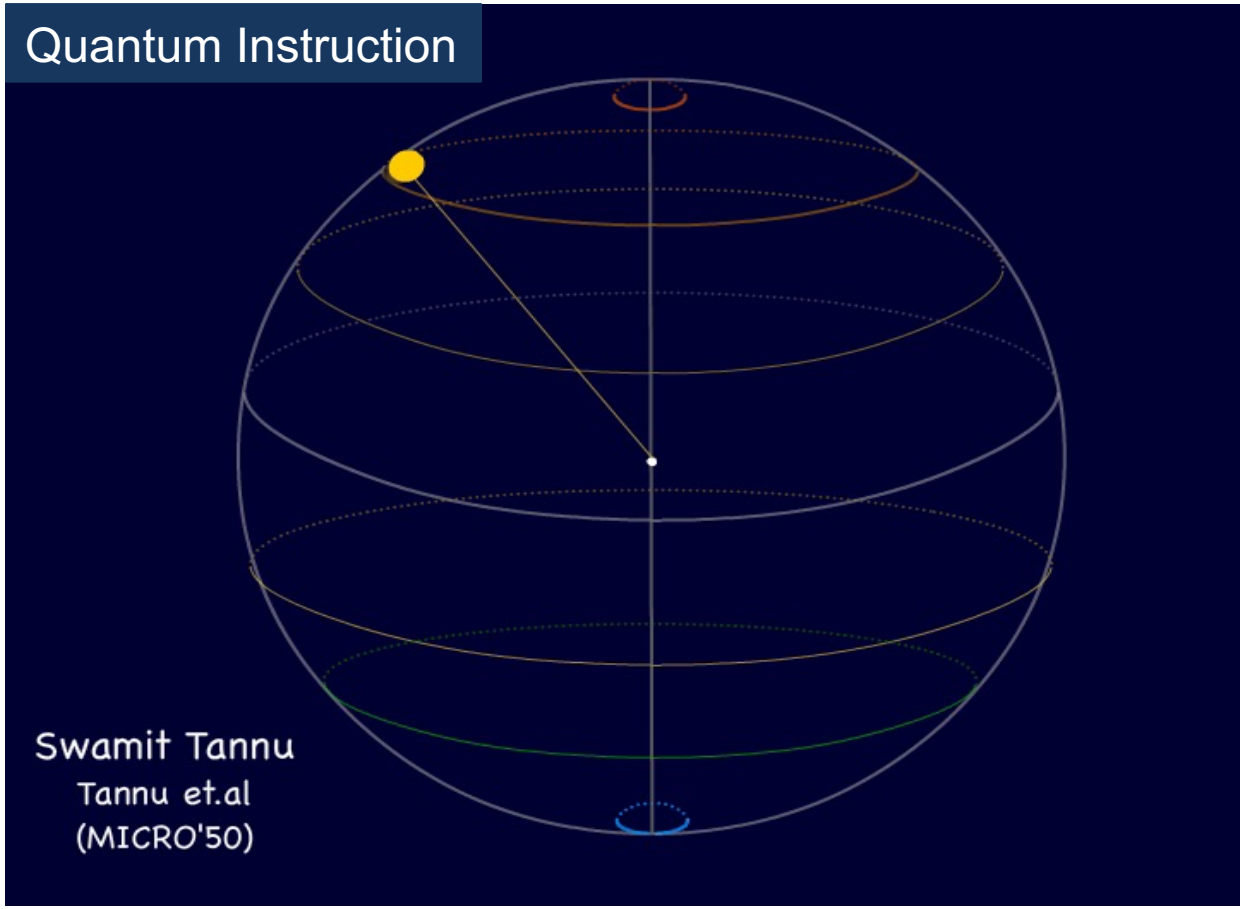

Organization of Quantum Computer

Qubits

Quantum Computer

Organization of Quantum Computer

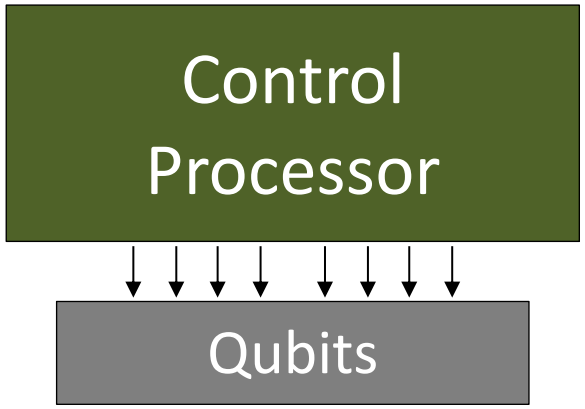

Quantum Computer

Organization of Quantum Computer

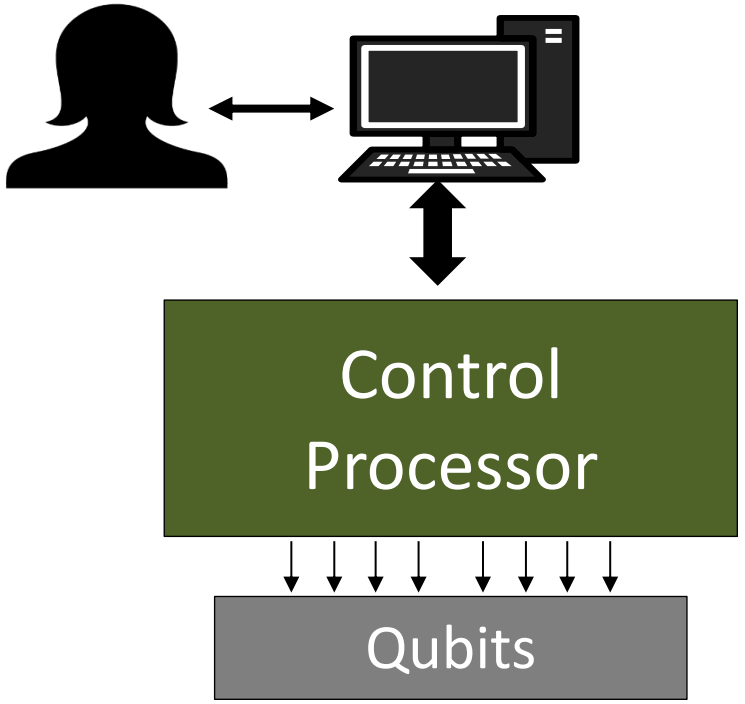

Quantum Computer

Organization of Quantum Computer

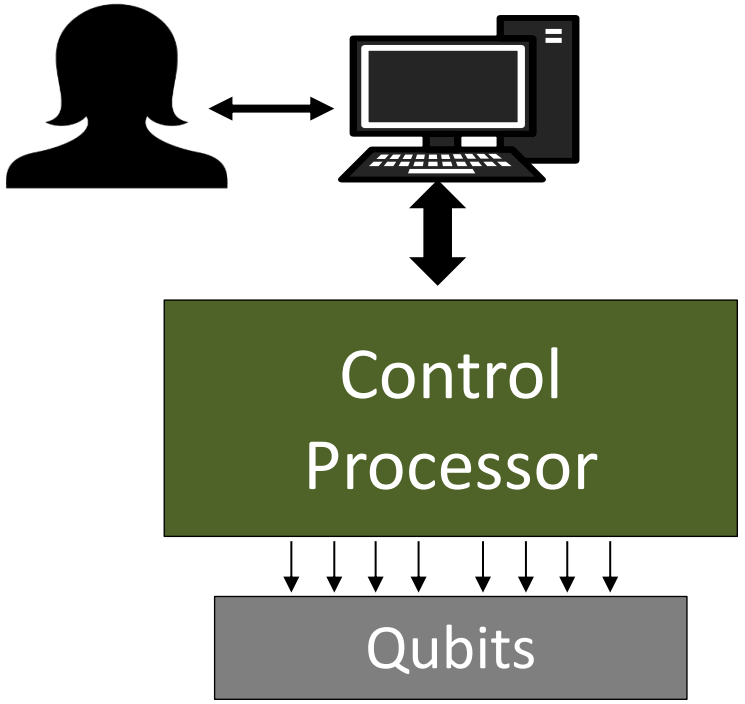

Quantum Computer

Control Processor -- Interface between Qubits & Programmer

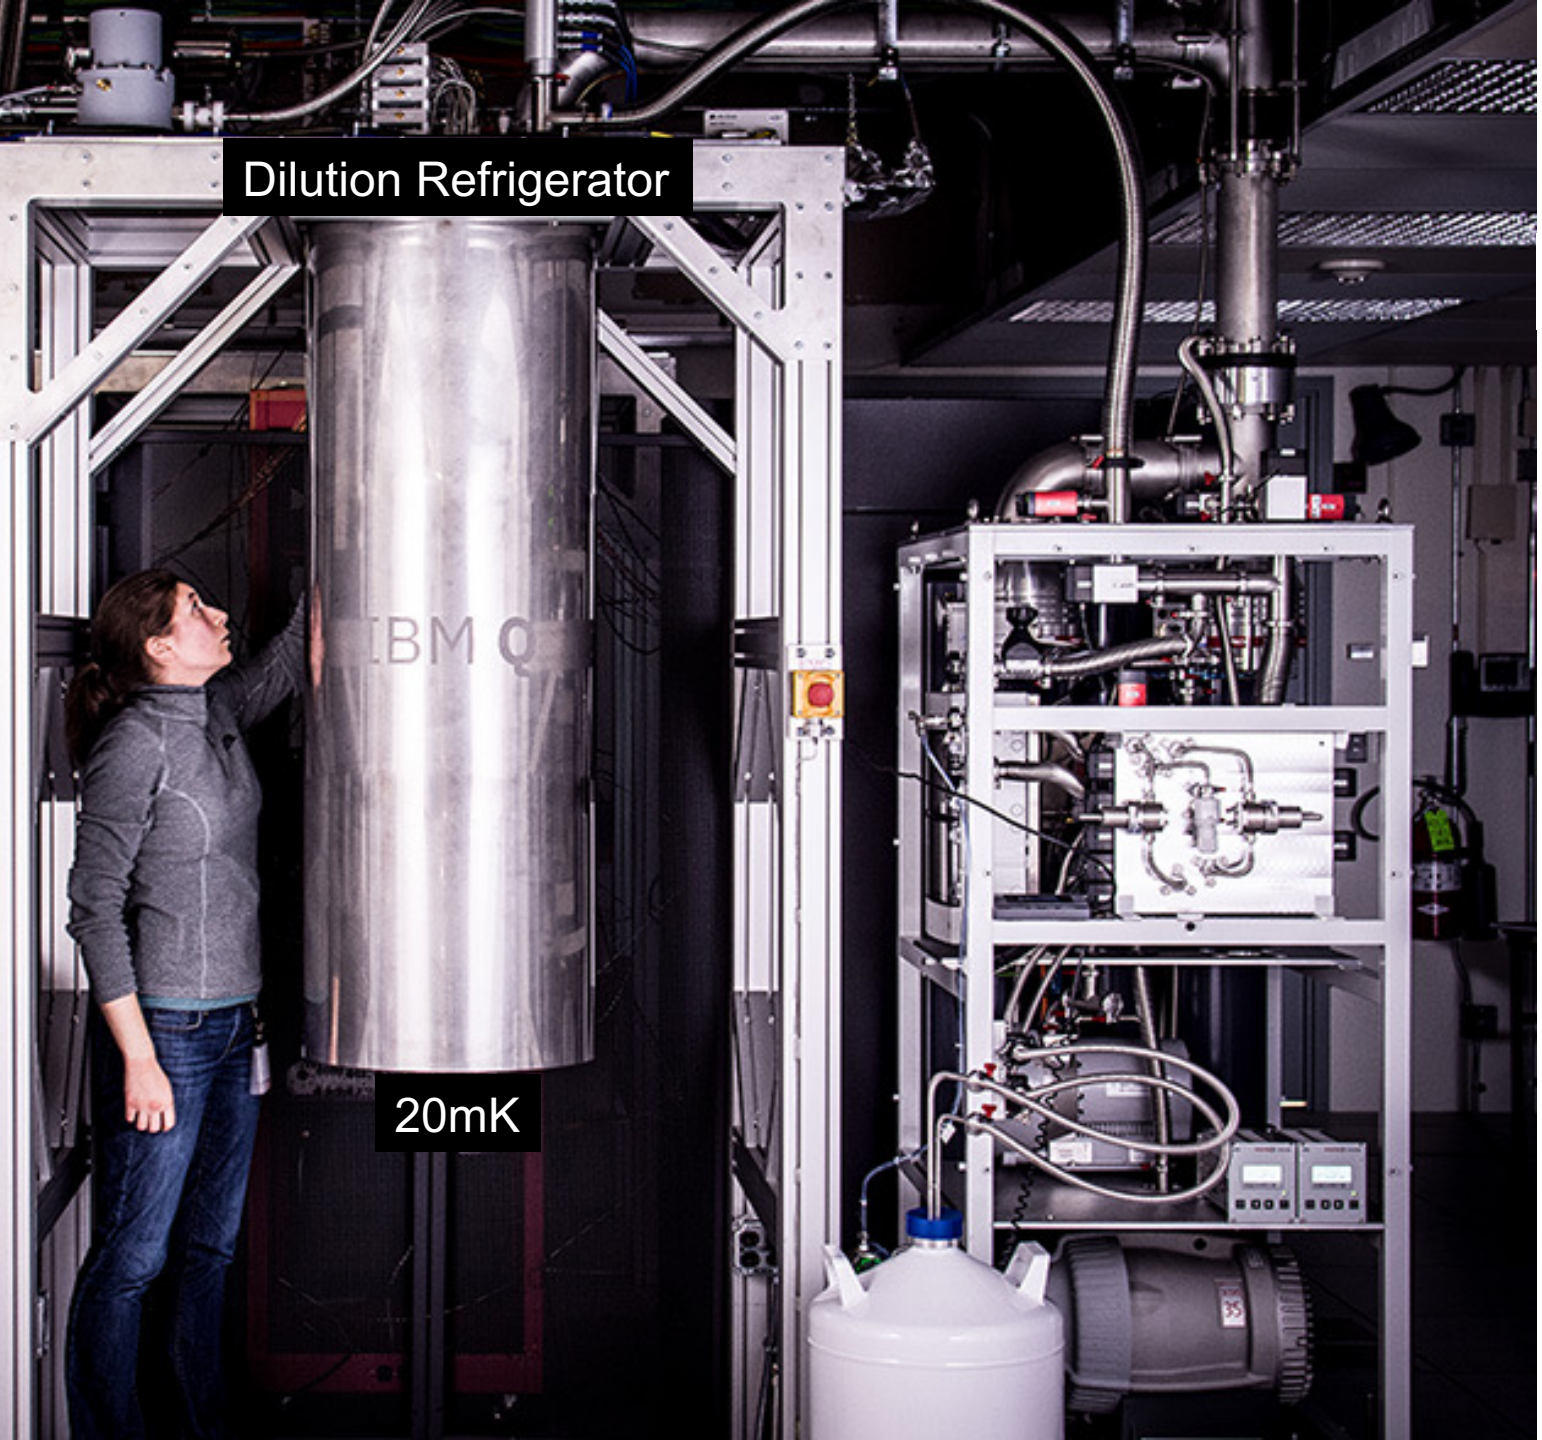

Ref: IBM Quantum Experience

Dilution Refrigerator

Qubits

20mK

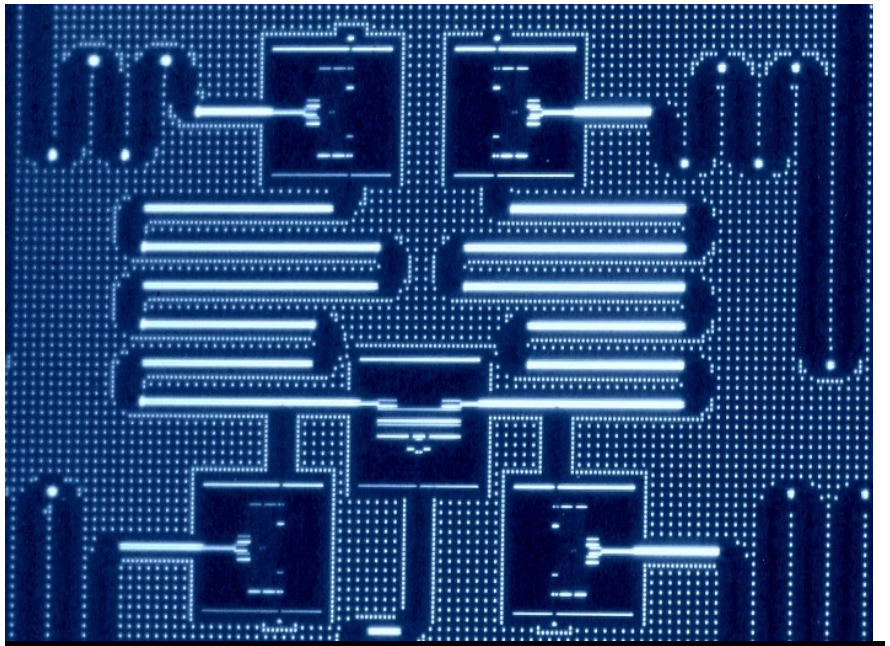

5 Qubit Chip (IBM)

Ref: IBM Quantum Experience

Cryogenic Control Processor

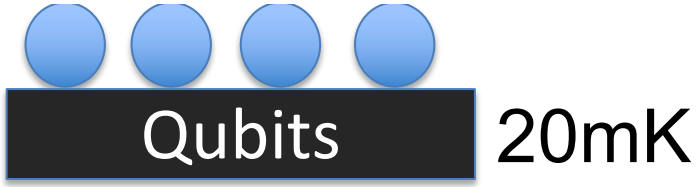

Cryogenic Control Processor

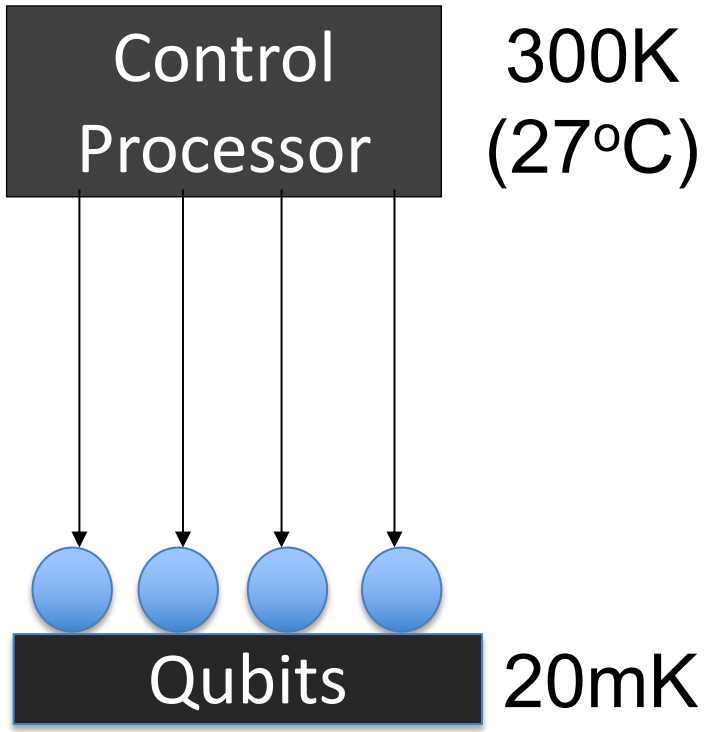

Cryogenic Control Processor

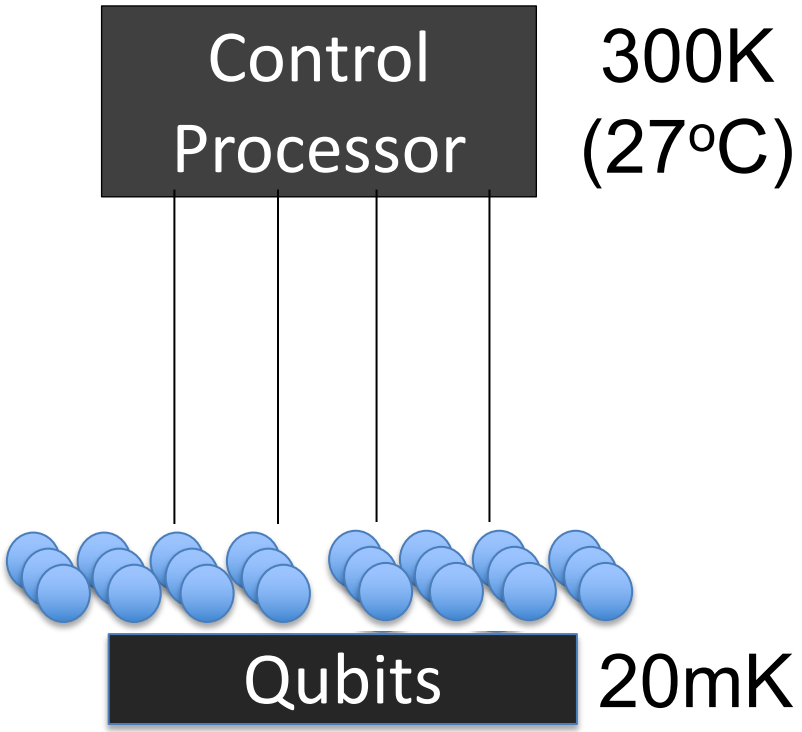

Cryogenic Control Processor

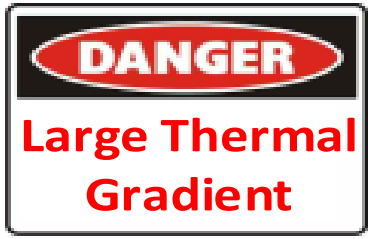

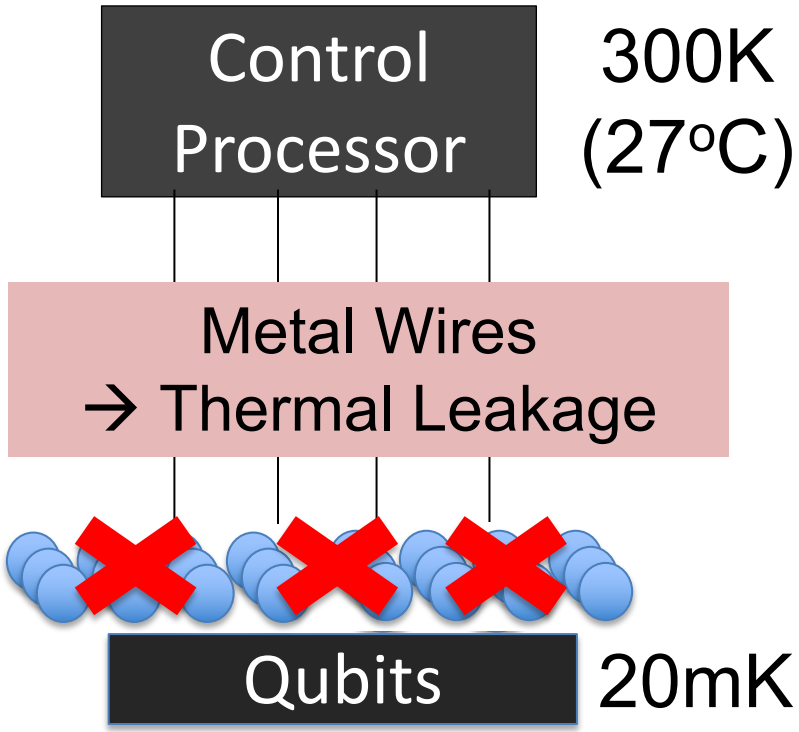

Cryogenic Control Processor

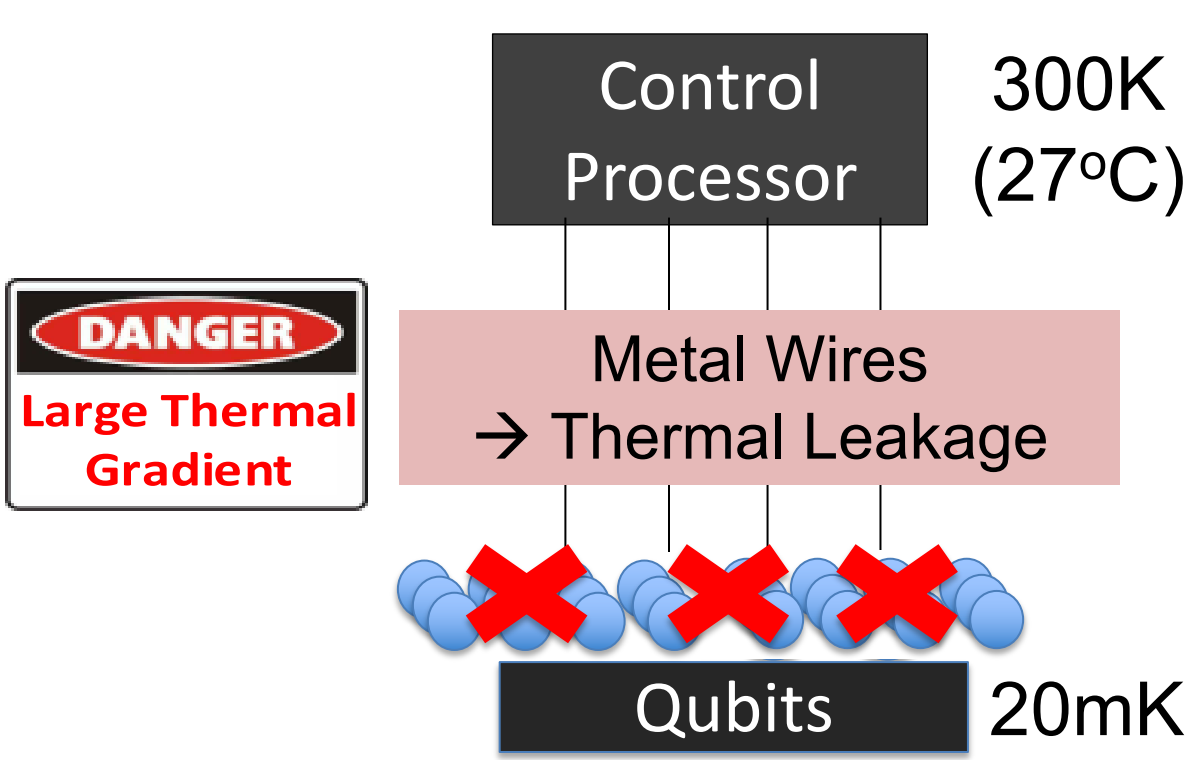

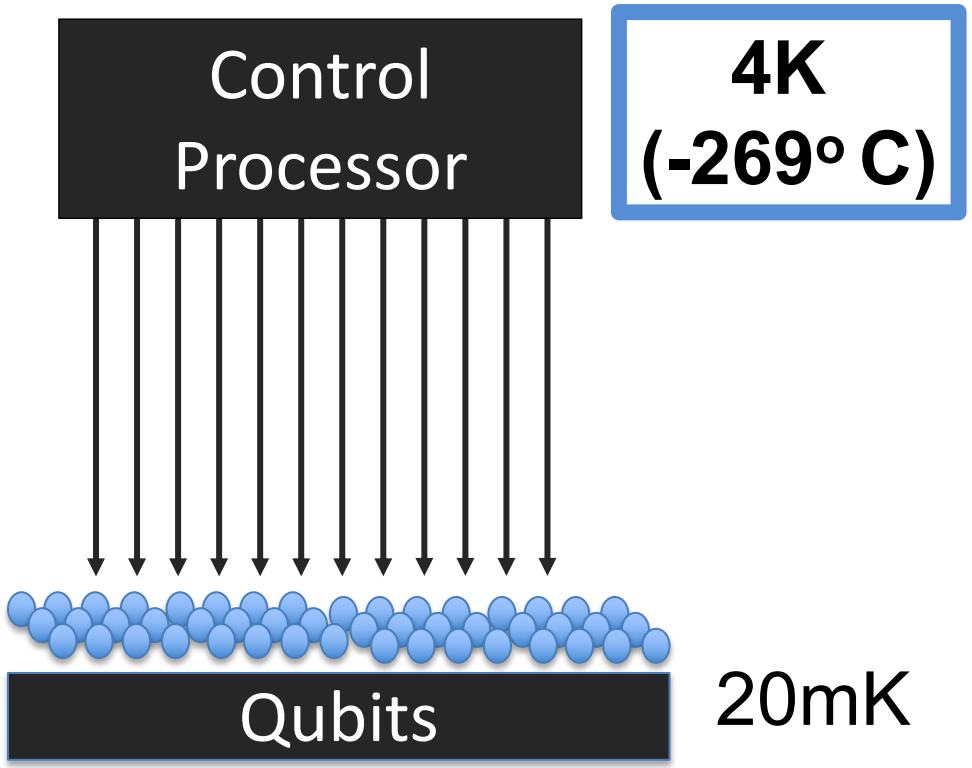

Cryogenic Control Processor is essential for scalable Quantum Computer

(Ref: Cryogenic Control Architecture for Large-Scale Quantum Computing by Hornibrook et. al)

Cryogenic Control Processor

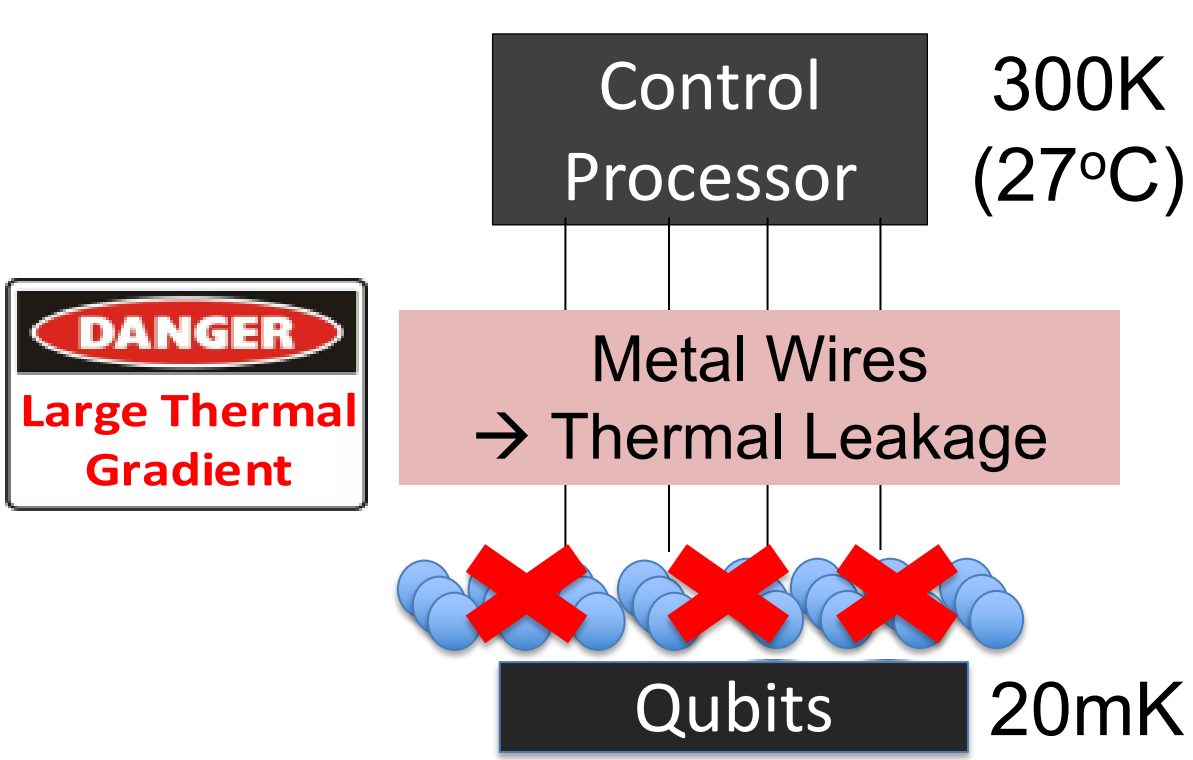

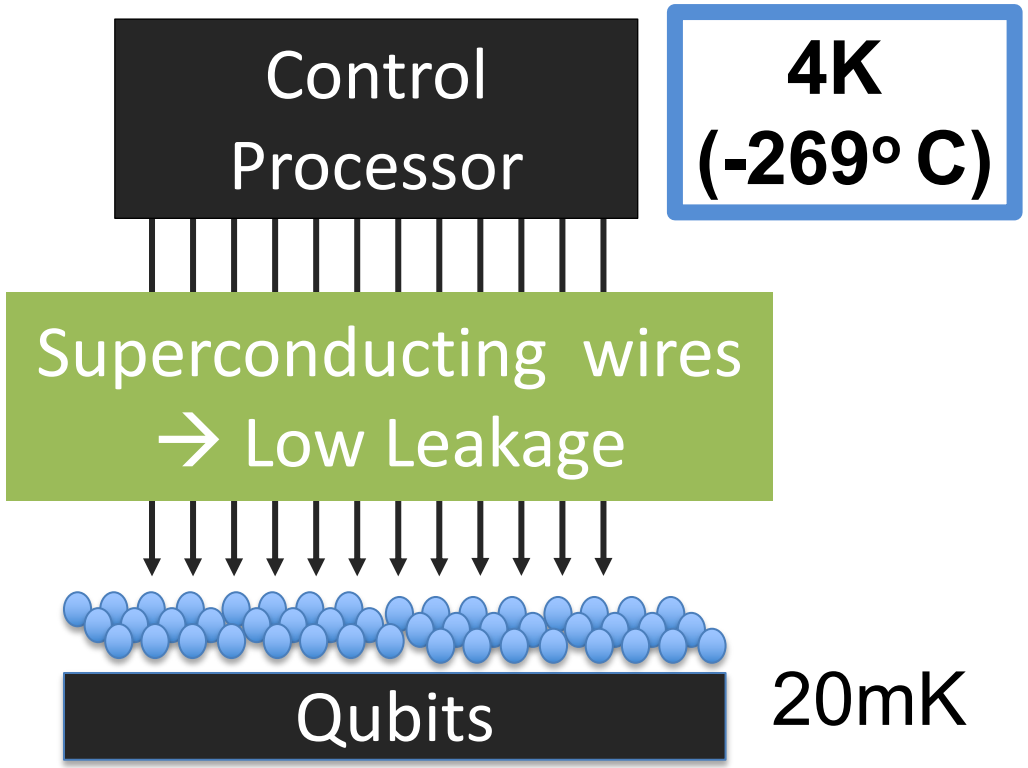

Cryogenic Control Processor is essential for scalable Quantum Computer

(Ref: Cryogenic Control Architecture for Large-Scale Quantum Computing by Hornibrook et. al)

Scalable Organization

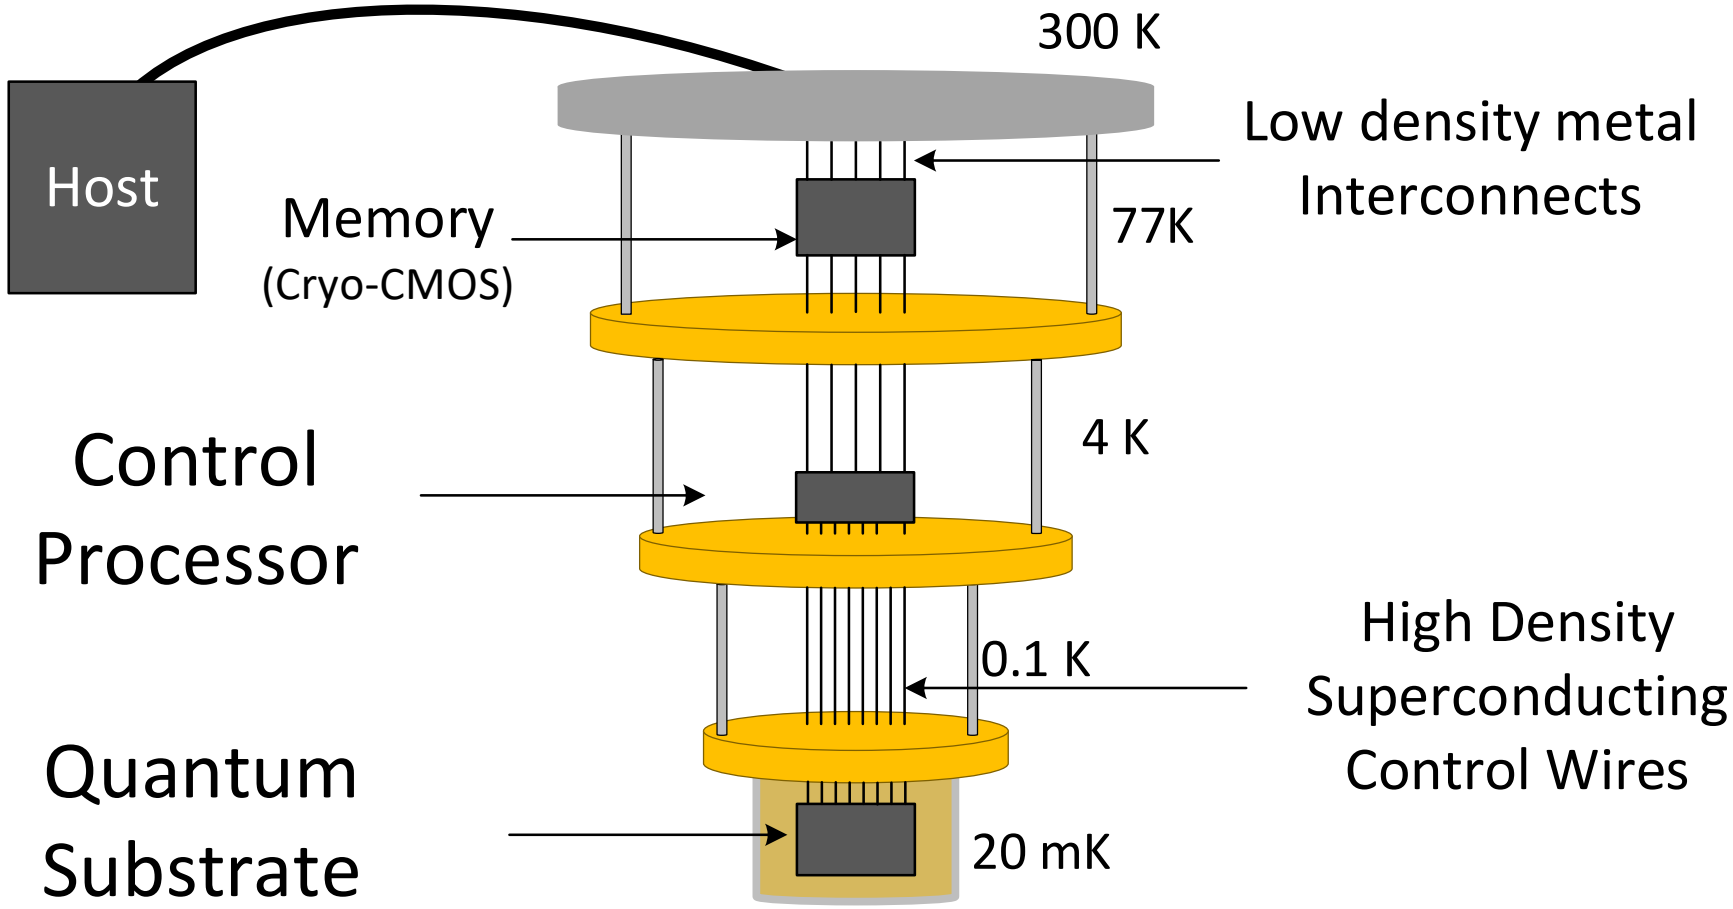

Outline

- Background
- Problem → Instruction Bandwidth Bloat due to Error Correction
- Design → HW managed error correction
- Challenges
- Evaluation & Conclusion

Quantum bits are fickle

Even at *20mK,* qubits can lose state

Classical Bit

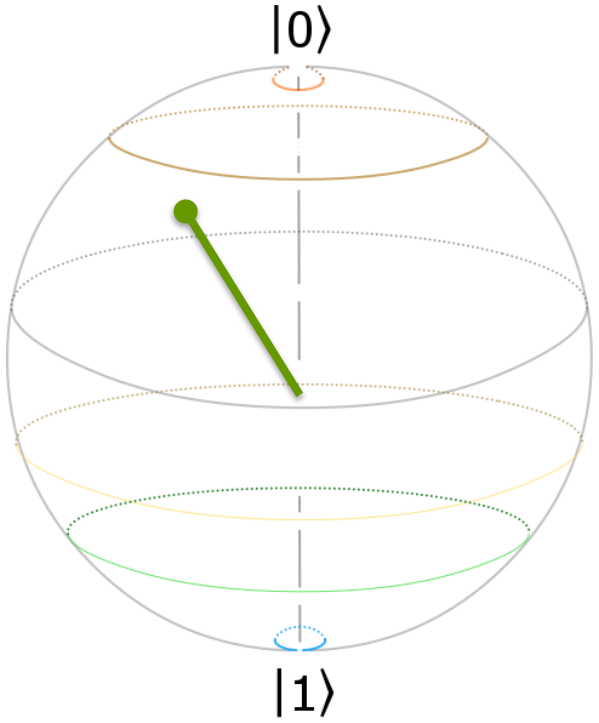

Quantum Bit

Quantum bits are fickle

Even at *20mK,* qubits can lose state

Classical Bit

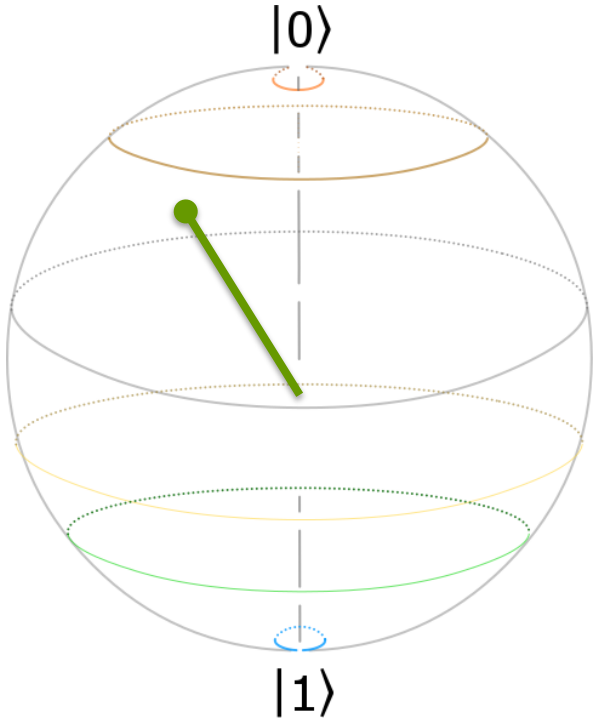

Quantum Bit

Quantum bits are fickle

Even at *20mK,* qubits can lose state

Classical Bit

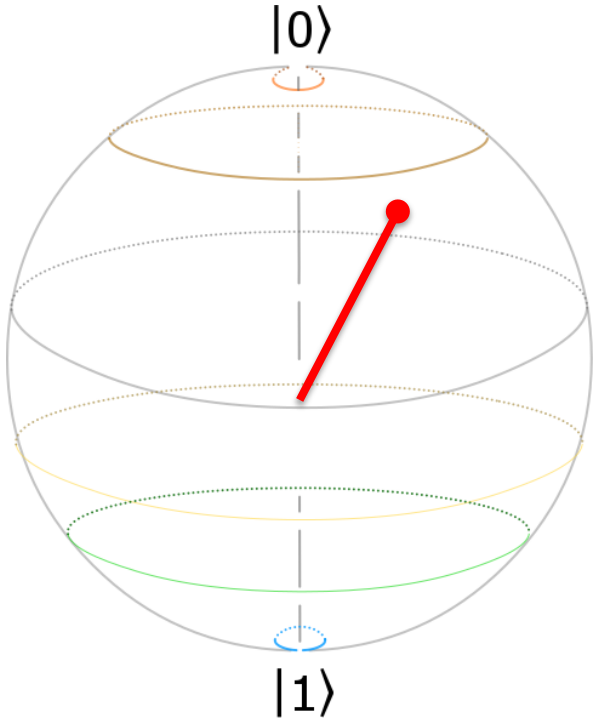

Quantum Bit

Need Error Correction to protect Quantum bits

Quantum Error Correction

- Copying qubits is not allowed

- Measurement destroys the qubit state

Quantum Error Correction

- Copying qubits is not allowed
- Measurement destroys the qubit state
- Quantum Error Correction can protect qubits

Continuous Quantum Error is essential to protect the state

Quantum Error Correction

- Copying qubits is not allowed
- Measurement destroys the qubit state
- Quantum Error Correction can protect qubits

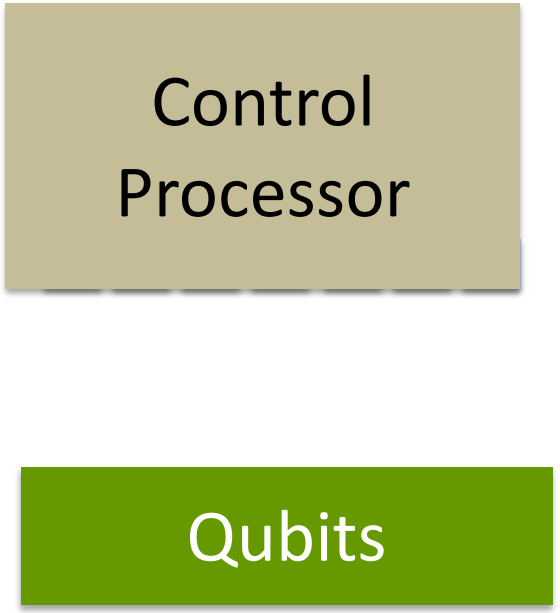

Continuous Quantum Error is essential to protect the state

Quantum Error Correction

- Copying qubits is not allowed
- Measurement destroys the qubit state
- Quantum Error Correction can protect qubits

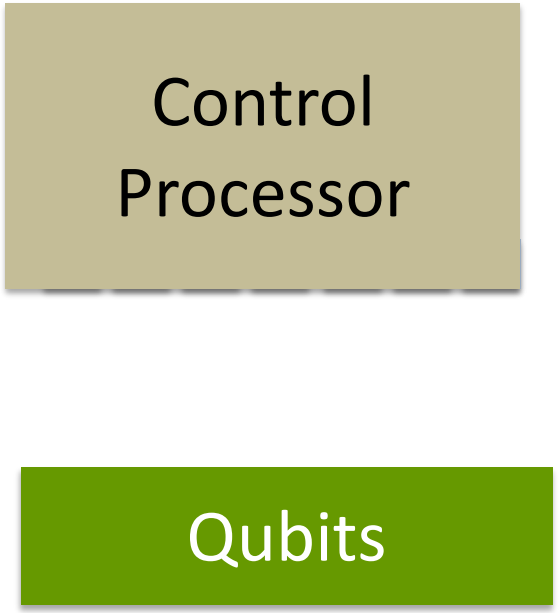

Continuous Quantum Error is essential to protect the state

Programmable Quantum Error Correction

❖Different QECC Designed

❖ New error correction codes → QECC needs to be **programmable**

Different Error Correction Designs

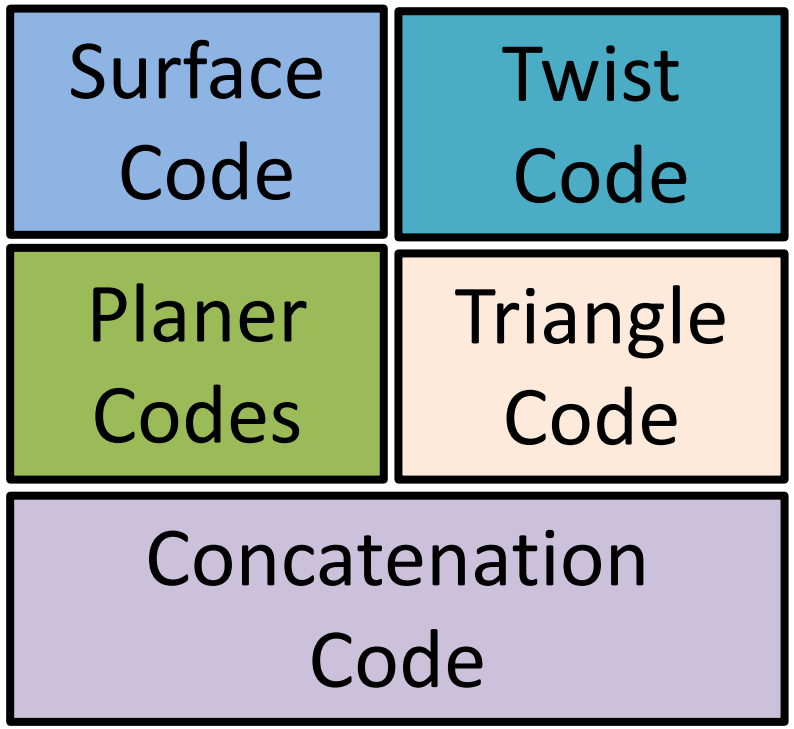

Programmable Quantum Error Correction

❖Different QECC Designed

❖ New error correction codes → QECC needs to be **programmable**

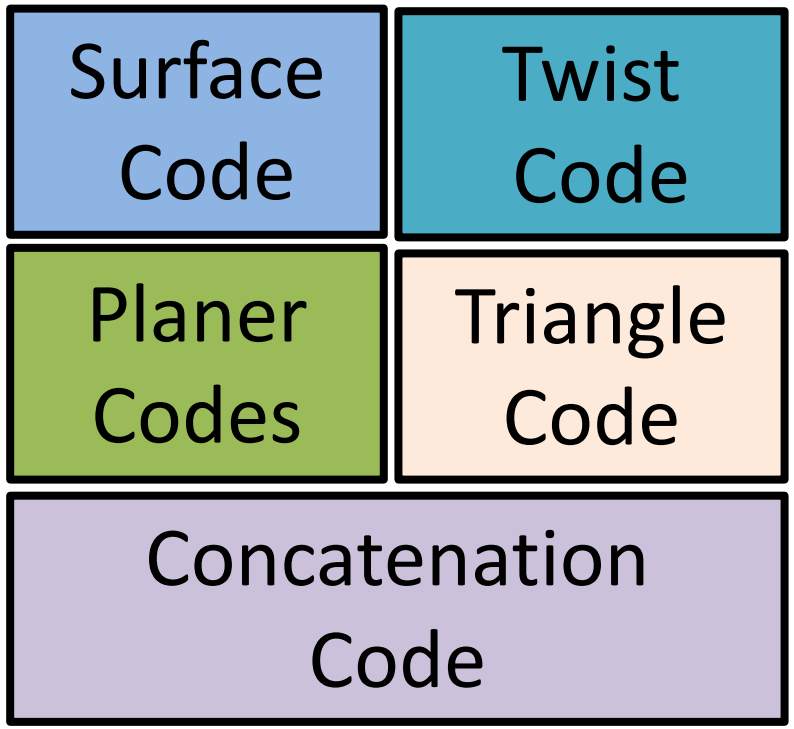

Programmable Quantum Error Correction

❖Different QECC Designed

❖ New error correction codes → QECC needs to be **programmable**

Surface Code

Concatenation Code

Programmable Quantum Error Correction

❖Different QECC Designed

❖ New error correction codes → QECC needs to be **programmable**

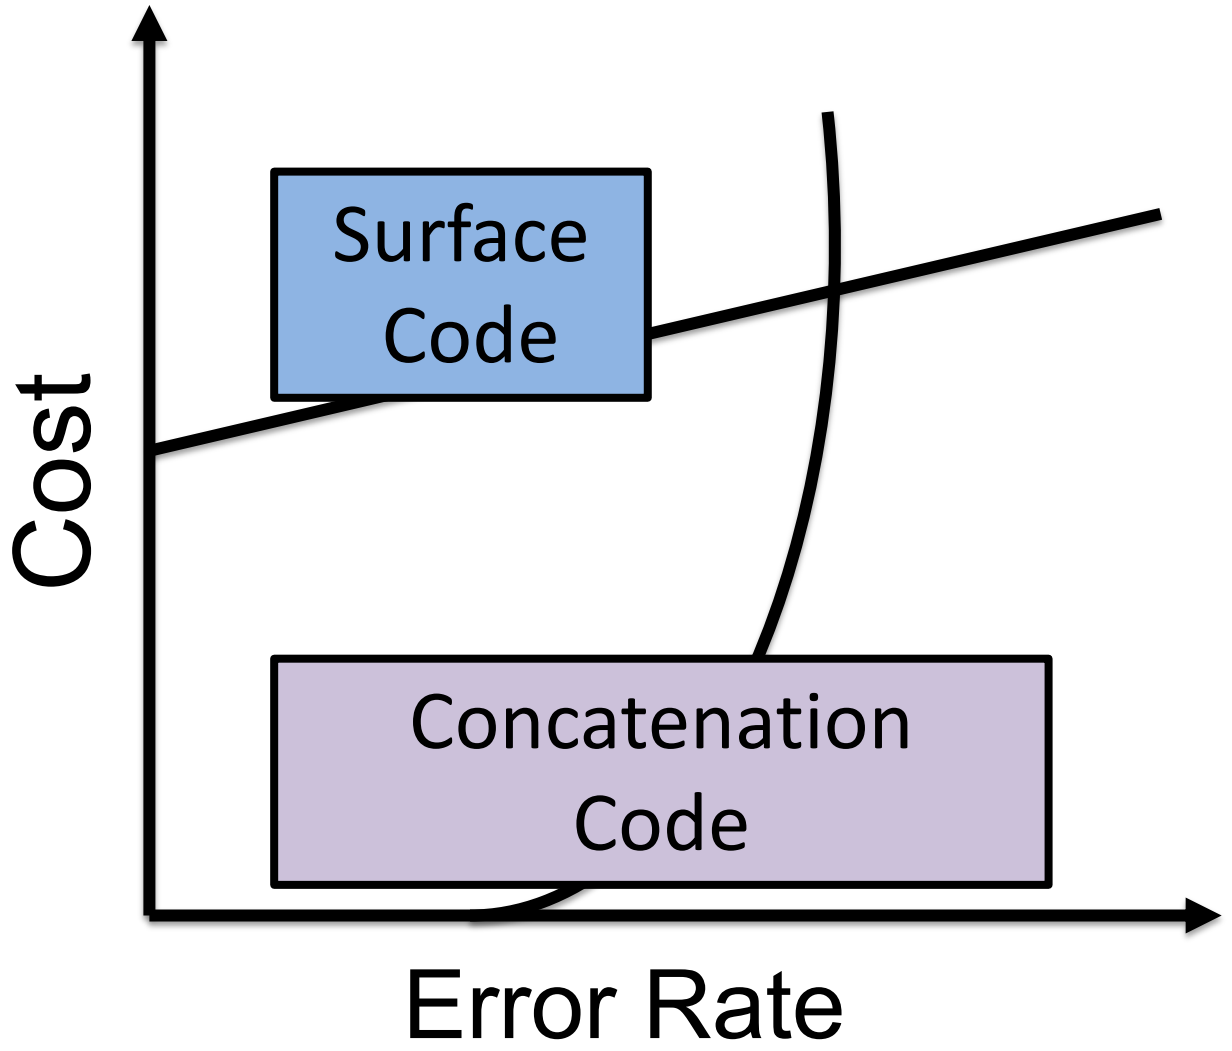

Programmable Quantum Error Correction

❖Different QECC Designed

❖ New error correction codes → QECC needs to be **programmable**

❖ Software managed QECC → Compiler inserts QECC instructions in regular instruction stream

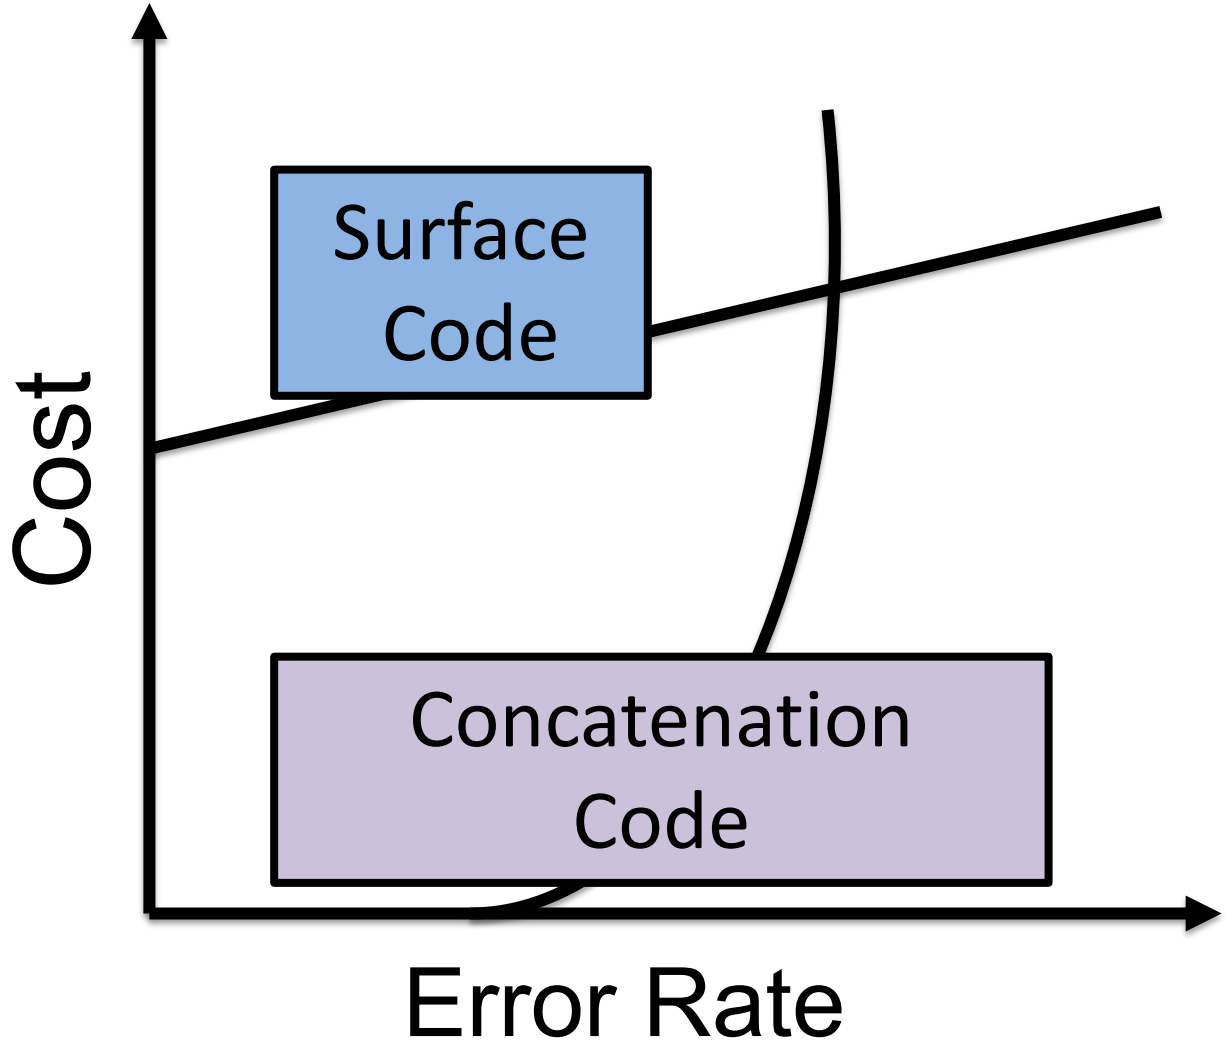

Baseline Architecture

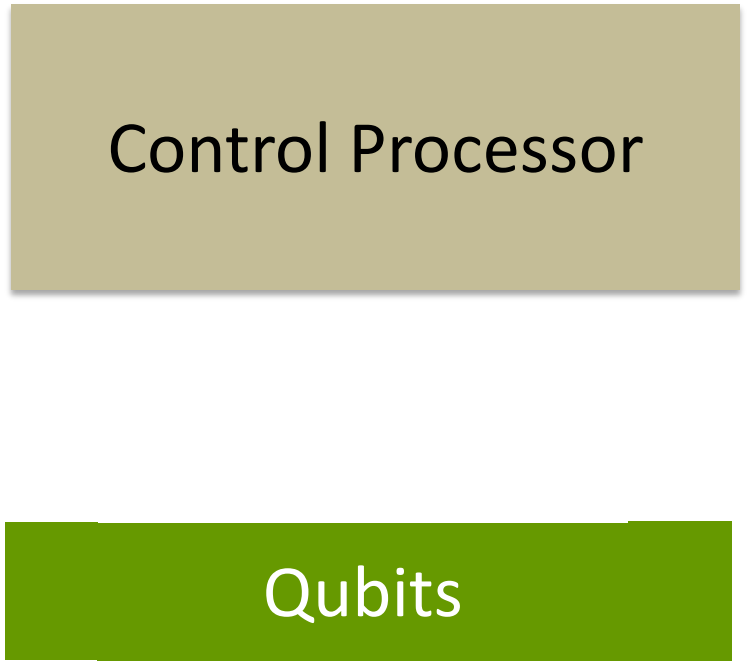

Baseline Architecture

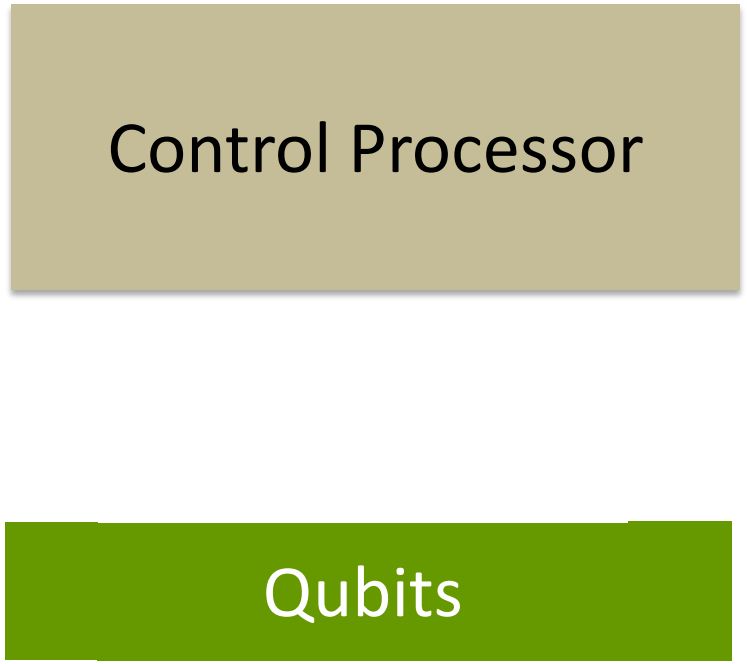

Baseline Architecture

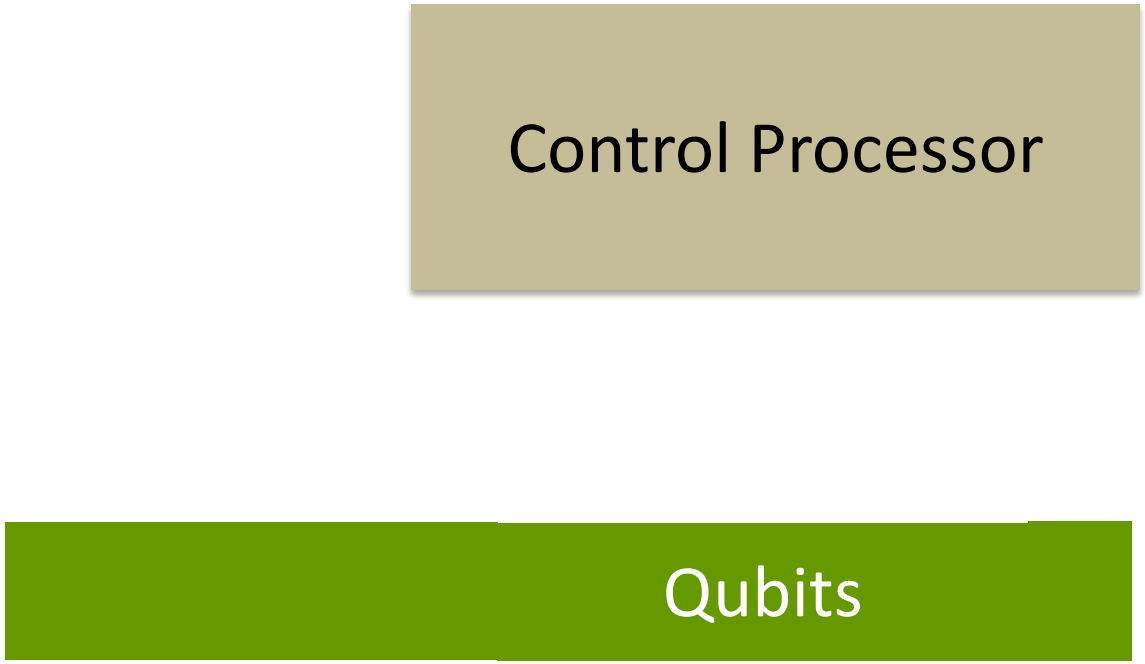

Baseline Architecture

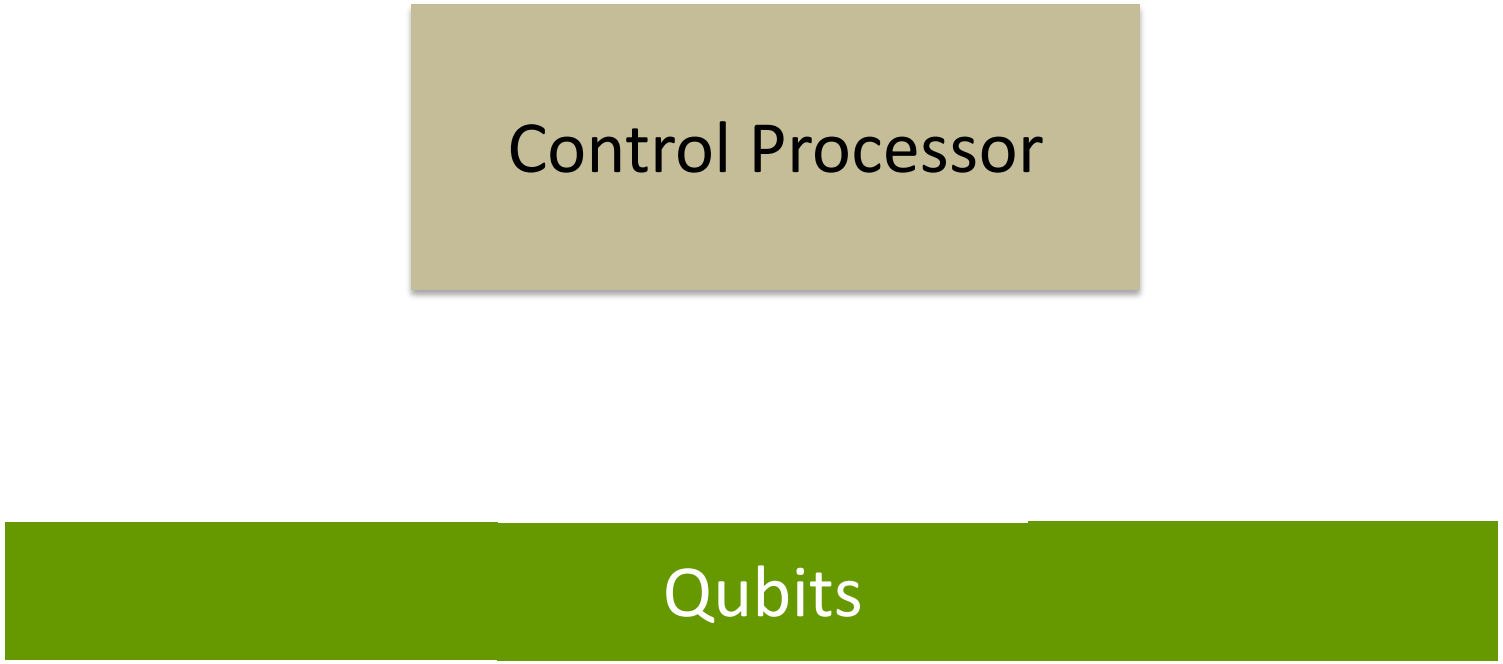

Baseline Architecture

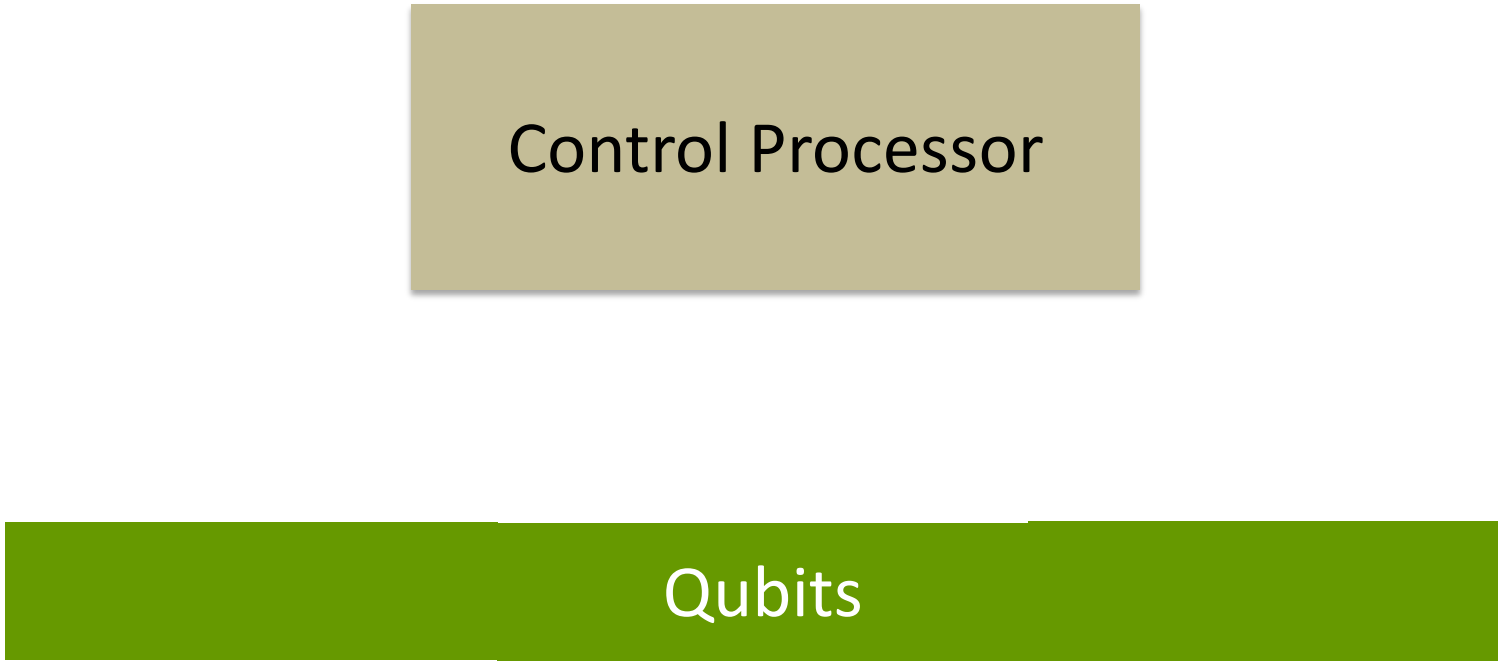

Baseline Architecture

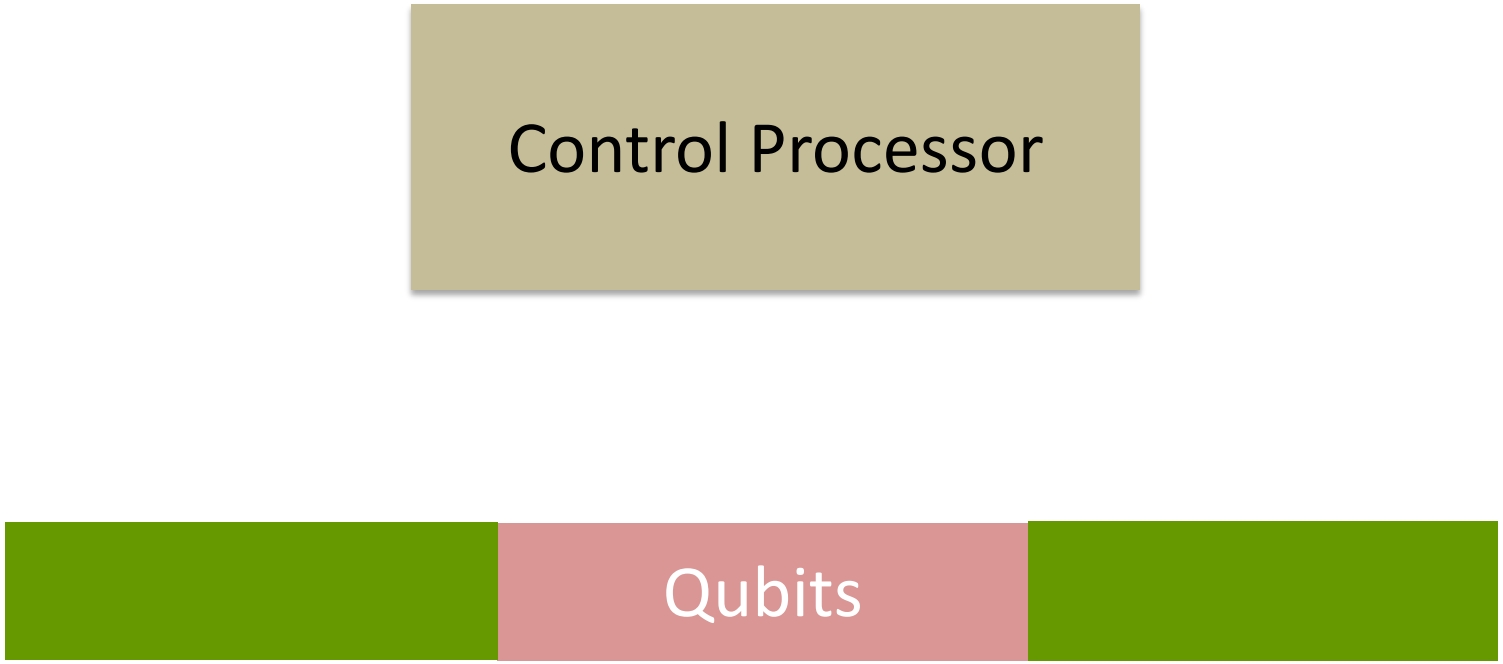

Baseline Architecture

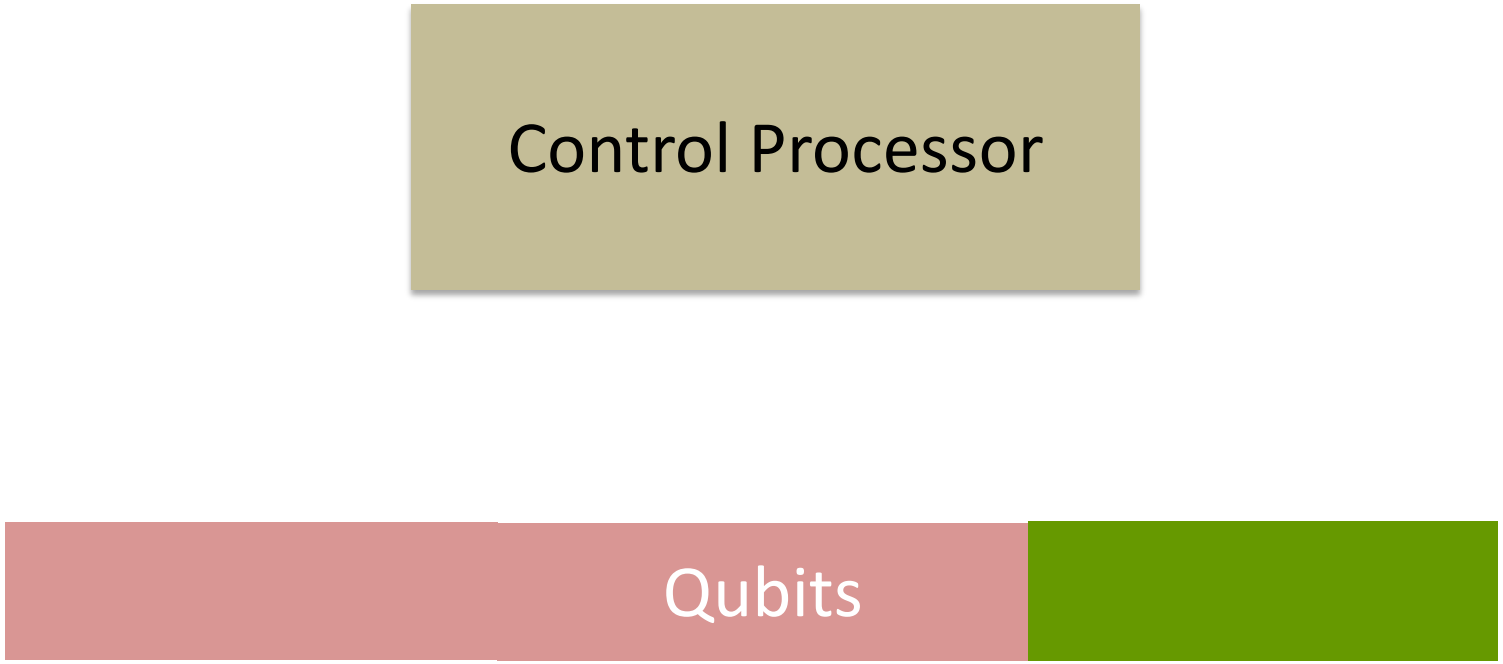

Baseline Architecture

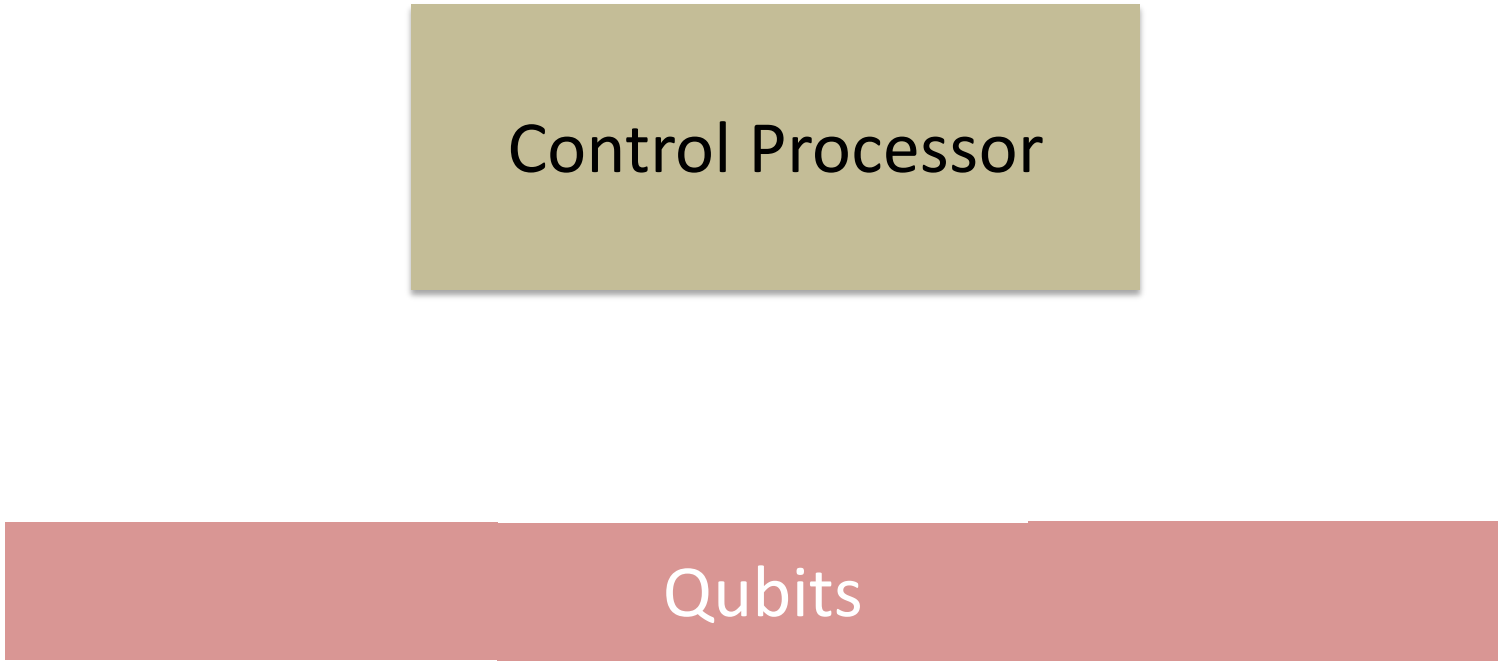

Baseline Architecture

Control Processor

Qubits

Problem: Instruction Bandwidth Bottleneck

- Can’t use i-cache -- delay in delivery of QECC instructions results in error

Problem: Instruction Bandwidth Bottleneck

- Can’t use i-cache -- delay in delivery of QECC instructions results in error

- SW-QECC + no i-cache→ instruction bandwidth **must** scale linearly with number of qubits

Problem: Instruction Bandwidth Bottleneck

- Can't use i-cache -- delay in delivery of QECC instructions results in error

- SW-QECC + no i-cache→ instruction bandwidth **must** scale linearly with number of qubits

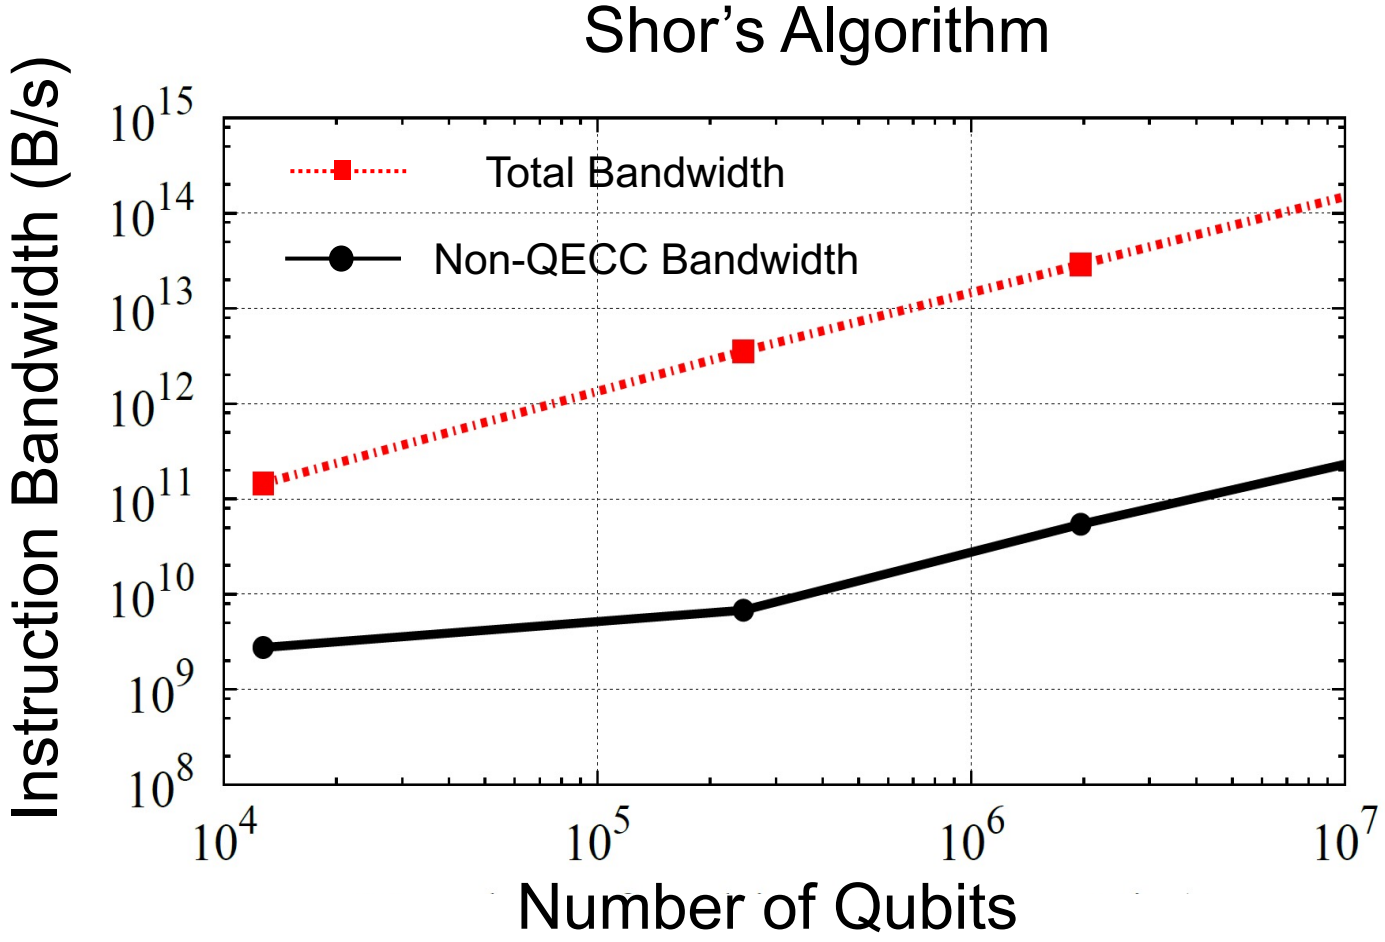

Problem: Instruction Bandwidth Bottleneck

- Can't use i-cache -- delay in delivery of QECC instructions results in error

- SW-QECC + no i-cache→ instruction bandwidth **must** scale linearly with number of qubits

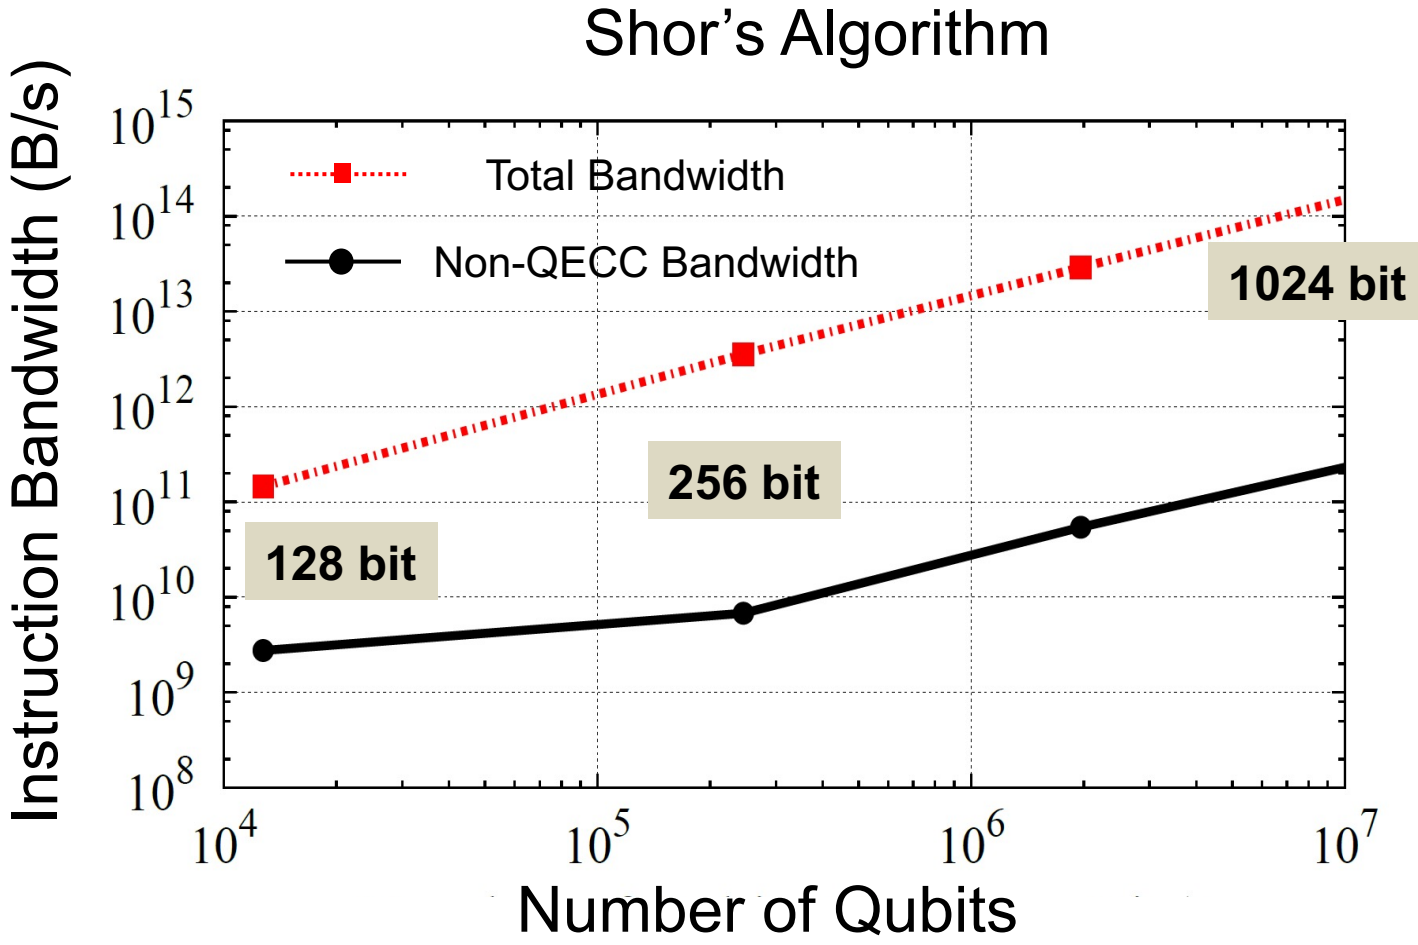

Problem: Instruction Bandwidth Bottleneck

- Can't use i-cache -- delay in delivery of QECC instructions results in error

- SW-QECC + no i-cache → instruction bandwidth **must** scale linearly with number of qubits

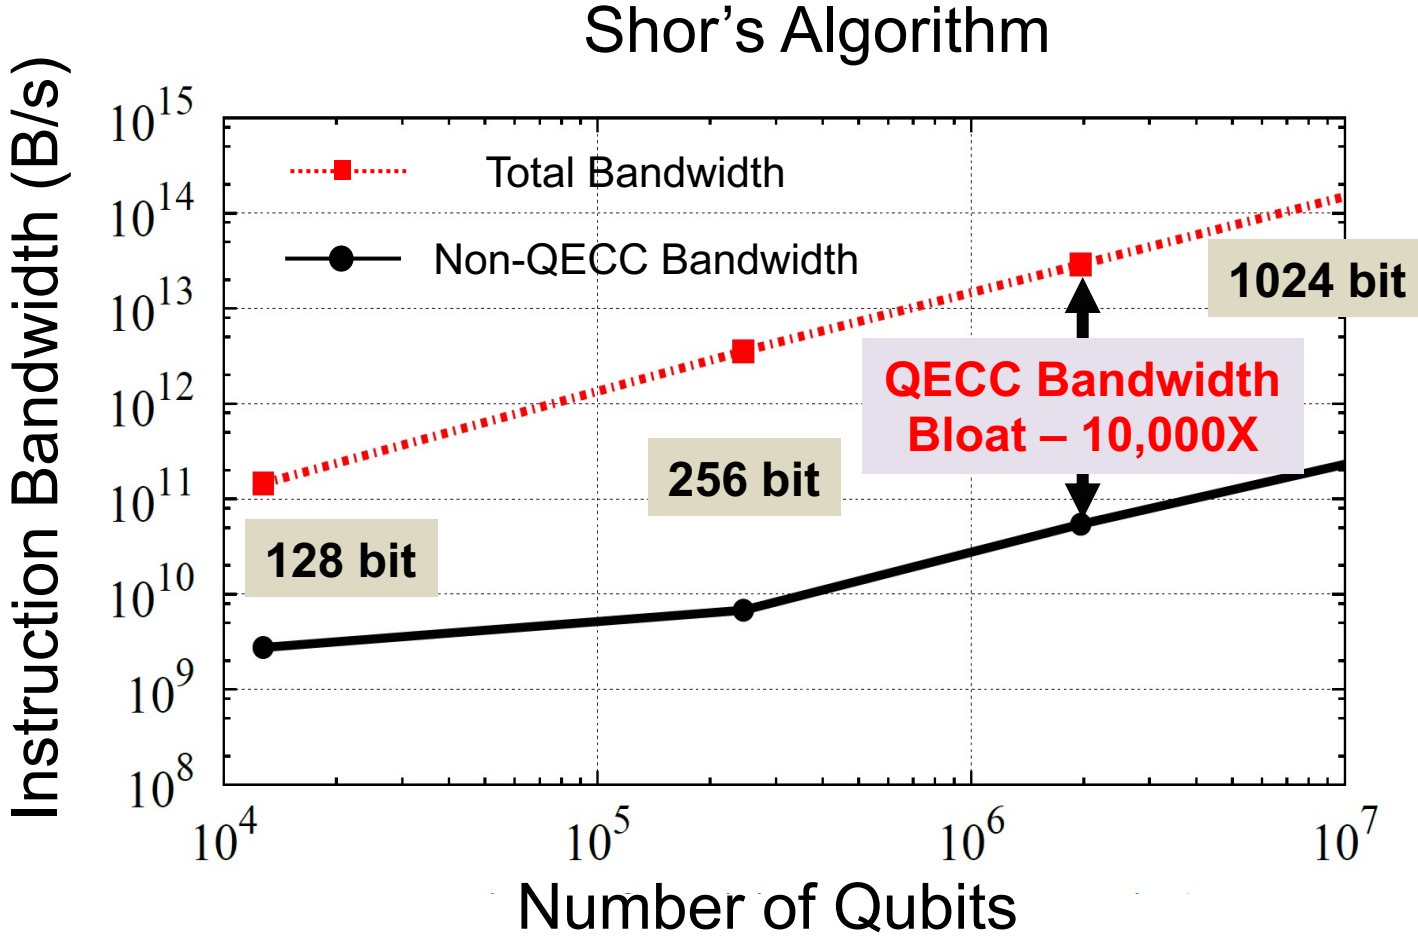

Problem: Instruction Bandwidth Bottleneck

- Can't use i-cache -- delay in delivery of QECC instructions results in error

- SW-QECC + no i-cache → instruction bandwidth **must** scale linearly with number of qubits

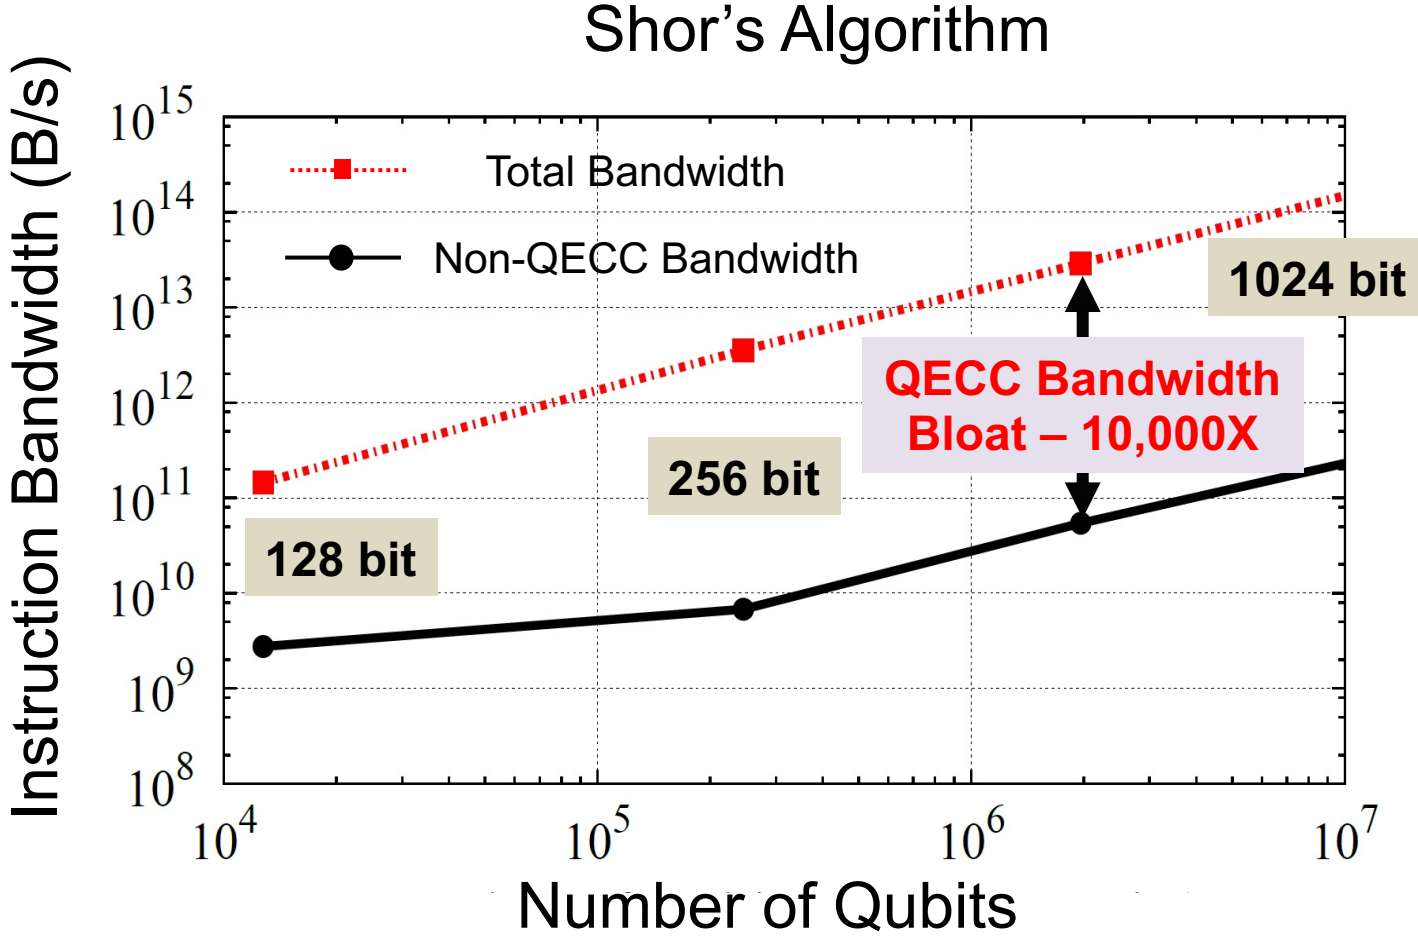

Instruction Bandwidth bottleneck → 99.99% of instructions are QECC

QECC Instruction Bandwidth Bloat

- Realistic Quantum workloads require substantially large number of qubits

- Large number of qubits → must support large instruction bandwidth

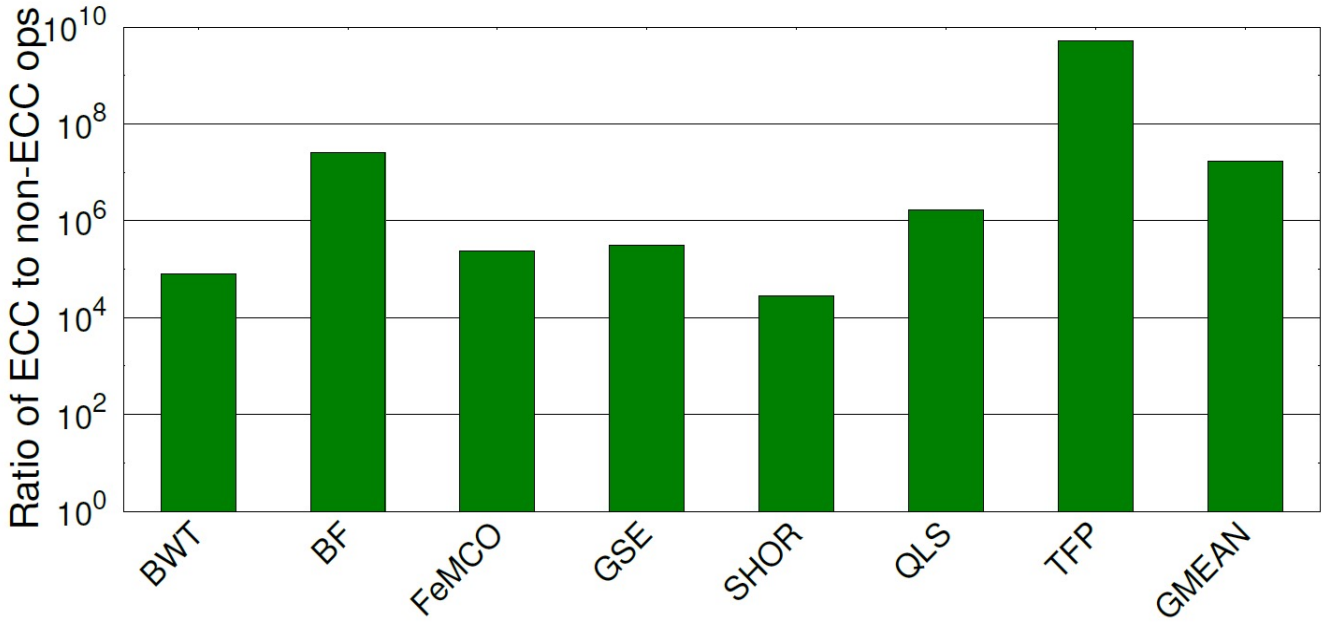

QECC instruction bloat is dominant in all large scale quantum workloads

	Programmability	BW Efficiency
Software-Managed QECC		
Hardcoded QECC		

	Programmability	BW Efficiency
Software-Managed QECC	✓	✗
Hardcoded QECC		

	Programmability	BW Efficiency
Software-Managed QECC	✓	✗
Hardcoded QECC	✗	✓

	Programmability	BW Efficiency
Software-Managed QECC	✓	✗
Hardcoded QECC	✗	✓

Goal: To enable programmable Quantum Error Correction without bandwidth bloat

Insight

QECC can be executed independently without any global synchronization

Insight

QECC can be executed independently without any global synchronization

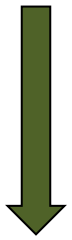

Insight

QECC can be executed independently without any global synchronization

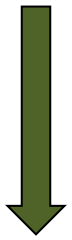

QECC is simple enough to manage in hardware using programmable microcode

QuEST Architecture

- MCE continuously issues QECC µops from local µcode

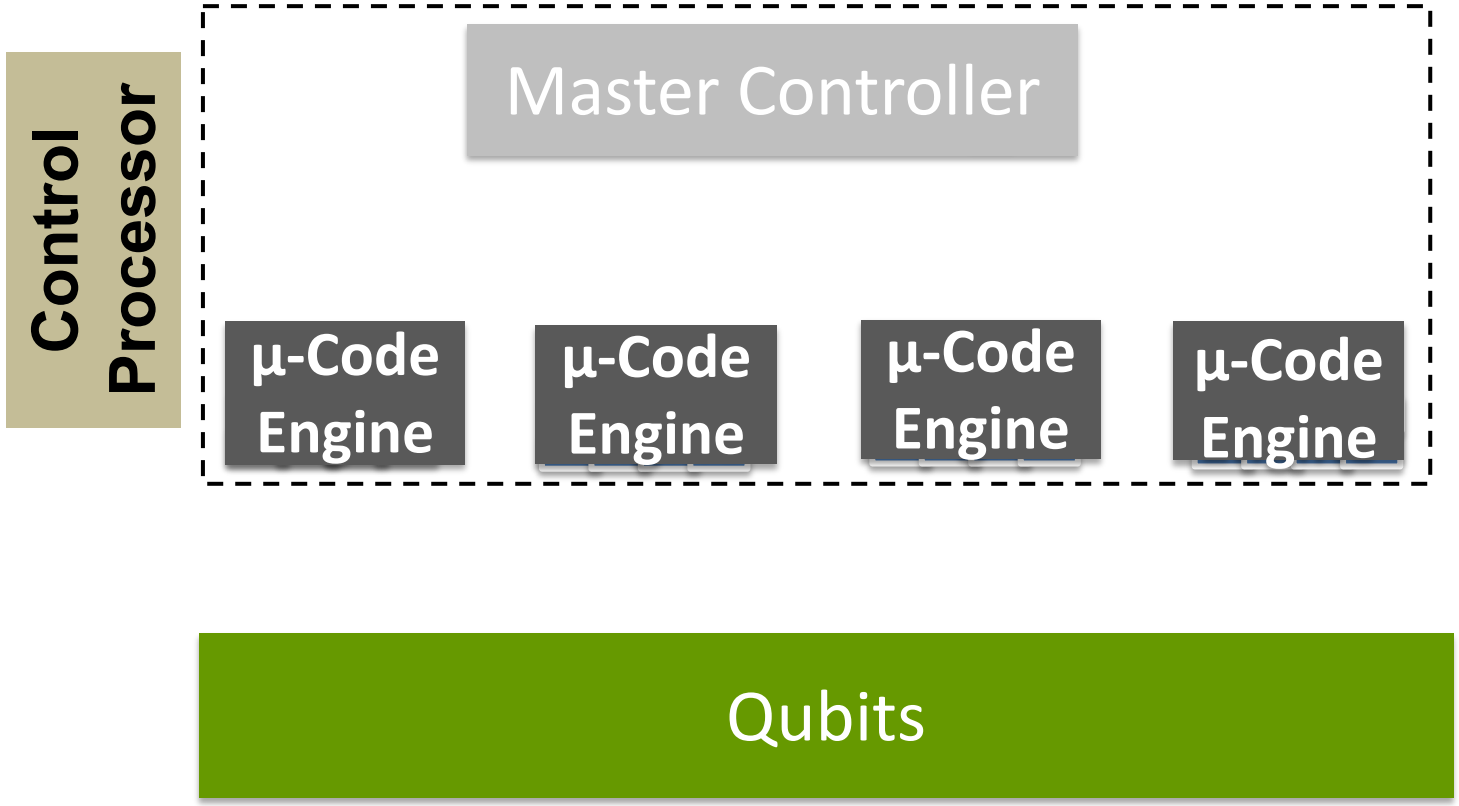

QuEST Architecture

- MCE continuously issues QECC µops from local µcode

- Master controller issues regular instructions. MCE decodes it to µops

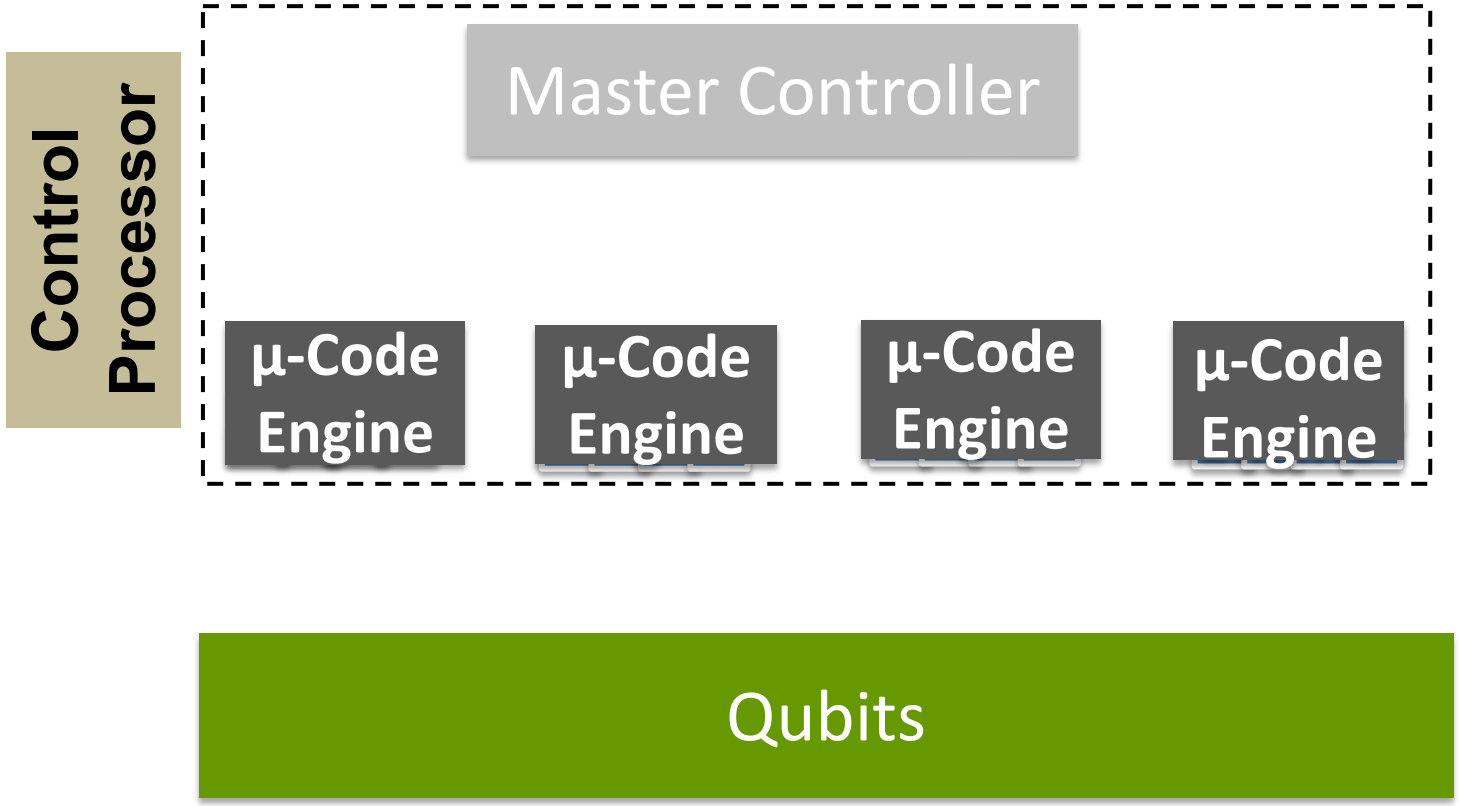

QuEST alleviates instruction bandwidth by issuing QECC ops locally

QuEST Architecture

- MCE continuously issues QECC µops from local µcode

- Master controller issues regular instructions. MCE decodes it to µops

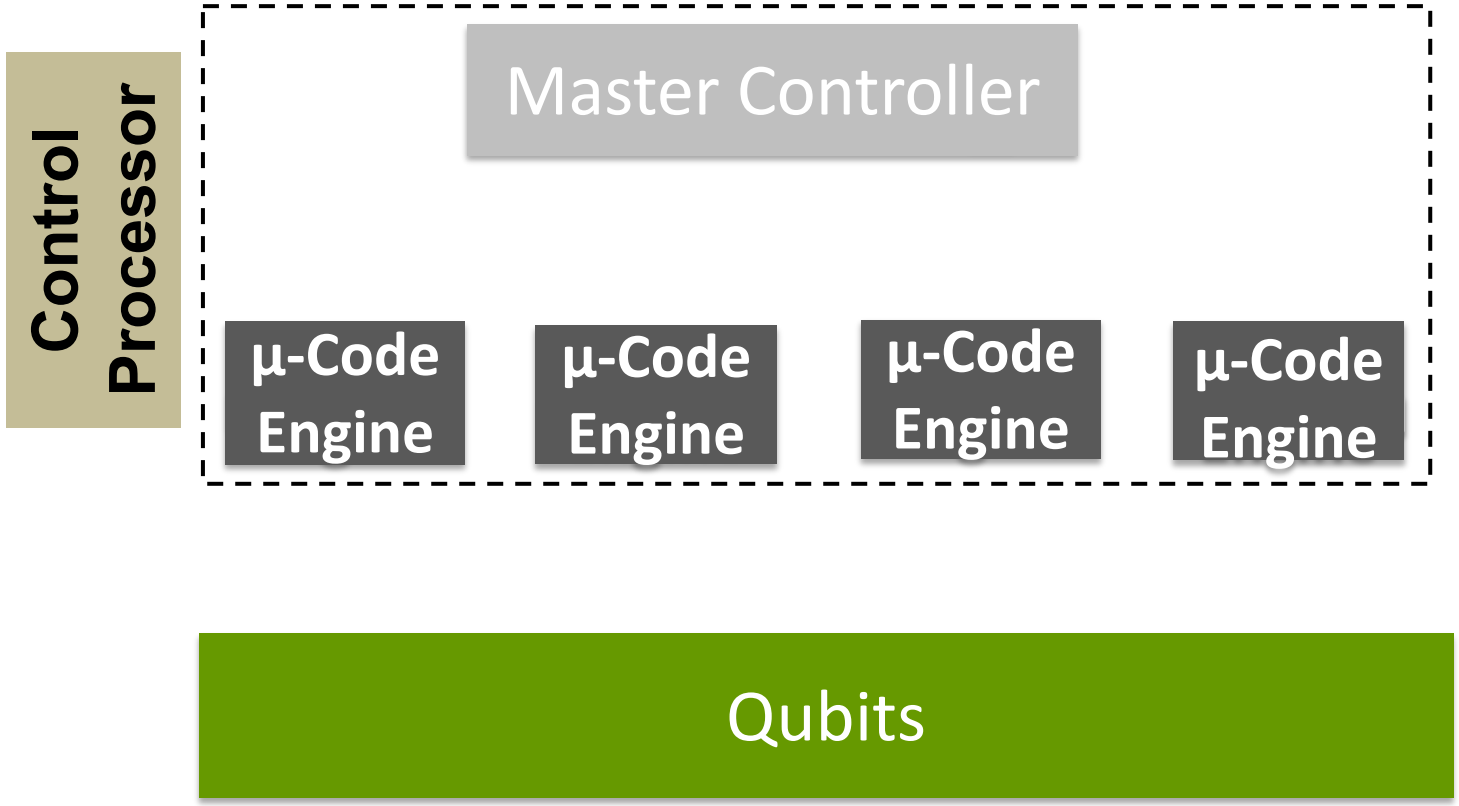

QuEST alleviates instruction bandwidth by issuing QECC ops locally

QuEST Architecture

- MCE continuously issues QECC µops from local µcode

- Master controller issues regular instructions. MCE decodes it to µops

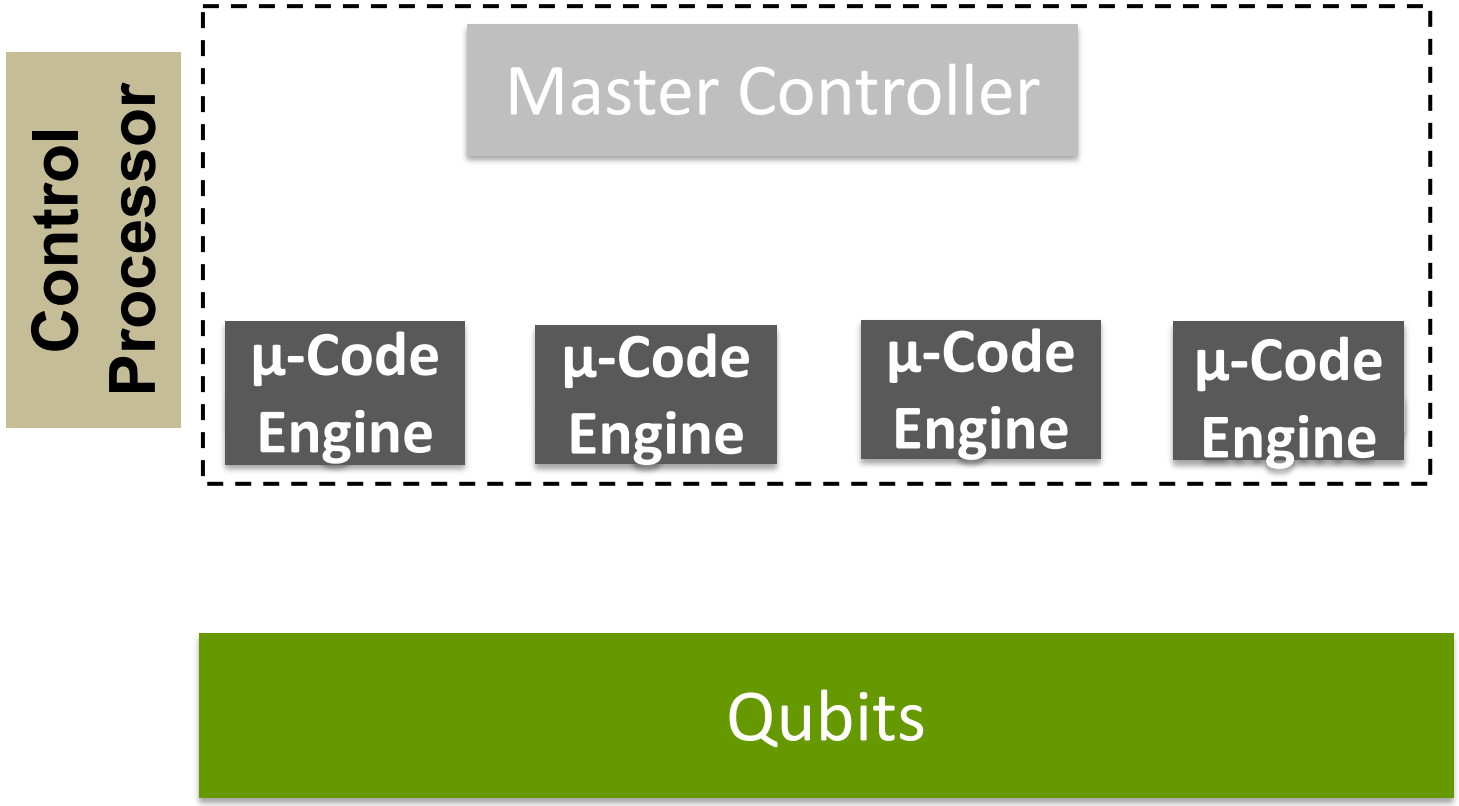

QuEST alleviates instruction bandwidth by issuing QECC ops locally

QuEST Architecture

❖ MCE continuously issues QECC µops from local µcode

❖Master controller issues regular instructions. MCE decodes it to µops

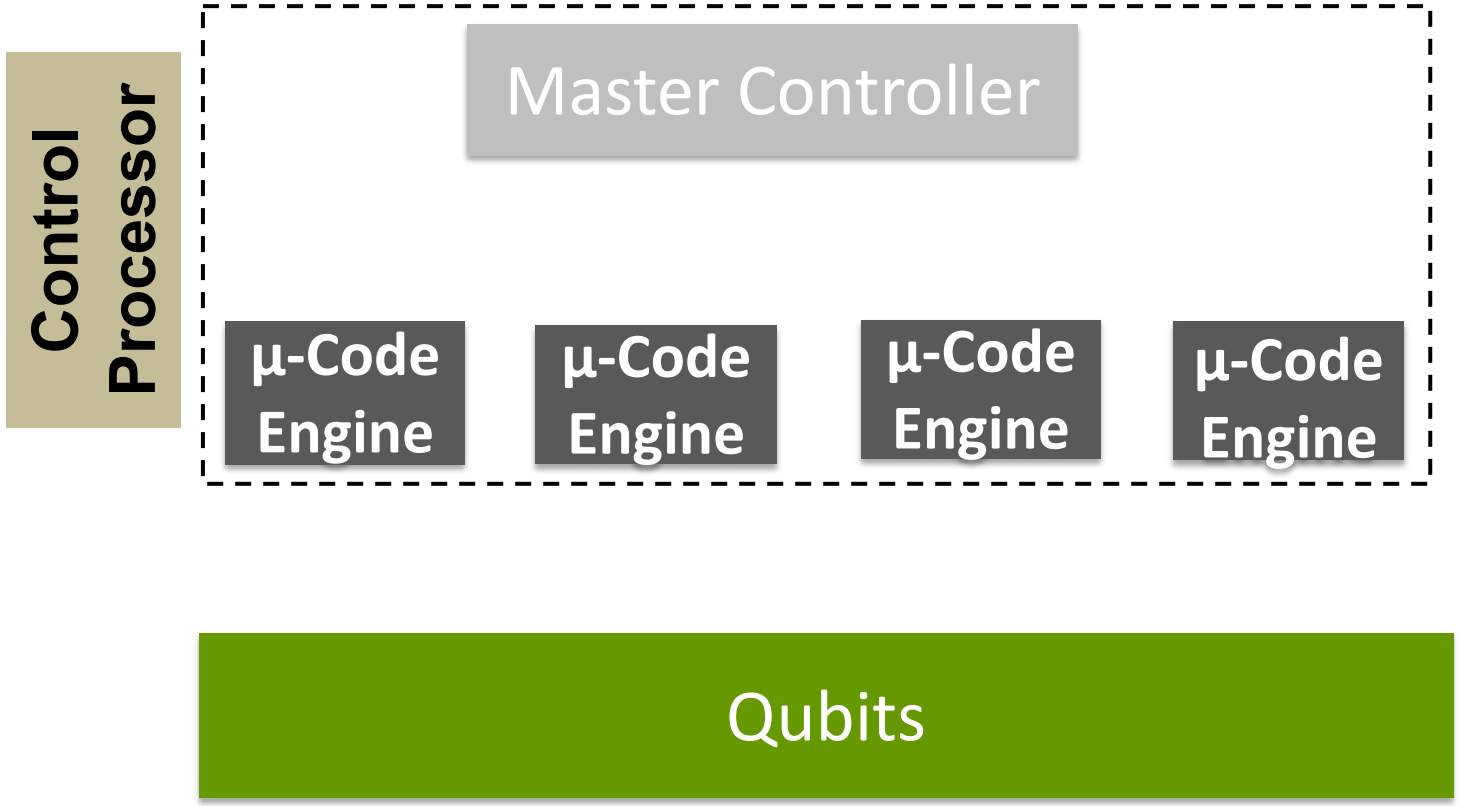

QuEST alleviates instruction bandwidth by issuing QECC ops locally

MCE Microarchitecture

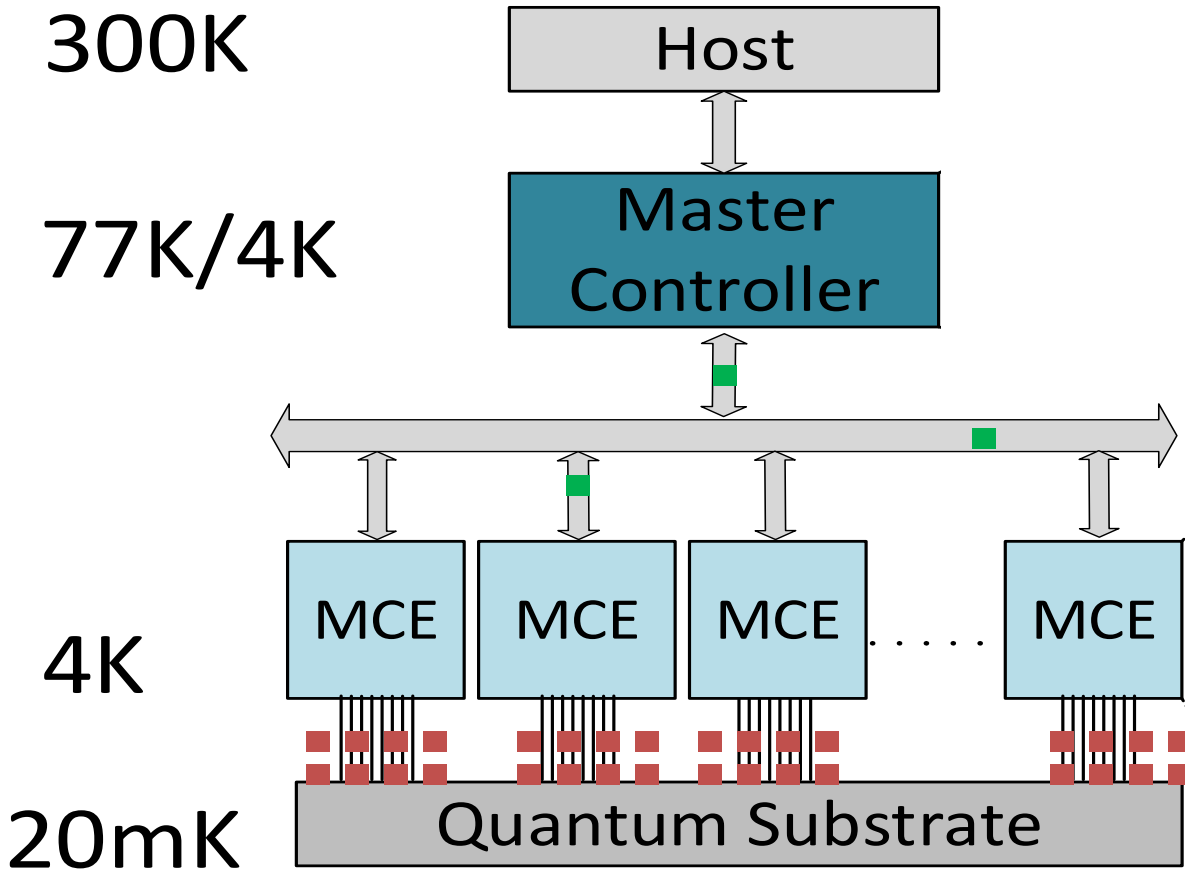

QuEST alleviates instruction bandwidth by issuing QECC ops locally

MCE Microarchitecture

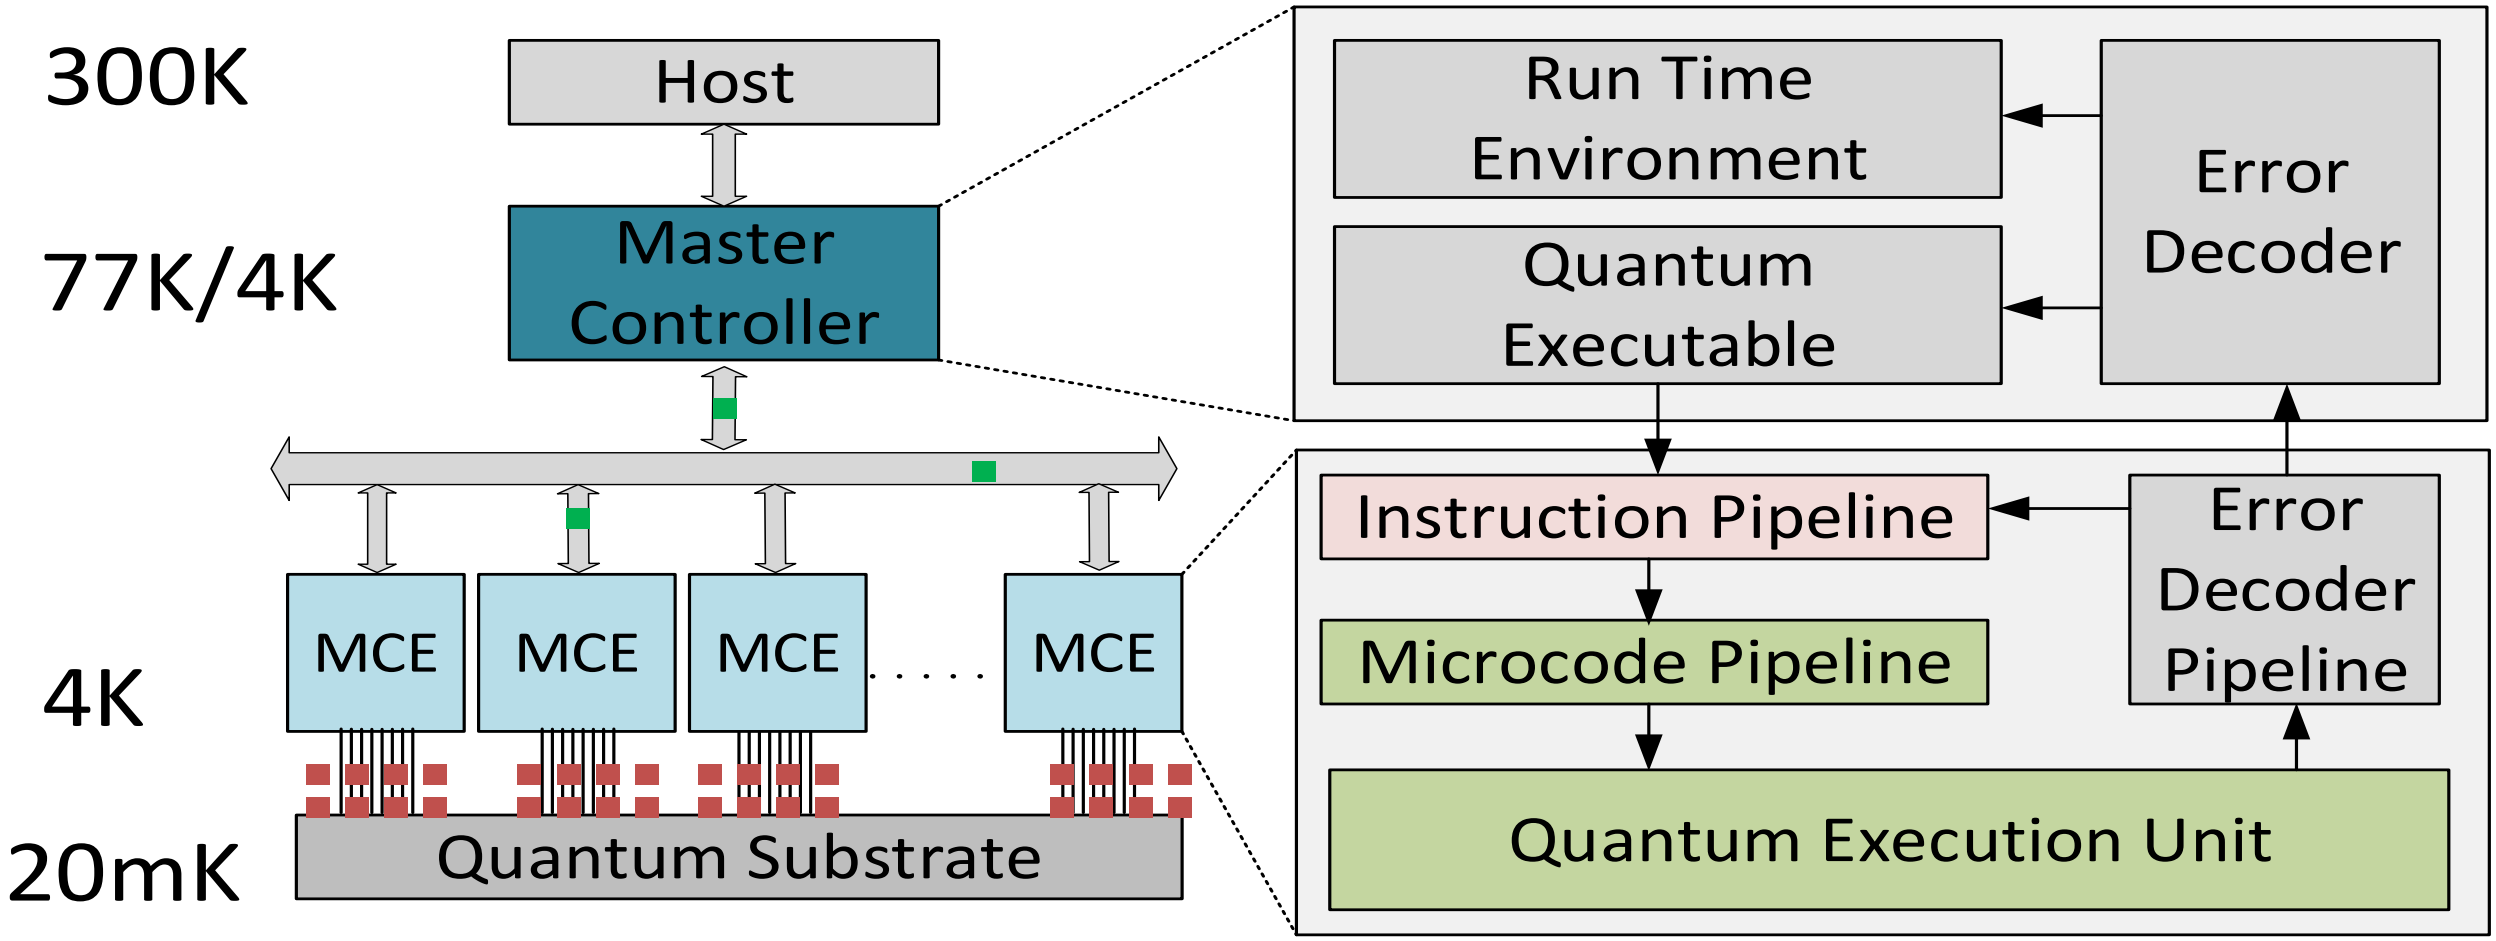

QuEST alleviates instruction bandwidth by issuing QECC ops locally

MCE Microarchitecture

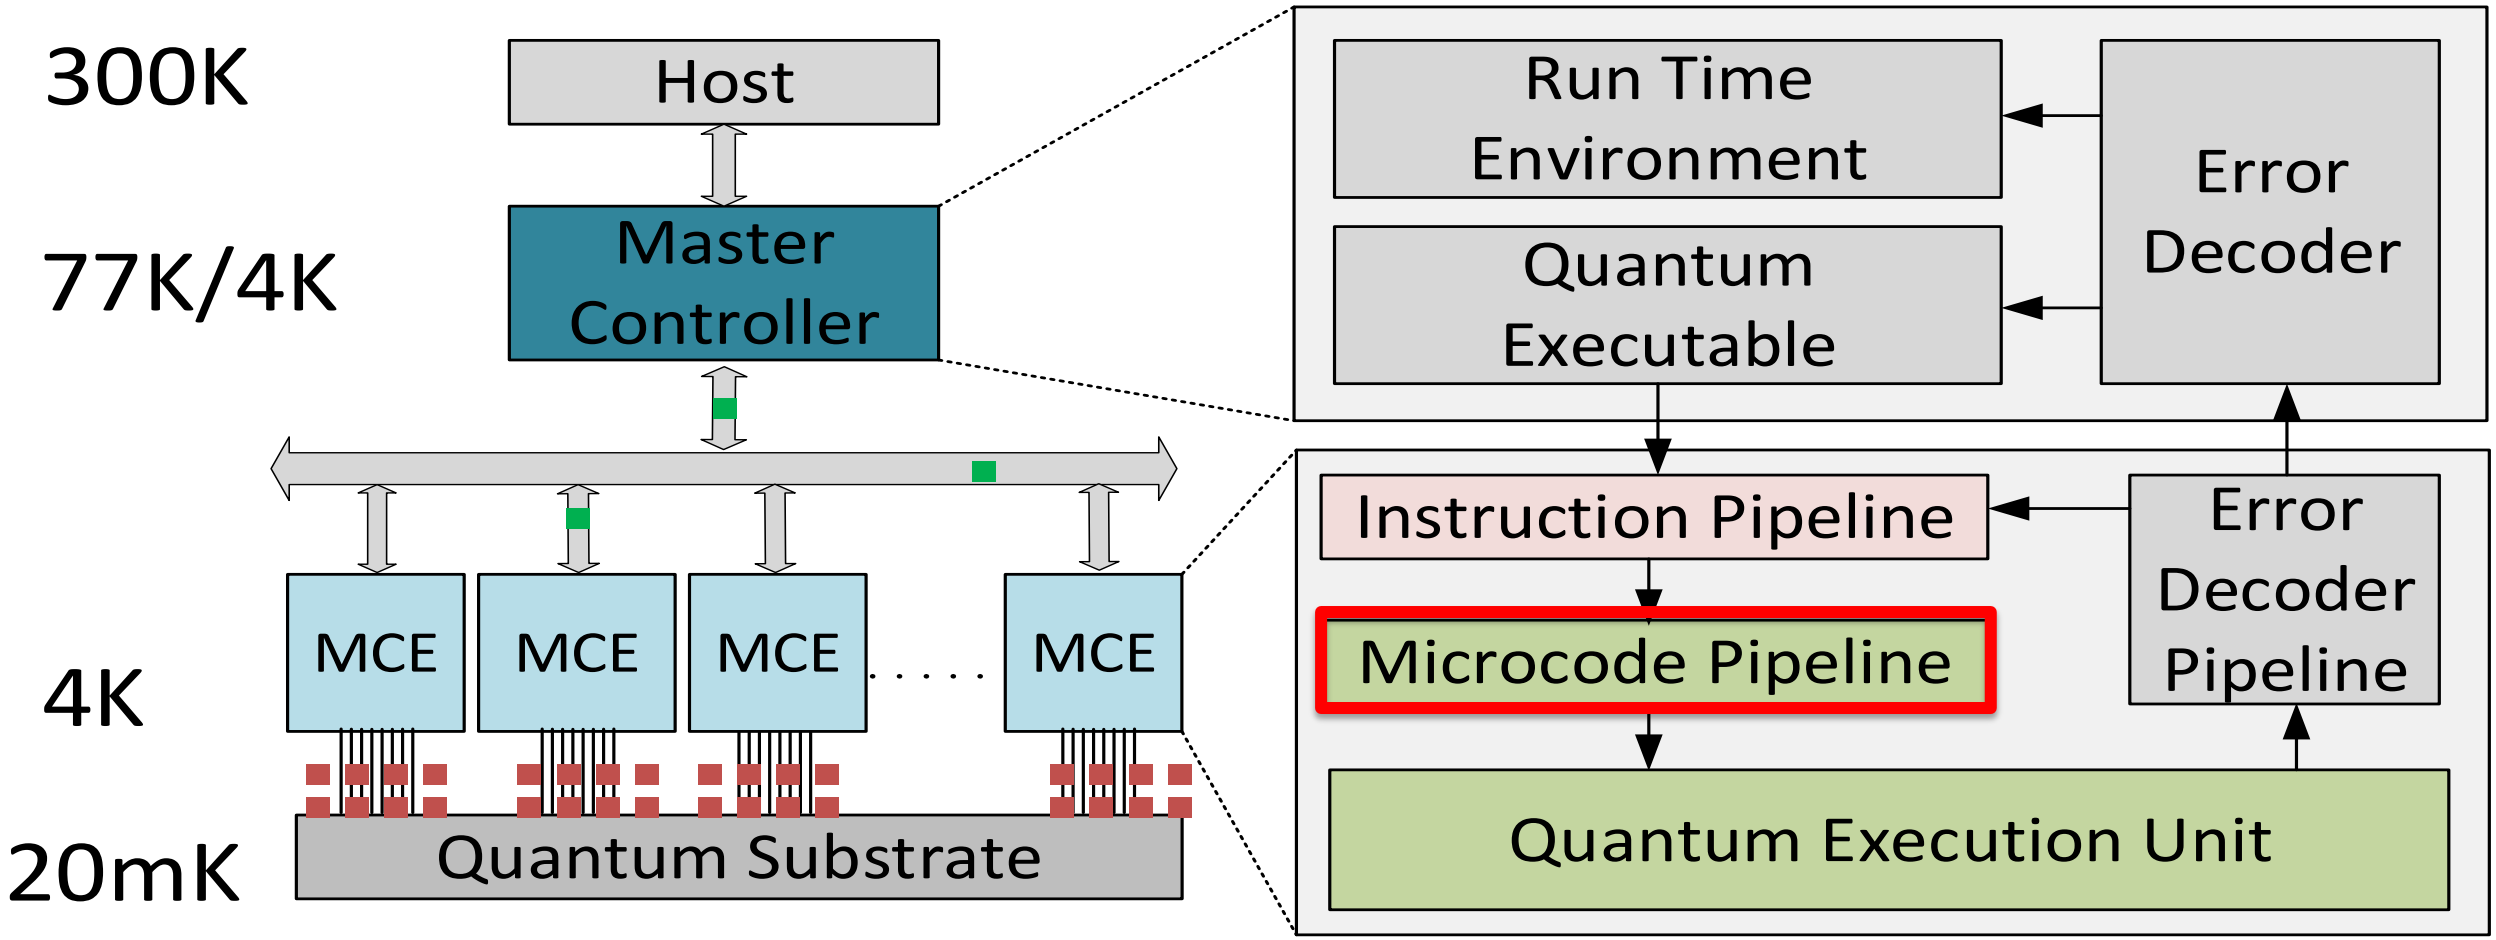

QuEST alleviates instruction bandwidth by issuing QECC ops locally

Microcode Design

- Logical Microcode → supplies logical µops
- QECC Microcode → supplies QECC µops to all qubits
- Mask Bits → selects between QECC µop and Logical µops

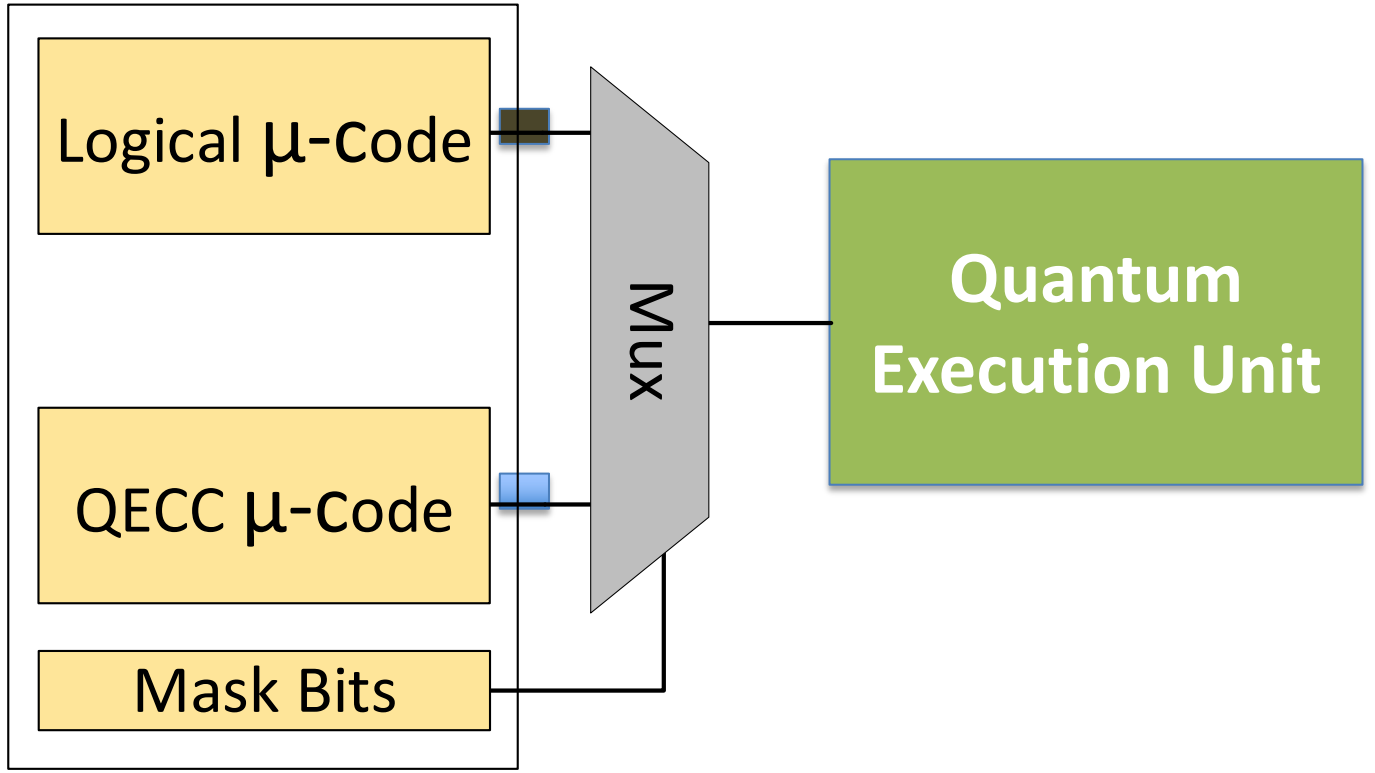

Micro-code Memory

Dedicated QECC microcode continuously run QECC

Microcode Design

- Logical Microcode → supplies logical µops
- QECC Microcode → supplies QECC µops to all qubits
- Mask Bits → selects between QECC µop and Logical µops

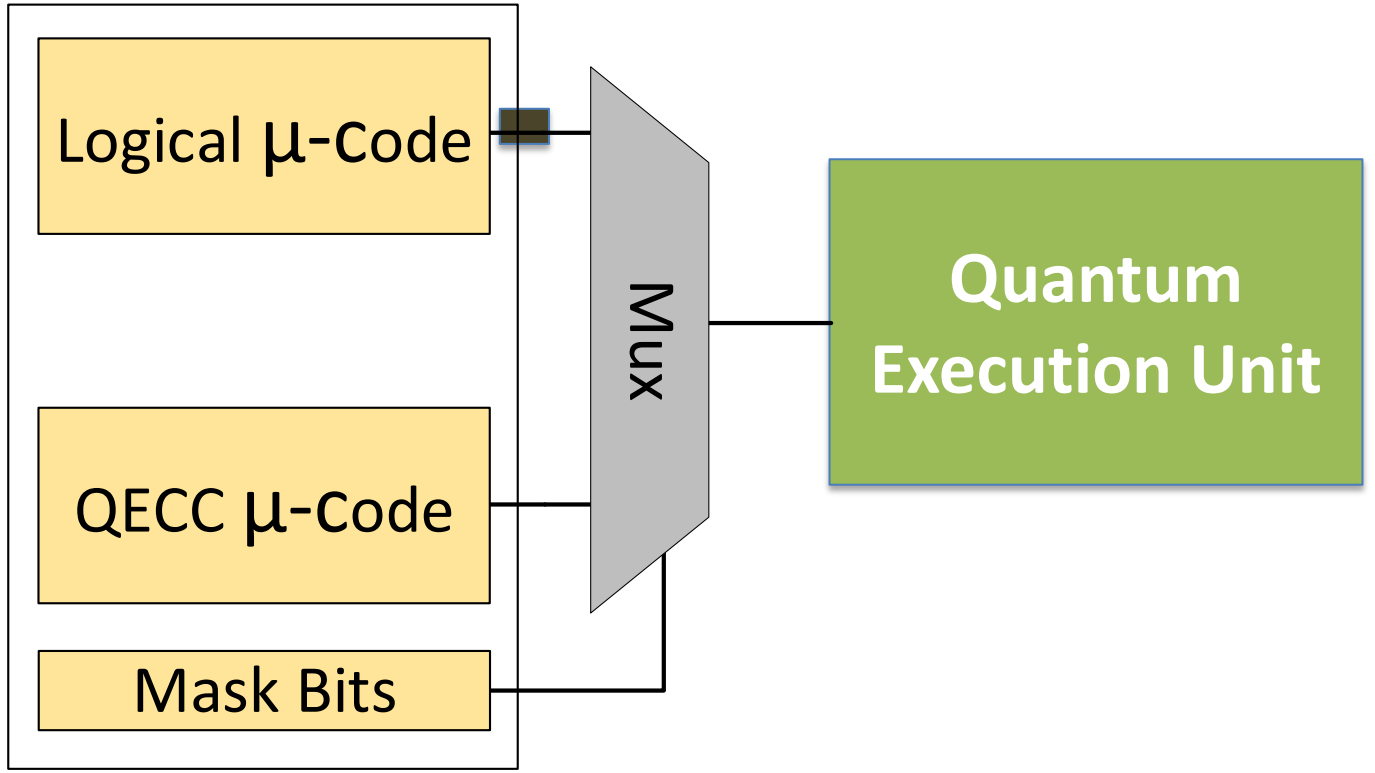

Micro-code Memory

Dedicated QECC microcode continuously run QECC

Microcode Design

- Logical Microcode → supplies logical µops
- QECC Microcode → supplies QECC µops to all qubits
- Mask Bits → selects between QECC µop and Logical µops

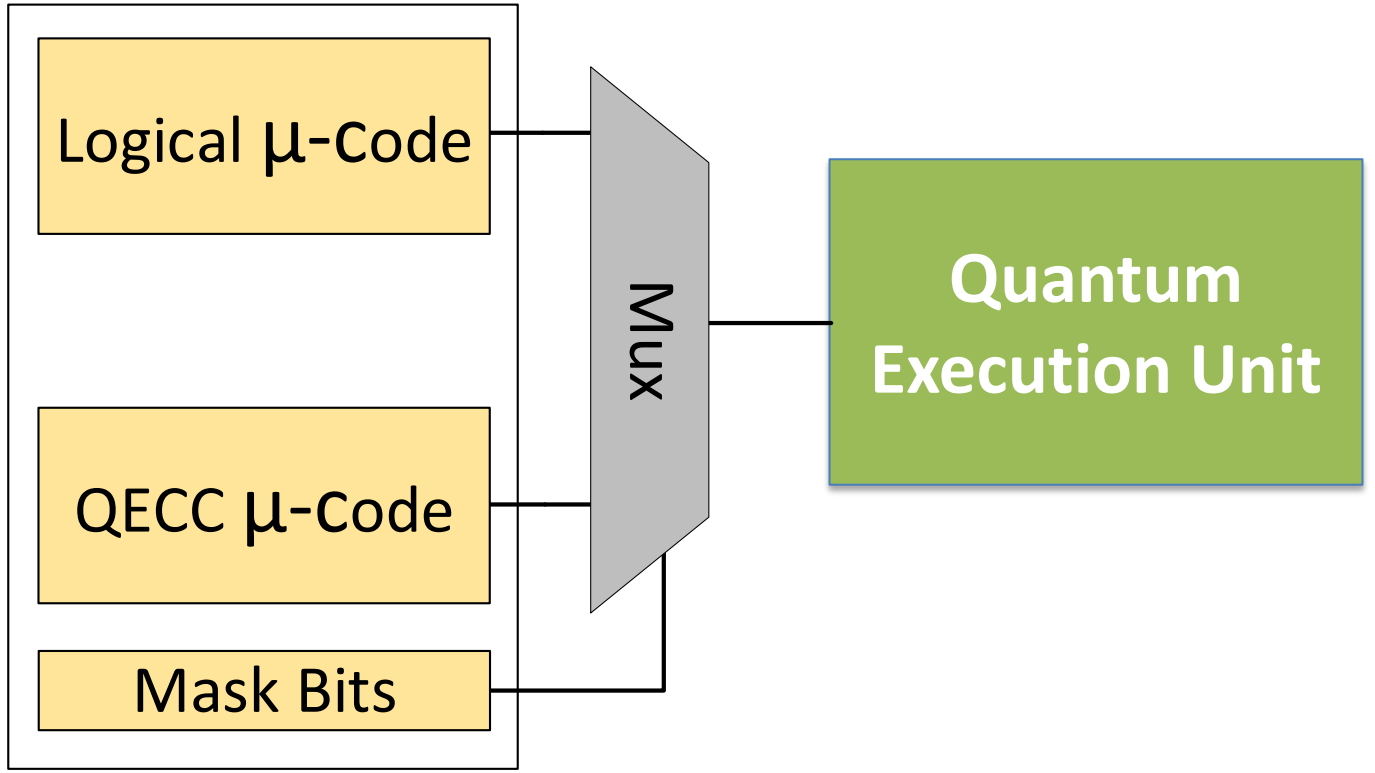

Micro-code Memory

Dedicated QECC microcode continuously run QECC

Microcode Memory Capacity Requirements

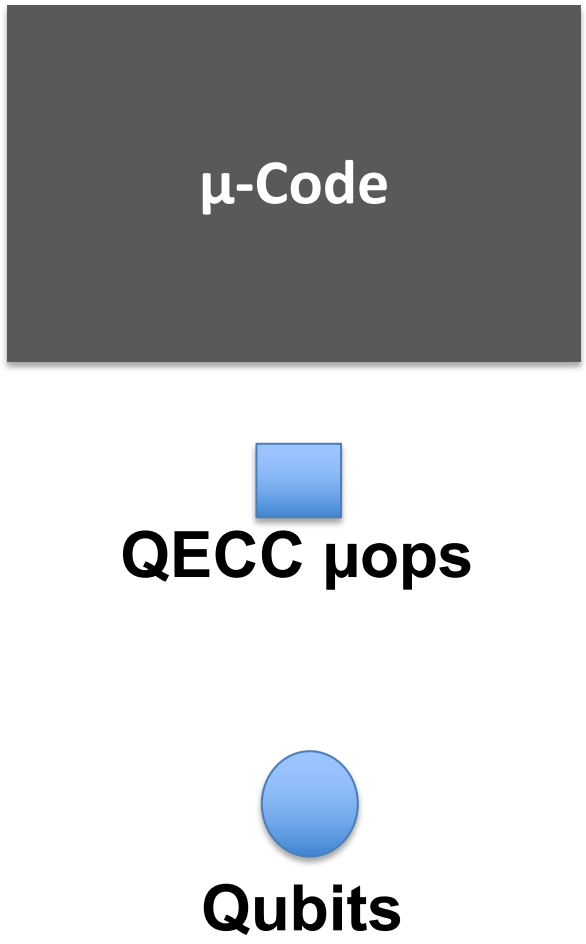

Microcode Memory Capacity Requirements

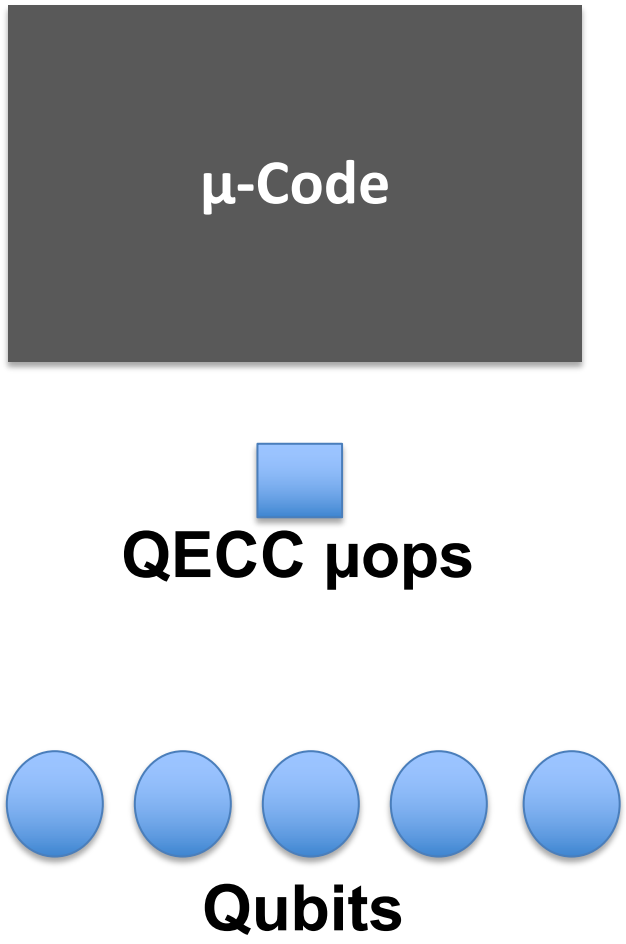

Microcode Memory Capacity Requirements

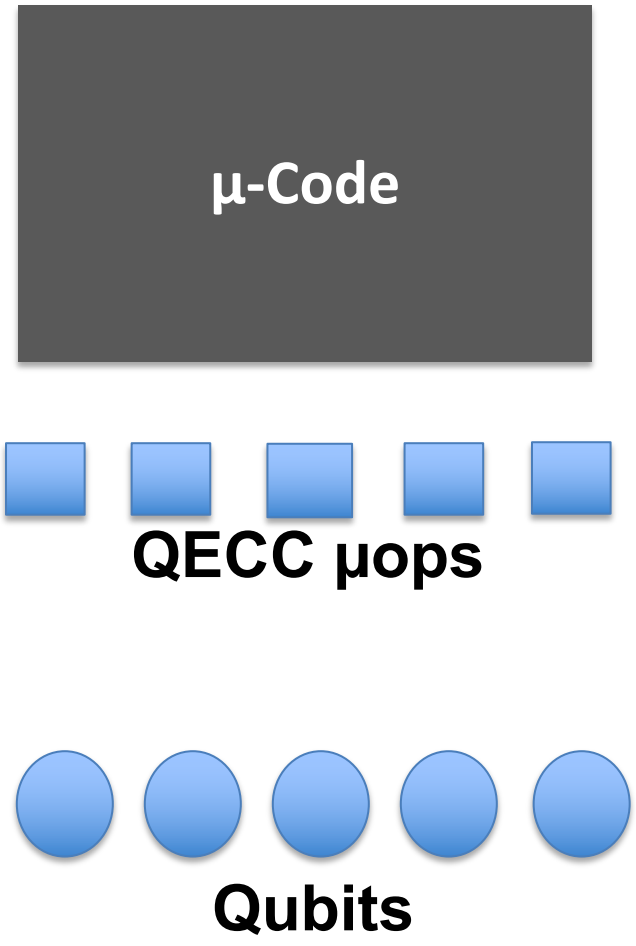

Microcode Memory Capacity Requirements

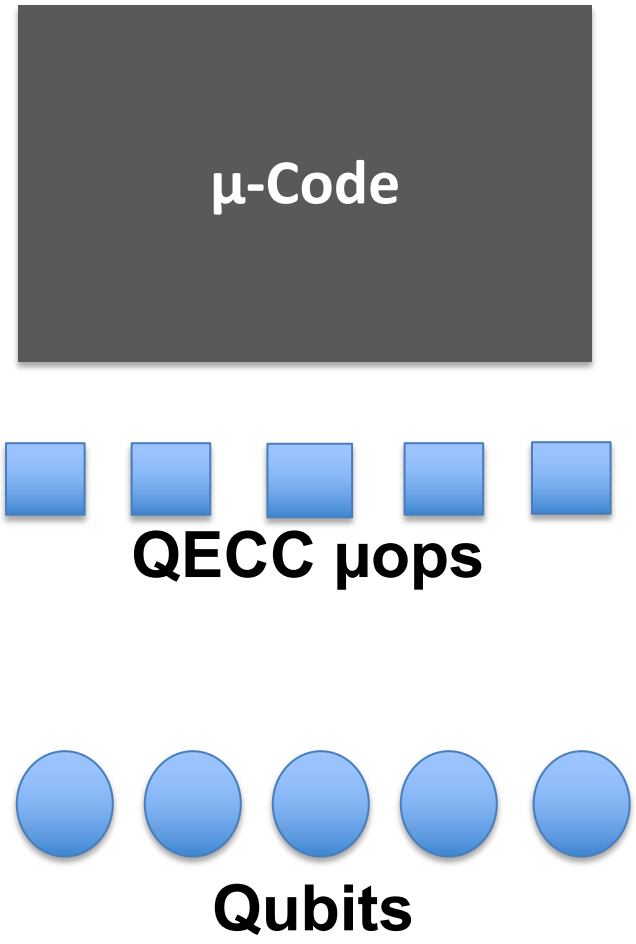

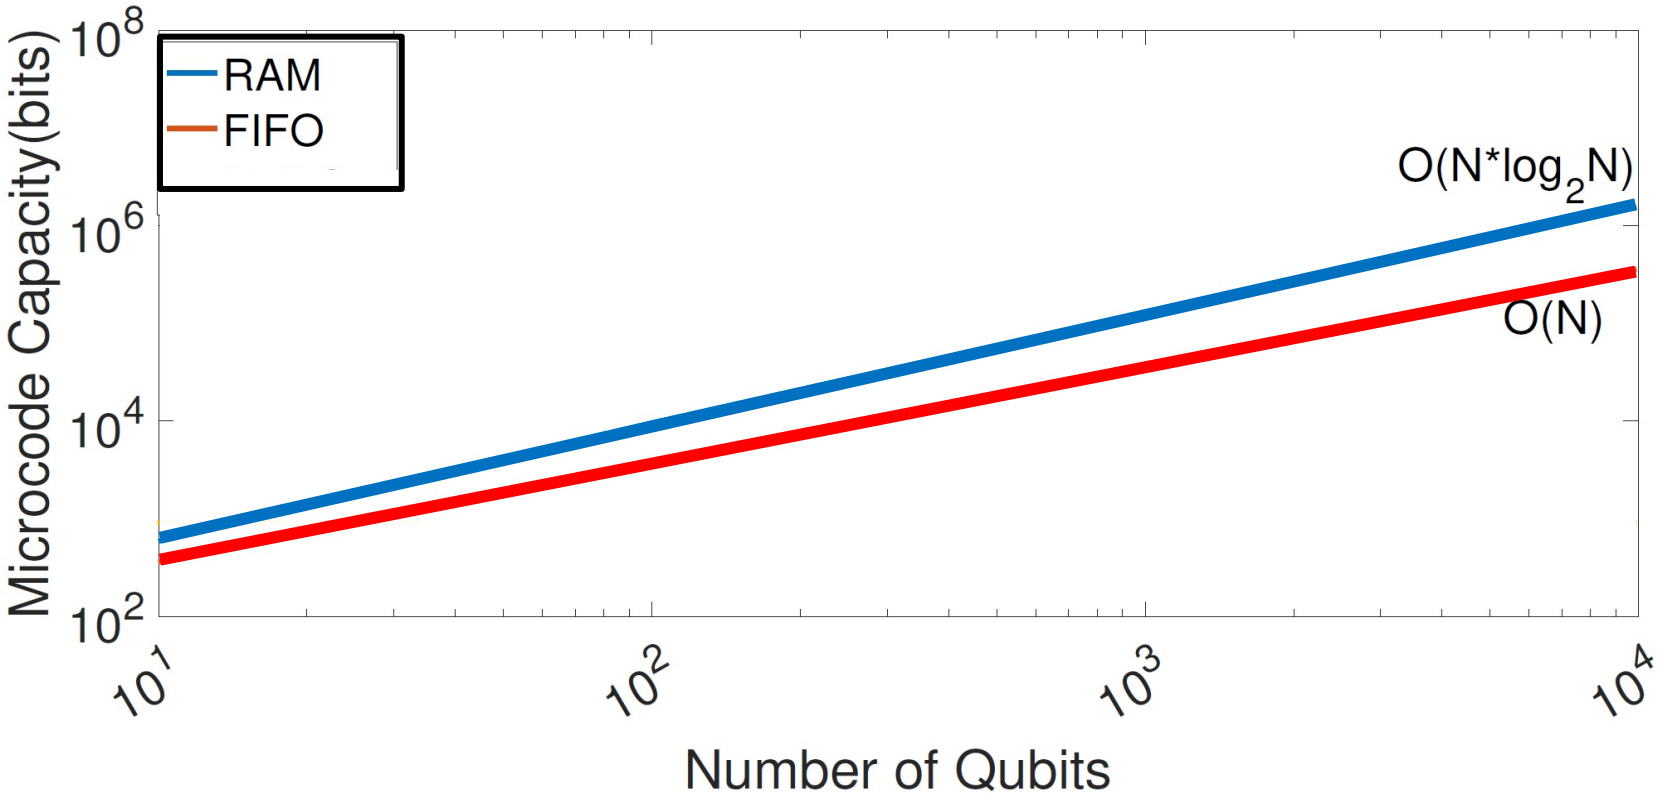

Microcode Memory Capacity Requirements

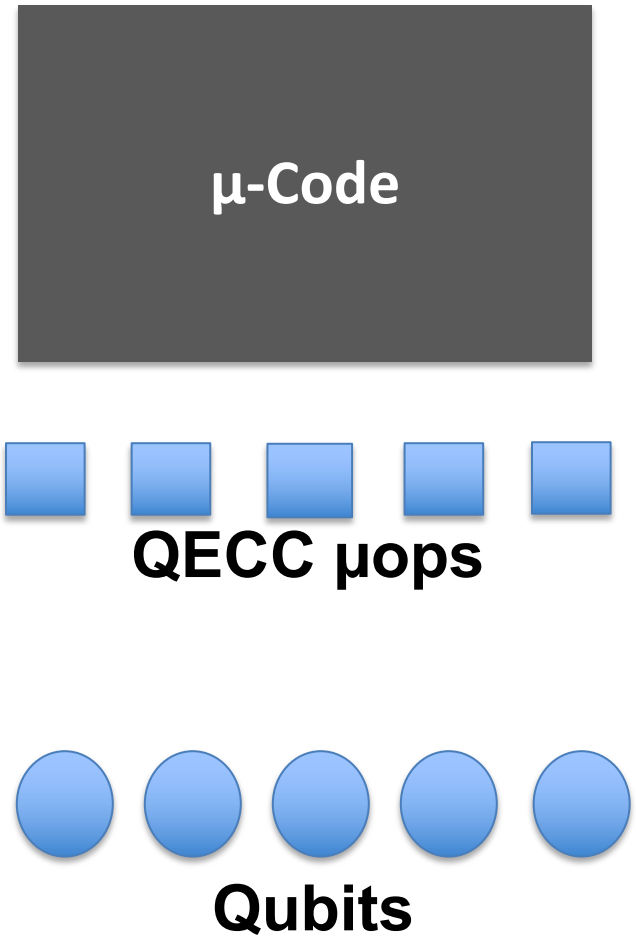

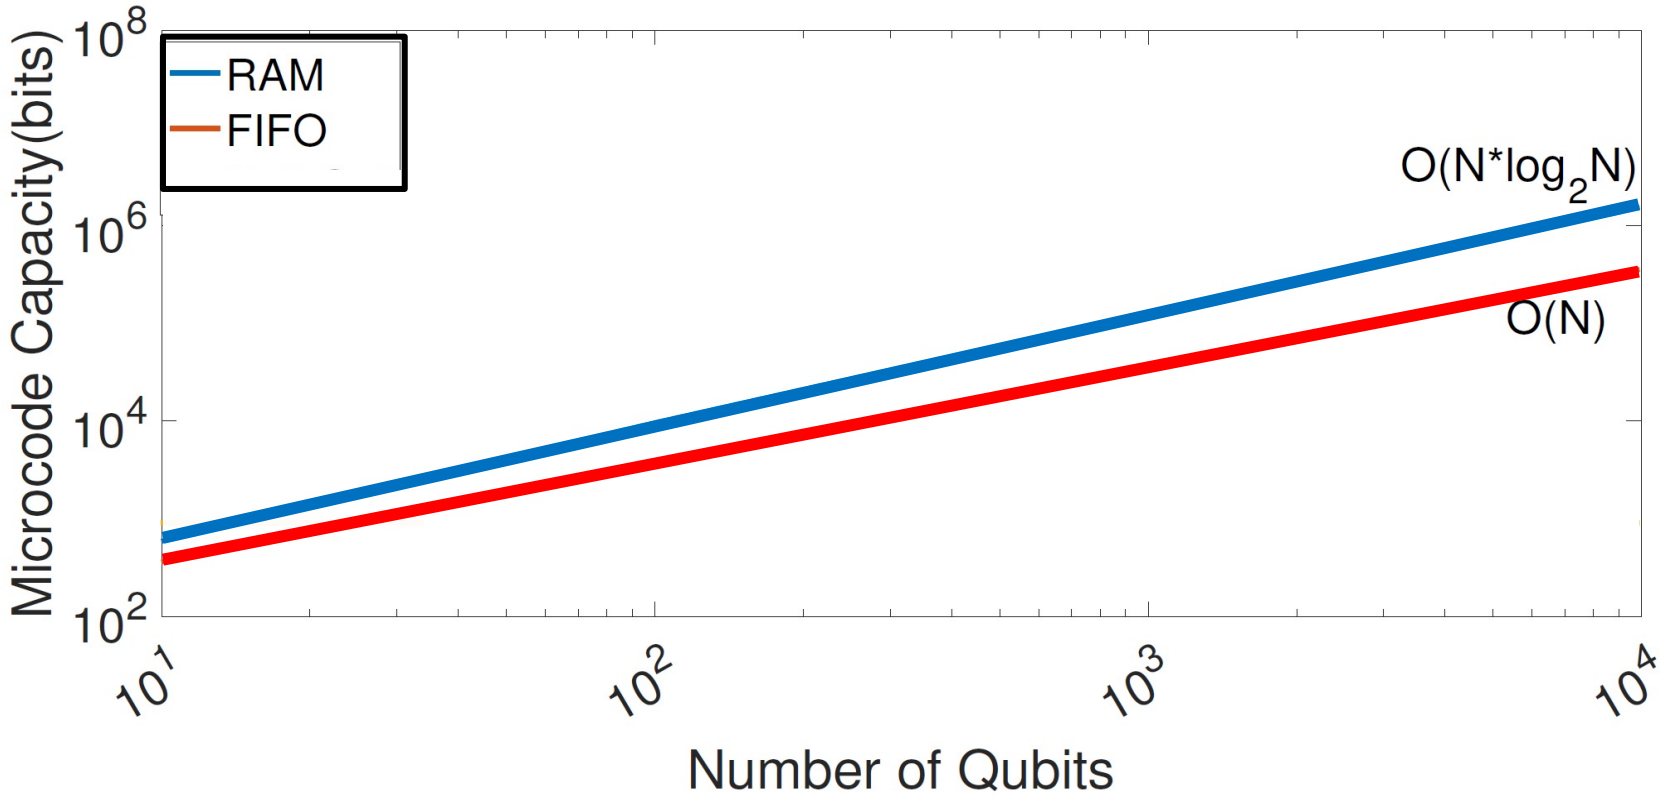

Microcode memory capacity requirement increase linearly with no. of qubits

Limited Memory Capacity at 4 Kelvin

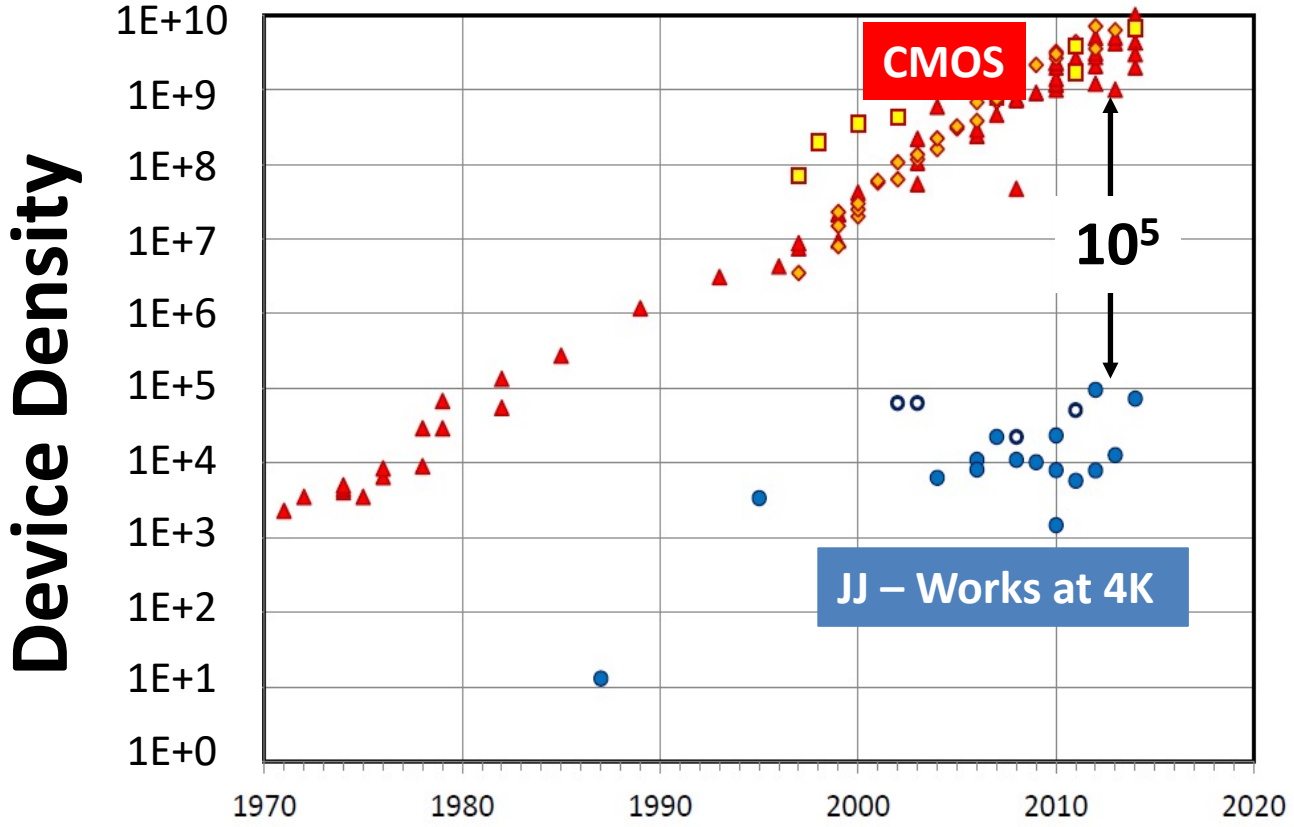

Ref: Beyond CMOS Superconducting Digital Circuits by MIT Lincoln Labs

Limited Memory Capacity at 4 Kelvin

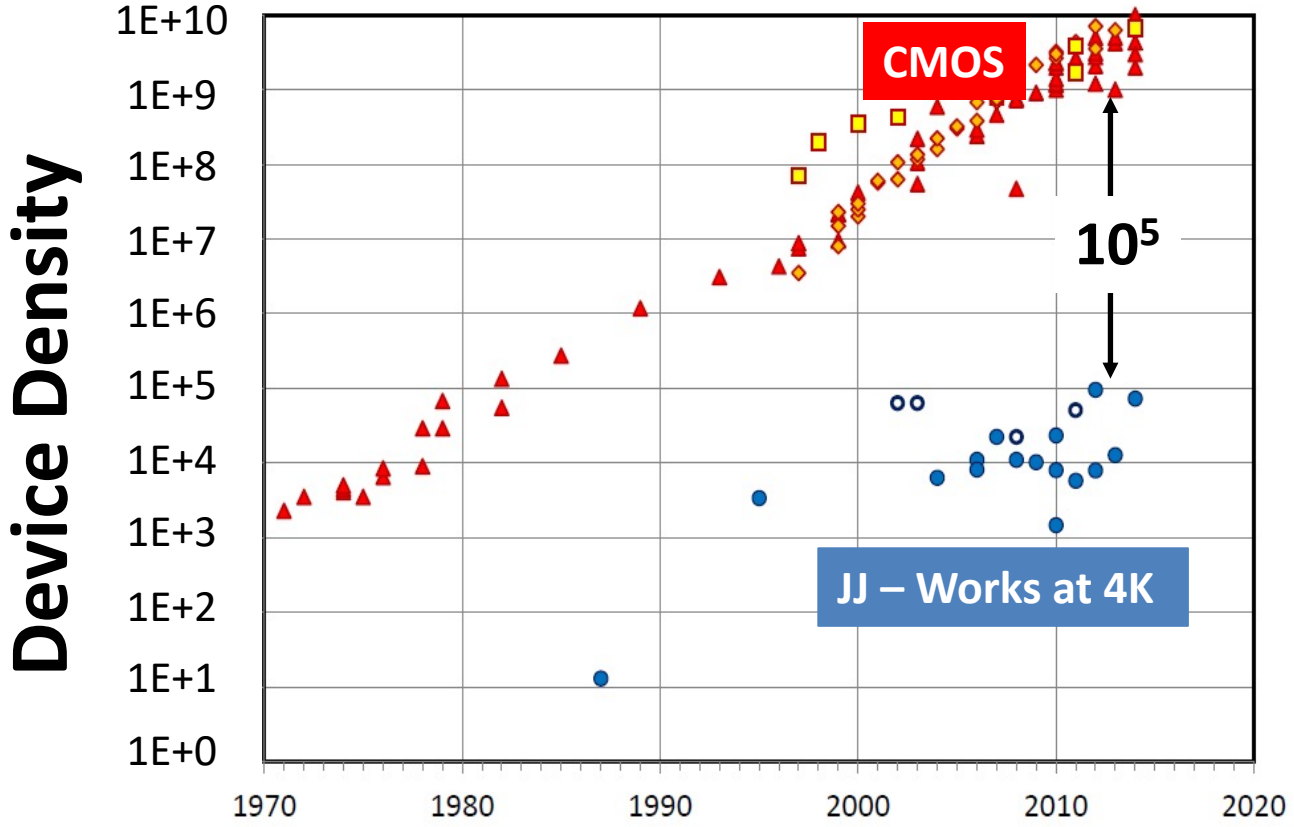

JJ-Memory Capacity

4 Kb/cm^2

Ref: Beyond CMOS Superconducting Digital Circuits by MIT Lincoln Labs

Limited Memory Capacity at 4 Kelvin

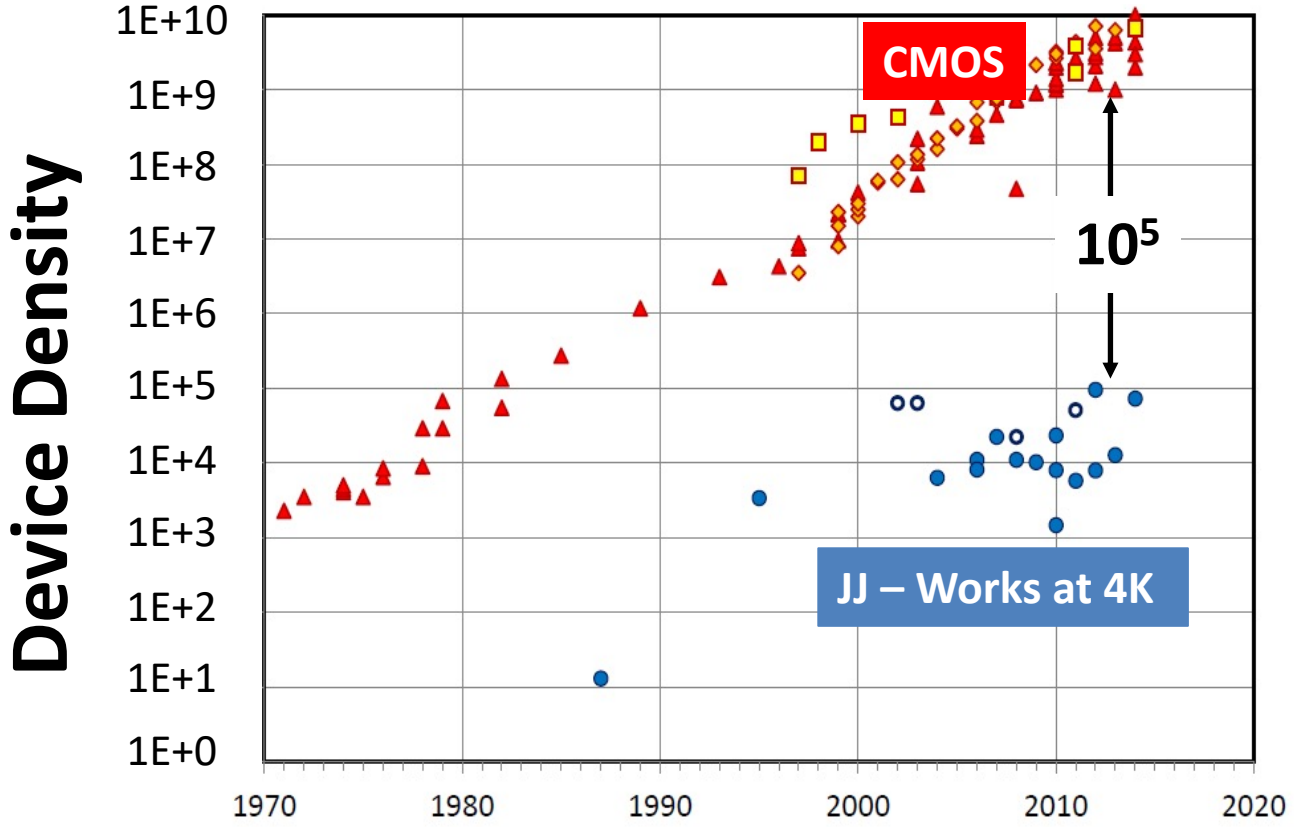

JJ-Memory Capacity
4 Kb/cm^2

Ref: Beyond CMOS Superconducting Digital Circuits by MIT Lincoln Labs

Capacity of JJ-based memory is limited

Memory Capacity Limits Number of Qubits

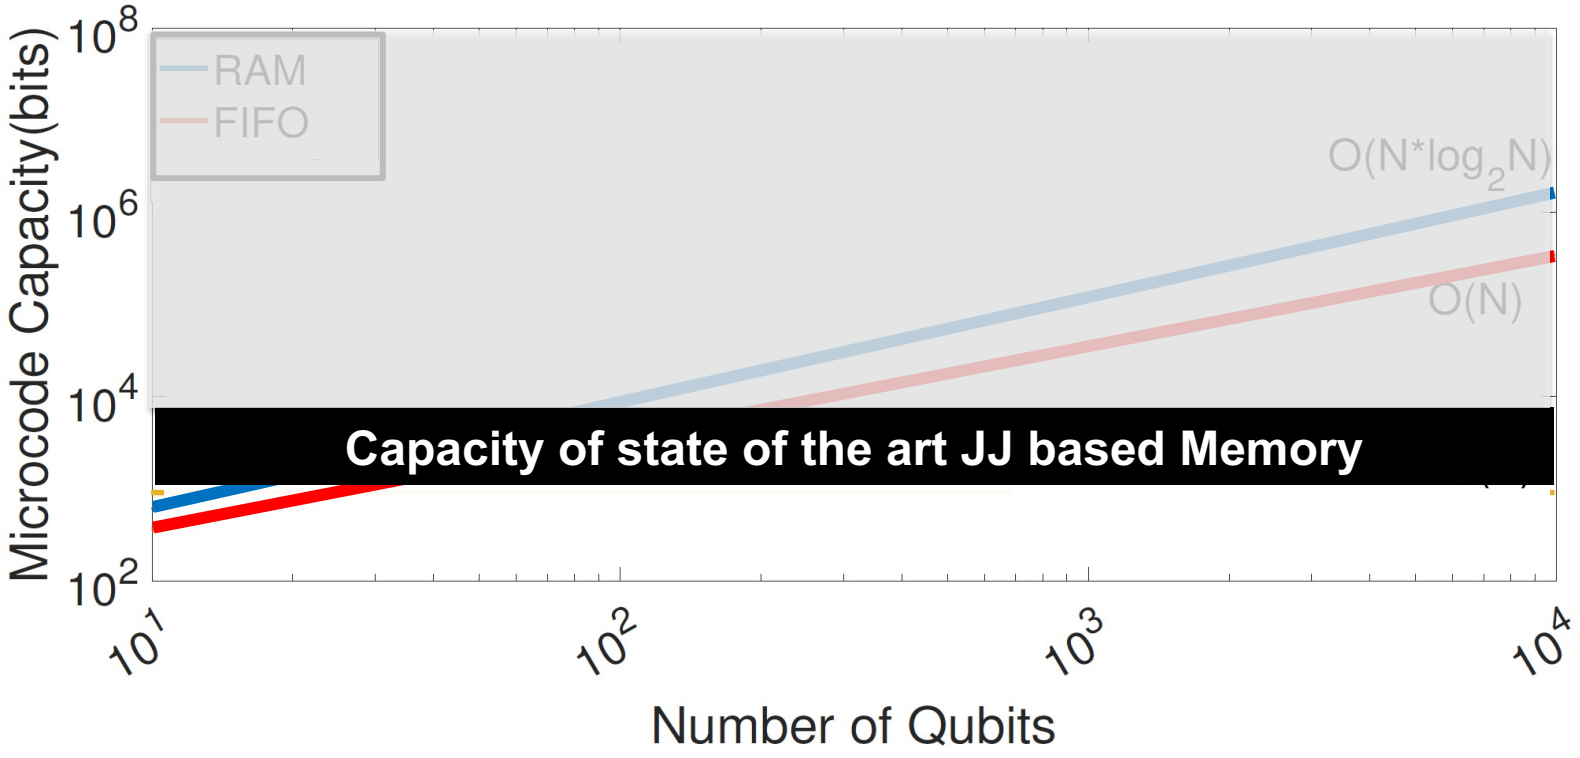

Memory Capacity Limits Number of Qubits

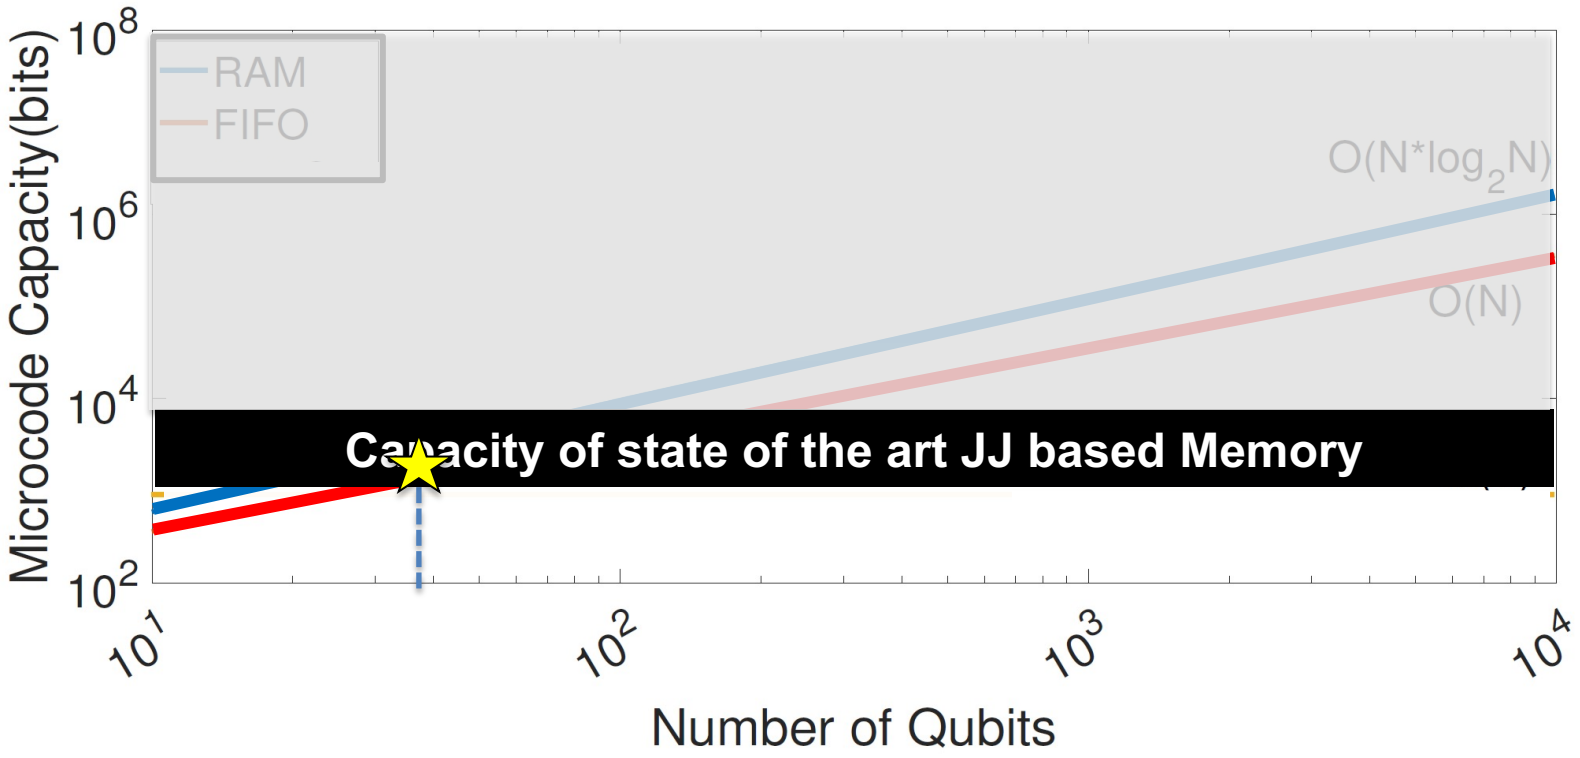

4Kb µcode memory can only support about 35 qubits

Insight- Repetition in QECC Instructions

QECC µops

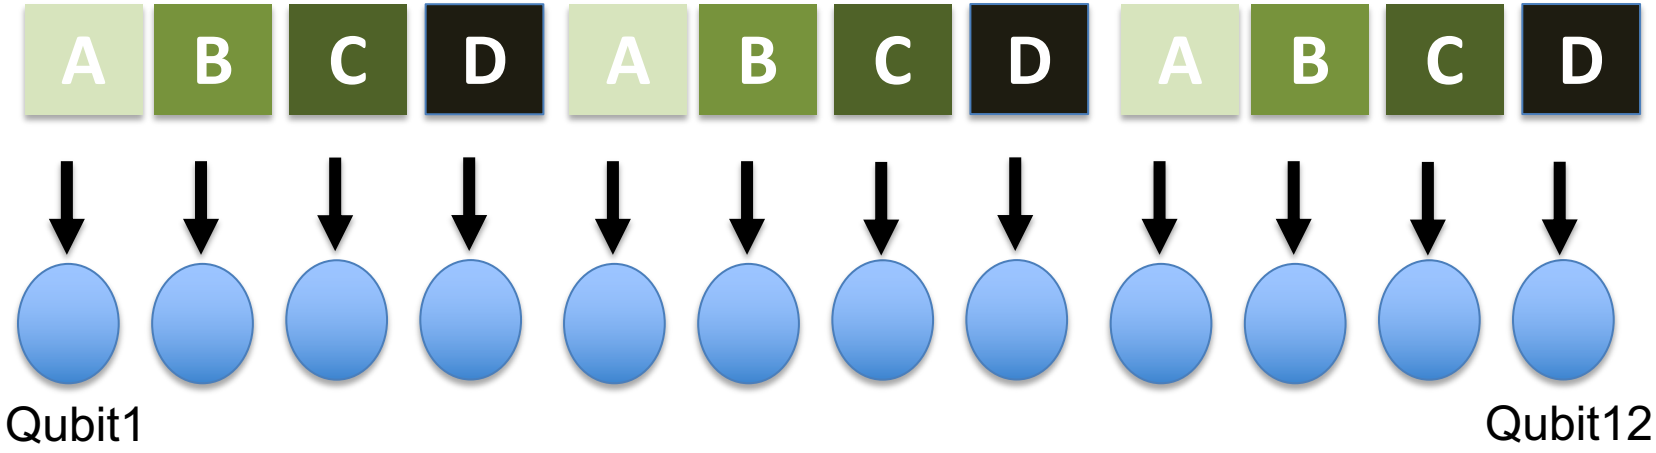

Insight- Repetition in QECC Instructions

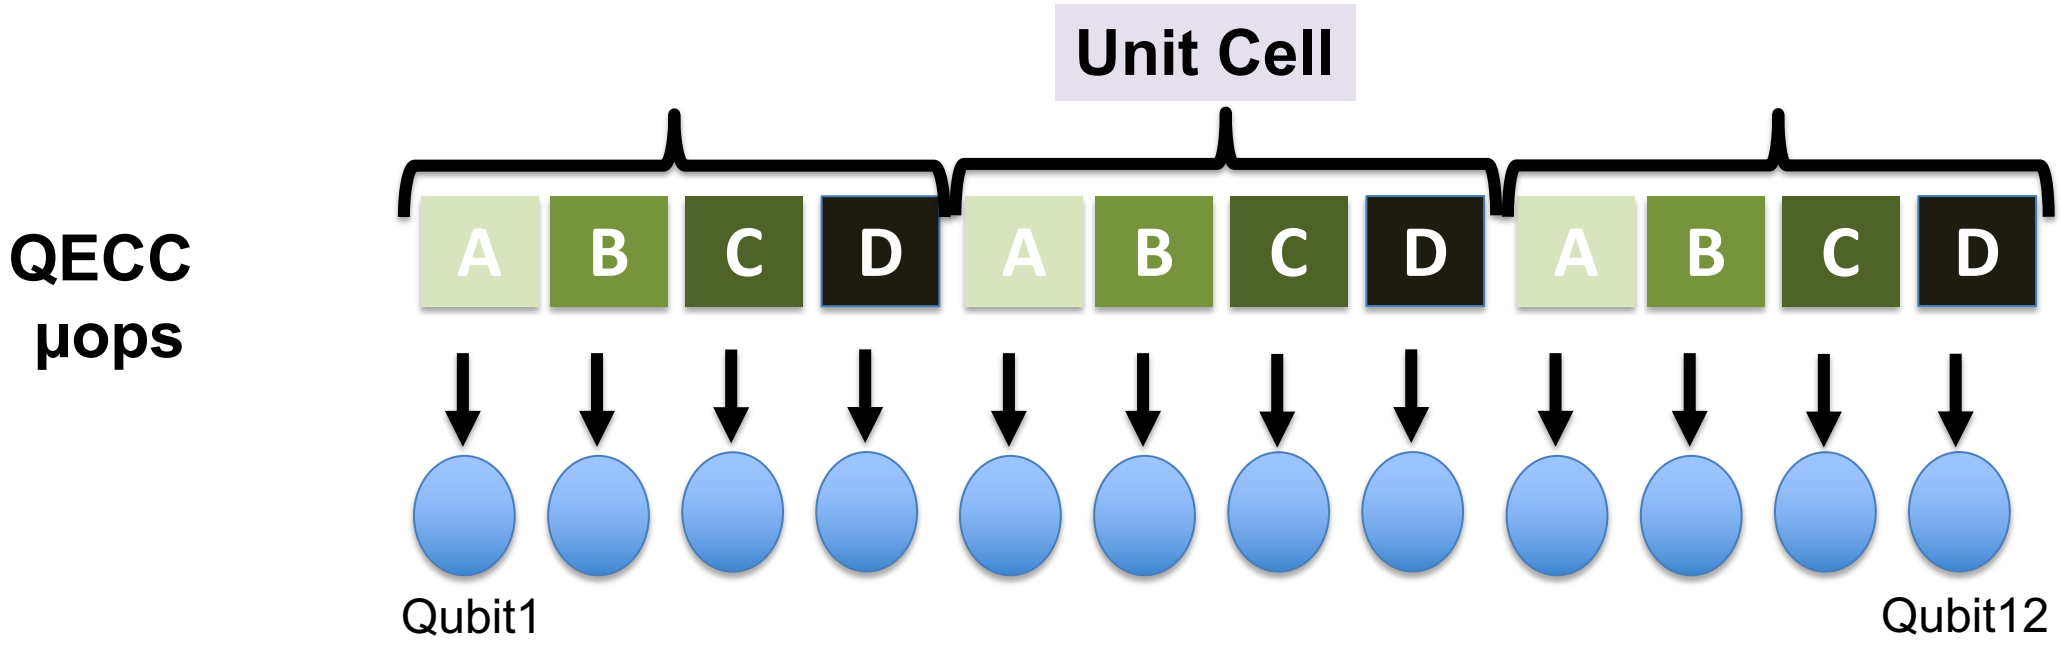

Insight- Repetition in QECC Instructions

QECC µops

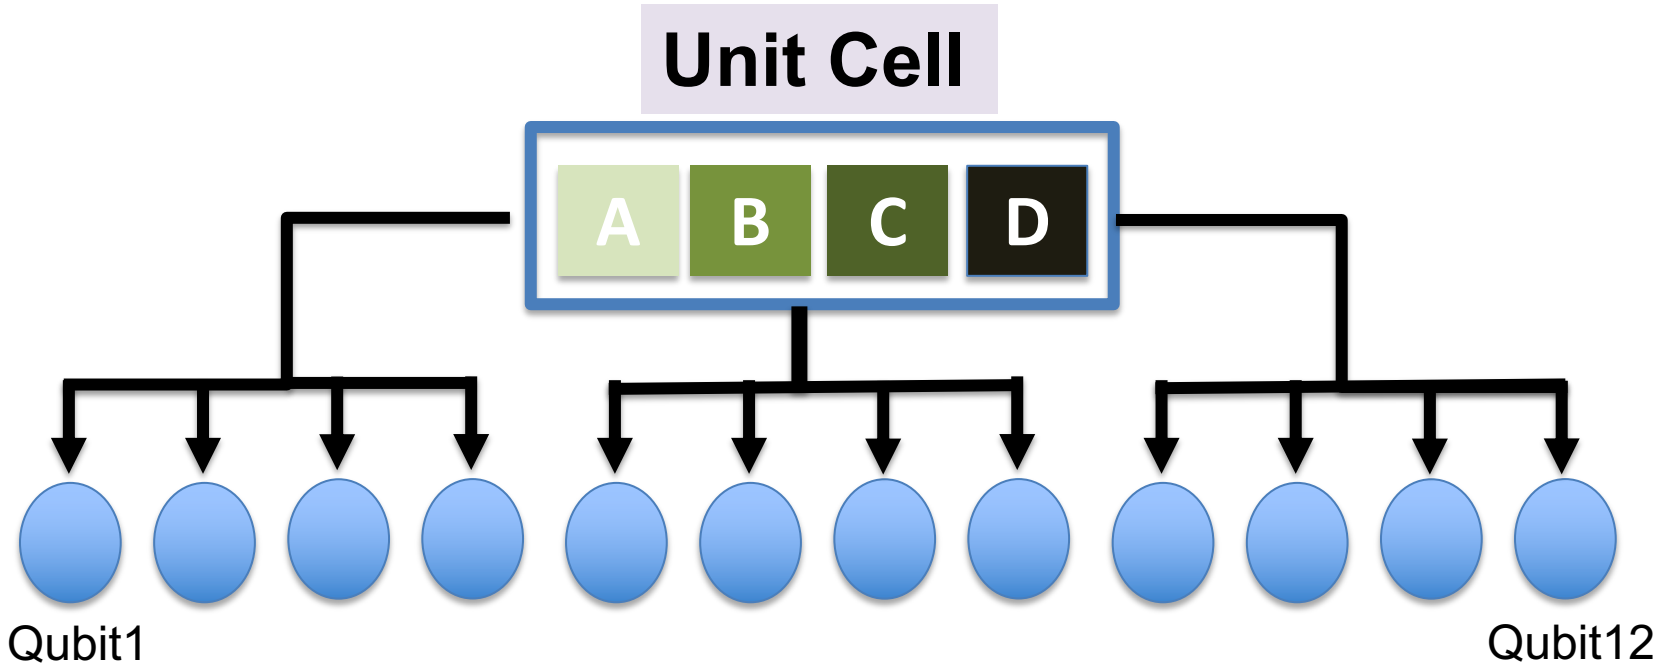

Insight- Repetition in QECC Instructions

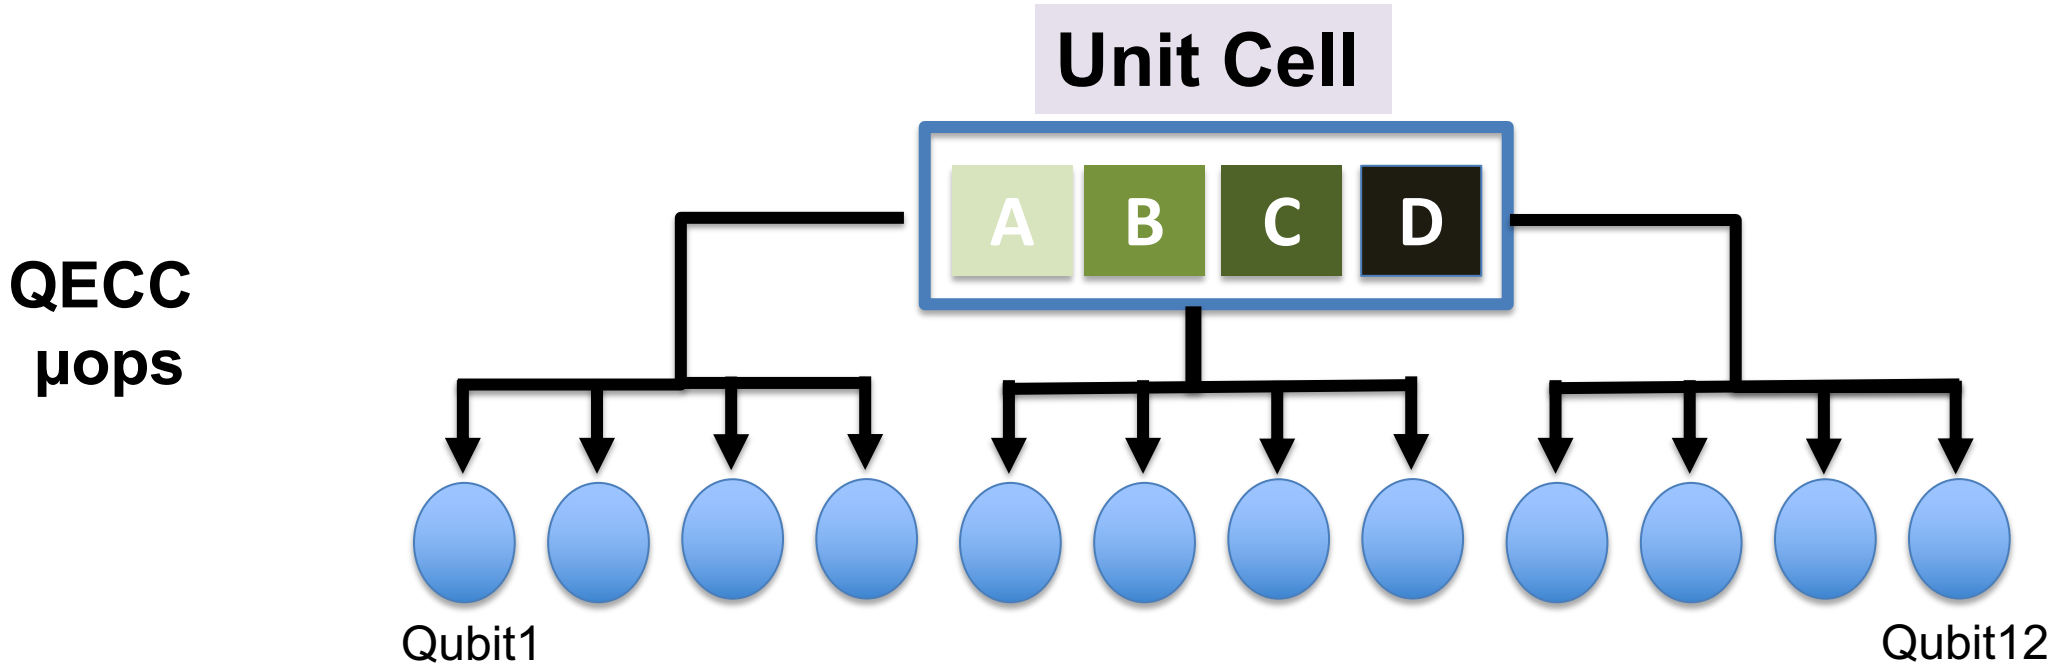

❖ QECC have spatially repeating instruction blocks -- Unit Cell

Insight- Repetition in QECC Instructions

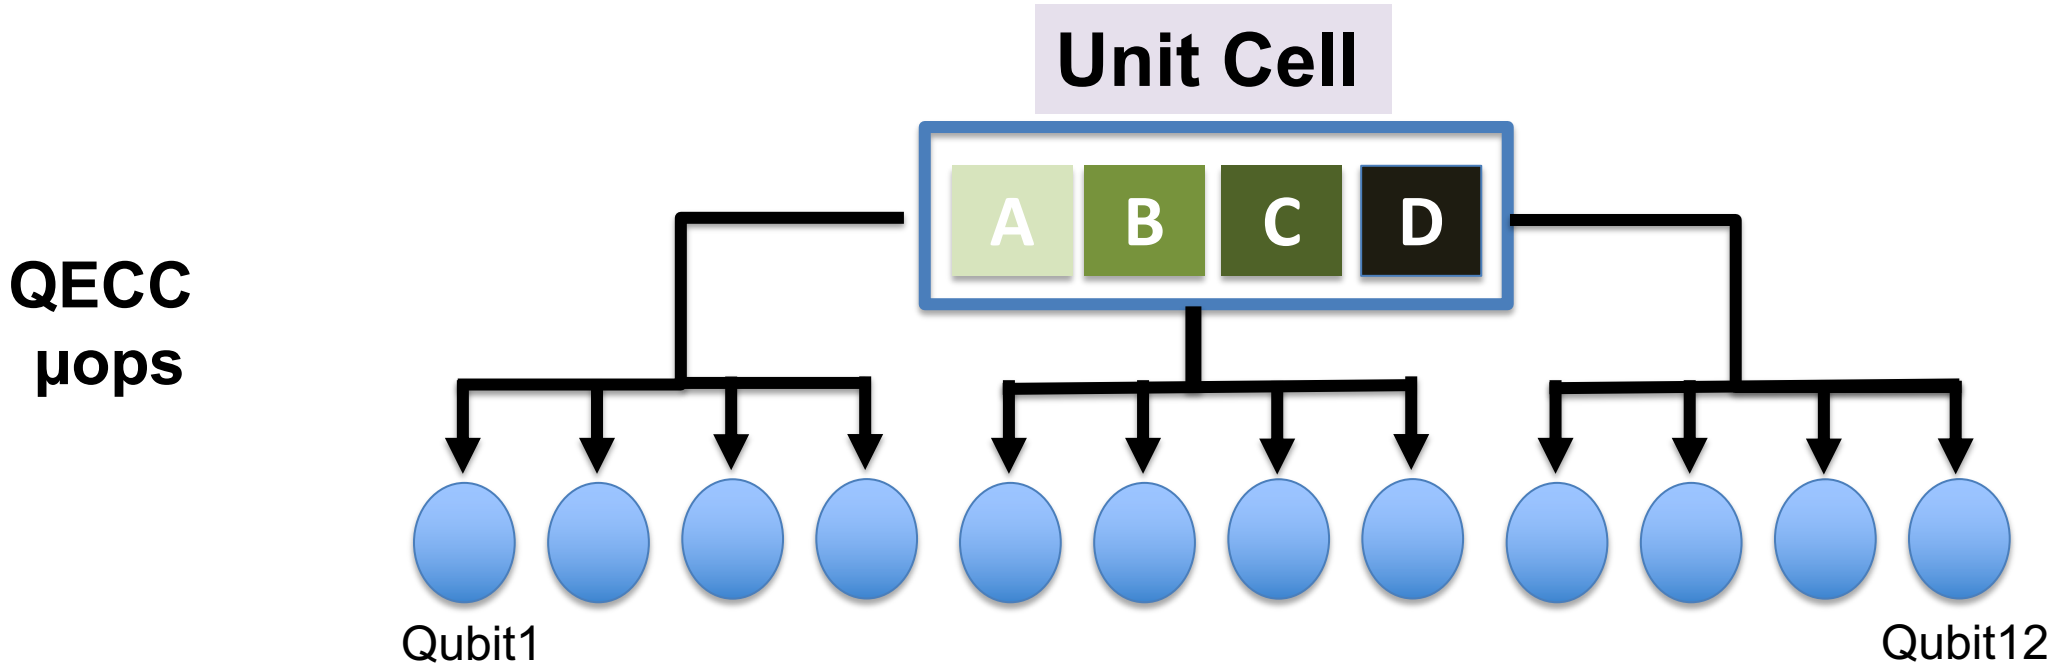

- QECC have spatially repeating instruction blocks -- Unit Cell
- Unit-Cell µops can generate QECC µops for all the qubits

Insight- Repetition in QECC Instructions

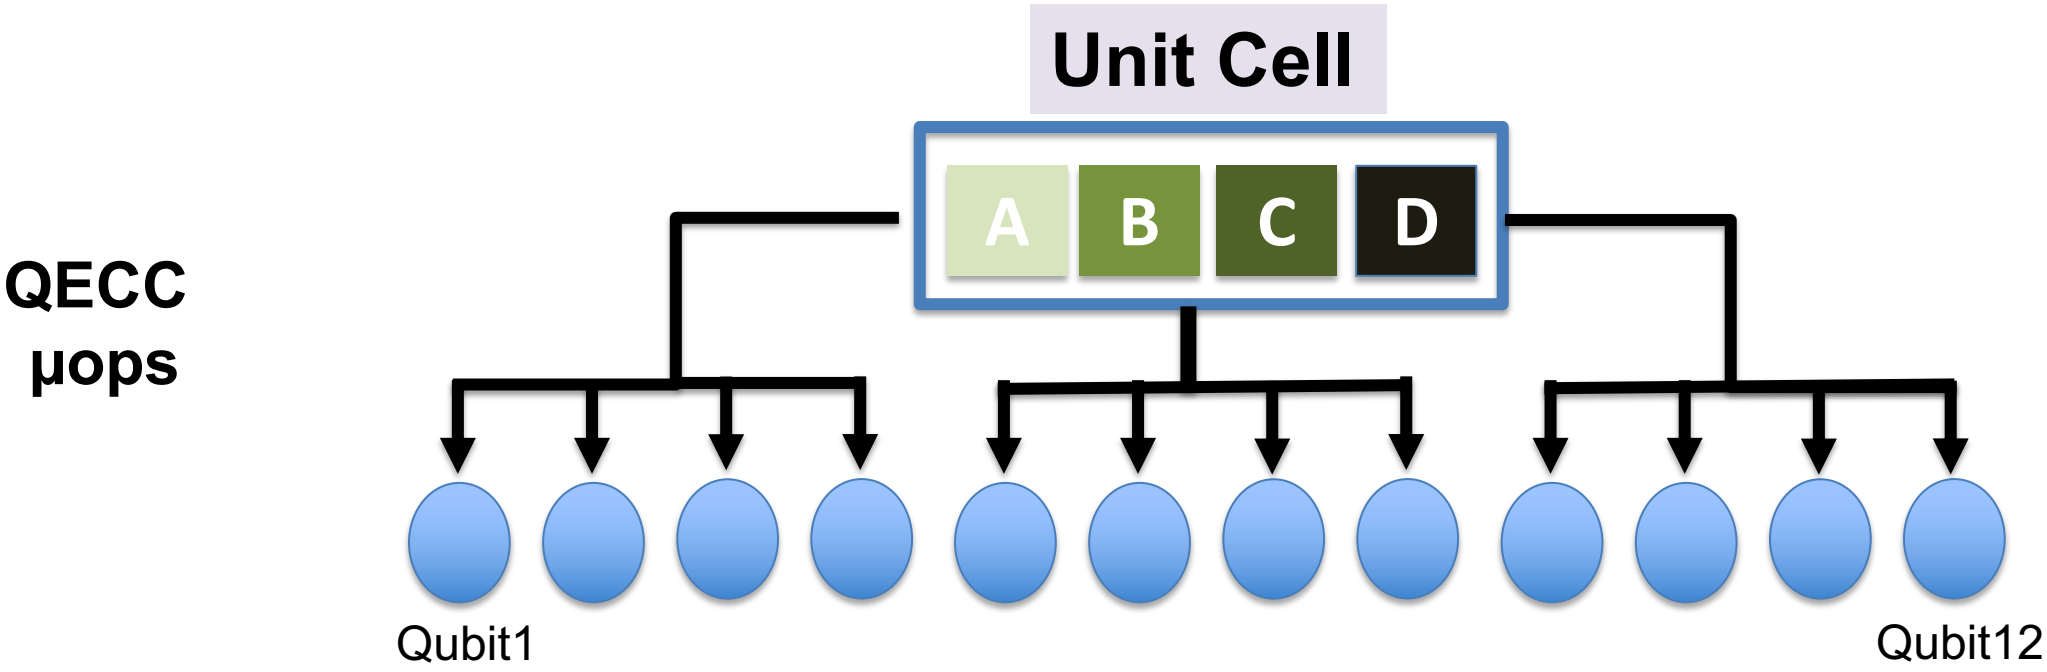

- QECC have spatially repeating instruction blocks -- Unit Cell
- Unit-Cell µops can generate QECC µops for all the qubits

QECC instruction patterns are leveraged to minimize µcode capacity

Unit-Cell Enables Constant Storage

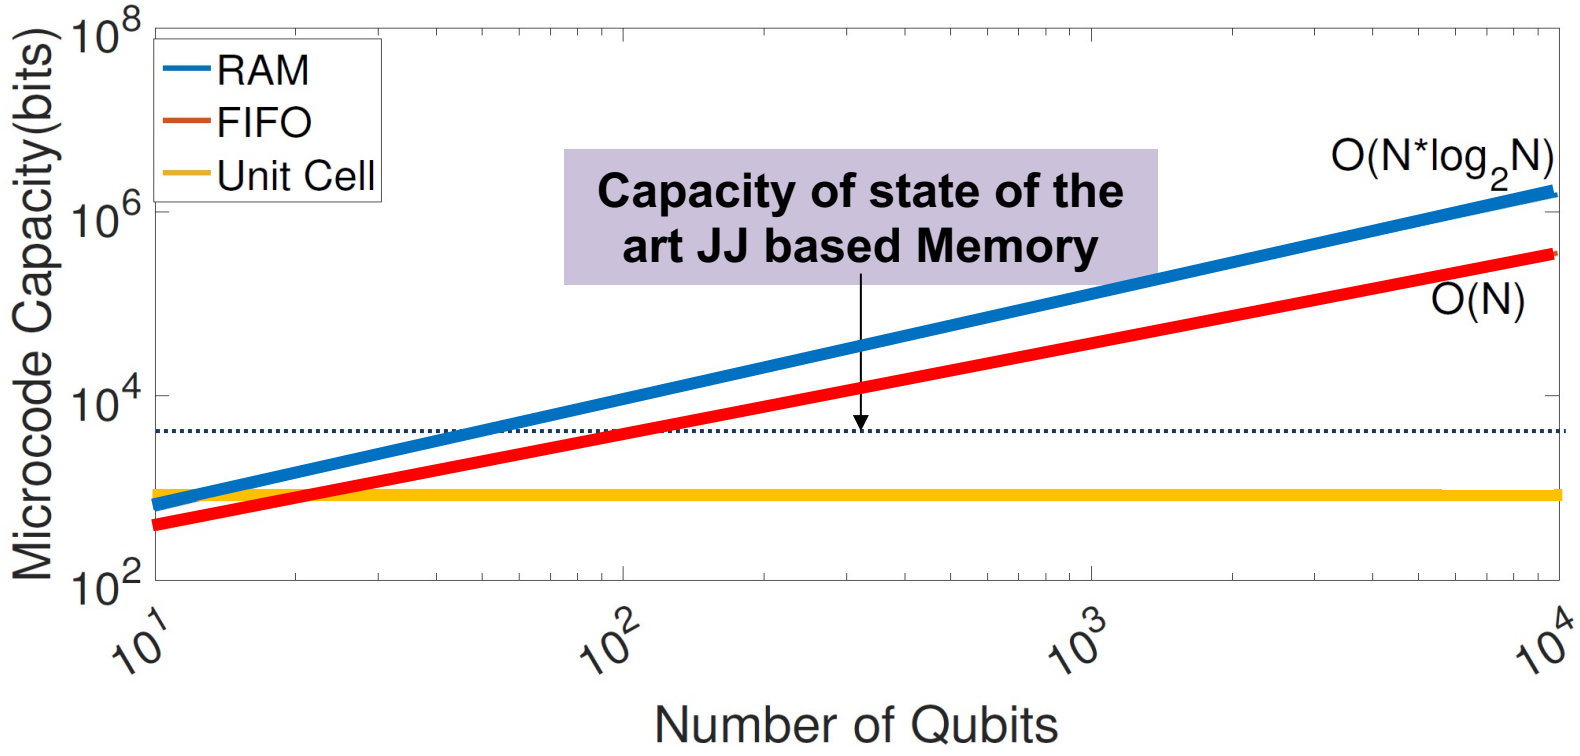

QECC instruction patterns are leveraged to minimize µcode capacity

Unit-Cell Enables Constant Storage

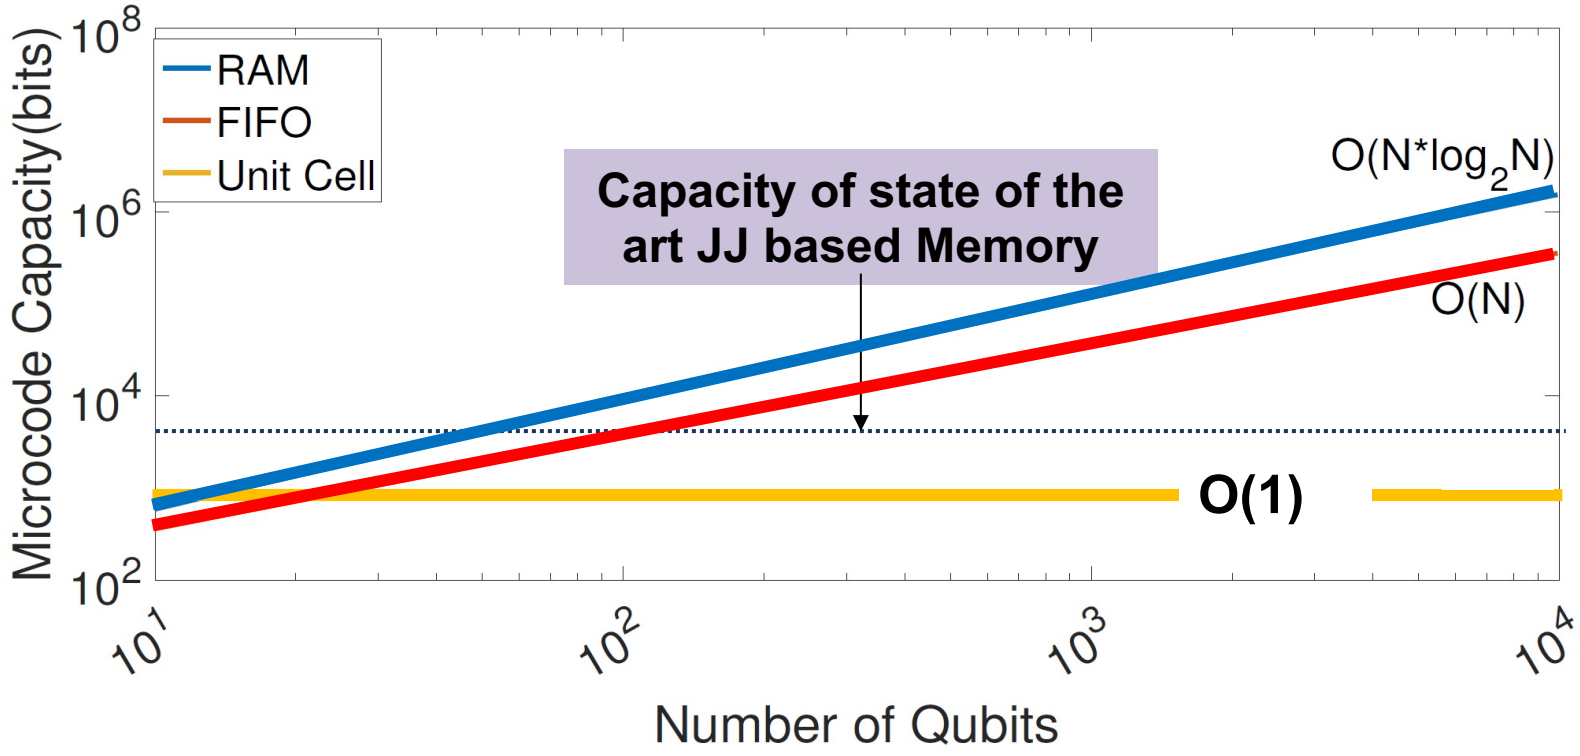

QECC instruction patterns are leveraged to minimize µcode capacity

Number of Qubits Serviced

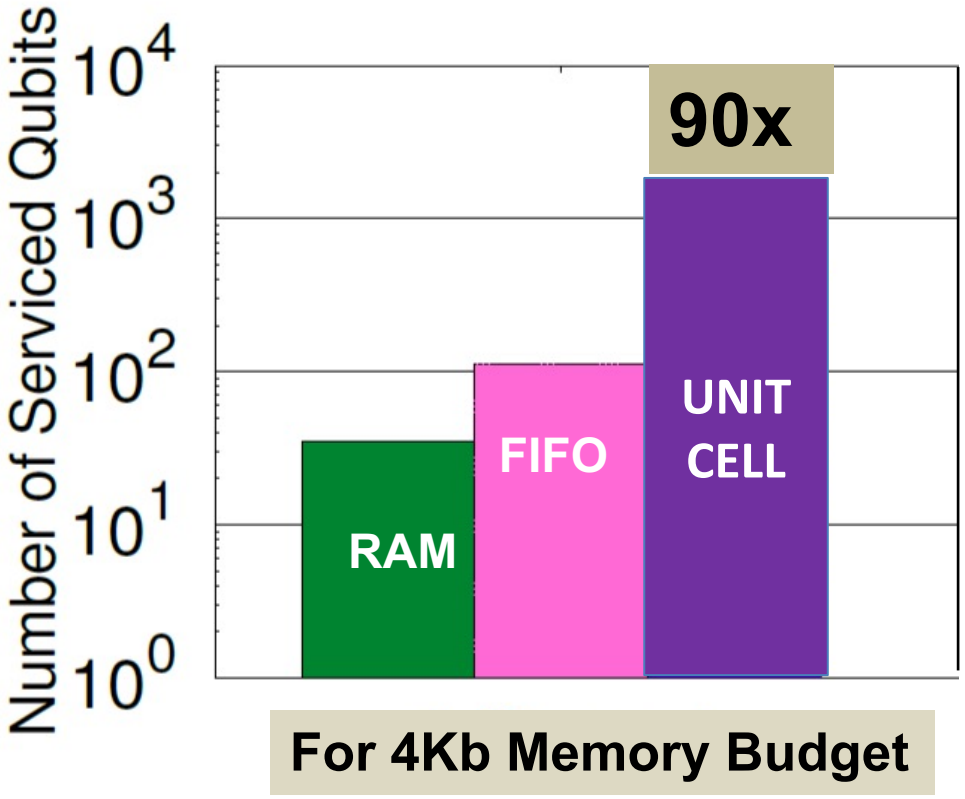

Microcode can service up to 2000 qubits with 4Kb of microcode

Evaluations

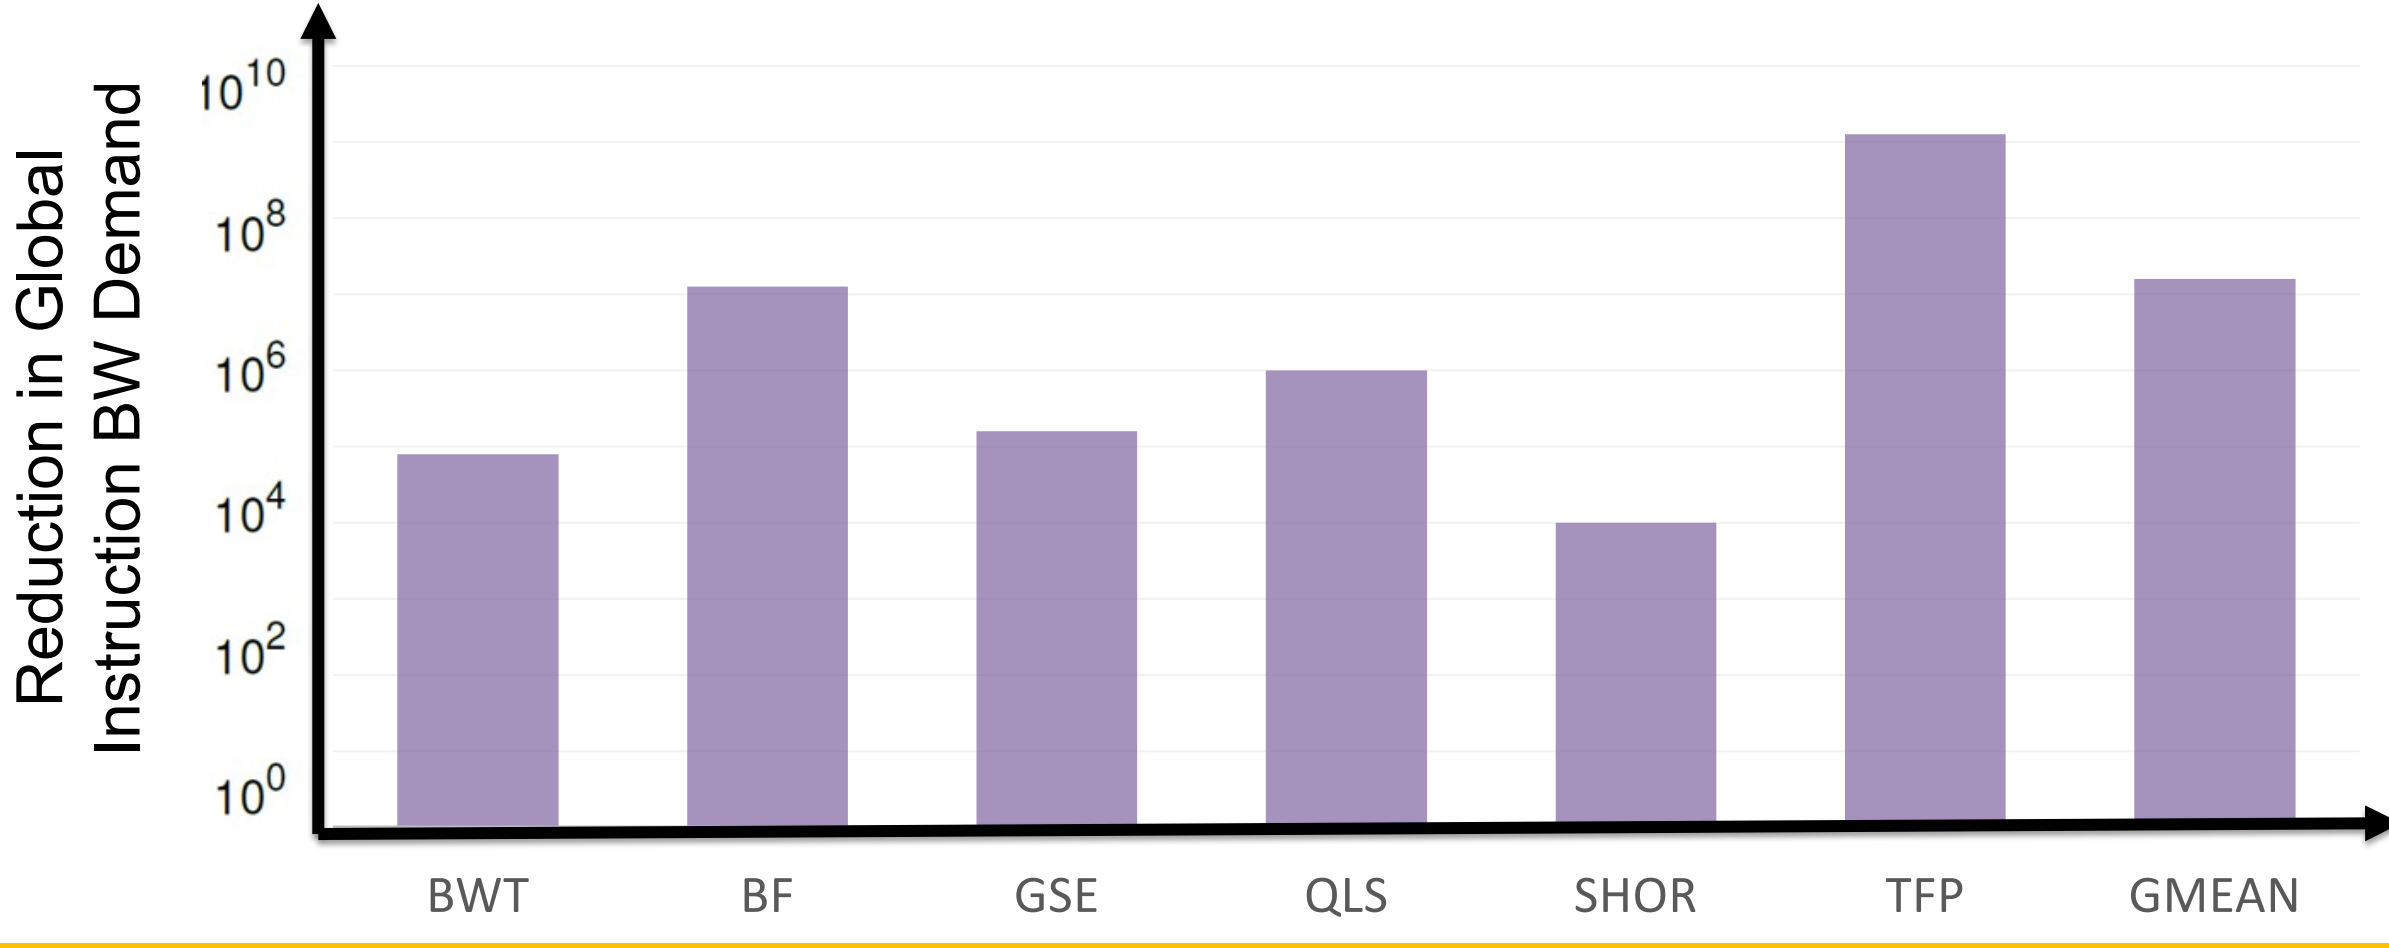

QuEST reduces global bandwidth demand by seven orders of magnitude

Evaluations

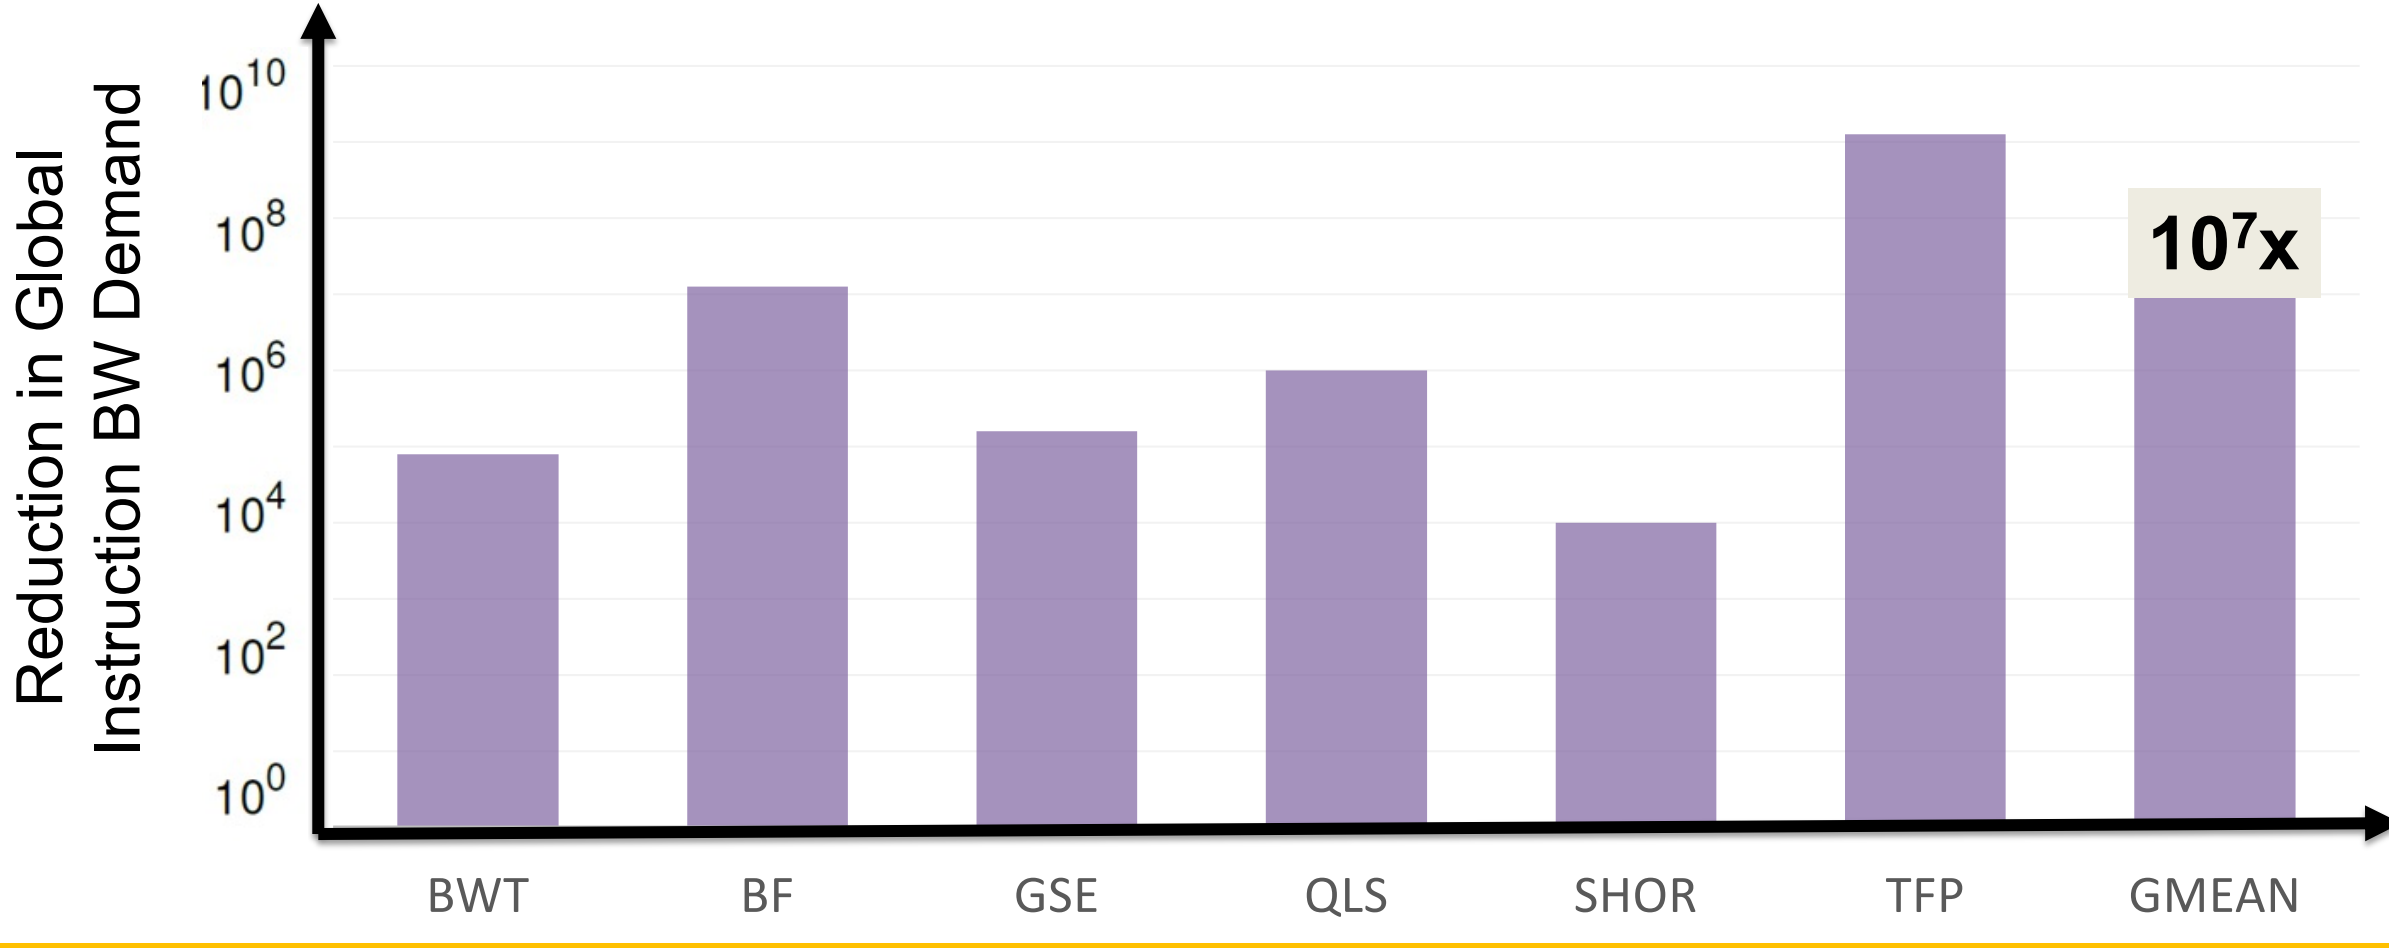

QuEST reduces global bandwidth demand by seven orders of magnitude

Conclusion

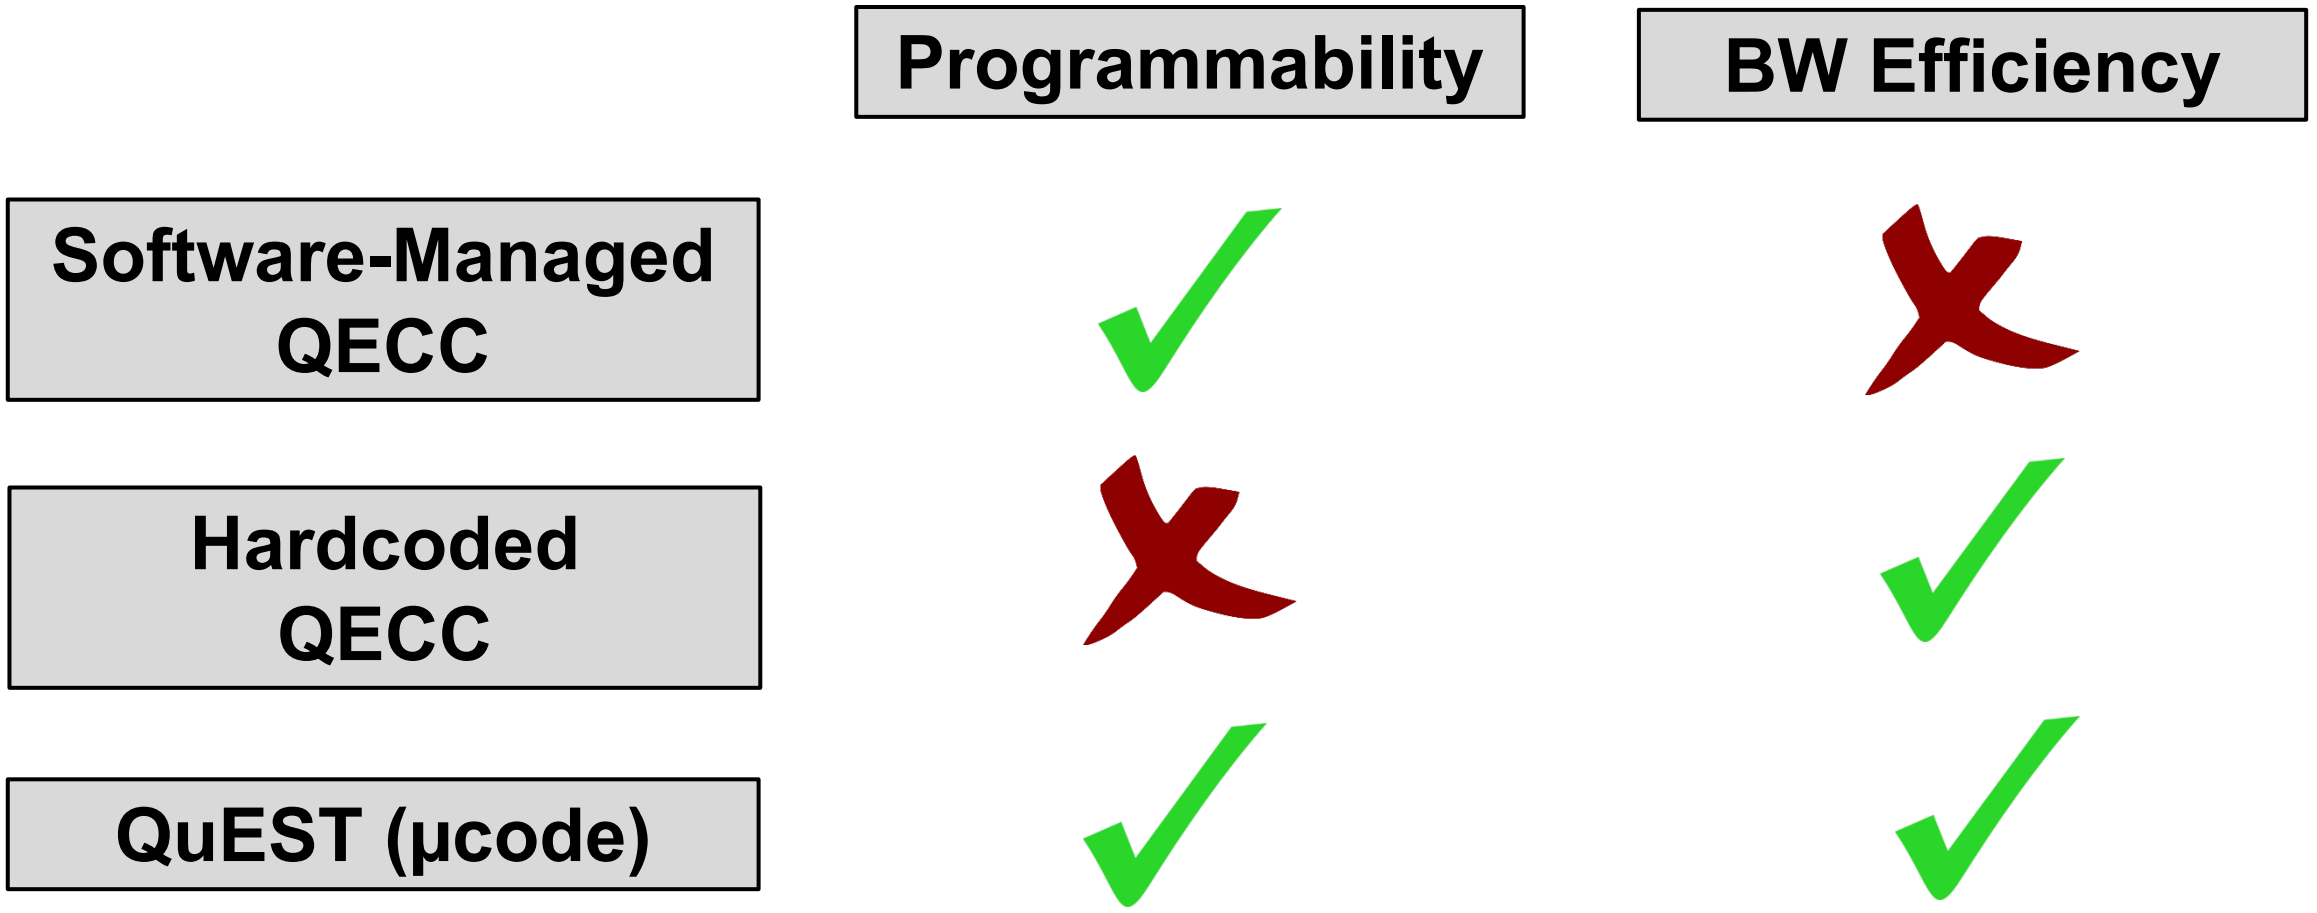
Programmability
BW Efficiency
Software-Managed QECC
Hardcoded QECC
QuEST (µcode)

Questions?


HOW'S YOUR QUANTUM COMPUTER PROTOTYPE COMING ALONG?
GREAT!
Dilbert.com DilbertCartoonist@gmail.com
THE PROJECT EXISTS IN A SIMULTANEOUS STATE OF BEING BOTH TOTALLY SUCCESSFUL AND NOT EVEN STARTED.
4-17-12 ©2012 Scott Adams, Inc./Dist. by Universal Uclick
CAN I OBSERVE IT?
THAT'S A TRICKY QUESTION.

Backup slides