Quantum Experiment Data Framework

Enabling Low Latency, High Throughput Data Handling for Quantum Control Systems

Kenji Kikuchi[1], Ranjith Kumar Ramachandran[2], Ahmet Bilgili[2], Arnaud Carignan-Dugas[3], Eric Holland[4]

Keysight Technologies

[1]Tokyo, Japan; [2]Singapore, Singapore; [3]Kitchener, Canada; [4]Boston, United States

Background

- Superconducting qubits, a leading quantum computing architecture, are controlled and read out using microwave pulses managed by a room temperature system. Keysight's Quantum Control System (QCS), introduced in 2022, now powers large scale platforms such as AIST's 1000 qubit and RIKEN RQC Fujitsu Collaboration Center's 256 qubit systems.
- As qubit counts grow, measurement data scales rapidly, creating challenges in visualization, analysis, and ultra low latency transfer. Traditional pipelines struggle with scalability, and current QCS software lacks low latency distributed analysis.
- This work proposes an architecture for low latency, high throughput data handling from QCS to analytics and evaluates its performance under concurrent workloads.

Methods

- To address these challenges, we designed and implemented a high-performance data pipeline and scalable web service architecture for quantum control systems. The following subsections describe the design principles and implementation details for each component.

Method 1: Data Pipeline Optimization

- In conventional implementations, relational databases such as a relational SQL database were used to store qubit measurement results, but this approach has several limitations: storage I/O bottlenecks under high data rates and limited support for distributed processing across clusters. These constraints become critical as qubit counts and measurement volumes scale.
- To overcome these challenges, we adopted a combination of a high-performance object store and an embedded analytical database engine. A high-performance object store provides scalable, high speed object storage, while an embedded analytical database engine's in-memory execution avoids storage I/O overhead and accelerates analytical queries. This design establishes the foundation for low-latency, high throughput data handling in subsequent stages of the pipeline.

Method 2: Low Latency Data Transport

- In the conventional software stack, database access for storing measurement results involved multiple layers, leading to repeated memory copies and substantial latency in real-time streaming.
- To address this, we streamlined the data path by reducing intermediate serialization and copy operations. Furthermore, we eliminated copying at critical boundaries by enabling memory object sharing between software layers. These optimizations form the foundation for ultra low latency data transfer in our architecture.

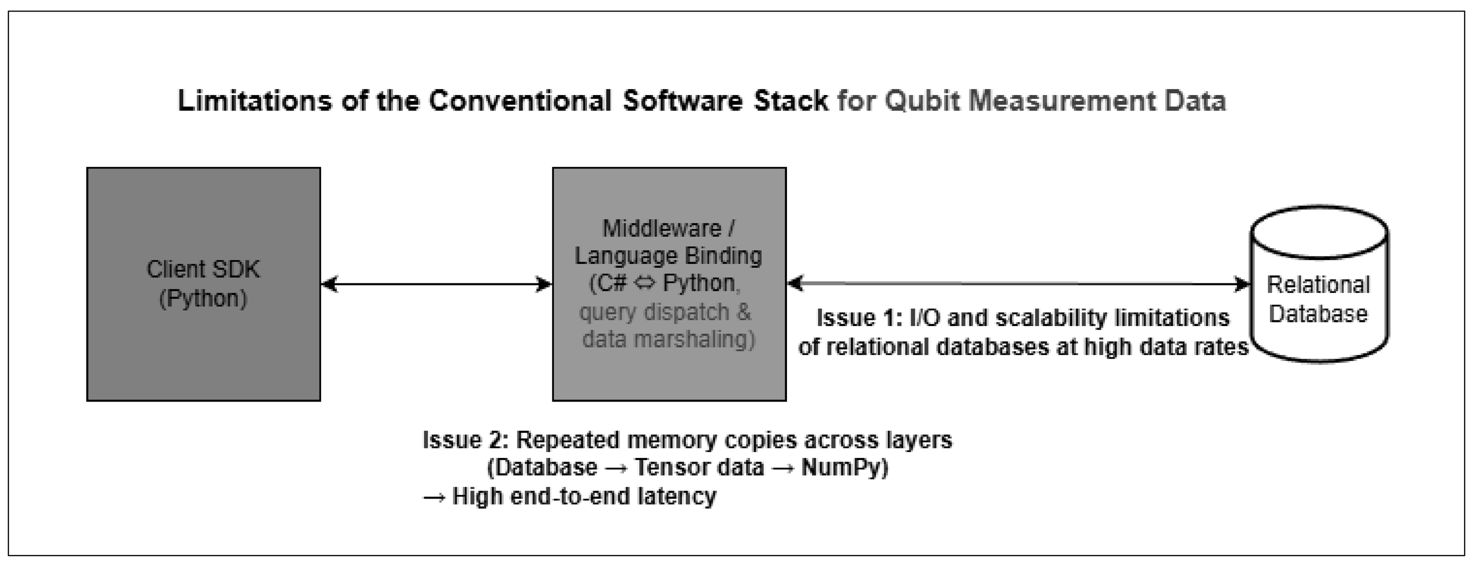

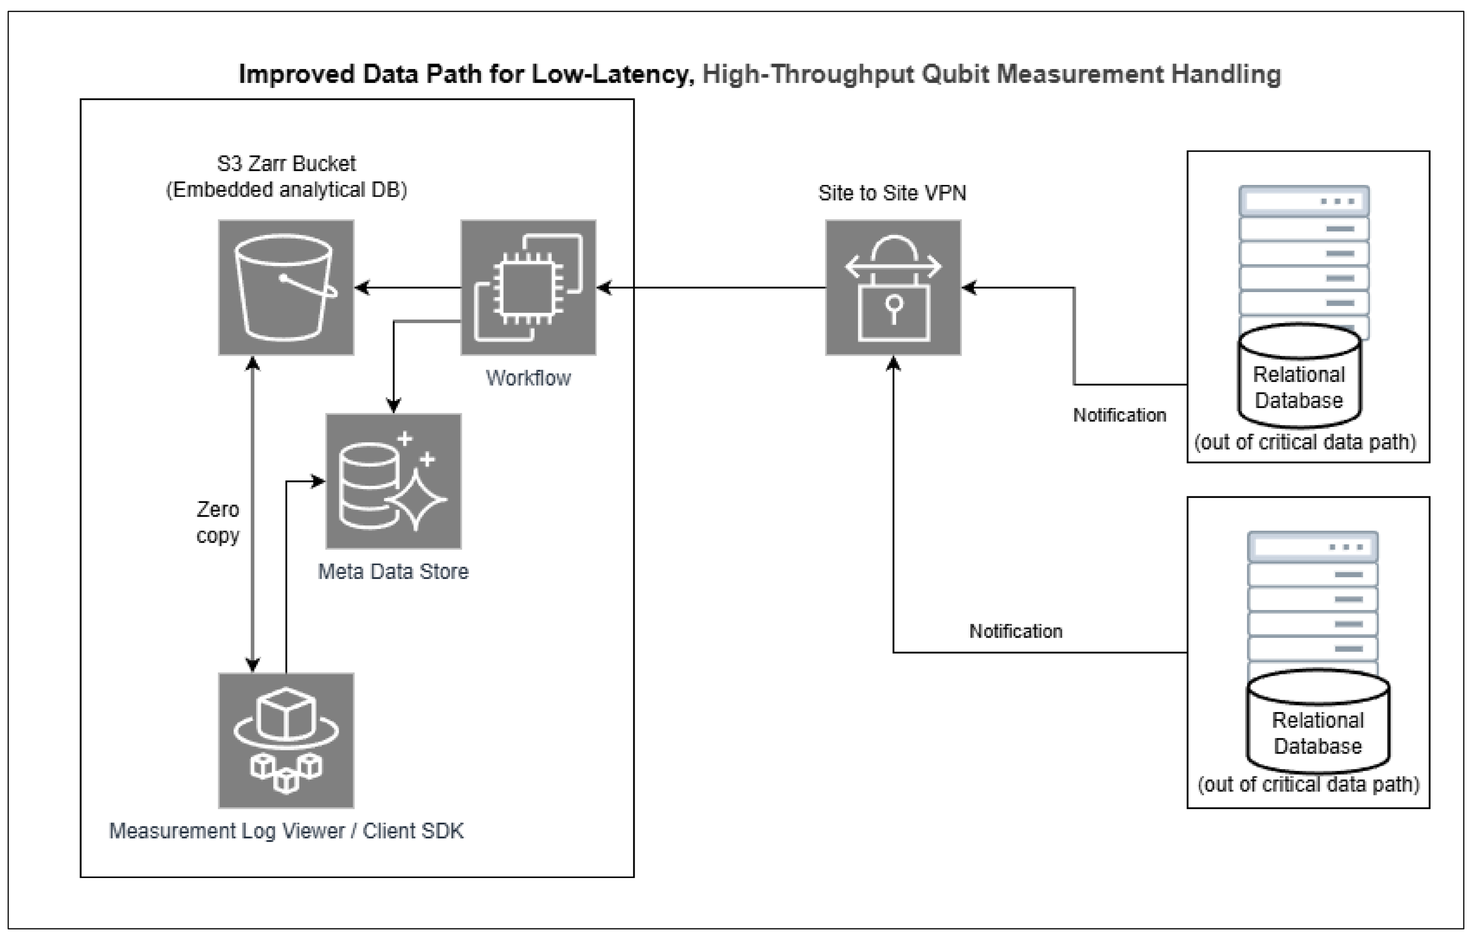

Method 3: Scalable Web Services

- Large scale quantum control systems are expected to handle concurrent access from multiple users, but a single server compute architecture suffers from throughput degradation under simultaneous workloads.
- To address this, we adopted Kubernetes for automatic scaling and load balancing, in addition to a distributed backend built on a high-performance object store and an embedded analytical database engine. This architecture enables HPC-style, data-parallel analysis of quantum experiments under concurrent, multi-user workloads.

Method 4: Integration with Cloud Infrastructure

- To enhance user convenience, we integrated the newly developed Quantum Experiment Data Framework (QEDF) functionality into the existing QCS Cloud platform, which already supports multi-user execution of quantum programs using QCS.

Results

- Prototype evaluations of the Quantum Experiment Data Framework (QEDF) demonstrated gains in data transfer efficiency and visualization responsiveness. High throughput data handling and array optimized storage enabled up to **611x** faster data access compared to conventional SQL-based data access in QCS pipeline. n = 30 concurrent runs.

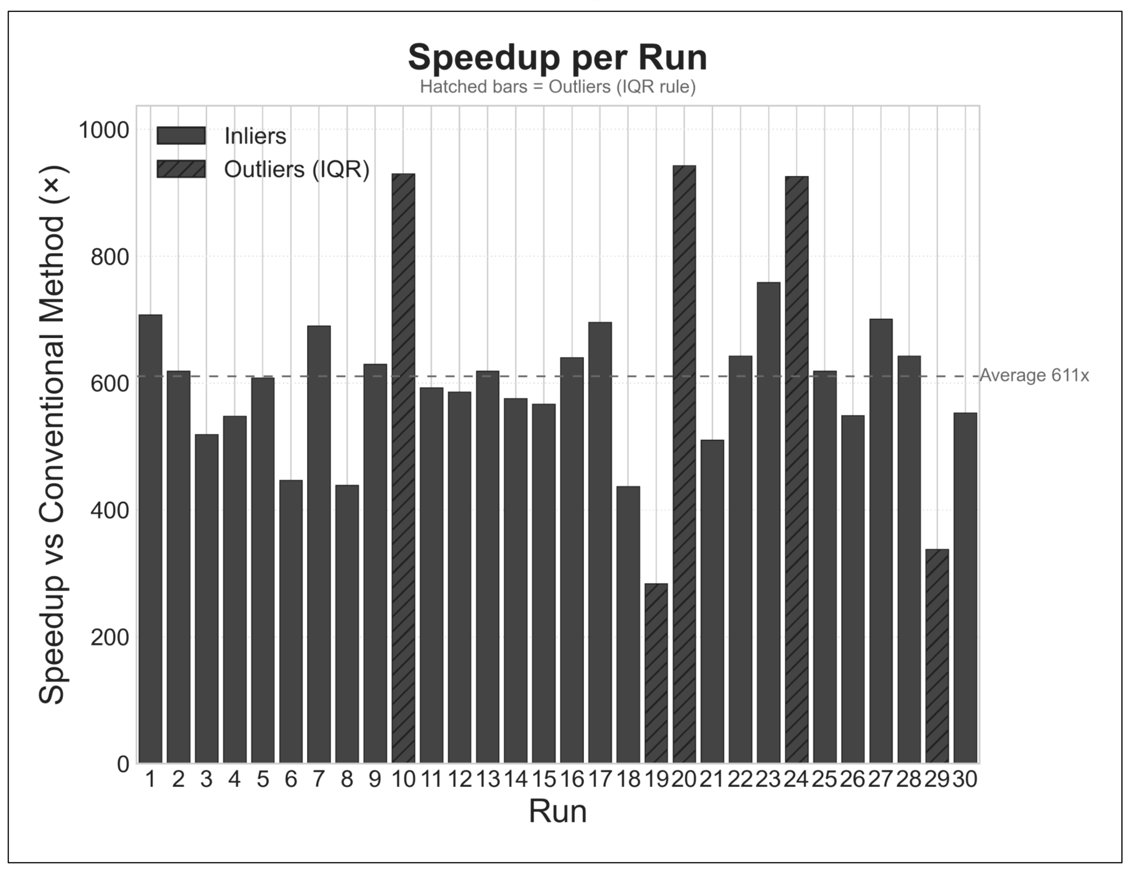

n = 30; Outliers: Runs 10, 19, 20, 24, 29

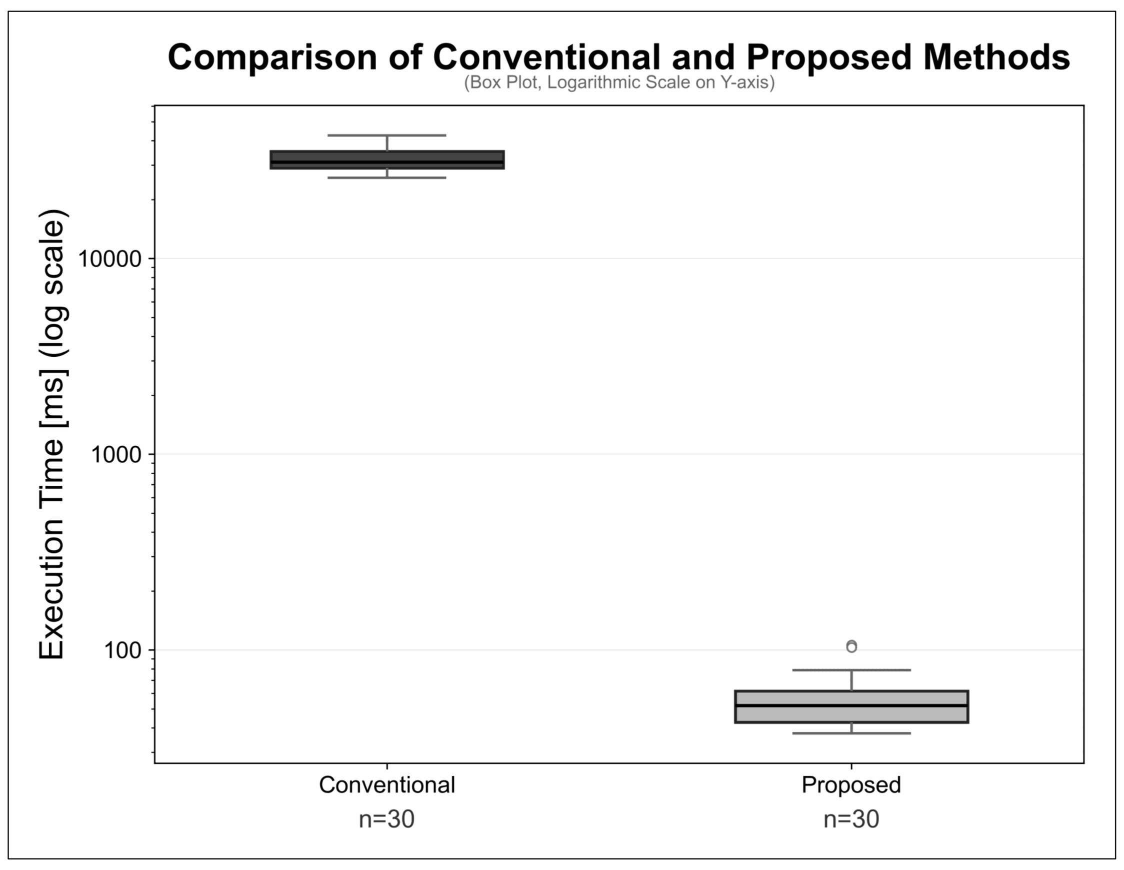

- Tests with representative quantum measurement datasets showed reduced end-to-end latency, real-time dashboard updates with sub second refresh rates, and consistent multi-user performance across concurrent client sessions.

Conclusion

- We developed the Quantum Experiment Data Framework (QEDF) to enable low latency, high throughput data visualization and analysis for quantum control systems. The key contribution of QEDF is a unified, low latency data path that combines object-based storage, in-memory analytics, and zero-copy software integration for quantum control workloads.
- The foundations are being laid for the integration of Remote Direct Memory Access (RDMA) to enable inter-host zero-copy data transfer, through which we plan to achieve microsecond-level latency (compared to tens of microseconds with TCP) and line-rate throughput exceeding 100 Gb/s. Future work will incorporate machine learning and reinforcement learning techniques for predictive modeling and adaptive experiment optimization.

KEYSIGHT

Kenji Kikuchi
Kenji.kikuchi@keysight.com