A Two-level Control Framework for Quantum Networks

Se-young Yu*†, Elia Perego+†, Justin Phillips+†, You-Wei Cheah*†, Prathwiraj Umesh$,
Guangqi Gao+, Jiarui Liu+, Ezra Kissel*, Michael Bregar$, Ke Sun+, Qiming Wu+, Raju Valivarthi$
Erhan Saglamyurek*, Wenji Wu*, Maria Spiropulu$, Hartmut Häffner+, Inder Monga*
* *Lawrence Berkeley National Laboratory*, + *University of California, Berkeley*, $ *California Institute of Technology*

***Abstract*—Quantum network control is a major research area of the QUANT-NET project. We strive to build a quantum network control plane to orchestrate and manage all the physical-layer technologies, and to explore what a quantum network control plane should look like in the future, so as to automate high-rate and high-fidelity entanglement generation, distribution, and storage in an efficient, reliable, and cost-effective way. To these ends, we have designed a two-level control framework for quantum networks. Within such a framework, a two-level scheduler has been implemented to support synchronous time slot scheduling, network-wide non-real-time control, and node-wide real-time control. This two-level control framework and the scheduler are being deployed and evaluated in the QUANT-NET testbed. Enabled by this two-level control framework and the scheduler, several basic quantum network operations have been automated in the testbed, which include automated quantum node calibration and on-demand single photon generation. In this paper, we present this two-level control framework, the scheduler, and the automated quantum network operations that have been enabled in the QUANT-NET testbed.**

***Index Terms*—quantum network architecture, quantum network control, quantum network real-time control, trapped-ion quantum networking, quantum network calibration**

I. INTRODUCTION

Quantum networks are envisioned to achieve novel capabilities that are provably impossible using classical networks and be transformative to science, economy, and national security [1]. These novel capabilities range from cryptography [2], sensing and metrology [3], distributed systems [4], [5], to secure quantum cloud computing [6]. To realize these great potentials, quantum networks are actively researched and developed worldwide across different physical platforms, from solid-state systems such as superconducting circuits [7] and nitrogen vacancy centers in diamond [8], [9], to neutral atoms [10], [11], and also to trapped ions [12]–[15]. Today, quantum networks are in their infancy. Like the Internet, quantum networks are expected to undergo different stages of research and development until they reach a level of practical functionality. Existing quantum network testbeds and prototypes are typically implemented in laboratory settings with limited functionality and managed in a manual and unstructured manner. However, as the size and complexity of quantum networks increase, it is evident that manual control of quantum networks is not a sustainable model. The issues of quantum network control and software automation are becoming increasingly important.

The U.S. Department of Energy (DOE) currently supports Berkeley Lab (LBNL), UC Berkeley (UCB), California Institute of Technology (Caltech), and the University of Innsbruck in constructing a testbed for quantum networking technologies (QUANT-NET) [16]. The goal is to establish a testbed network between LBNL and UCB, where both sites are connected with an entanglement swapping substrate over ∼5 km optical fiber and managed by a quantum network protocol stack. On top of this entanglement swapping substrate the research team will implement the most basic building blocks of distributed quantum computing and quantum repeaters by teleporting a controlled-NOT gate between two far trapped-ion quantum computation nodes [17], [18]. Quantum network control is a major research area of the QUANT-NET project. We strive to build a quantum network control plane to orchestrate and manage all the physical-layer technologies, and to explore what a quantum network control plane should look like in the future, so as to automate high-rate and high-fidelity entanglement generation, distribution, and storage in an efficient, reliable, and cost-effective way. To these ends, we have designed a two-level control framework for quantum networks. Within such a framework, a two-level scheduler has been implemented to support synchronous time slot scheduling, network-wide non-real-time control, and node-wide real-time control. Enabled by this two-level control framework and the associated scheduler, several basic quantum network operations have been automated in the QUANT-NET testbed, which include automated quantum node calibration and on-demand single photon generation. It should be noted however, that this two-level control framework and the associated scheduler are not only designed specifically for the QUANT-NET testbed, but can also be extended to support other types of quantum networks. This paper presents the two-level control framework and the associated scheduler. The rest of the paper is organized as follows. Section 2 presents related work in quantum networking and background on the ARTIQ real-

†These authors contributed equally.

This research is supported by the Quantum Internet to Accelerate Scientific Discovery ASCR Research Program funded through Berkeley Lab FWP FP00013429.

979-8-3315-5736-2/25/$31.00 ©2025 IEEE
DOI 10.1109/QCE65121.2025.00145

time control system. Section 3 discusses the basic needs of a quantum network control plane. Section 4 presents the design and implementation of the two-level control framework and the scheduler. Section 5 discusses the application of this two-level control framework and the scheduler in the QUANT-NET testbed. Lastly, section 6 concludes this paper.

II. Related works and Background

A. Quantum network testbeds

To date, a number of quantum network testbeds have been, or are being built around the world to demonstrate the building blocks of quantum networks [12], [15], [17], [19]–[25]. Each testbed has its own unique design, architecture, and underlying quantum technologies. Most of the research focus by leading academic institutions are focused on the physics and device related aspects of quantum communication, which includes implementation of quantum teleportation, entanglement, and quantum memory with transduction to make the wavelengths compatible with telecom frequencies to enable transport of photons over longer distances using telecom fiber. We do not discuss QKD related quantum communications efforts.

B. Quantum network control

Today, quantum networks are in their infancy, and quantum network control is still a new research area. Until now, demonstrations of quantum network control have been performed using ad-hoc software that is specific to each network's experimental setup, programmed to perform functions directly into low-level control devices using expertise in experimental physics. Wehner's group has been working on "an operating system for executing applications on quantum network nodes" [26], which is focused on a single quantum node control. In contrast, our work focuses on the control of an overall quantum network aka. a system of nodes, and not of the control for an individual node.

C. ARTIQ real-time control system

ARTIQ is a leading-edge control system for quantum information experiment [27]. It features a high-level programming language, based on Python, that helps describe complex experiments. ARTIQ software consists of real-time sections, called "kernels", and non-real-time sections. Kernels are compiled and executed on dedicated FPGA hardware with nanosecond timing resolution and sub-microsecond latency. The non-real-time sections are executed on CPU with no tight timing constraints.

ARTIQ software controls and manages Sinara hardware to handle real-time operations in experiments. The Sinara hardware family consists of several commonly used generic devices of an atomic physics experiment, such as Transistor-transistor logic (TTL), Direct Digital Synthesiser (DDS)-based frequency synthesizer, Digital-to-Anlogue Converter (DAC), and Analogue-to-Digital Converter (ADC) devices.

III. Quantum network control plane goals

In this section, we discuss the basic design goals for a quantum network control plane.

Goal 1: Support high-rate and high-fidelity entanglement generation, distribution, and storage in an efficient, reliable, and cost-effective way. Entanglement is the fundamental building block of a quantum network. Key entanglement-related functions include entanglement generation & distribution, entanglement routing, entanglement swapping, and entanglement distillation [28]. These functions typically involve in-network operations and rely on the maintenance of state inside the network, which is a divergence from the traditional end-to-end principle for the Internet design [29]. Therefore, a control plane that supports such a paradigm is required.

Goal 2: Schedule synchronized operations in a distributed environment. Entanglement generation, distribution, and storage in a quantum network typically involve multiple quantum objects (e.g., qubits) located at physically dispersed locations, and a set of operations that are to be executed on these quantum objects in a distributed manner. Because quantum objects in the Noisy Intermediate-Scale Quantum (NISQ) era are fragile, sensitive to environmental noise, and have limited coherence time, it is necessary to execute the set of operations in a tightly synchronized manner before decoherence destroys the quantum information [30].

Goal 3: Schedule calibrations that are to be executed regularly in a quantum network. Realizing a practical and useful quantum network requires long coherence-time qubits, high-fidelity quantum gate operations, and high-fidelity quantum communications. To enable such capabilities, each quantum network device and channel must be actively monitored, calibrated, and optimized in a periodical manner [17].

Goal 4: Support real-time control in a distributed environment. Timing constraints to perform entanglement-related functions in a quantum network are stringent. In addition, the accurate control and manipulation of quantum objects in a quantum network requires a time resolution of μs scale or less [12], [15], [17]. This means that the control plane requires tight timing control, which imposes hard real-time constraints at the lowest levels.

Goal 5: Ability to discover network resources, build and maintain network topology, make routing decisions, manage and configure quantum network devices, and monitor network status. These are basic functions crucial for any network control, and involves network-wide non-real-time operations. The challenge is therefore to develop an efficient, extensible, robust, and reliable quantum network control architecture to provision these functions.

IV. A two-level control framework for quantum networks

A. Major design choices

We present major design choices in this section to realize the goals described in Section III.

Design choice 1: Logically centralized control. Considering that remote entanglement generation, distribution, and storage in a quantum network typically involves in-network operations and rely on the maintenance of state inside the network, centralization of control state and logic inside a logically centralized control plane will support faster and more efficient entanglement operations. This design choice helps to realize Goal 1.

Design choice 2: Synchronous time slot scheduling (SynTS). With such a scheme, the total duration of quantum network operations is divided into fixed interval time-slots that run continuously. The control plane allocates available time-slots and schedules quantum network operations for specific quantum devices in allocated time-slots . Such a scheduling scheme brings two benefits: (i) supporting synchronized operations in a distributed environment and (ii) allowing calibration tasks to be executed regularly in a quantum network. This design choice helps to realize Goal 2 and 3.

Design choice 3: Real-time control in a distributed environment. A systematic approach is applied to meet tight time constraints: (i) time-critical functions run on dedicated real-time control systems (e.g., ARTIQ real-time control systems); (ii) dedicated control and clock channels are established for real-time communication between relevant devices; (iii) non-real-time functions such as monitoring proper functioning of all components, re-initiating calibration, and data analysis are executed in software in the CPU domain; (iv) non-real-time communications run through TCP/IP channels. This design choice helps to realize Goal 4.

Design choice 4: An open plugin framework. Quantum network control with a plugin framework allows for dynamic insertion of new services and functionalities in the control plane. This design helps to realize Goal 5.

B. A two-level control framework

1) The big picture: Logically, a quantum network's control plane is split into network-wide (global) non-real-time control and node-wide (local) real-time control (see Fig. 1). A two-level scheduler is employed to support such a framework. *(i) At network level.* A logically centralized Quantum Network (QN) server coordinates all activities in the network. This server controls and manages the underlying quantum network nodes (i.e., devices) through a general message bus implementation and broker service. Each class of quantum network nodes has well-defined functions and interfaces. The QN server has a global SynTS scheduler to handle network-wide (global) non-real-time functions such as quantum network topology discovery, quantum network monitoring, and quantum link calibration etc. The server also handles user requests on an event-driven basis. When serving a user request, the server performs alternating cycles between calibration and operational modes on the quantum network so that sustained high-fidelity is achieved. *(ii) At node level.* Each quantum network node is built using a dedicated real-time control system. Each node has a local SynTS scheduler to handle node-wide (local) real-time functions that have tight time constraints, such as qubit control and quantum protocol implementation etc.

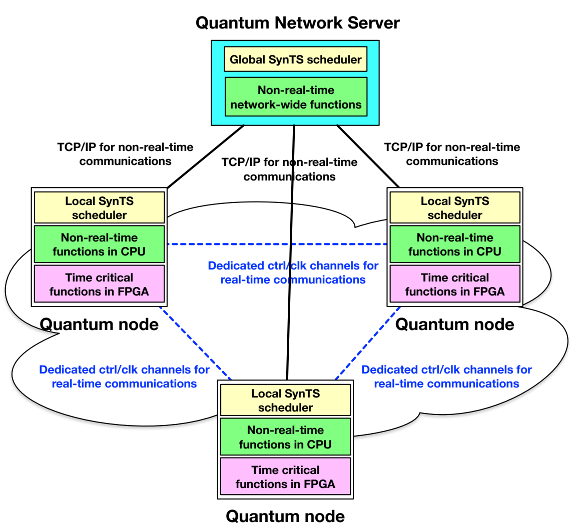

Fig. 1: A two-level control framework for quantum networks.

2) The two-level scheduler: The two-level scheduler adopts synchronous time slot scheduling. The total duration of quantum network operation is divided into fixed interval time-slots that run continuously. Each quantum network function is designed and implemented with a maximum runtime duration. Therefore, a quantum network function is allocated with one or multiple time slots, depending on its maximum runtime duration. Each local scheduler generates and manages a schedule composed of time slots allocated for locally-attached real-time devices, responsible for the execution of quantum network functions at its own node. The global scheduler manages all local schedulers in the network. At request by the global scheduler, multiple local schedulers can allocate available time-slots and schedule quantum network operations for specific quantum devices in allocated time-slots.

3) Support for regular calibration: The control framework supports quantum network calibration. *(i) Quantum channel calibration.* The global scheduler periodically schedules quantum channel calibration on a per-quantum-channel basis to minimize quantum channel loss, reduce background noise, and compensate for polarization drifts for high-fidelity quantum communication. *(ii) Quantum node calibration.* A quantum node typically has one or multiple control parameters that need to be calibrated for proper operations. For example, in a trapped-ion quantum node where a single $^{40}Ca^{+}$ is trapped, a laser is applied to the ion to perform single photon generation. Since the laser frequency drifts over time, the exact frequency shift must be known in order to correct for drifts when running single photon generation. A calibration of the laser frequency must therefore be executed periodically to monitor the frequency drifts and maintain up-to-date frequency corrections. Quantum node calibrations are typically specific to each individual node, and are scheduled periodically by the node's local scheduler.

4) Support for synchronized operations: The control framework supports high-fidelity and high-rate entanglement generation, distribution, and storage in a quantum network. An

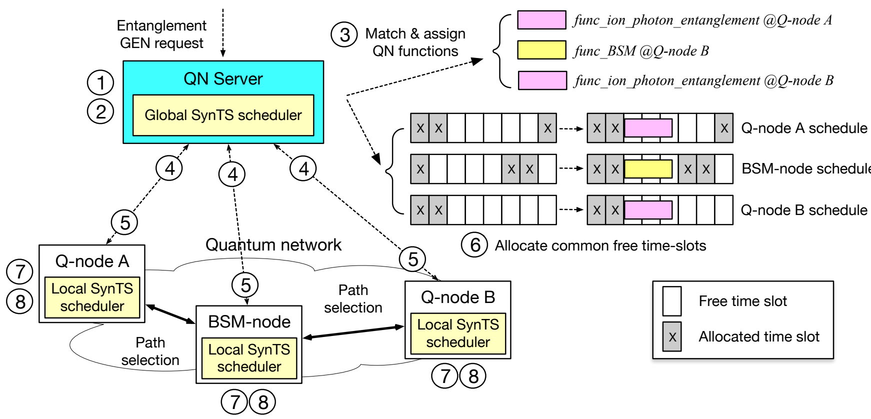

Fig. 2: The two-level scheduler scheme.

ion-ion entanglement establishing process involving Q-nodes and a BSM-node was previously presented in this paper [17]. Here, we illustrate how the control framework supports such a process. In Fig. 2, we assume a trapped-ion based quantum network as illustrated. At time T_0, the QN server receives an entanglement generation request between Q-node A and B with a specified Quality of Service (QoS) requirement (e.g., rate). Afterwards, the control framework works as follows.

1) The QN server performs a request analysis. The request is admitted only if its requirements can be fulfilled.
2) If the request is admitted, the QN server proceeds to call a routing routine to select a path, *Q-node A* $\leftrightarrow$ *BSM-node* $\leftrightarrow$ *Q-node B*, along which to establish entanglement.
3) The QN server matches and assigns functions required for remote entanglement generation for each node along the path. We assume a single function is required at each node, which are *func_ion_photon_entanglement* at Q-node A, *func_BSM* at BSM-node, and *func_ion_photon_entanglement* at Q-node B, respectively. The control framework has full knowledge of these functions, which are well-characterized and pre-configured.
4) The QN server reaches out to Q-node A, BSM-node, and Q-node B for the time-slot allocation status of their local SynTS schedulers.
5) Q-node A, BSM-node, and Q-node B all reply with their local schedules.
6) The QN server proceeds to call the global scheduler to allocate common available time-slots for *func_ion_photon_entanglement* at Q-node A, *func_BSM* at BSM-node, and *func_ion_photon_entanglement* at Q-node B, such that these functions can be scheduled simultaneously at time T_1. Afterwards, the QN server updates Q-node A, BSM-node, and Q-node B with the newly assigned time-slots.
7) Q-node A, BSM-node, and Q-node B accept and confirm the newly assigned time-slots.
8) At time T_1, the local schedulers run *func_ion_photon_entanglement* at Q-node A, *func_BSM* at BSM-node, and *func_ion_photon_entanglement* at Q-node B, respectively.

Therefore, as requested by the global scheduler, multiple local schedulers can schedule synchronized operations for entanglement generation, distribution, and storage.

V. Testbed Evaluation

We implement and evaluate this two-level control framework and the scheduler in the QUANT-NET testbed. As illustrated in Fig. 3, the testbed consists of two trapped-ion quantum nodes, and one BSM-node, which are located in two labs separated by $\sim$5 km in Berkeley, and are connected through telecom fibers. Each of these nodes is built by using a dedicated ARTIQ-based real-time control system. Each node runs a local Synchronous Time Slot (SynTS) scheduler to handle local real-time functions that have tight time constraints. In addition, a Quantum Network Server (QN Server) is deployed to control and manage the underlying trapped-ion quantum nodes and BSM-node, and the server runs a global SynTS scheduler to handle network-wide non-real-time functions.

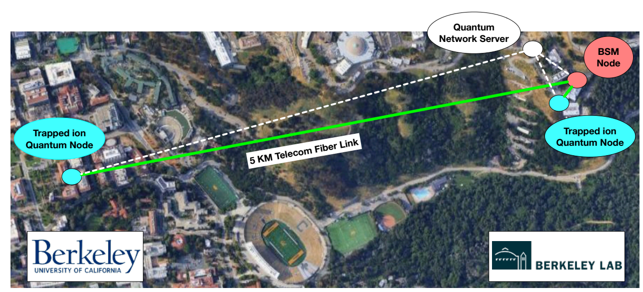

Fig. 3: Overview of the QUANT-NET testbed in Berkeley.

The two trapped-ion quantum nodes are able to generate single photons. Enabled by this two-level control framework and the scheduler, automated quantum node calibrations and on-demand single photon generation (SPG) are supported in our testbed. This system has been lively demonstrated in OFC'25 [31]. In the following subsections, we discuss the design and implementation of the trapped-ion quantum nodes, and how automated quantum node calibration and on-demand single photon generation are implemented in our testbed.

A. The design and implementation of trapped-ion quantum nodes

The underlying technologies in the QUANT-NET testbed are trapped-ion quantum processors. One of the ultimate goals of the QUANT-NET project is to demonstrate remote entanglement between two trapped ion systems at LBNL and UCB using a photonic Bell-State measurement (BSM) scheme. To this end, we built two experiment setups. Recently, we are able to confine and manipulate single ions and perform single photon generation, which is a crucial step towards enabling remote ion-ion entanglement generation.

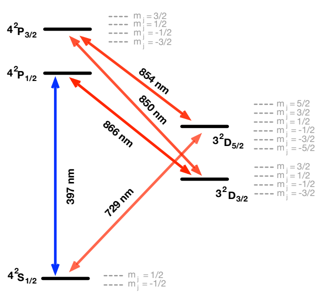

Fig. 4: Level scheme of a ^{40}Ca$^+$ ion. A magnetic field splits each fine structure into several Zeeman sublevels.

1) Single photon generation using ^{40}Ca$^+$ ions: In each ion trap setup, a single ^{40}Ca$^+$ ion is confined in a surface-electrode radiofrequency ion trap by a combination of static (DC) and RF fields. The relevant energy levels of ^{40}Ca$^+$ are shown in Fig. 4. A static magnetic field is applied to split the Zeeman sublevels of the relevant states (e.g., $4S_{1/2}$, $3D_{3/2}, 3D_{5/2}, 4P_{1/2}$, and $4P_{3/2}$). Two types of qubits can be encoded in each ^{40}Ca$^+$ ion, a Zeeman qubit and an optical qubit. The former can be encoded in the D-manifold, with $|0\rangle = |3D_{5/2}, m_j = 5/2\rangle$ and $|1\rangle = |3D_{5/2}, m_j = 3/2\rangle$, and will be used for next stage ion-photon entanglement generation. The latter can be encoded in the S-D optical transition between the states $|0\rangle = |3D_{5/2}, m_j = 5/2\rangle$ and $|1\rangle = |4S_{1/2}, m_j = 1/2\rangle$.

The trapped ion is laser-cooled via the strong $4S_{1/2} \leftrightarrow 4P_{1/2}$ transition, driven by light at 397 nm [32]. The mechanical effect of the laser on the ion motion depends on the frequency detuning of the laser relative to the transition frequency, i.e. if the laser frequency is lower than the $4S_{1/2} \leftrightarrow 4P_{1/2}$ transition, the laser removes motional quanta, thus leading to ion cooling ("red-detuning"). Conversely, if the laser frequency exceeds the transition's resonance frequency, the ion is heated ("blue-detuning").

This transition scatters a high rate of photons (on the order of 10^7 photons per second), allowing ion detection via a photo-multiplier tube (PMT). Since the ion can decay from $4P_{1/2}$ to the $3D_{3/2}$ state with a probability of approximately 6%, application of light at 397 nm would optically pump the population into the metastable state D, trapping the ion in a so-called "dark state" for the entire experiment duration (usually shorter than the level lifetime). To prevent population trapping in the $3D_{3/2}$ state, a repump laser at 866 nm drives any population which decays to $3D_{3/2}$ back to the $4P_{1/2}$ manifold.

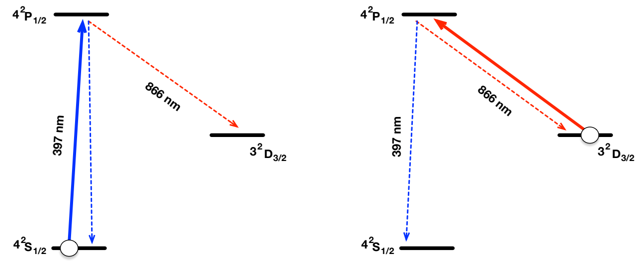

Fig. 5: Atomic level diagram for SPG using ^{40}Ca$^+$ ion.

As illustrated in Fig. 5, single photons are initially produced in the $4S_{1/2}$ manifold by first applying only resonant 397 nm light (blue solid arrow) on the $4S_{1/2} \leftrightarrow 4P_{1/2}$ transition in order to optically pump population to the $3D_{3/2}$ manifold (the "pump" pulse), and subsequently applying only resonant 866 nm light (red solid arrow) on the $3D_{3/2} \leftrightarrow 4P_{1/2}$ transition, causing the ion to scatter 397 nm photons as it returns to the ground manifold (the "drive" pulse). Dashed arrows represent spontaneous decay pathways.

2) Trapped-ion Q-node setup and ARTIQ-based real-time control: An ARTIQ-based trapped-ion Q-node is illustrated in Fig. 6. Specifically, the major components and their interactions in this Q-node are the following: (1) An RF ion trap and an atom oven are enclosed in a vacuum chamber. When fed with electric current, the oven will be heated and will begin generating neutral Ca atoms. (2) RF and DC are applied to trap the ions. (3) CW lasers are used for various controls, such as photoionization, cooling, state readout, and coherent operations like optical pumping, quantum gates, ion-photon generation, etc. Each laser is fine-tuned by an AOM. By controlling each AOM's driving RF signals, the corresponding laser light's amplitude, frequency, and attenuation can be accurately controlled. In addition, the laser used to manipulate the optical qubit is stabilized using an ultra-low expansion (ULE) cavity. (4) Magnetic coils (at the UCB Q-node) or permanent magnets (at the LBNL Q-node) generate magnetic fields to split the Zeeman sublevels of the relevant states. (5) The use of a photo-multiplier tube (PMT) detects the fluorescence of ions for qubit readout and quantum state tomography. And (6) other components such as deflectors, and shutters enables tasks such as shifting the direction of laser beams, blocking/unblocking lasers during ion control operations.

As of the writing of this paper, DC trap voltages, AOMs, and PMTs are controlled by an ARTIQ-based real-time control system, as well as various shutters and switches. The ARTIQ system consists of a control node, and several Sinara hardware devices.

The control node is responsible for node-wide operations that do not require tight timing constraints. Major functions

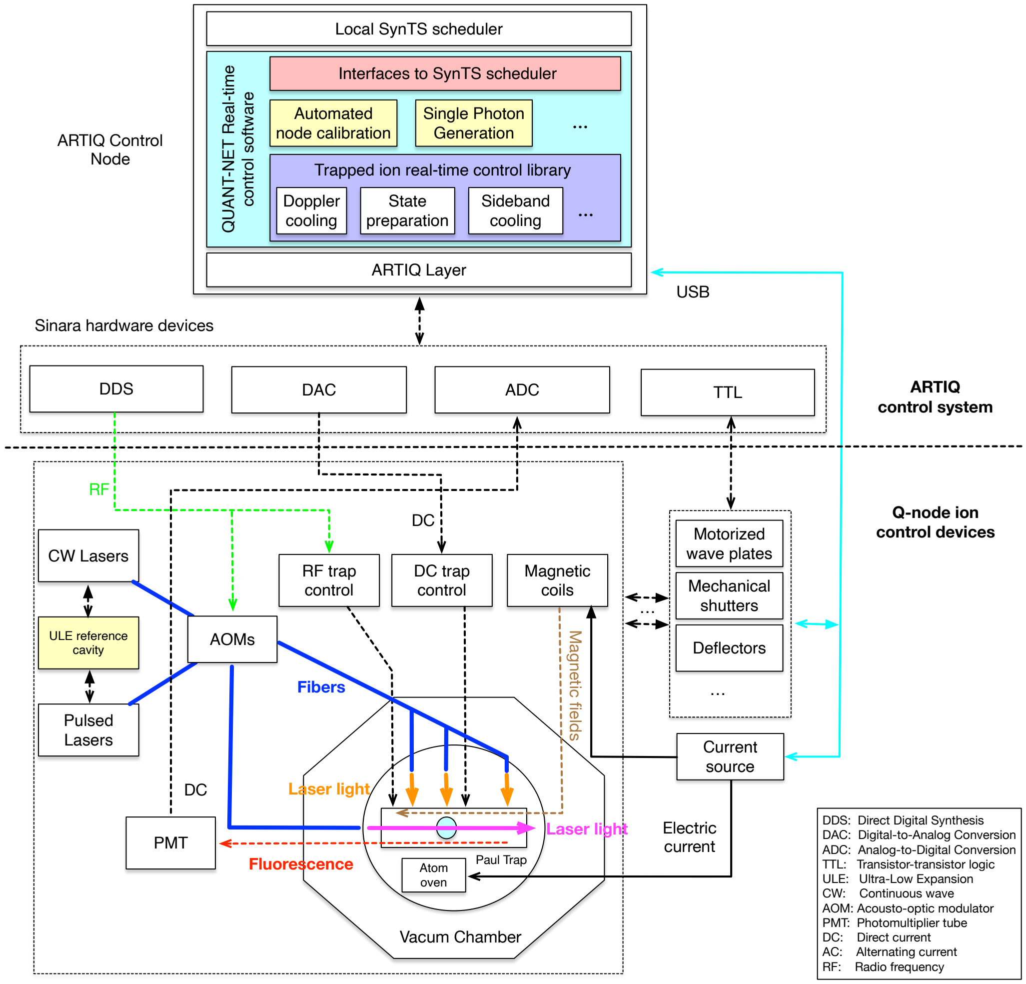

Fig. 6: ARTIQ-based trapped-ion Q-node control.

Real-time control APIs	Purposes
Exp_SPG	Generate single photon generation experiment pulse sequence
Cal_freq_397	Generate 397 nm laser frequency calibration pulse sequence
Cal_power_397	Generate 397 nm laser power calibration pulse sequence
Cal_freq_866	Generate 866 nm laser frequency calibration pulse sequence
Cal_power_866	Generate 866 nm laser power calibration pulse sequence

TABLE I: ARTIQ-based real-time control APIs developed for single photon generation and dependent calibration routines.

include: (1) Preparing and building up the ion control environment such as setting parameters, initialization of devices, and system self-checking etc; (2) Coordinating control functions such as triggering calibration and control sequences; (3) Collecting, recording, and analyzing data; (4) Managing Sinara hardware devices; and (5) Controlling non-real-time devices via USB or network connections.

At the command of the control node, each Sinara hardware device interfaces and interacts with the underlying ion control devices, executing dedicated real-time operations:

- RF signal generation and control. Sinara DDS cards provide sub-Hz frequency resolution, controlled phase steps, and accurate output amplitude control. The generated RF signals are used to drive AOMs to accurately control the amplitude, frequency, and phase of laser light for ion/qubit control.
- DC generation and control. Sinara DAC cards provide high-resolution and high-speed DAC, with a maximum update rate of 2Msps. The generated DCs are used for the DC trapping voltages.
- ADC processing. This is not currently used in the testbed.
- TTL I/O. Used for controlling devices such as mechanical shutters and deflectors, notifying events, and detecting input events, such as PMT counts.

3) The implementation of quantum network operations: As discussed earlier (see subsection V-A1), the quantum physics operations for $^{40}Ca^+$ ions/qubits initialization, control and readout, and single photon generation, correspond to tailored laser-pulse or voltage-control sequence applied to ions and qubits. The amplitude, frequency, phase, timing, and duration of laser light as well as the voltages used to manipulate the internal as well as motional degrees-of-freedom of the trapped ions must be accurately controlled in order to execute the desired operations with a high success rate. To this end, real-time control functions at the node level are purposely designed to program the ARTIQ-control systems to output laser-pulse or voltage-control sequences, collect input events, and record data. These functions realize the quantum physics operations required for trapped-ion quantum networking.

At each Q-node, real-time control functions are implemented by QUANT-NET's own real-time control software framework, which interfaces with the underlying ARTIQ system. The software framework is designed in a way such that fundamental building blocks in trapped-ion quantum networking can be developed in the ARTIQ framework, and that these building blocks are modular, composable, and flexible, providing the base for building more complex operations such as automated node calibration and single photon generation. Table I lists the ARTIQ real-time control functions that we developed for single photon generation and the required calibrations.

To execute a quantum network operation, the QN server typically prepares the related quantum nodes (e.g., Q-node and BSM-node) in the ready state with resources assigned and the related real-time control functions scheduled in a distributed manner. When the real-time control functions run, pulse sequences will be generated for underlying real-time quantum control to achieve the desired operations.

B. Automated quantum node calibration in the testbed

As earlier described (see subsection V-A1), the experimental sequence for generating single photons depends strongly on the power and frequency of the 397 nm and 866 nm lasers, which can be predetermined by using some experimental sequences as calibration routines.

In general, during the calibration of a particular parameter, a number of experiments are repeated while varying the value of that parameter by scanning through a preset range, and keeping all other parameters fixed. By analyzing the results as a function of the scanned parameter, the optimal value can be found and used for subsequent experiments or sequences, such as the single photon generation (SPG) experiment.

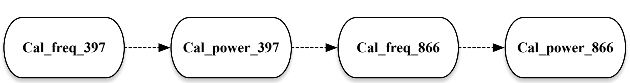

Fig. 7: Calibration routines for frequency and power of the 397 and 866 nm lasers organized as a Directed-Acyclic Graph (DAG).

A set of calibration sequences are organized using a directed acyclic graph (DAG) as shown in Fig. 7 (order of calibrations are performed from left to right), and invoked by the local SynTS scheduler at configurable time intervals, as scheduled by agents running locally on each quantum node. For the SPG experiment described, the full set of laser power and frequency calibrations for the 397 nm and 866 nm frequency and power was performed every 15 minutes.

Fig. 8 presents the 397 nm laser's AOM frequency calibration at the LBNL Q-node over an 8-hour period, within which the SPG data presented in Fig. 12 was taken. The vertical axis reports the difference from the initial calibration value $f_{397}(0)$ as a percentage of $f_{397}(0)$. For this data, $f_{397}(0) = 222.15$ MHz. The apparent discretization of the calibrated frequency is due to the resolution of the calibration scan, which spanned a 40 MHz range in 50 steps.

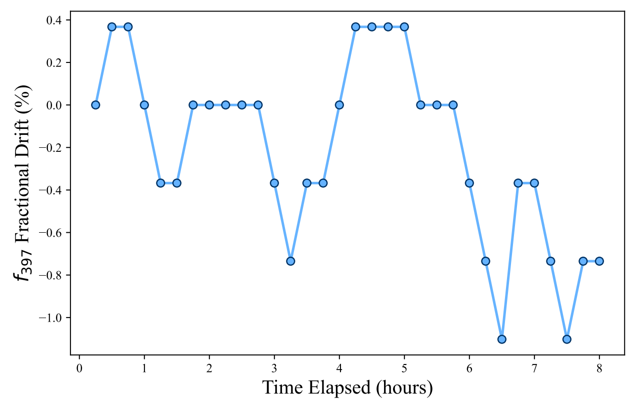

Fig. 8: Tracking of the 397 nm laser's calibrated AOM frequency f_{397} over a period of 8 hours, with periodic frequency calibration invoked by the scheduler every 15 minutes.

At the Q-node, when a calibration routine is invoked, the routine will read in stored parameters, and perform scans for the frequency or power that is to be calibrated. For example, when a frequency calibration routine for the 397 nm laser is executed, the scan for the frequency in question is scanned based on a scanning range with predetermined step sizes. Frequencies and powers for other sequences, such as ion storage and Doppler cooling, are read in from the parameter storage, such as those for 866 nm and 854 nm lasers. For each step of the 397 nm laser frequency scan, PMT photon counts are captured and used as an evaluation metric to determine the optimal frequency to be used for dependent sequences or experiments, and stored back in to the Q-node's local parameter storage.

The base experimental sequence for these calibrations involves collecting ion fluorescence for a set duration (in our case 1 ms), with 397 nm and 866 nm lasers active. Frequency scans are effectively spectroscopic experiments, measuring the transition lines ($S_{1/2} \leftrightarrow P_{1/2}$ and $D_{3/2} \leftrightarrow P_{1/2}$). Fig. 9 shows a typical 397 nm frequency calibration result, with totals counts plotted against laser frequency.

As the frequency increases, fluorescence rises until reaching a maximum at the atomic transition frequency. Theoretical calculations predict optimal Doppler cooling when the detuning equals $\Gamma/2$, where Γ is the natural linewidth of the transition

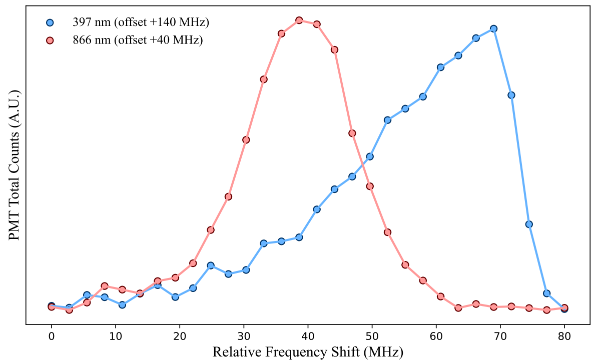

Fig. 9: Example of laser frequency calibration, taken from the UCB quantum node. An offset (shown in the legend) has been applied to the frequencies to combine the 397 nm and 866 nm data in a single plot. For the 397 nm laser, the transition from red-detuning (cooling) to blue-detuning (heating) is evident from the sharp drop in counts around 70 MHz.

[32]. An automated procedure determines the best frequency as the one yielding the midpoint between the peak fluorescence and the light background level. At the UCB Q-node, this value is saved for use in both Doppler cooling and state readout 397 nm controls. At the LBNL Q-node, this value is used for Doppler cooling, while the resonant frequency (with a 500 kHz red detuning) is used for state readout. At both Q-nodes, the resonant frequency is used for SPG attempts.

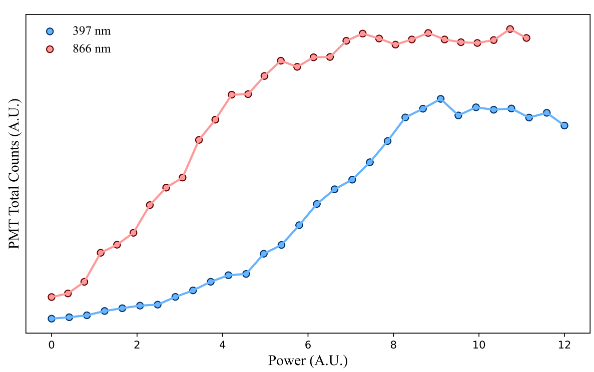

Fig. 10: Example of laser power calibration. In the lab, attenuation is scanned rather than the power directly; however the plot displays power for improved readability.

Fig. 10 reports a typical 397 nm power scan, where total PMT counts are plotted against the power attenuation at the optimal frequency. The automated procedure saves two key attenuation values: one giving the maximum photon count for readout and SPG attempts, and the other achieving roughly $1/3$ of the maximum photon counts (accounting for background) for Doppler cooling and ion storage.

For the repump laser, analogous sequences are run, covering the entire transition line. The ideal repump frequency and power settings for both Doppler cooling and state readout are those giving the highest photon counts.

These calibrations are autonomous, and can be launched either sequentially or individually. Results are automatically evaluated, tested with custom checks, stored, and shared among all the experiments running on that node.

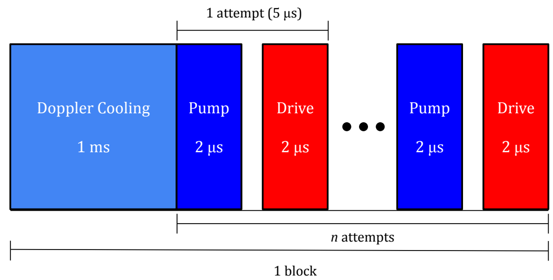

Fig. 11: Single photon generation experiment pulse sequence.

C. On-demand single photon generation in the testbed

Single photon generation requests can be submitted by users on-demand via the QN server. Next, the global SynTS scheduler of the QN server prepares the related quantum nodes (LBNL Q-node, UCB Q-node, or both) in the ready state. At the command of the global scheduler, each relevant local SynTS scheduler schedules and invokes *Exp_SPG - ARTIQ real-time function for single photon generation*. The pulse sequence for the SPG experiment is depicted in Fig.11. SPG attempts are composed of a pump pulse followed by a drive pulse, each lasting 2 µs, separated by a 1 µs delay. Attempts are further organized into blocks, each of which consists of 1 ms of Doppler cooling before n back-to-back attempts. For our experiments presented in this paper, we used 1000 blocks, each with $n = 5000$. A total of 17 experiments were executed intermittently over the course of a 15 hour period, for a total of $17 \times 10^3 \times 5 \times 10^3 = 85 \times 10^6$ distinct SPG attempts. During this period, the 397 and 866 nm lasers were automatically calibrated at 15-minute intervals, facilitated by the scheduler.

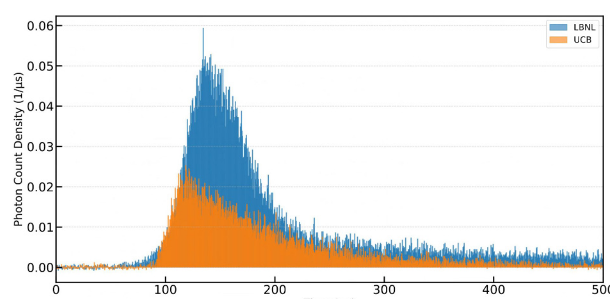

Fig. 12: Background-subtracted photon count densities for both LBNL (blue) and UCB (orange) nodes performing SPG intermittently over a 15 hour period.

Fig. 12 presents data from the SPG collection attempts that were performed intermittently over a 15 hour period of continuous and autonomous calibration, where manual SPG were requested by users and performed simultaneously at both LBNL and UCB quantum nodes. Some manual calibrations were also performed during the data acquisition period. Following the drive pulse, a high numerical aperture objective lens (NA = 0.6) directs scattered photons through free space towards a PMT, which sends voltage pulses to a pulse counter that records the relative time difference between a trigger pulse (ARTIQ triggers this at the same time that the 866 nm pulse begins), and the arrival of a pulse from the PMT.

At each Q-node, we record a histogram of the difference between arrival times of photon pulses and trigger pulses. These histograms are background subtracted and normalized such that integrating the background-subtracted histogram yields the cumulative collection efficiency of a single SPG attempt. The time axes are arbitrarily shifted such that the wave packet onsets roughly overlap. Currently, the implemented scheduler does not guarantee that SPG attempts are made at each Q-node such that the photon wavepackets, which are of order 100 ns long, overlap temporally. Due to a variety of experimental factors, the performance of the LBNL and UCB quantum nodes is not identical. The SPG efficiencies for data described in Fig. 12 are 0.53% and 0.22% for the LBNL and UCB Q-nodes, respectively. The known loss mechanisms for the Q-nodes (with transmissions in parentheses) include: solid-angle coverage of the objective (10%), polarization filtering caused by orientation of collection axis with respect to magnetic field (67%), transmission through a pellicle beam splitter (50% at UCB, 70% at LBNL), transmission through a bandpass filter (80%), and the efficiency of the PMT at 397 nm (20%), yielding an upper limit on collection efficiency of 0.54% at UCB and 0.75% at LBNL. The signal-to-noise ratio (SNR) is computed for the full-width half-max of the data in Fig. 12 to be 56.2 for the LBNL node and 32.6 for the UCB node.

VI. Conclusion

This paper presents a two-level control framework for quantum networks. We discuss design goals and choices influencing our framework, and we implement a two-level scheduler to support synchronous time slot scheduling, network-wide non-real-time control, and node-wide real-time control. Preliminary results from the use of this two-level scheduler is discussed in this paper and showcases the success of this implementation in managing single-photon generation as well as auto-calibration. Further evaluation is being done in the QUANT-NET testbed as new experiments are incorporated. This framework will be released to the community upon completion of internal testing.

Acknowledgements

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

References

[1] Koji Azuma, Sophia E Economou, David Elkouss, Paul Hilaire, Liang Jiang, Hoi-Kwong Lo, and Ilan Tzitrin. Quantum repeaters: From quantum networks to the quantum internet. *Reviews of Modern Physics*, 95(4):045006, 2023.

[2] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:7–11, 2014.

[3] Peter Komar, Eric M Kessler, Michael Bishof, Liang Jiang, Anders S Sørensen, Jun Ye, and Mikhail D Lukin. A quantum network of clocks. *Nature Physics*, 10(8):582–587, 2014.

[4] Michael Ben-Or and Avinatan Hassidim. Fast quantum Byzantine agreement. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 481–485, 2005.

[5] Vasil S Denchev and Gopal Pandurangan. Distributed quantum computing: A new frontier in distributed systems or science fiction? *ACM SIGACT News*, 39(3):77–95, 2008.

[6] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 517–526. IEEE, 2009.

[7] Youpeng Zhong, Hung-Shen Chang, Audrey Bienfait, Étienne Dumur, Ming-Han Chou, Christopher R Conner, Joel Grebel, Rhys G Povey, Haoxiong Yan, David I Schuster, et al. Deterministic multi-qubit entanglement in a quantum network. *Nature*, 590(7847):571–575, 2021.

[8] Peter C Humphreys, Norbert Kalb, Jaco PJ Morits, Raymond N Schouten, Raymond FL Vermeulen, Daniel J Twitchen, Matthew Markham, and Ronald Hanson. Deterministic delivery of remote entanglement on a quantum network. *Nature*, 558(7709):268–273, 2018.

[9] SLN Hermans, Matteo Pompili, HKC Beukers, Simon Baier, Johannes Borregaard, and Ronald Hanson. Qubit teleportation between non-neighbouring nodes in a quantum network. *Nature*, 605(7911):663–668, 2022.

[10] Severin Daiss, Stefan Langenfeld, Stephan Welte, Emanuele Distante, Philip Thomas, Lukas Hartung, Olivier Morin, and Gerhard Rempe. A quantum-logic gate between distant quantum-network modules. *Science*, 371(6529):614–617, 2021.

[11] Manuel Brekenfeld, Dominik Niemietz, Joseph Dale Christesen, and Gerhard Rempe. A quantum network node with crossed optical fibre cavities. *Nature Physics*, 16(6):647–651, 2020.

[12] V Krutyanskiy, M Galli, V Krcmarsky, S Baier, DA Fioretto, Y Pu, A Mazloom, P Sekatski, M Canteri, M Teller, et al. Entanglement of trapped-ion qubits separated by 230 meters. *Physical Review Letters*, 130(5):050803, 2023.

[13] LJ Stephenson, DP Nadlinger, BC Nichol, S An, P Drmota, TG Ballance, K Thirumalai, JF Goodwin, DM Lucas, and CJ Ballance. High-rate, high-fidelity entanglement of qubits across an elementary quantum network. *Physical review letters*, 124(11):110501, 2020.

[14] Matthias Bock, Pascal Eich, Stephan Kucera, Matthias Kreis, Andreas Lenhard, Christoph Becher, and Jürgen Eschner. High-fidelity entanglement between a trapped ion and a telecom photon via quantum frequency conversion. *Nature communications*, 9(1):1998, 2018.

[15] Victor Krutyanskiy, Marco Canteri, Martin Meraner, James Bate, Vojtech Krcmarsky, Josef Schupp, Nicolas Sangouard, and Ben P Lanyon. Telecom-wavelength quantum repeater node based on a trapped-ion processor. *Physical Review Letters*, 130(21):213601, 2023.

[16] Berkeley Quantum Network Testbed — quantnet.lbl.gov. https://quantnet.lbl.gov. [Accessed 12-04-2025].

[17] Inder Monga, Erhan Saglamyurek, Ezra Kissel, Hartmut Haffner, and Wenji Wu. QUANT-NET: A testbed for quantum networking research over deployed fiber. In *Proceedings of the 1st Workshop on Quantum Networks and Distributed Quantum Computing*, pages 31–37, 2023.

[18] Damian Schon, Prathwirag Umesh, You-Wei Cheah, Se-Young Yu, Ezra Kissel, Raju Valivarthi, Erhan Saglamyurek, Lavanya Ramakrishnan, Wenji Wu, Alp Sipahigil, Maria Spiropulu, Hartmut Haffner, and Inder Monga. The quant-net testbed development and preliminary results. *submitted to IEEE International Conference on Quantumm Computing and Engineering - QCE24*, 2024.

[19] Joaquin Chung, Gregory Kanter, Nikolai Lauk, Raju Valivarthi, Wenji Wu, Russell R Ceballos, Cristián Peña, Neil Sinclair, Jordan Thomas, Si Xie, et al. Illinois Express Quantum Network (IEQNET): metropolitan-scale experimental quantum networking over deployed optical fiber. In *Quantum information science, sensing, and computation XIII*, volume 11726, page 1172602. SPIE, 2021.

[20] Chaohan Cui, William Horrocks, Lauren McCaffrey, Vijay Nafria, Ivan B Djordjevic, and Zheshen Zhang. Entanglement distribution and routing in a multi-node quantum network testbed. In *2021 Conference on Lasers and Electro-Optics (CLEO)*, pages 1–2. IEEE, 2021.

[21] Eric Bersin, Matthew Grein, Madison Sutula, Ryan Murphy, Yan Qi Huan, Mark Stevens, Aziza Suleymanzade, Catherine Lee, Ralf Riedinger, David J Starling, et al. Development of a Boston-area 50-km fiber quantum network testbed. *Physical Review Applied*, 21(1):014024, 2024.

[22] Can M Knaut, Aziza Suleymanzade, Y-C Wei, Daniel R Assumpcao, P-J Stas, Yan Qi Huan, Bartholomeus Machielse, Erik N Knall, Madison Sutula, Gefen Baranes, et al. Entanglement of nanophotonic quantum memory nodes in a telecom network. *Nature*, 629(8012):573–578, 2024.

[23] Momtchil Peev, Christoph Pacher, Romain Alléaume, Claudio Barreiro, Jan Bouda, Winfried Boxleitner, Thierry Debuisschert, Eleni Diamanti, Mehrdad Dianati, James F Dynes, et al. The SECOQC quantum key distribution network in Vienna. *New journal of physics*, 11(7):075001, 2009.

[24] PAN Jianwei. Progress of the quantum experiment science satellite (QUESS) Micius project. *Chinese Journal of Space Science*, 38(5):604–609, 2018.

[25] Yu-Ao Chen, Qiang Zhang, Teng-Yun Chen, Wen-Qi Cai, Sheng-Kai Liao, Jun Zhang, Kai Chen, Juan Yin, Ji-Gang Ren, Zhu Chen, et al. An integrated space-to-ground quantum communication network over 4,600 kilometres. *Nature*, 589(7841):214–219, 2021.

[26] C Delle Donne, M Iuliano, B van der Vecht, GM Ferreira, H Jirovská, TJW van der Steenhoven, A Dahlberg, M Skrzypczyk, D Fioretto, M Teller, et al. An operating system for executing applications on quantum network nodes. *Nature*, 639(8054):321–328, 2025.

[27] M-labs. Sinara open hardware project. https://m-labs.hk/experiment-control/artiq/.

[28] Stephanie Wehner, David Elkouss, and Ronald Hanson. Quantum internet: A vision for the road ahead. *Science*, 362(6412):eaam9288, 2018.

[29] Jerome H Saltzer, David P Reed, and David D Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems (TOCS)*, 2(4):277–288, 1984.

[30] Kenneth R Brown, John Chiaverini, Jeremy M Sage, and Hartmut Häffner. Materials challenges for trapped-ion quantum computers. *Nature Reviews Materials*, 6(10):892–905, 2021.

[31] Optica. Optical Fiber Communications Conference and Exhibition. https://www.ofcconference.org/en-us/home/.

[32] H. J. Metcalf and P. van der Straten. Laser cooling and trapping of atoms. *Journal of the Optical Society of America B Optical Physics*, 20(5):887–908, May 2003.