第３部　特集3 Quantum Internet
AQUA (Advancing Quantum Architecture) Annual Report 2021

Rodney Van Meter

Abstract

The AQUA (Advancing Quantum Architecture) working group extended the breadth, depth, importance, and reach of its research in 2020. We now divide our activities into four major areas: Quantum Community, Quantum Education, Quantum Computing, and Quantum Internet. WIDE members contributed to operation of the IBM Qiskit Challenge; provided our online course (MOOC) for free during the pandemic and acquired a large grant for quantum online curriculum development in Japanese from the Japanese government's Q-LEAP office; continued research on algorithms, compilers and programming tools for noisy, intermediate scale quantum (NISQ) computers; and moved toward bringing the Quantum Internet to fruition.

Our primary Quantum Internet activities included continuing leadership of the Quantum Internet Research Group (QIRG), which is part of the Internet Research Task Force (IRTF); the open source release of QuISP, our Quantum Internet Simulation Package; and participation in the formal establishment of the Quantum Internet Task Force (QITF) as a research consortium of the Keio Research Institute at SFC, providing a funding organ for work toward a Quantum Internet testbed we are planning for the Kanto area. This year's report focuses primarily on these Quantum Internet activities.

第１章　Introduction

WIDE, through the AQUA working group, is well positioned to participate in and help guide the field in the exciting area of quantum computing and quantum networking, particularly as it moves from theoretical papers and small laboratory technology demonstrations toward actual systems.

This report focuses on the Quantum Internet activities of AQUA over the last year. We begin with a short introduction to the Quantum Internet itself (Sec.2), including a very brief overview of some of AQUA's contributions and our key architectural ideas. AQUA members informally created the Quantum Internet Task Force in 2019, and turned it into a formal research consortium in 2020 (Sec.3). We co-chaired the IRTF's Quantum Internet Research Group, and contributed as authors to an Internet Draft that is approaching publication as an RFC (Sec.4). Our most important technical contribution of the year was the open-source
release of the Quantum Internet Simulation Package (Sec.5).

AQUA activities extend well beyond Quantum Internet. Our work on quantum computing, quantum education and quantum community are also briefly addressed in this report. This report closes with a summary of major publications over the year. An introduction to the AQUA group and work areas is included as Appendix A. A brief introduction to the field of quantum information is included as Appendix B.

第２章　The Quantum Internet

2.1　The Value of the Quantum Internet

The Quantum Internet is coming, and it will make your world more secure, increase the sensitivity and accuracy of sensors, and help quantum computers share data and cooperate to solve problems -- not to mention make the world much more interesting [20, 41, 36]. Testbed networks of increasing flexibility and functionality are being built or are on the drawing boards in the Netherlands, the U.S., Japan, and throughout the world.

Today's Internet, which we can call the classical Internet, connects computers that are in distant locations, and share data and/or computational services. That data may be located in a particular place either because that is where the data is originally collected (e.g., environmental measurements), or where the data curator has provided facilities. Likewise, although computational services are increasingly virtual, part of the rationale driving the development of the original ARPANET was access to mainframes (and later supercomputers), expensive computational resources that only make economic sense when shared among a large group of users.

The same rationale applies to quantum computers. Combining the power of multiple quantum computers and sharing specialized capabilities are especially attractive. However, there is a catch: quantum computers cannot connect to each other using the classical Internet. They require a specialized *Quantum Internet* in order to talk to each other and share quantum data.

Technically, the purpose of a true Quantum Internet is to create quantum entanglement, which Einstein famously called *spukhafte Fernwirkungen*, usually translated as "spooky action at a distance". If you have heard about the weirdness of quantum mechanics, you likely have heard of Schrodinger’s cat, which captures two important quantum phenomena: superposition and measurement. Entanglement is the third important phenomenon. Two or more quanta, which might be individual electrons or photons or any of dozens of natural and artificial systems that display similar characteristics, can be entangled with each other by making them interact in special ways. That entanglement remains even if they are moved far apart. When one member of the entangled pair is measured, the other is instantaneously affected, regardless of distance.

A common misconception is that entanglement allows us to communicate faster than the speed of light, violating Einstein’s principle of relativity, but that’s not possible. Consider a researcher in Japan sharing entanglement with a researcher in the Netherlands. When the Japanese researcher examines, or measures, her entangled particle, something happens to the state the particles share, but it’s immediately clear what has happened. In order for the American researcher to interpret the meaning of the change to the entangled particle he holds, she must send some classical information from Japan to the U.S. That information, of course, travels no faster than the speed of light, saving the day for relativity.

The very first quantum networks were developed in the early to mid 2000s [13, 28]. These were highly specialized networks with a single purpose: to securely create shared, secret random numbers between (classical) computers, numbers that can be used to secure communications. Those numbers are used as a "key" for encryption, hence the term quantum key distribution [4, 12, 6]. A true Quantum Internet can do this, too, but is qualitatively different and moves far beyond this single function.

Perhaps the most exotic use of quantum entanglement is to improve the sensitivity and precision of sensors. High-precision sensing forms the basis of our GPS system, geological surveys such as oil exploration, and high-resolution radio telescopes using many radio dishes. Intriguingly, LIGO, the famous gravitational wave observatory, uses special quantum states created using lasers already; it’s an open question whether the Quantum Internet can help by creating entanglement between the LIGO facilities in Washington and Louisiana.

The biggest payoff, though, may come via the most obvious use. A Quantum Internet will allow us to connect multiple quantum computers together, first within a single laboratory and eventually around the planet. Entangling distant quantum computers lets us attack larger programs than any single computer. It also lets us perform blind, or secure, quantum computation, safer than classical cloud computing.

Today, quantum computers are moving out of the lab and becoming a competitive industry, aiming to solve problems in material science, chemistry, finance, and AI. Standalone machines are always less useful than connected ones. Tomorrow, we will connect those computers together, setting in motion the next phase of the Second Quantum Revolution.

2.2 Quantum Repeaters

The quantum equivalent of the Internet's routers are called *quantum repeaters* [11]. Despite the name, they are not repeaters in the classical sense, but rather cooperate over quantum channels to create end-to-end quantum entanglement useful for applications. A quantum repeater's work consists of four tasks:

1. generation of base-level entanglement with its nearest neighbors, using fiber or free space links;
2. managing errors (via error detection or error correction);
3. coupling the single-hop entanglement into longer-distance entanglement, e.g. via a method known as entanglement swapping [18]; and
4. participating in management of the network.

Experimental and theoretical physicists have worked hard on the physical layer mechanisms for generating entanglement, and theorists have studied means of managing errors while building entanglement along a chain of repeaters, but little energy has been invested so far in designing *networks* of quantum repeaters. A network diagram showing the elements of a network with Quantum Byzantine Agreement as an application is shown in Fig.1.

2.3 Architectural Work by AQUA

Building on the work done over the last ten years, the work done by AQUA has completed our list of provisional technical proposals for almost every aspect of creating a true entanglement-based Quantum Internet above the physical layer.

Our work builds on years of prior work quantum network and quantum internetwork architecture [36]. Almost every major network design issue has now been touched upon, though the design choices will evolve continually even after initial deployments. Our overall goal, therefore, is to build a future-proof, flexible architecture that allows indefinite innovation. A few key results are highlighted here.

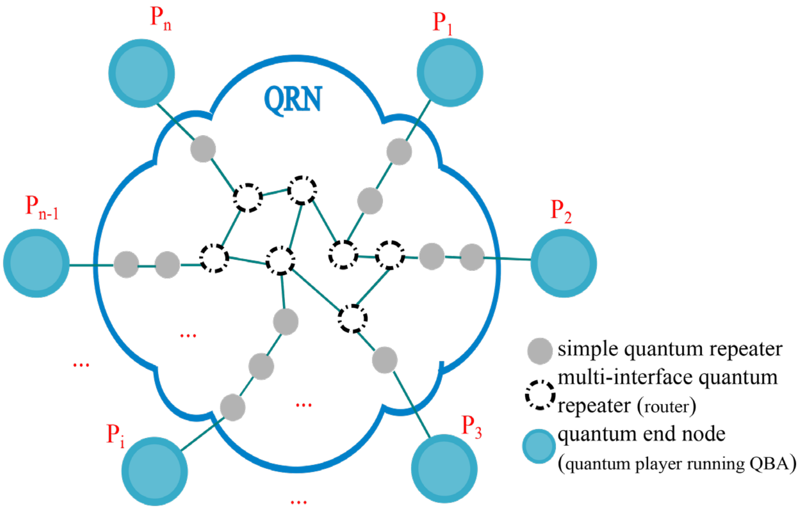

Fig 1 Required elements in Quantum Repeater Networks (QRN) for running scalable quantum distributed applications such as Quantum Byzantine Agreement. From [32].

Internetwork architecture: The Quantum Recursive Network Architecture (QRNA), building on the classical Recursive Network Architecture [34], will scale management of connections, work across heterogeneous networks [26], and retain autonomy and privacy of network operations [40].

Routing: A modified form of Dijkstra's algorithm will work for intermediate-sized networks [38].

Multiplexing: Assignment of resources to quantum network connections will likely follow circuit switching principles due to the need to continually execute distributed operations, but moment-by-moment allocation of entangled states may work acceptably even when performed in Internet-style best-effort, first-come, first-served fashion [1].

2.4 Our Key Architectural Ideas

The core ideas in our architectural proposal are worth examining more closely. These ideas represent over a decade of thought and simulation, but of course to date do not represent real operational experience.

The first issue is to define the *semantics* of the network: what service does the network provide? We have settled on end-to-end delivery of entangled states, primarily Bell pairs. The network is not forward-and-forget, as in the classical Internet, but instead requires distributed cooperation among the nodes along the path to build entanglement. The semantics at the connection end points can vary according to the application class [39].

As noted above, recursion in the network provides scalability. Our current plan is only a two-level hierarchy, equivalent to the autonomous systems (ASes) of the Internet. Border routers hide the internal topology and technology of a given network, allowing wider-area routing protocols to treat the entire network as a single hop. This dramatically simplifies the creation of RuleSets for a connection as well.

Our original idea for protocols was to build state machines, as is common in network protocols [2, 36]. However, the complex set of tasks involved in purification and entanglement swapping, as well as the complexity introduced by recursion, led us in recent years to develop RuleSets. Each rule in a RuleSet consists of a Condition Clause and an Action Clause. These rules are similar to the match and action clauses in software-defined networking (SDN) [23].

Naturally, we must set up the connection by creating and distributing those RuleSets. A network connection setup is started by the Initiator, and terminates at the Responder. Our current proposed approach is to collect information about the links on an outbound pass from the Initiator. The Responder takes that information and plans the connection, then distributes the RuleSets on a return pass toward the Initiator, after which nodes operate asynchronously within the bounds of the RuleSets to participate in the distributed state creation, as shown in Fig.2.

Perhaps the most important open architectural question is resource allocation and routing. In the Internet, these two issues are separate, but in the original, circuit-switched network they were a combined system, as analog and later digital switch ports had to be allocated. While we

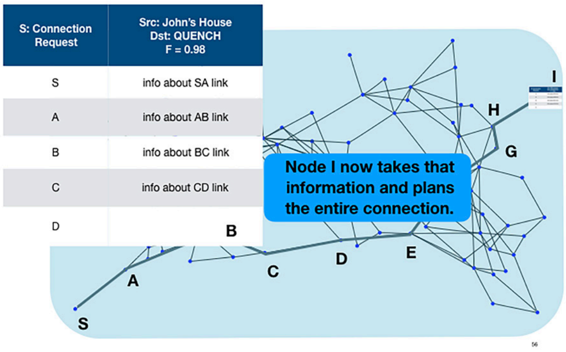

S: Connection Request	Src: John's House Dst: QUENCH F = 0.98
S	info about SA link
A	info about AB link
B	info about BC link
C	info about CD link
D	

Fig 2 Setting up a connection. Node S is the Initiator, and node I is the Responder. Information about individual links is collected hop by hop from Initiator to Responder, then processed at the Responder, after which RuleSets are distributed.

have created a preliminary proposal for routing that builds on shortest path first and conducted simulations that suggest that statistical multiplexing can be effective, these issues still need to be addressed at Internet scale.

第3章　Quantum Internet Task Force

In May of 2019, Shota Nagayama initiated the Quantum Internet Task Force as an organization dedicated to creating a Quantum Internet in Japan. QITF will focus on development of a quantum repeater testbed in the Kanto area, helping to bring the technology out of the physics laboratory and eventually into the commercial market.

In July 2020, QITF was formally recognized as a research consortium underneath the Keio Research Institute at the Shonan Fujisawa Campus, providing a formal status that allows it to accept research funds. Similar to the WIDE Project, QITF will both leverage existing and new projects run by individual members, and attract large-scale funding to support the development of the testbed network. This crucial effort will take place over a number of years and will require the use of dark fiber between locations, as well as custom-building substantial amounts of experimental infrastructure.

A photo of the first meeting is shown in Fig.3. Since that time, QITF has grown to include almost all of the research groups in Japan that are working on quantum repeaters. As of January 202, membership of QITF includes:

- Shota Nagayama, Board member, Representative
- Rikizo Ikuta, Board Member
- Toshihiko Sasaki, Board Member
- Hiroki Takahashi, Board Member
- Tomoyuki Horikiri, Board Member
- Rekishu Yamazaki, Board Member
- Hideo Kosaka, Advisory Board
- Kae Nemoto, Advisory Board
- Rodney Van Meter, Advisory Board
- Takashi Yamamoto, Advisory Board
- Nobuyuki Imoto, Special Advisor
- Takao Aoki, Member
- Masato Koashi, Member
- Takahiko Satoh, Member
- Masahiro Takeoka, Member
- Akihisa Tomita, Member
- Kenji Toyoda, Member
- Yasunobu Nakamura, Member
- Kazuhiro Hayasaka, Member
- Keisuke Fujii, Member

The expertise of this team covers the theoretical physics of individual devices, design of quantum repeaters using those devices, quantum error correction, quantum information theory, applications of the Quantum Internet (notably quantum key distribution), network architecture, and network protocols. The full spectrum of capabilities will ensure the success of the overall project.

第4章　Quantum Internet Research Group (QIRG)

One of our most far-reaching efforts is leadership of and participation in QIRG. QIRG will be a critical forum for not

Fig 3　The first meeting of the Quantum Internet Task Force, May 17, 2019. WIDE members Jun Murai and Takahiko Satoh are present. Shota Nagayama (front) is the founder and leader of QITF.

only reaching consensus on the architecture and eventually protocols for a Quantum Internet, but also serves as an opportunity to educate participants in Internet systems and architecture through the IETF and IRTF. This RG will also utilize the accumulated know-how of the IETF and IRTF in network design and operations, and in that sense is critical. QIRG has begun writing Internet Drafts intended to reach RFC status, described below.

4.1 History and Activities

On March 21, 2018, Rod Van Meter presented a vision for a Quantum Internet Research Group at the Internet Research Task Force (IRTF) Open Meeting held in London at IETF 101. 106 participants signed the bluesheets (attendance list) and additional people followed the irtfopen session online. As the presentation was scheduled on short notice, Rod presented remotely, as shown in Fig.4.

QIRG began when Stephanie Wehner (TU Delft) approached Rod Van Meter about creating such a group. Rod contacted Allison Mankin (IRTF Chair at the time), who was enthusiastic about the idea. Provisional meetings were held at IETFs 103 (115 attendees as tracked by bluesheet signatures), 104, and 106, following which QIRG was approved as a full IRTF Research Group in March 2020. At IETF 104, in addition to the regularly scheduled QIRG meeting, Rod Van Meter and Tracy Northup held a two-hour tutorial on quantum repeaters. Further meetings have been held at IETFs 107, 108 and 109.

As of this writing (Jan. 2021), there are 401 subscribers to the QIRG mailing list, which as with all IRTF lists is open to anyone[1]

Plans for 2021 include an open seminar series as well as meetings during IETF.

4.2 Architectural Principles Draft

Wojciech Kozlowski of TU Delft is leading the writing of an Internet Draft titled, "Architectural Principles for a Quantum Internet"[2]. The I-D describes some of the key ideas in quantum networking, including teleportation, and begins

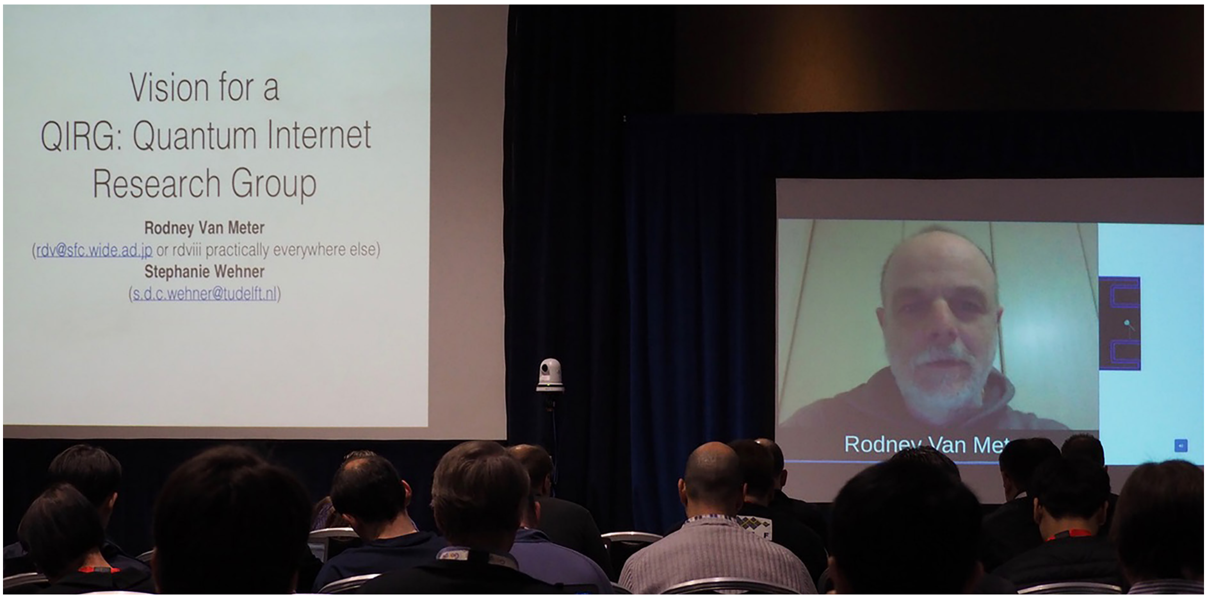

Fig 4 Rod Van Meter proposing the Quantum Internet Research Group at IETF 101 in London, via remote connection from Japan. Photo courtesy of Colin Perkins.

---

[1] WIDE members interested in subscribing can do so at https://www.irtf.org/mailman/listinfo/qirg.

[2] https://datatracker.ietf.org/doc/draft-irtf-qirg-principles/

the process of discussing the elements that will be needed to complete a network architecture and provide quantum services to applications, as well as how those applications will integrate with existing Internet services. Rod Van Meter and Shota Nagayama are coauthors on the draft.

The abstract of the draft is as follows:

> The vision of a quantum internet is to fundamentally enhance Internet technology by enabling quantum communication between any two points on Earth. To achieve this goal, a quantum network stack should be built from the ground up to account for the fundamentally new properties of quantum entanglement. The first realisations of quantum networks are imminent, but there is no practical proposal for how to organise, utilise, and manage such networks. In this memo, we attempt to lay down the framework and introduce some basic architectural principles for a quantum internet. This is intended for general guidance and general interest, but also to provide a foundation for discussion between physicists and network specialists.

4.3 Other Drafts

In addition to the architecture draft, a draft titled "Applications and Use Cases for the Quantum Internet" is in development. Two other drafts have expired, one on the design of a link layer that was also from Delft, where a quantum repeater testbed is in development.

Table 1 Summary of QIRG meeting attendees, as evidenced by bluesheet signatures.

IETF 101 (Mar. 2018)	IRTFopen	106
IETF 103 (Nov. 2018)	meeting	115
IETF 104 (Mar. 2019)	tutorial	157
IETF 104 (Mar. 2019)	meeting	61
IETF 106 (Nov. 2019)	meeting	59
IETF 107 (Mar. 2020)	mtg. (online)	N/A
IETF 108 (Jul. 2020)	mtg. (online)	86
IETF 109 (Nov. 2020)	mtg. (online)	73

The other expired draft is our own draft on connection setup in a quantum repeater network, formalizing the process described in Sec.2.4. We expect to revive that draft in 2021 as our development advances.

第５章 QuISP: Quantum Internet Simulation Package

QuISP was released as open source on April 5, 2020, and is version controlled and distributed via github[＊3]. It is licensed under the 3-Clause BSD License.

QuISP represents our fourth-generation quantum network simulator. The first generation was written in Matlab, by Thaddeus Ladd. The second generation was written by Thaddeus Ladd and Rod Van Meter in C++ as a standalone program, capable of moderate configurability and several choices with respect to purification protocols [37]. Both of those simulators work on density matrices and were primarily designed to assess performance and tuning parameters for a specific form of physical entanglement generation [21]. The latter simulator was also used for some analysis of a quantum network routing protocol [38]. Our third-generation simulator represented the first known simulation of a quantum repeater *network*, as opposed to a simple chain of repeaters, and was used for the first analysis of multiplexing for repeaters [1].

5.1 QuISP README

This subsection contains the README for the QuISP software package.

The Quantum Internet Simulation Package (QuISP) is an event-driven simulation of quantum repeater networks, which will be the ultimate foundation of the coming Quantum Internet. QuISP's goal is to simulate a full Quantum Internet

---

[＊3] https://github.com/sfc-aqua/quisp

consisting of up to 100 networks of up to 100 nodes each. Its focus is on protocol design and emergent behavior of complex, heterogeneous networks, while keeping the physical layer as realistic as possible.

QuISP is a product of the Advancing Quantum Architecture (AQUA) research group headed by Prof. Rodney Van Meter, at Keio University's Shonan Fujisawa Campus, Fujisawa, Japan. `http://aqua.sfc.wide.ad.jp`

5.1.1 Research questions

A simulator is one or more of three things: a time machine, an X-ray machine, or a telescope.

Research questions we hope to answer:

- Emergent behavior
  - Classical networks exhibit *congestion collapse*; are quantum networks subject to the same thing?
  - Will the dynamics of large networks prevent us from making end-to-end connections under realistic scenarios, even when a naive model suggests it should be possible?
  - Are there other unexpected behaviors in large-scale networks?
- Protocol design
  - Testing of detailed protocol designs to validate correct operation.
  - Are there interactions between the classical and quantum portions of the network?
- Connection architecture and performance prediction
- All three proposed network generations exhibit complex behavior that makes analytic prediction of performance difficult with realistic parameters; simulation, of course, will require the best effort we can make at validation, as well.

5.1.2 Simulation goals

We have a number of long-term goals for the simulator:

- Complex network topologies, including the notion of network boundaries and heterogeneity at the physical and logical levels
- support 1G, 2G and 3G quantum networks, utilizing either purify-and-swap (1G) or quantum error corrected (QEC) (2G and 3G) protocols for managing errors
- Distinct link architectures: memory-to-memory (MM), midpoint interference (MIM), and midpoint source (MSM)
- Internetworking protocols for connecting different types of networks
- Various applications running in complex traffic patterns

Because these protocols can result in hundreds of qubits in a single entangled state, and the entire system may consist of up to a million qubits, simulation at the physical Hamiltonian level or even just above that at the unitary (gate, e.g. CNOT) level is infeasible. We cannot calculate and store full density matrices for such states. Instead, like simulators for large-scale error correction, QuISP operates primarily in the *error basis*, in which we maintain a description of errors the states have incurred rather than the full state. However, unlike QEC simulators, QuISP supports non-Pauli errors, in a limited fashion.

QuISP is almost endlessly configurable; for example, it is possible to set not only different lengths for different links in the network, but also different gate error rates and memory lifetimes on individual qubits. Non-Pauli errors that are at least partially supported in the current release include qubit loss, relaxation to ground state, excitation to excited state, and complete mixing.

If you are unfamiliar with the research literature or the terminology above, see "Learning more", below.

In addition, we aim to make simulations run on QuISP *completely reproducible*, to the extent humanly possible. It will be possible for others to verify work done using QuISP if they have the name of the QuISP release, version numbers for supporting software, the `.ini` file, any changed `.ned` files, and the seed for the pseudo-random number generator.

5.1.3 Current status

Most of the *infrastructure* is up and running, though the sets of actual experiments (interesting simulations) you can do are fairly limited still.

Working infrastructure:

- All of the basic OMNeT++ functionality, such as events, visualization, logging, analysis, etc.
- Complex network topologies, defining per link parameters including length, channel error rates, numbers of qubit memories, gate error rates, etc. (You will see the included demonstration networks in the documentation linked below.)
- A complete internal model of the software architecture for a repeater, including a connection manager, the RuleSet execution engine, real-time tracking of quantum states, etc.

Besides the obvious joys of the endless network configurability, here are the key quantum protocols that are implemented:

- basics of RuleSet creation & distribution
- various purification protocols: Single round of X purification, alternating X/Z purification, etc. Extending these to test your own custom purification protocol is pretty straightforward.
- tomography: when the simulation boots, it assumes that the software at each end of each link knows *nothing* about the link, so it begins by performing tomography on the links. This is actually problematic, because it turns out to take a long time for tomography to converge, which means a lot of boot-up time in the simulation before other interesting things start to happen. We are working on a way to pre-calculate this, so that you can choose to either include tomography or not; it sort of works, but may be a bit kludgy.

Current generic networking-level status:

- fully blocking circuit switching
- random pairwise traffic pattern (flat distribution)

In-progress work:

- entanglement swapping is a relatively new feature, and taking data using it is still not fully implemented
- although the connection setup protocol works, the teardown on completion of a connection still needs a little love
- the set of demo networks is still being polished

Upcoming features in near-term releases:

- more general resource allocation & multiplexing
- more general mechanism for establishing traffic patterns
- MSM links
- graph states at the link level
- multi-party states at the application level

Mid-term to long-term release features:

- 2G networks, esp. Jiang style
- full quantum internetworking

5.1.4 Installation requirements

The full installation process is described in `https://github.com/sfc-aqua/quisp/blob/master/doc/INSTALL.md`. The main software tools you will need are:

- QUISP requires OmNET++ `https://omnetpp.org/` and
- an external C++ library, Eigen `http://eigen.tuxfamily.org/`, to work.
- To contribute to QuISP development, you will also need to be familiar with at least the basics of git `https://git-scm.com/`.
- We recommend the use of Dozygen `http://www.doxygen.nl/` for source code comments, but the Doxygen tools are not required unless you want to build the source code documentation.

Depending on your local setup and how you intend to use QuISP, you may also need various tools (a C++ compiler, make, an X Windows server, Docker, etc.), documented in the installation notes.

5.1.5 Building and running

There are two main ways of working with QUISP. You can either use the Eclipse-like graphical interface of OmNET++, for which you will find instructions in `https://github.com/sfc-aqua/quisp/blob/master/doc/USAGE-omnetpp-gui.md`, or you can use the `Makefile` and GNU make, by looking at instructions in `https://github.com/sfc-aqua/quisp/blob/master/doc/USAGE-makefile.md`. Some operations are implemented in the Makefile and not explained for the graphical user interface.

5.1.6 Moving into useful work

Once you have gotten this far, you should be able to [run some of the most basic demos (`https://github.com/sfc-aqua/quisp/blob/master/doc/running-demos.md`). Next, you'll want to learn how to create your own test networks, and how to extend the source code for your own uses.

When you are ready to start contributing, you can start reading the code, as we have done (`https://github.com/sfc-aqua/quisp/blob/master/doc/code-spelunking.md`).

You will also want to look at some of the software design documents (`https://github.com/sfc-aqua/quisp/blob/master/doc/software-design.md`).

5.1.7 Development tools

A few tools (mainly scripts) can be used to make development a bit easier. Look around in the `bin` folder of this project.

5.1.8 Is QuISP right for me?

Fundamentally, the point of QuISP is that *networks are much more than point-to-point connections*.

If you want to know about the behavior of systems and networks, to study behavior of links that are too complex for simple analytic equations (esp. those with multiple qubits per link) or to contribute to the design of network protocols, QuISP is the simulator for you. If you're trying to adjust detector window timing v. entanglement fidelity, or figure out what Q factor your cavity needs to be, or understand dispersion in a fiber, it might not be.

5.1.9 Learning more

See the references `https://github.com/sfc-aqua/quisp/blob/master/doc/References.md`.

5.1.10 Contributing

First, join the QuISP Slack team (`https://aqua-quisp.slack.com`).

Please also refer to the [code of conduct] (`https://github.com/sfc-aqua/quisp/blob/master/CODE_OF_CONDUCT.md`).

5.1.11 License

QuISP was initially released on April 5, 2020, and is licensed (`https://github.com/sfc-aqua/quisp/blob/master/LICENSE`) under the 3-Clause BSD License (`https://opensource.org/licenses/BSD-3-Clause`).

QuISP builds on OMNeT++. OMNeT++ itself is a custom license (`https://omnetpp.org/intro/license`), open source and free for academic use, but a license fee required for commercial organizations. QuISP also requires the linear algebra library Eigen, where license is MPL2, and so probably not an issue (`http://eigen.tuxfamily.org/index.php?title=Main_Page#License`).

5.2 Quantum Router Software Architecture (QRSA)

One of the key outcomes of simply building a sophisticated quantum network simulator is the need to add realism to the behavior of the network nodes. We have created a five-subsystem structure, as shown in Fig.6.

The five major software subsystems are as follows:

- Routing daemon: As in an Internet router, responsible for exchanging classical messages concerning the status of links and nodes in order to enable distributed calculation of paths to connect nodes. This is an ongoing, background task.

- Connection manager: We have determined that end-to-end connections require the establishment of state information at each node along the chosen path; at least for 1G repeaters, and likely for all generations of repeaters, creation of end-to-end entangled states is a *distributed computation*, rather than a forward-and-forget stateless action. The CM is responsible for sharing RuleSets for connections. The CM acts only on connection setup requests, with moderately low latency preferred, but no specific real time requirements.

- RuleEngine: the heart of the software stack, matching the Condition clauses and executing the Action clauses of rules. The RE has soft real-time constraints, primarily to minimize waiting time for operations on qubits held in memory. The RE can send and receive classical messages as well as be triggered by quantum events.

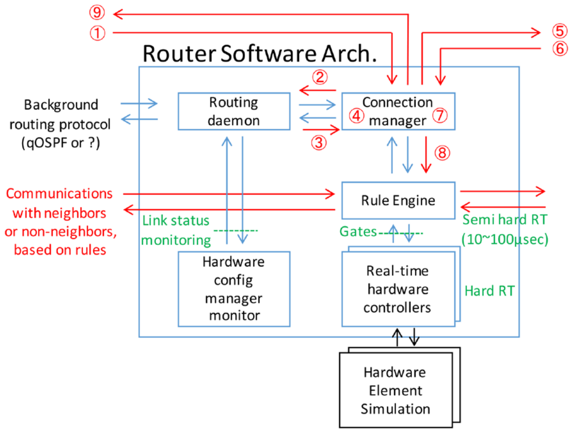

Fig 6 The Quantum Router Software Architecture. QuISP uses this structure within simulated routers and repeaters, and we expect that real-world systems will use a very similar structure.

- Hardware monitor: tracks the condition of memories and links for the routing daemon to use in reporting to other nodes.

- Real-time hardware controllers: the equivalent of device drivers, with hard real time constraints.

5.3 QuISP Hackathons

In support of direct QuISP development, attracting new developers and broadening the impact of our research, AQUA conducted several hackathons in 2020.

- February 8 -— SFC AQUA internal
- February 15 -— SFC AQUA internal
- March 14-15 -- SFC AQUA with a small number of external participants, preparing for the 1.0 release in April
- June 5-6 -- two-day virtual hackathon at WIDE Kenkyuukai

At the June hackathon, participants connected from France (2), Thailand (2), Korea (1), India (1), Italy (1), Netherlands (1), and Japan (12). Those connecting from Japan were Japanese (7), and one each from India, Argentina, Egypt, Korea and the U.S.

In October, IEEE held its first Quantum Week, intended to be a large-scale quantum conference with workshops, a technical program and perhaps a trade show. Rod Van Meter and Ryosuke "cocori" Satoh conducted a one-day workshop titled "Tutorial on Hands-on Simulation of a Quantum Network". Fifteen participants from around the world joined us. For this tutorial, we developed a detailed spreadsheet giving the individual tasks to be conducted, and we discussed live what the purpose of the tasks was and what was to be learned. We estimated the total time for the tutorial, excluding the preparatory tasks of learning key concepts and preparing the platform, at six hours. The steps were:

1. Watched IETF 104 tutorial from Prague (or have equivalent background)
2. Watched APS March Meeting 2020 talk

3. Have C++ compiler (Windows, Mac, Linux native, Linux on Docker)
4. Have git
5. Joined the AQUA-QuISP Slack
6. My platform is:
7. Downloaded OMNeT++ 5.6 (268-685MB, depending on platform)
8. Downloaded Eigen 3.3
9. Cloned QuISP from github
10. Built OMNeT++ (n.b.: this takes a long time! Worse, getting your system configured with the right libraries, etc. can be hard.)
11. OMNeT++ itself runs! Congratulations!
12. Configured OMNeT++ to find Eigen library
13. Configured OMNeT++ to find the QuISP source
14. QuISP builds! Congratulations!
15. Can find `quisp_tutorial.ini` and run `Tutorial_MM network` (any parameters) `Tutorial_MM` with tomography
16. Add purification
17. With modified parameters
18. With connection setup and data traffic instead of tomography
19. Try one or more of the same actions with `Tutorial_MIM` network
20. Try one or more of the same actions with the Y-shaped network Try one or more of the same actions with the core-and-arms network Try one or more of the same actions with the ISP network
21. Begin "code spelunking"

第6章　Ongoing Networking Research

AQUA's ongoing Quantum Internet research covers a broad range of architecture and protocol topics. The most important focus at this time is the creation and use of *graph states* on quantum repeater networks.

Repeater graph states: One of the biggest drawbacks to first-generation quantum repeaters dependent on purification and entanglement swapping is the need for an *acknowledged link layer*, resulting in link-level round-trip latencies and limiting throughput as buffer memories are occupied awaiting classical messages. The use of repeater graph states has the potential to *dramatically* increase performance, perhaps by several orders of magnitude, by eliminating these waits. However, the loss of photons is a severe problem, handled

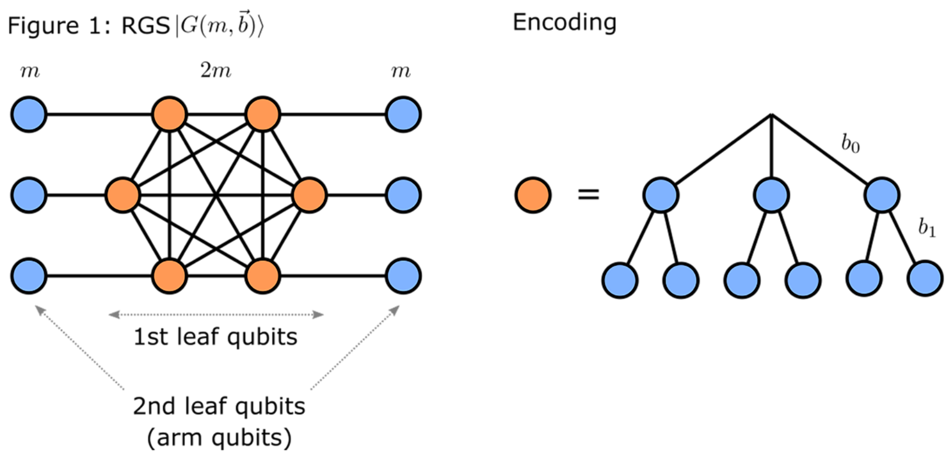

Fig 7　The multi-layer repeater graph state encoding for photons. Each blue circle represents a physical photon, each orange circle is a logical qubit encoded in multiple photons as shown on the right. Loss of individual photons, normally fatal to a quantum state, can be mitigated by adaptive choice of the measurement basis for further photons as they arrive at the network node that performs measurements.

with the clever use of an encoding scheme that allows us to remove subtrees of a large entangled state from that entangled state by choosing measurement bases based on on which qubits have successfully been received. The first proposal for repeater graph states (see Fig.7) [3] was wildly inefficient, but an improvement from Virginia Tech may bring it into the realm of practical implementation [9, 17]. In consultation with VaTech, we are working to implement variants of this scheme in QuISP to evaluate potential performance and engineering challenges.

Multiparty graph states: The above work is a system-level use of graph states, targeting the creation of end-to-end Bell pairs for ordinary application uses. In similar fashion, graph states that span multiple nodes and are provided directly to applications can be used for distributed quantum computation, including blind quantum computation [8, 10, 24]. We are working on both network methods for distributing these states and methods for dividing desired graph states into pieces for distribution based on the infidelity of links and qubit memories.

2G repeater simulation: In conjunction with Mahidol University, Thailand, we are working to simulate quantum repeater networks that utilize quantum error correction directly on encoded qubits [19].

Quantum network security: Work on understanding how to make quantum networks operationally robust against misbehavior (whether or not intentional) is needed as we begin to plan quantum repeater networks. We submitted one paper on this topic in 2020, and as of this writing it is under revision for resubmission to the journal (see the list below).

Impact of tomography on network operations: The tomography work and hijacking framework described above represent the first steps in understanding how the network is monitored in real time operation. A key concern is doing so in Internet-scale interconnected systems, where latencies are high and heterogeneity forces technology-independent data representations and communications.

Connection establishment methods: We now understand reasonably well how quantum connections want to use the network and at the abstract level how to select the resources to use. However, actually identifying and reserving the resources requires a good deal of engineering. A key concern is doing so in Internet-scale interconnected systems, where privacy and autonomy of operation are paramount.

Application analysis: Many more prospective applications of distributed quantum entanglement exist in the literature, but the operational demands they make of quantum networks is still poorly understood.

第７章　#QuantumNative: Online Education and Research for the Next Generation

For a number of years, the term *digital native* has been common:

> **digital native** (noun) a person born or brought up during the age of digital technology and so familiar with computers and the Internet from an early age.
> *Apple Dictionary 2.2.1, 2014*

Similarly, we can define a *quantum native*:

> **quantum native** (noun) a person born or brought up during the age of quantum technology and so familiar with quantum computers from an early age; a person whose first serious study of algorithms involved quantum algorithms, whether exclusively or in conjunction with classical algorithms.

We are using the hashtag `#QuantumNative` to describe such people. An important goal of the AQUA working group is to find and nurture quantum native talent; because the thought process for creating quantum algorithms is very different than that of classical algorithms, we believe it is important that potential quantum programmers are exposed to the concepts as early in their education as possible. For some years, WIDE member Rodney Van Meter has been teaching students as young as first-year bachelor's students, and in 2020 we extended our reach substantially.

Table 2 Summary of FutureLearn MOOC learners through Jan. 25, 2020.

Joiners	12,064
Comments	7,824
Countries	159

7.1 FutureLearn MOOC

In October 2017, WIDE members Rodney Van Meter and Takahiko Satoh, working with WIDE member Keiko Okawa, brought online a massive online open course (MOOC) titled, "Understanding Quantum Computers" (UQC). The course was presented online through the platform FutureLearn. Keio University has partnered with FutureLearn since 2016. UQC has attracted more than 12,000 learners from 159 countries and territories A screenshot of the trailer is shown in Fig.8.

In 2020, the biggest advance in the MOOC was the introduction of Indonesian subtitles and translations for the articles, bringing the total number of languages to four: English, Japanese, Thai and Indonesian.

Many MOOC platforms are ultimately passive, composed of little more than hour-long lectures by a professor or other expert, often recorded in a lecture hall and with no more attention to the clarity and attractiveness of visual aids than pointing a camera at a projection screen. FutureLearn, in contrast, emphasizes interactively for learners and high production values for audio and video. The FutureLearn philosophy features three points: (1) tell stories; (2) provoke conversation; and (3) celebrate progress. Each learning Step is targeted at approximately ten minutes. Several types of materials can be presented; we used video; articles with text and graphics; quizzes; semi-moderated discussion boards, in which the educators and a team of assistants participated; and two types of external materials. Because the Futurelearn discussion boards do not allow learners to upload video or images, we linked to a separate site that allows this. We also included plans for 3-D printable objects for explaining some concepts.

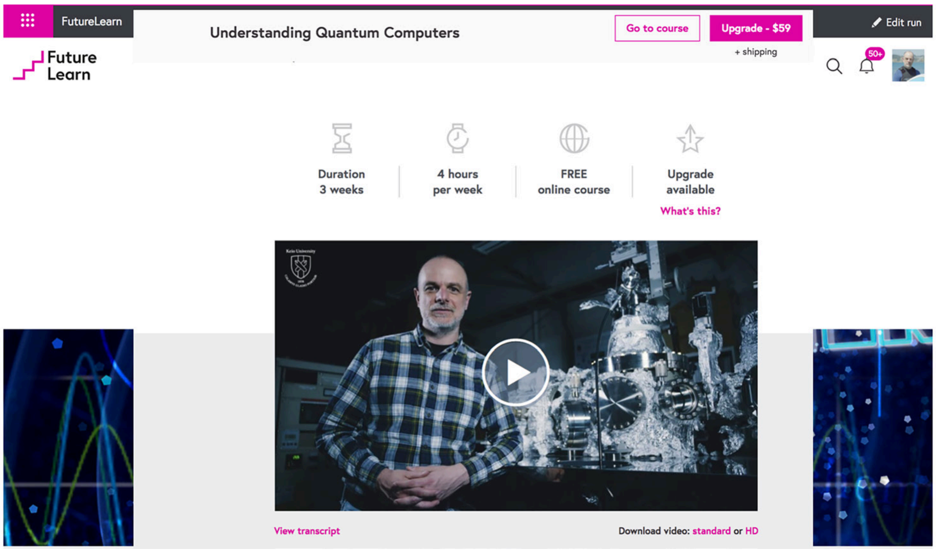

Fig 8 The trailer on the front page of the "Understanding Quantum Computers" MOOC.

7.2 FutureLearn MOOC

In addition to the MOOC, learners now have access through the web to several superconducting quantum computing systems of different architecture and capabilities.

IBM, with one of the world's leading industrial research efforts, has produced numerous systems, some public and some restricted access, beginning with a five-qubit system [33]. These systems can be accessed via a web interface[*4]. A screen shot of the front end is shown in Fig.9. Some WIDE members are members of the Keio Quantum Computing Center (KQCC) with access to the best available systems, and have published papers using these machines [27, 30, 25].

7.3 Q-LEAP Education

In fall 2020, The Q-LEAP office awarded to a group led by Prof. Kae Nemoto of NII a large, six-year contract for development of an online quantum information curriculum for Japanese universities. Rod Van Meter is a member of the group, and he is responsible for the quantum communications and networking portion of the curriculum. A proposed four-year curriculum for students concentrating in this area is shown in Fig.10.

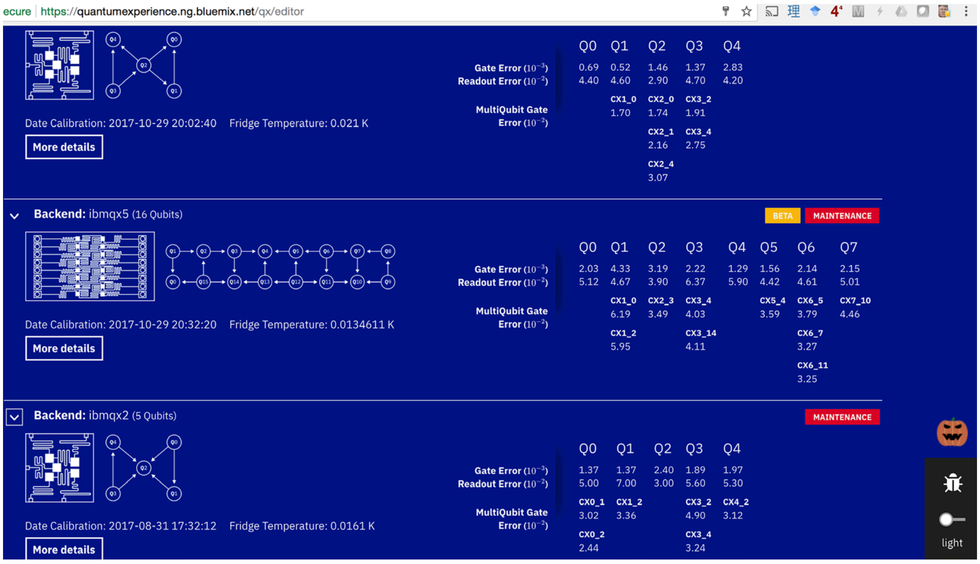

	Q0	Q1	Q2	Q3	Q4
Gate Error (10^{-3})	0.69	0.52	1.46	1.37	2.83
Readout Error (10^{-2})	4.40	4.60	2.90	4.70	4.20

MultiQubit Gate Error (10^{-2}): CX1_0 1.70, CX2_0 1.74, CX3_2 1.91, CX2_1 2.16, CX3_4 2.75, CX2_4 3.07

Date Calibration: 2017-10-29 20:02:40 Fridge Temperature: 0.021 K

Backend: ibmqx5 (16 Qubits)

	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
Gate Error (10^{-3})	2.03	4.33	3.19	2.22	1.29	1.56	2.14	2.15
Readout Error (10^{-2})	5.12	4.67	3.90	6.37	5.90	4.42	4.61	5.01

MultiQubit Gate Error (10^{-2}): CX1_0 6.19, CX2_3 3.49, CX3_4 4.03, CX5_4 3.59, CX6_5 3.79, CX7_10 4.46, CX1_2 5.95, CX3_14 4.11, CX6_7 3.27, CX6_11 3.25

Date Calibration: 2017-10-29 20:32:20 Fridge Temperature: 0.0134611 K

Backend: ibmqx2 (5 Qubits)

	Q0	Q1	Q2	Q3	Q4
Gate Error (10^{-3})	1.37	1.37	2.40	1.89	1.97
Readout Error (10^{-2})	5.00	7.00	3.00	5.60	5.30

MultiQubit Gate Error (10^{-2}): CX0_1 3.02, CX1_2 3.36, CX3_2 4.90, CX4_2 3.12, CX0_2 2.44, CX3_4 3.24

Date Calibration: 2017-08-31 17:32:12 Fridge Temperature: 0.0161 K

Fig 9 The GUI element showing the IBM quantum computers available for online use as of Oct. 30, 2017.

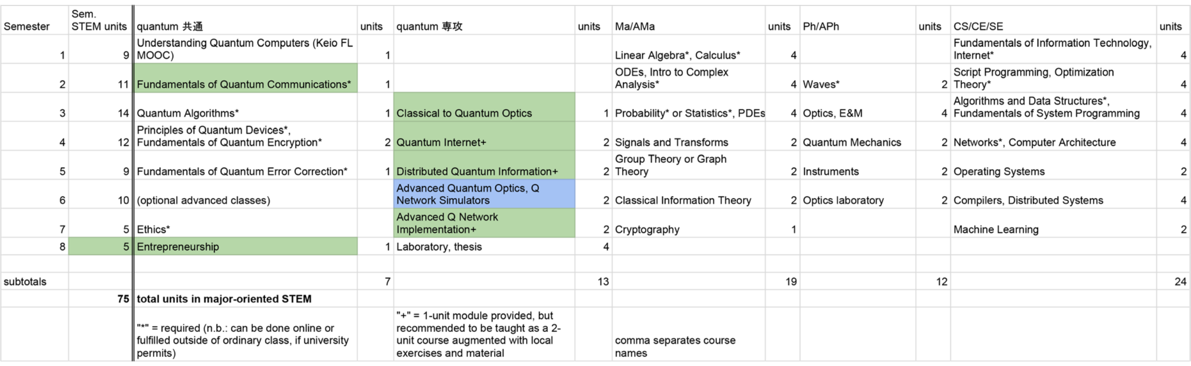

Semester	Sem. STEM units	quantum 共通	units	quantum 専攻	units	Ma/AMa	units	Ph/APh	units	CS/CE/SE	units
1	9	Understanding Quantum Computers (Keio FL MOOC)	1			Linear Algebra*, Calculus*	4			Fundamentals of Information Technology, Internet*	4
2	11	Fundamentals of Quantum Communications*	1			ODEs, Intro to Complex Analysis*	4	Waves*	2	Script Programming, Optimization Theory*	4
3	14	Quantum Algorithms*	1	Classical to Quantum Optics	1	Probability* or Statistics*, PDEs	4	Optics, E&M	4	Algorithms and Data Structures*, Fundamentals of System Programming	4
4	12	Principles of Quantum Devices*, Fundamentals of Quantum Encryption*	2	Quantum Internet+	2	Signals and Transforms	2	Quantum Mechanics	2	Networks*, Computer Architecture	4
5	9	Fundamentals of Quantum Error Correction*	1	Distributed Quantum Information+	2	Group Theory or Graph Theory	2	Instruments	2	Operating Systems	2
6	10	(optional advanced classes)		Advanced Quantum Optics, Q Network Simulators	2	Classical Information Theory	2	Optics laboratory	2	Compilers, Distributed Systems	4
7	5	Ethics*		Advanced Q Network Implementation+	2	Cryptography	1			Machine Learning	2
8	5	Entrepreneurship	1	Laboratory, thesis	4						
subtotals			7		13		19		12		24
	75	**total units in major-oriented STEM**									
		"*" = required (n.b.: can be done online or fulfilled outside of ordinary class, if university permits)		"+" = 1-unit module provided, but recommended to be taught as a 2-unit course augmented with local exercises and material		comma separates course names					

Fig 10 A proposed four-year curriculum for those interested in quantum networking. Modules labeled in green will be developed by the SFC team led by Rod Van Meter.

[*4] https://www.research.ibm.com/ibm-q/

第８章　Quantum Community

In addition to the QuISP hackathons described in Sec.5.3, WIDE member Sara Metwalli is active in Women Who Code, where she has conducted hands-on tutorials introducing quantum computing, and WIDE members helped to design and conduct the IBM Qiskit Challenge hackathons, shown in Fig.11.

Two Qiskit Challenges have been conducted out of Japan to date. The fall 2019 edition guided complete beginners through the basics and concluded with a competition for efficiently solving a graph coloring problem, colorfully couched as placement of convenience stores in Tokyo[＊5]. The fall 2020 edition focused on a satisfiability problem[＊6]. Both editions were created by WIDE members Takahiko Satoh and Shin "parton" Nishio, working with Atsushi Matsuo (IBM). WIDE member Samanvay Sharma worked to organize and promote the 2020 edition.

Rod Van Meter has joined the editorial board of the new journal *IEEE Transactions on Quantum Engineering* as an associate editor. He also participates in the new IEEE interest group on quantum education[＊7].

第９章　Quantum Computing

The SFC AQUA group continues to devote about 50% of its effort to quantum computing, in part by participating in the Keio Quantum Computing Center and the IBM Q Network. This year's WIDE annual report focuses on Quantum Internet research, so this section will only briefly introduce our current research and results.

9.1　Quantum Software: IBM Q Network

In December 2017, IBM announced an organizational network of hubs around the world, where customers can gather to study quantum computing and develop algorithms suited to their own businesses. Members of the hubs have access to the fifty-qubit system and larger systems as they come online. Keio University was the first such hub in Asia, and WIDE Member Rodney Van Meter is the Vice Chair.

In the short run, many researchers are focusing on the development of *hybrid* algorithms, using noisy, intermediate-scale quantum computers [29] to execute specific subroutines, and augmenting the quantum computation with significant amounts of classical computation. An important element in the success of such an approach is error mitigation, before

#QuantumNative の育成

IBM Quantum Challenge 2019
量子プログラミングコンテストによるエキスパート養成

✔ IBM社との共催（問題作成・ジャッジを担当）
✔ 大会の登録者は700人以上

✔ 優勝者が今秋からKQCCに所属（博士課程）

東京を模した地図の塗り分けを量子コンピュータで解決

IBM Quantum Challenge 2020 Fall

✔ 大会の登録者は3000人以上
✔ 国別最多は**インド（50%超）**

✔ インド初のQiskit Advocatorが今秋からKQCCに所属（修士課程）

問題を解いてProf. Ryokoを量子世界から救出

Fig 11　The 2019 and 2020 IBM Qiskit Challenge hackathons.

[＊5] https://github.com/quantum-challenge/2019

[＊6] https://github.com/qiskit-community/IBMQuantumChallenge2020

[＊7] https://ed.quantum.ieee.org/

complete quantum error correction becomes technically feasible [14].

9.2 Research Accomplishments

Our most important publications this year focused on the use of NISQ machines and how algorithms will grow from small scale, noisy systems to large scale, error corrected systems.

Many quantum algorithms will need digital arithmetic, e.g. for cryptography and for counters in graph theory problems, but existing systems are still too noisy to execute arithmetic effectively. Our *subdivided oracle* concept uses the analog phase of a qubit as an analog counter. Grover's lgorithm [16] uses a phase flip to *mark* solutions to a given problem. We applied the subdivided oracle as a counter, using partial rotations instead of a full π phase flip in a graph theory problem [30]. This dramatically reduces gate count and increases success probability, but its full theoretical implications are still not fully understood. This may be an important technique for years to come.

A common AQUA working approach is to take an algorithm defined abstractly by theorists and work to understand how quantum computers will execute the algorithm in practice. This is far more than simple programming, as new techniques must often be discovered and assumptions made in theoretical work prove to be impractical. *Clique finding* is an important graph theory problem, and a theoretical algorithm for it has been proposed for quantum computers [22]; we implemented this algorithm and assessed its prospective performance in both the near and long term [25]. We proposed the use of constrained quantum states that improve the Grover search in this case, as well as examined the impact of the assumed used of QRAM, a quantum computing memory for large classical datasets that does not yet exist.

9.3 Ongoing Research

Discrete logarithms: Shor's algorithm for factoring large numbers [31] has a variant that applies to discrete logarithms, useful for attacking certain classical encryption algorithms. While our own focus is not on the use of quantum computers for cracking cryptography, cryptographers must work on this topic in order to understand the capabilities of quantum computers as they attempt to design post-quantum cryptographic techniques. Joint work with NICT, Mitsubishi UFJ Financial Group, and Mizuho Information
& Research Institute through KQCC.

Hybrid algorithms: The current most important topic in the quantum computing community is how to combine noisy intermediate-scale quantum computers [29] with large-scale classical computers to solve problems more quickly than classical systems alone. These *hybrid algorithms* must take into account noise, and decompose larger problems into small sub-problems that can be effectively handled on existing and near-future quantum computers.

Error mitigation: The companion problem to developing hybrid algorithms is developing *error mitigation* techniques that will allow us to execute quantum algorithms effectively in the presence of noise, without the overhead and technical challenges of full-scale quantum error correction and fault tolerance.

Quantum machine learning: Several AQUA students are working on machine-learning related projects, such as a quantum form of the k-means clustering algorithm, used for analyzing audio data, and image processing algorithms.

Quantum chemistry and quantum mechanics: Some students are working on quantum chemistry, which was the original driving vision behind quantum computers, and simulation of quantum systems that are open to the environment.

第10章　2020 Publications

In 2020, AQUA completed five peer-reviewed journal publications:

1. SA Metwalli, FL Gall, R Van Meter, **Finding Small and**

Large k-Clique Instances on a Quantum Computer, *IEEE Transactions on Quantum Engineering*, to appear; arXiv preprint arXiv:2008.12525

2. T Satoh, Y Ohkura, R Van Meter, **Subdivided phase oracle for NISQ search algorithms**, *IEEE Transactions on Quantum Engineering* 1, 1-15, 2020

3. AS Cacciapuoti, M Caleffi, R Van Meter, L Hanzo, **When Entanglement meets Classical Communications: Quantum Teleportation for the Quantum Internet**, *IEEE Transactions on Communications*, 68(6), 3808-3833, 2020

4. S Nishio, Y Pan, T Satoh, H Amano, R Van Meter, **Extracting Success from IBM's 20-Qubit Machines Using Error-Aware Compilation**, *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 16 (3), 1-25, 2020

5. P Pathumsoot, T Matsuo, T Satoh, M Hajdušek, S Suwanna, R Van Meter **Modeling of Measurement-based Quantum Network Coding on IBMQ Devices**, *Physical Review A*, 101(5), 052301, 2020

One blog posting for the IETF was picked up and published in an IEEE magazine:

1. W Kozlowski, R Van Meter, **Schrödinger's Internet at the IRTF**, *IEEE Communications Standards Magazine* 4 (3), 4-6, 2020

We had several conference presentations that were not peer reviewed:

1. Yoshinori AONO, Sitong LIU, Tomoki TANAKA, Shumpei UNO, Rodney VAN METER, Naoyuki SHINOHARA, and Ryo NOJIMA, **Executing discrete logarithm problem on superconducting quantum processors**, 43rd Quantum Information Technology Symposium (QIT43), Dec. 2020

2. SA Metwalli, FL Gall, R Van Meter, **Finding Small and Large k-Clique Instances on a Quantum Computer**, IPSJ Quantum Software Kenkyuukai, Oct. 2020

3. Satoh Ryosuke, Van Meter Rodney Hajdusek Michal, **Federated Graph State Preparation on Noisy, Distributed Quantum Computers** IPSJ Quantum Software Kenkyuukai, Oct. 2020

4. Ohkura Yasuhiro, Van Meter Rodney Hajdusek Michal, **Crosstalk-aware NISQ Multi-programming** IPSJ Quantum Software Kenkyuukai, Oct. 2020

5. SA Metwalli, FL Gall, R Van Meter, **Finding Small and Large k-Clique Instances on a Quantum Computer**, IEEE Quantum Week, Oct. 2020

We also have several papers currently under review:

1. K Oonishi, T Tanaka, S Uno, T Satoh, R Van Meter, N Kunihiro, **Efficient Construction of a Control Modular Adder on a Carry-Lookahead Adder Using Relative-phase Toffoli Gates**, arXiv preprint arXiv:2010.00255 2020

2. T Satoh, S Nagayama, S Suzuki, T Matsuo, R Van Meter, **Attacking the Quantum Internet**, arXiv preprint arXiv:2005.04617 2020

Appendix A What is AQUA?

A.1 Goals

The primary goal of AQUA is to advance the deployment of quantum technologies in the real world, principally by applying known techniques from classical computer architecture, networking and distributed systems to the problems of scalability in quantum systems. This work will both bring new computational capabilities and help ensure that the progress of information technology does not end when the size of transistors can no longer be reduced.

The physical technology on which modern computing systems are built will change dramatically over the course of the next several decades. Beyond the research goals, AQUA also aims to expose the current generation of students to the principles that drive the evolution of computing technology, and the underlying physics of computation, preparing the students for forty-year careers in which they will work with applied physicists and electrical engineers to drive the coming technological revolutions.

A.2 Work Areas

AQUA works in five technical areas contributing to distributed quantum computing systems:

- Devices: In conjunction with researchers at Stanford University, RIKEN, and the University of Tokyo we worked on the design of semiconductor-based chips using optically-controlled *quantum dots* and *superconducting flux qubits*.

- Workloads: Although AQUA does not focus primarily on the creation of new quantum algorithms, we do work on how to implement known quantum algorithms efficiently on realizable architectures. We also perform the reverse analysis: to implement a given algorithm, how large and how accurate a quantum system is required?

- Tools: Proper analysis of new ideas in architecture and networks requires software tools for compiling programs and optimizing their mapping to particular systems, as well as physical simulation of quantum devices and effects.

- Principles: We are searching for new principles in quantum architecture and networking, as well as applications of known principles.

- Networks: Large systems must combine multiple devices into one system that can compute collaboratively, as well as share information; we are investigating both system-area and wide-area quantum networks.

Underlying all of these is the critical issue of error management in quantum systems; quantum data is far too fragile to store or compute upon without continuous, active correction. Our primary focus is on the promising surface code error correction, looking for ways to makes its implementation resource-friendly and robust in the face of various system constraints.

Beyond these specific technical areas, we work in two equally vital areas for the creation of a quantum industry that is healthy and supplied with well-trained engineers and scientists:

- Education: We are consciously working on educational materials and methods for bringing new people into the field, focusing on early undergraduate education.

- Community: Building a community that is broad based and independent of any given university, company or technology is vital for both the short and long term.

Appendix B Quantum Concepts

The following is a brief summary of the key aspects of quantum communication and computation that impact network and system architecture.

Qubits. Quantum information is most often discussed in terms of *qubits*. A qubit, like a classical bit, is something with two possible values that we can label zero and one. Unlike a classical bit, a qubit can occupy both values simultaneously, known as *superposition*.

Superposition and measurement. A qubit can represent multiple values in different proportions at the same time, e.g., two-thirds of a "one" and one-third of a "zero". This *superposition* determines the relative probability of finding each value when we *measure* the state. When we measure the qubit, we get only a single classical bit of information (the "one" or "zero") with 100% probability, and the superposition *collapses*.

Entanglement and Bell pairs. Some groups of qubits exhibit strong correlation between the qubits that cannot be explained by independent probabilities for individual qubits. Instead, the group must be considered as a whole, with interdependent probabilities. This phenomenon is known as *quantum entanglement*. A special entangled state known as a *Bell pair* or *EPR pair*, consisting of two quantum bits, figures prominently in quantum communication. Each qubit in the pair has a 50% probability of having a value of 1 and a 50% probability of having a value of 0 when we measure it. Although we cannot predict which will be found, when we measure one member of the pair, the value of the other is immediately determined. This happens independent of the distance between the two members of the Bell pair.

Interference. Quantum algorithms use some building blocks derived from classical concepts, such as adder designs, but the overall thrust of a quantum algorithm is very different from that of a classical algorithm. Rather than attempting to solve a problem and checking for the answer, a quantum algorithm's goal is to create *interference* between the elements of a superposition quantum state. The basic concept is shown in Figs.12 and 13 in the main text, while in quantum computers the interference happens across a much larger space. Constructive interference reinforces desirable states, increasing the probability of finding a desirable outcome on measurement, while destructive interference reduces the probability.

No cloning. As mentioned above, a key restriction of quantum systems is that we cannot make *independent* copies of an unknown state [42]. This makes error correction exceedingly difficult.

Fidelity. The quality of a quantum state is described by its *fidelity*, which is, roughly, the probability that we correctly understand the state -- if we ran the same experiment many times and measured the results, how close to our desired statistics would we be? Unfortunately, any physical operation results in a loss of fidelity, gradually degrading the state as we manipulate or even store it. We can counter this by using a form of error correction or detection.

Purification. The form of error detection historically favored in quantum repeater networks is *purification*, which uses minimal resources [7]. It sacrifices some quantum states to test the fidelity of others. There are various purification mechanisms, with different purification algorithms and different methods for determining which states are sacrificed, each with particular tradeoffs.

Quantum error correction (QEC). QEC may be based on classical codes or purely quantum concepts. The primary difficulties are extraction of errors without damaging quantum state, avoiding error
propagation, and the increased resources required. (See references contained in [35], [19] and [15].)

Teleportation. Teleportation destroys the state of a qubit at the sender and recreates that state at the destination, teleporting information rather than matter, as explained in Figure 14 [5]. The process uses a Bell pair's long-distance correlation, followed by transmission of a pair of classical bits.

With these basic concepts, we can begin to construct networks. Bell pairs are consumed by teleportation, so one way to organize a network is to create a continuous stream of Bell pairs between source and destination -- as long as we identify those sources and destinations, choose paths to get there, and manage the resources along the way.

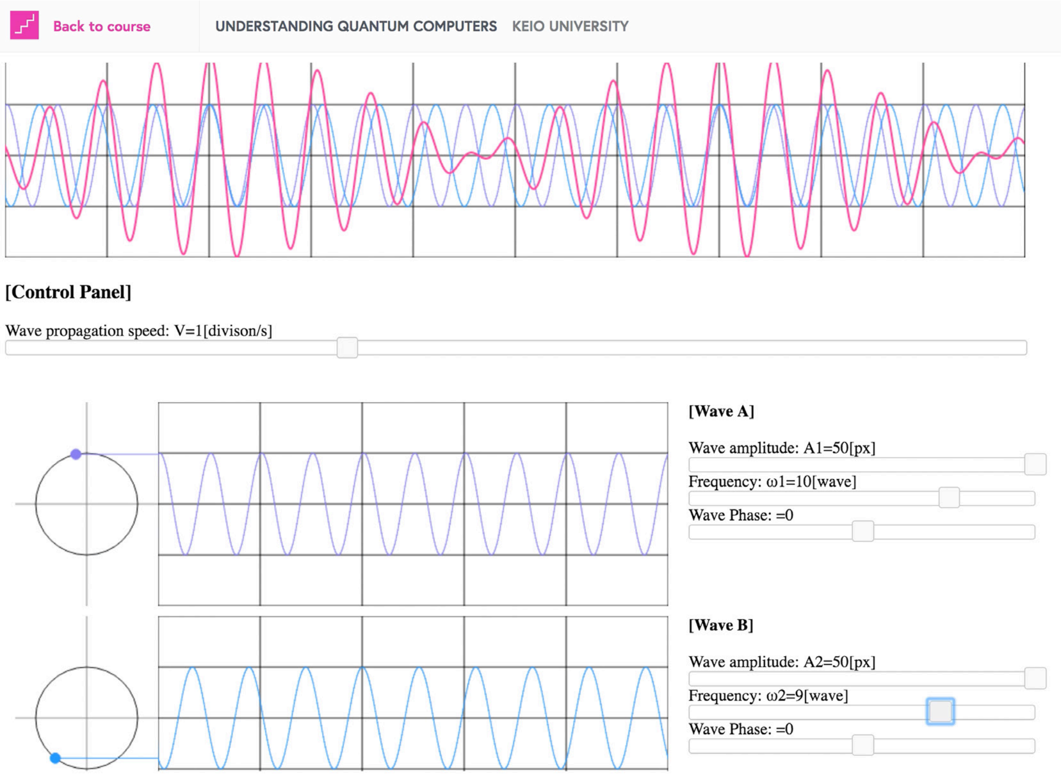

[Control Panel]

Wave propagation speed: V=1[divison/s]

[Wave A]

Wave amplitude: A1=50[px]

Frequency: ω1=10[wave]

Wave Phase: =0

[Wave B]

Wave amplitude: A2=50[px]

Frequency: ω2=9[wave]

Wave Phase: =0

Fig 12 JavaScript app allowing the learner to adjust various parameters to learn about one-dimensional interference. Constructive and destructive interference are the key to quantum algorithms.

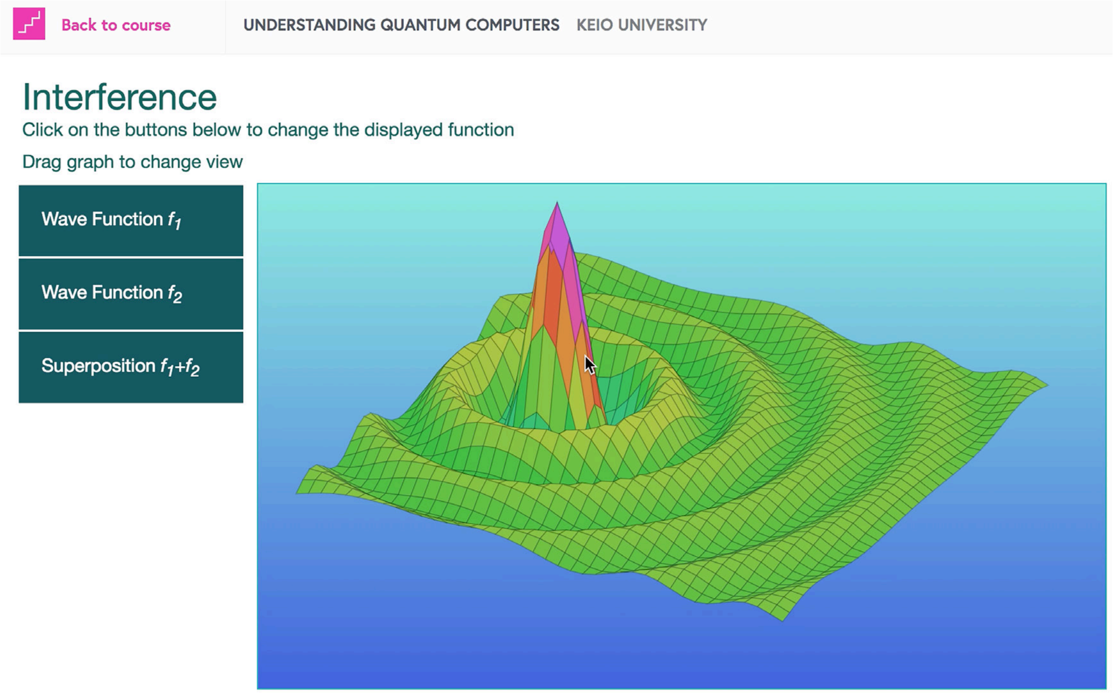

Interference

Click on the buttons below to change the displayed function

Drag graph to change view

Fig 13 JavaScript app for understanding two-dimensional interference, created using the D3 library.

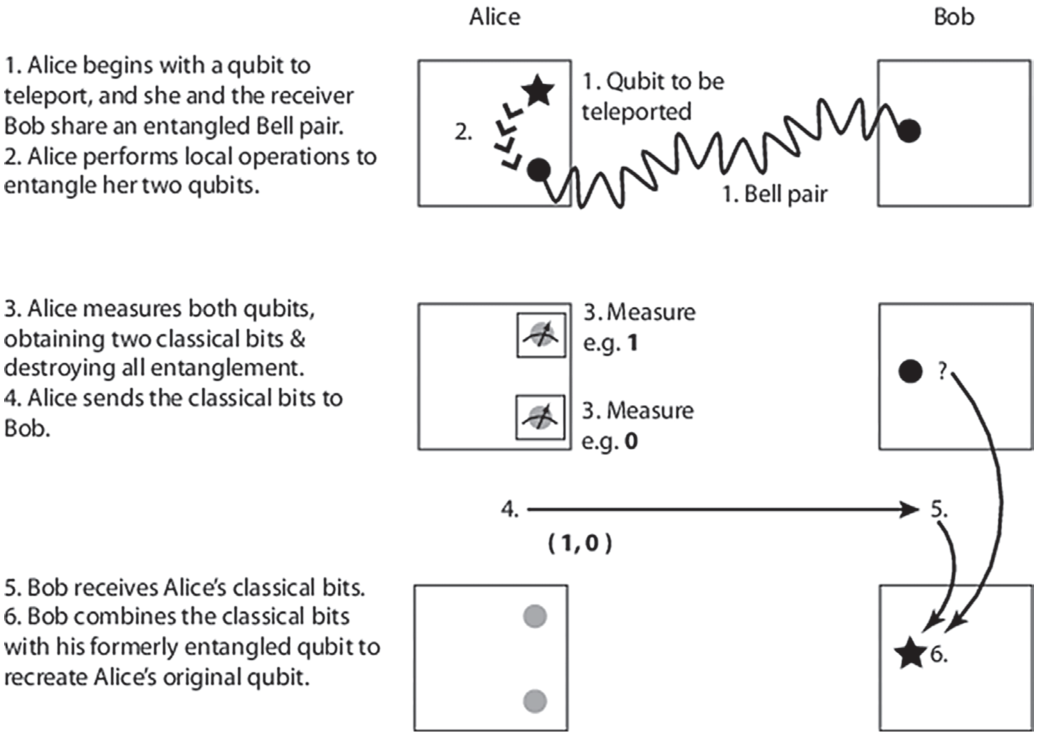

1. Alice begins with a qubit to teleport, and she and the receiver Bob share an entangled Bell pair.
2. Alice performs local operations to entangle her two qubits.

3. Alice measures both qubits, obtaining two classical bits & destroying all entanglement.
4. Alice sends the classical bits to Bob.

5. Bob receives Alice's classical bits.
6. Bob combines the classical bits with his formerly entangled qubit to recreate Alice's original qubit.

Fig 14　Operations in teleporting a qubit from Alice to Bob.