RFI response

Raytheon
BBN Technologies

Organization Submitting	Raytheon BBN Technologies Corp.
Type of Business	Large Business
Title	Response to *Request for Information on Quantum Information Science and the Needs of U.S. Industry*
Technical Point of Contact	Dr. Zachary Dutton 10 Moulton Street Cambridge, Massachusetts 02138 Phone: 617-873-5122 Email: sguha@bbn.com
Date	May 8th, 2015
RFI No.	R15009-BBN

The *Quantum Information Processing* group at Raytheon BBN Technologies has been pursuing research and development towards applications of quantum information science (QIS) in *computation*, *communication*, and *sensing.* Since our group's founding in 2009, we have worked closely with various government agencies in identifying such applications and mapping out technology roadmaps to realize systems capable of implementing them. We are pleased to offer our perspective on the important questions posed by the government in this request for information and look forward to continuing to engage the government on these topics. As the lead organization for QIS within Raytheon we can speak to both our specific perspective as researchers in this field and some of the motivation and interest within Raytheon for the eventual benefits of developing this technology.

1. Opportunities

From our perspective, the United States is ahead of the other industrialized countries in terms of the quantum computation technology. This advantage has been established and maintained largely due to the IARPA MQCO program as well as several additional programs within IARPA and LPS. This position is extremely valuable and should absolutely be maintained. Within this research, several qubit technologies have emerged as particularly strong candidates for development into useful, scalable quantum processors and this has guided some of our own investment and prioritization.

Superconducting base qubits have made massive improvements in coherence times (with the top groups going from routinely achieving 10 ns-100 ns to 10 μs-100 μs in the span of several years). Being in this regime has qualitatively changed the research activity as the emphasis has moved beyond simply exploring limitations to coherence time and improving them. Rather, we now are approaching a regime in which effects beyond coherence times such as control electronics will have an non-negligible impact on gate fidelities. Furthermore even coherence times are often understood at a microscopic level (for example due to mode couplings, which can be calculated with HFSS and other EM software). For this reason, the qubit design is beginning to require a more engineering-like approach and we believe continued investment into this is going to pay huge dividends. In this activity, industrial research organizations such as Raytheon BBN have and can continue to play a key role. Going forward, superconducting qubits have several properties that make it a strong candidate for eventual scalable quantum processors:

- **Fast gates.** Many of the gates have been demonstrated at high fidelity in ~10s of ns and all of them have now been demonstrated in ~100s of ns. Thus, the fundamental physical gate clock rate is among the fastest of any candidate qubit technology.
- **Ability to fabricate multi-qubit chips.** There are many challenges associated with eventually fabricating the controlling qubit systems at large scale, however, current methods can fabricate large numbers of qubits on a single chip.

Another extremely promising qubit technology are ions. They have been demonstrated with gate fidelities in excess of that for superconducting qubits, we have an even better understanding of microscopic mechanisms for decoherence and gate errors in these systems, and their coherence times are extremely long – in excess of seconds. They do suffer two drawbacks: (1) the slowest gates are much slower than that in superconducting qubits and will thus impose a bottleneck on clock speed; (2) there are questions regarding whether traps can be scaled to support sufficiently large qubit systems. Despite these drawbacks, investment in ion qubit technology is still highly valuable in our view. The two technologies have different advantages and disadvantages and it is too early to fully bet on one or the other as the eventual winner for a scalable processor.

Perhaps more importantly, it is extremely likely that these two qubit technologies will play complimentary roles in the eventual implementation of quantum networks, which we believe eventually enable the highest impact applications of QIS. We discuss these applications in Section 2 below. However, we note here that a quantum network should contain some nodes with a high degree of processing power (for which superconducting qubits may be best suited), while all nodes should be able to distribute, route and purify entanglement. For this latter task, a qubit technology with a longer coherence time (to match the longer time scales associated with transit time over the network and wait times for successful entanglement connection events) will be more important than qubits optimal for high fidelity gates at large-scale. Thus a system which uses a hybrid of these two technologies will likely be the most powerful one when we reach the scale of large scale quantum networks.

A key capability to enable quantum networks with hybrid qubit technology will be interconversion of quantum information among qubits in very different frequency regimes. In particular ionic and optical qubits are generally defined by optical frequencies whereas superconducting and quantum dot qubits are at microwave transition frequencies. There are several intriguing approaches to achieve this, including using mechanical oscillators coupled to optical cavities as an intermediary, as well as more direct conversion utilizing nonlinear materials with high electro-optic coefficients (such as LiNbO3). This technology will be highly impactful to future quantum networks. In the shorter term, we believe it will additionally be a key enabler for quantum computation. Entanglement distribution would massively increase the power of quantum processors via distributed quantum processing analogous to current classical multi-core super-computers. Even as qubit densities and performance increases, networking capability will always be an important ingredient to circumvent practical limits on the scale of single processors (for example a single DR or a single ion trap). We envision such quantum super-computers will be implemented over short links even within a single facility, much as current super-computers currently are networked with optical links.

Moving beyond the scope of quantum computing specifically, we believe that highly integrated and power efficient hardware to implement quantum communication will be an important piece to realization of large-scale quantum networks. Power efficient and coherent nonlinear processing in the optical domain will be important as well as low-loss and high channel count linear devices such as waveguides, switches, and filters. To this end research in quantum communications in integrated nano-photonics (both in traditional materials such as silicon and interesting alternatives such as SiC, diamond, and GaN) will be an important step in realizing scalable and robust devices for networks.

While the above points to specific hardware development directions, some of the government's funding of theoretical analysis of such systems has been extremely valuable and should be expanded. For example, the ARO/LPS verification program is tackling important questions regarding how to characterize quantum processors even of modest scale (~10's of qubits, which we expect to be routinely achievable in the next few years). This program and related work is getting at important questions regarding how to verify fault tolerance in the presence of more realistic noise assumptions (*e.g.* non-Markovian, correlated, and non-Clifford noise models). Without particular focus on issues of this type, quantum computing development runs the risk of either over-specifying hardware requirements (that is, not take advantage of characteristics in the correlations and types of errors that may in reality be the dominant ones) or conversely, be overly optimistic and not accounting for the logical fault paths which can be enhanced by errors present in real systems not properly described by models assumed in calculations and simulations.

2. Market Areas and Applications

Interest in quantum computing has been largely motivated by Shor's algorithm, due to the importance of its application and the exponential speed-up it promises over any existing techniques. Progress in discovering other algorithms with this degree of impact has been slow. However, there now exists an interesting collection of results that point to some important directions for application and algorithm work to go. We note three such directions in the following paragraphs that could lead to a more robust and concrete set of 'killer apps' that the quantum computation community could point to.

Grover's algorithm has captured less attention for two reasons: (1) the speed-up offered is only polynomial in the problem size; (2) a common criticism is that search of an unstructured list is not a very common problem. However, recent work has indicated that the speed-up carries over to structured list searches as well and is more robust than first thought. Furthermore, the general technique is applicable to a wider class of graph optimization problems with an application space much broader than that offered by Shor. We would encourage a more serious investment in studying graph optimization problems which offer even polynomial speed-up on the grounds that the range of applications could be extremely broad and the speed-ups offered could still be of significant interest for those applications. There are multiple examples of provable speed-up now established in this field. Adding to this motivation, we note that recent work has also begun to explore more subtle questions than the usual computer science standard of scaling of the worst-case. For example, average case performance is often of more practical interest than the worst case, especially if those worst cases are driven by specific pathological examples. A second important approach is that of approximation algorithms or, more precisely, algorithms to find approximate solutions. These algorithms can potentially run faster and some recent progress has been made on them.

Hamiltonian simulation is another area which merits much tighter focus in the coming years. While it is likely that calculations on large interesting simulations (such as of large, complex molecules, material science systems, and high energy physics models) will still require large scale quantum computers with error correction, the economic and technological impact of this capability could be far-reaching within the science, technology and engineering communities. It is imperative that the community converge on some new conclusions about this application space to increase the motivation for quantum computation research and development and the government can play a key role in helping this process.

Finally, it is worth noting that algorithms to break cryptographic schemes, of which Shor's algorithm is a prime example, have been extended to other crypto systems. For instance, Hallgren's efficient algorithm for the solution of Pell's equation can crack the Buchmann-Williams cryptosystem. This has led the cryptographic community to focus on post-quantum cryptography, which would be robust against quantum attacks. One potential example of post quantum cryptography is lattice-based cryptography. Interestingly, there exist subexponential time quantum algorithms to crack lattice-based schemes, though no efficient algorithms are known. This points to additional more subtle questions, which can and should be studied in more detail. There are many open fundamental and practical questions regarding the competition between quantum computers and alternative cryptographic schemes in a world where quantum computers are assumed to exist.

Now turning our attention to applications which go beyond quantum computation specifically, we return to our point made in Section 1 above regarding the importance of quantum networks. Quantum Key Distribution (QKD) is the most mature area of quantum information science but it is of limited application in current instantiations due to the exponential loss of rate with increasing distance, limiting practical systems to ~50 km range. Fortunately, a promising path forward is offered by quantum repeaters. This approach is much more nascent but differs from QKD in two important ways. First, it can surpass fundamental limits on secret key generation rates with distance, significantly altering the scaling with distance and enabling continental scale quantum communication. Second, they offer a way of not just sharing key but rather sharing entanglement among nodes. Shared entanglement can be used for much more general purposes, including distributed quantum computing and a host of interesting privacy protocols.

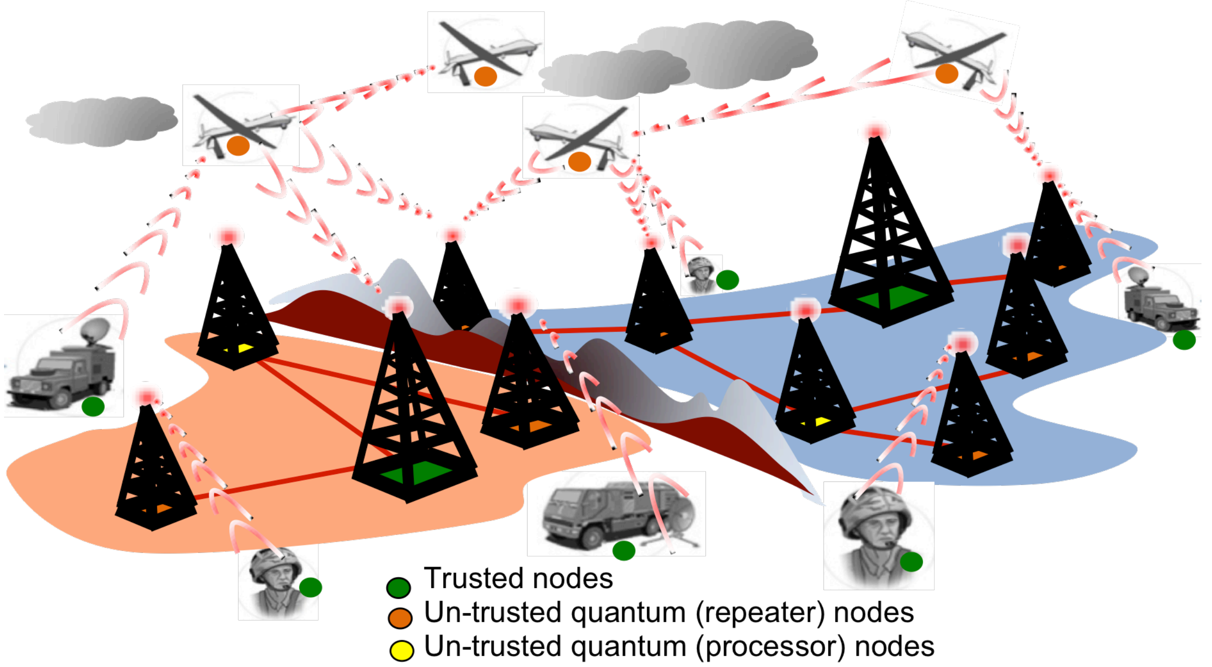
🟢 Trusted nodes
🟠 Un-trusted quantum (repeater) nodes
🟡 Un-trusted quantum (processor) nodes

Figure 1. Notion of a future quantum network.

The notional diagram in Fig. 1 demonstrates how a future quantum network could look. In it, one has several different flavors of nodes. Un-trusted ones are ones in which tampering can be detected (*e.g.* repeater nodes) while trusted ones are assumed to be secure against an adversary. The presence of un-trusted repeater nodes allows the distance between trusted nodes to go far beyond the short ~50 km distances imposed by photon loss. Non-processing repeater nodes would only require a scale of processing that could necessary entanglement purification and routing, and could be implemented with ions and integrated photonics with gate error rates commensurate with these protocols. Additionally, we envisage some of the nodes (both trusted and untrusted) would contain significant processing power, *i.e.* full scale fault-tolerant quantum computers. Processing nodes could be implemented with superconducting quantum processors with error rates supporting universal fault tolerant quantum computing. In addition to the hardware development needed to realize this vision, there is significant research needed to consider entanglement distribution and routing within such a network. The optimal protocols in this context will likely be quite different than that for point-to-point quantum communication and will need to borrow from well-established techniques in classical network theory to develop such protocols.

A quantum network such as in Fig. 1 could obviously support distributed quantum computing as well as long-distance and networked key distribution. However, these two things can and should be combined in a new powerful way to enable new privacy applications. For example, a protocol for quantum auctions was proposed several years ago which enables an auction over a network in which the losing bids are hidden from the auctioneer and all other bidding parties. Realizing this protocol requires both quantum communication between the parties as well as moderate quantum processing power at the nodes. As another example, a protocol for quantum private queries was proposed which enables one to probe a public data-base without the proprietor of the data-base knowing the content of the question (more specifically, the party asking the question could detect if the proprietor learned the content, in much the same way the parties in QKD can detect an evesdropper stealing secret key). These are but two examples of a large class of protocols combine quantum processors linked on quantum communication networks. They share a common theme that they enable a completely new physically based security on top of the mathematical complexity techniques which are used for similar tasks today. Given the unproven status of the unbreakability of mathematical techniques and the eventual implementation of quantum computers, this physics based approach is bound to be a fundamental ingredient to security in future systems and eventually of extreme economic importance. We believe focus on this class of protocols will be another strong motivator for investment in quantum computation and communication hardware in the coming years.

Finally, we would like to note quantum information has allowed some interesting new explorations into low power communication and imaging and possible enhancements to existing systems. As one example, we have discovered some new fundamental results and performed an initial demonstration regarding low probability to detect (LPD) communications, in which the very presence (rather than the content of) of a message is hidden. In particular, one can use the tools of quantum information to understand the detectability of a communication system making minimal assumptions on the specific hardware of the adversary (instead using basic fundamental physics laws to impose limits). Using this approach we have seen that within certain environments, undetectable reliable communication can be accomplished within a noisy environment. This line of research is a good example of a case where tools of QIS can be used to potentially enhance current commercial systems produced by Raytheon in the areas of communication and remote sensing. We believe further research in this direction will be an important complementary part of research towards privacy and security more generally.

3. Barriers

As discussed in Section 2, there is a strong need for more focus on algorithms beyond than the current “famous” examples of Shor and Grover algorithms. This activity is naturally picking up in the community, with the increased interest in quantum simulation and quantum network privacy protocols. This is a critical to increasing the interest in QIS beyond the research community which is already engaged and efficiently identifying some of the impacts QIS could have which are currently not known or understood. This becomes an increasingly important issue each year, as more time passes and the relatively small number of ‘killer apps’ becomes more apparent.

ARO’s funding towards quantum algorithms has been an important driver in moving this field forward. Many of the interesting angles referred to in Section 2 have been enabled by this and related programs. However, we see benefit in programs of larger scale and greater technical focus (akin to a typical DARPA or IARPA program) being very beneficial at this stage. We have been encouraged by the appearance of two BAAs out of NRO this year which explicitly mention interest in quantum processing applications to optimization problems of interest to that agency. Programs which are motivated by specific applications and problem scales can drive creative solutions which will likely uncover surprising benefits. In particular, they will provide a technical basis on which to resolve some of the issues referenced above, such as whether a polynomial speed-up is of sufficient benefit to justify implementation with quantum processors, or whether there is a set of constraints that would motivate us to consider the average rather than worst case performance within those constraints. In fact, we are somewhat alarmed that such a large fraction of the public discourse on quantum computing is focused on the question of quantum annealing to solve optimization problems. While this is a perfectly reasonable question to pose, it gains an outsize amount of public attention relative to the substantial amount of more rigorous analysis being done elsewhere. Government programs can be a key part of increasing the attention focused on important work going on in algorithms and focusing this work in appropriate ways towards appropriately specific problems and metrics.

Similarly, we would very much encourage programs focused on protocols for large-scale entanglement based quantum networks, towards the kinds of applications mentioned in Section 2 above.

4. Workforce Needs

Robust funding in focused programs in QIS can have a measureable positive impact on the trained work-force available for industrial research. As an example, the IARPA MQCO program had a broad impact on labs across the country working on various experimental qubit technologies. Students and post-docs coming out of the labs on this program have provided a strong pool of talent for Raytheon BBN and other industrial research institutions. The pool is still limited due to strong competition for that talent among several such organizations, but the impact of this program has still been very clear. In addition, funding some of these industrial research institutions directly has had a positive impact in giving them the resources to hire and further train students and post-docs professionally, particularly ones that are inclined to a career aimed more at engineering of large scale and robust systems rather than more academically focused research projects. We strongly encourage continued activity in this regard.

We have observed a significantly less robust pool of talent with respect to quantum information theory and quantum computer science. This issue seems to be partially driven by the current organization and scope of many academic departments, which often don't have a natural home for researchers engaged in the computer science and information theory side of QIS. The government could help this situation in two ways. First, borrowing from the successful example of MQCO, we believe government programs that are crafted with the issue of training an industrial workforce in mind could have a real, tangible impact. Second, the formation of funded centers within academic institutions (for example, through NSF and MURI programs) is a natural way to institutionalize academic homes which could help encourage and train young talent in this area. Accounting for the particular shortage in quantum computer science and quantum information theory talent in the U.S. when formulating priorities for such programs and centers would be very valuable.