Thinking Ahead to a World with Quantum Computers

The Landscape of Responsible Research and Innovation in Quantum Computing

Philip Inglesant, Mark Hartswood and Marina Jirotka
University of Oxford

NQIT Networked Quantum Information Technologies

UK NATIONAL QUANTUM TECHNOLOGIES PROGRAMME

NQIT (Networked Quantum Information Technologies) is funded by the Engineering and Physical Sciences Research Council and is a part of the UK National Quantum Technologies Programme

EPSRC
Engineering and Physical Sciences Research Council

UK NATIONAL QUANTUM TECHNOLOGIES PROGRAMME

Authors: Philip Inglesant, Mark Hartswood and Marina Jirotka

Design & layout: Hannah Rowlands, based on a design by Hunts

DEPARTMENT OF COMPUTER SCIENCE

UNIVERSITY OF OXFORD

Table of Contents

Executive Summary

Responsible Research and Innovation (RRI) is a core work package in the Networked Quantum Information Technologies (NQIT) project, one of four research hubs funded as part of the UK National Quantum Technologies Programme. RRI aims to promote research and innovation which is socially desirable and undertaken in the public interest, stimulating creativity not only in the science but also in innovation arising from it, and ultimately leading to products and services that are accepted by the public.

Quantum technologies are advancing rapidly, in some cases already in the market, or becoming close to market. The main strand of NQIT, and the main focus of this report, is on developments towards quantum computing, which, in its full sense, is still some years from availability.

This document does four things:

- Provides a background to RRI within NQIT (Chapter 1 and Chapter 2);
- Identifies the challenges for Responsible Research and Innovation in NQIT, from interviews with participants in the project, funders, industrial partners; and from a review of quantum computing and RRI literature (Chapter 3);
- Makes recommendations about how to handle these challenges (throughout Chapter 3 and Section 4.1 - Section 4.3);
- Describes a framework and pathway to implement a tailored RRI process in NQIT (Section 4.4).

Steering the way between overstating the potential benefits of quantum computing and too much caution about the potential risks, it sets a course for NQIT to become established as a trusted source of knowledge and innovation in quantum technologies.

Background to RRI in NQIT

The Landscape starts by describing some of the challenges for RRI in NQIT from the lessons of earlier technologies, to recognise the importance of being honest about scientific uncertainty, the risks as well as the benefits from quantum technologies, and taking seriously public concerns, to avoid any deep seated "quantum phobia" taking root.

Chapter 2 gives a brief introduction to the principles of Responsible Research and Innovation, and shows how these relate to modern technologies in our increasingly "digitally hybrid" lives. With newly emerging technologies such as quantum computing, the eventual outcomes of innovation are hard to predict, and may entail unwanted and unanticipated side effects; powerful technologies harnessing the fundamental fabrics of nature - nano-, synthetic biology, quantum - may raise fears of serious and irreversible consequences.

Against this backdrop, RRI has been developed to help to ensure that the products and outputs of research and innovation are socially acceptable and desirable. Embedding RRI in research projects is increasingly encouraged by funders, and particularly by EPSRC, and is an integral part of the UK National Quantum Technologies Programme.

RRI has been developed to help to ensure that the products and outputs of research and innovation are socially acceptable and desirable.

EPSRC has adopted the Anticipate – Reflect – Engage - and Act AREA framework (Figure 1) as a practical activity for RRI within funded research programmes: a Responsible Innovation approach should continuously seek to: Anticipate the impacts, intended or otherwise that might arise – Reflect on the motivations for the research, the assumptions, dilemmas and social transformations these may bring – Engage in inclusive deliberation, dialogue, and debate – and Act responsively where appropriate to influence the direction and trajectory of the research and innovation process.

Foundations for RRI in NQIT

In developing a framework for RRI in NQIT, we suggest the following general principles, which cut across the dimensions of the AREA framework and which form the basis for RRI in the context of quantum computing and for refining practices of RRI throughout the lifetime of the Hub:

- All researchers should have the opportunity to engage in RRI related activities and the support to do so. Supporting the broadest possible engagement will embed RRI more deeply throughout the Hub.
- Elements of this framework should be connected up and down through any hierarchy. Thus, a senior scientist giving advice to a government committee should have access to the digests of more local RRI activities. Similarly, scientists as a group should be kept informed about the committee work and its outcomes.
- RRI should be thought of as interconnected responsibilities that are distributed across various groups of stakeholders, not all on the shoulders of scientists. Some key partnerships will be with social scientists, who can help with understandings of how science and society interact. Others will be with RRI specialists who will play a facilitator role. They will have some technical knowledge and can help initiate, coordinate and join up RRI activities.
- An RRI framework has to span between science and research activities, communication, commercialisation pathways, policy formation and funding directions, reaching across the research and innovation ecosystem.
- Liaison across the Hubs will enable each Hub to benefit from the experiences of other Hubs. Many of the issues are common across the Hubs; others are unique, and the comparison is itself powerful. A cross-Hub strategy would maximise the impact and minimise duplication of effort.
- Dialogue with the public, with early adopters, with civil society, and other stakeholders is crucial to ensure that the research delivers results which will be widely recognised and welcomed. Public communication exercises should aim to encourage dialogue, in specifically tailored activities such as workshops, and also in outreach events of all kinds.
- In comparison with RRI in some other areas of science and technology, there has been less attention, to date, given to quantum technologies. NQIT will start to redress the balance and gives a unique perspective to RRI in quantum as an emerging technology.
- The RRI framework in its own turn must follow the precepts of responsible research, continually evolving, reflective and responsive to changes in technology and society and incorporating new insights.

RRI challenges within NQIT

The areas for focus of RRI in NQIT have been identified from a series of interviews with project members, workshops, and a review of quantum technology and related literature; these key findings are presented in Chapter 3.

We have classified the focus areas into those that:

- are "**live**" and need addressing immediately – with an overarching theme around the **narratives** – positive and less positive - of quantum computing
- are on the **horizon**, are areas of active research now, and will become live as quantum computing technologies come into being – with an overarching theme of the increasing need to **trust** the results of quantum computing, and different forms of trust enabled by it; and
- may arise in the **future** as quantum computing produces applications and reaches widespread deployment - and enables **social transformations**.

How can we talk about Quantum?
Finding a way to explain quantum phenomena to school children at the Cheltenham Science Festival
/ Hannah Rowlands

Live, immediate issues

Section 3.1: these relate to the ways in which quantum computing and quantum as a whole is perceived by the public and others, rather than being focussed on specific applications. Here we include narratives around quantum mechanics and quantum computing going back to their origins. At this stage, the overwhelming issue is simply a high level of uncertainty; it is all the more important to communicate effectively about quantum as a whole and quantum computing.

1. **The uncertainties associated with realising quantum computing.** There is still some theoretical uncertainty as to whether quantum computers can actually be realised. Experimental and theoretical work is increasingly overcoming these uncertainties, but scientists and practitioners still cannot be sure of the timescale, what the capabilities of quantum computers will be, and ultimately what applications and social implications these will lead to.
2. **The character of quantum computing.** Quantum computing is not strictly analogous to classical computing, but rather has a very specific mode of operation. Quantum Computers challenge the limitations of classical computing, but they are physical devices, with physical limitations, and so quantum computing is re-focussing research on the physical underpinnings of computing, with implications for classical as well as quantum.
3. **How we can talk about Quantum?** Is the easily-misunderstood quantum "spookiness", and "spooky action at a distance" from Einstein still useful? It captures the counter-intuitive nature of quantum phenomena but also contributes to the impression of quantum mysterious and hard to understand. Is there a risk that this will lead to negative perceptions of quantum?

The overarching theme is the narratives and discourses which surround quantum computing; narratives around uncertainty and the still-uncertain character of quantum computing combined with older narratives about quantum as strange and hard to understand. This opens a space for negative counter-narratives, but is also an opportunity to develop a positive dialogue with the public.

At this stage, the overwhelming issue is simply a high level of uncertainty; it is all the more important to communicate effectively about quantum as a whole and quantum computing.

Recommendations to address the uncertainty and emerging narratives of quantum computing

1. **Adopt a strategy of being honest about the uncertainties surrounding quantum technologies**, but at the same time to find ways to articulate an informed understanding of what is likely to happen and when. We suggest:
   - A quantum application "clock" showing a realistic timeline indicating the timescale between scientific advance and its application. For example, in medical advances a time of 10 years is often suggested between the lab and a new treatment.
   - Uncertainty should never be used an excuse to delay or avoid dialogue; this can create a vacuum in which others can shape the quantum narrative, perhaps in unhelpful ways. Take the opportunity for mutually informative dialogue with all stakeholders, throughout the lifecycle, not only closer to market.
2. **Articulate the very specific character of quantum computing**, both its powers and limitations, risks as well as benefits, as distinct from classical computing. This is an area ripe for misunderstanding and where powerful mythologies can take root – such as the notion that quantum computing can solve any problem currently too hard for a classical computer.
   - It may be useful to develop "Frequently Asked Questions" on points of difference between classical and quantum computing.
   - Design a poster or pamphlet or video describing in simple terms how a problem is represented in a special way before it can be run as a quantum program.
   - It might be useful to explore ways in which an artist in residence can help to express complex scientific ideas; this has proven effective in other RRI projects.
3. **Develop and adopt a language with which to discuss quantum phenomena** in a way which is faithful to quantum effects but which avoids introducing elements of mystification.
   - Use "counter intuitive" as a much less loaded term than "spooky"
   - Emphasise the effectiveness of quantum mechanics underpinning the engineering of existing everyday technologies.

Issues on the horizon

Section 3.2: these revolve around concrete implementations of the technology to support actual applications. There is much active research in these areas, while still not yet being applications in themselves. The timescale to achieve quantum computing stretches from computing-related technologies already starting to appear (for example, Quantum Key Distribution), technologies a few years downstream (for example, quantum simulation), to decades before high-profile algorithms such as Shor's become implementable beyond a small scale. Three key issues we identify relate to:

1. **Strong claims for quantum technologies**. The laws of quantum mechanics can in principle be harnessed to provide theoretically unbreakable communications, truly random numbers, or perfectly trusted forms of computation. Yet on the other hand implementation and engineering challenges limit how far any "strong guarantees" can be applied in practice.
2. **Verification.** How do we know that the results are correct, if quantum computing performs functions which cannot be replicated classically? How can we be sure of the correctness of secure communications and quantum servers against malicious interception? How can the claimed speed-ups for non-universal quantum computers such as D-Wave, and the "quantumness" of their operations, be verified?
3. **Other medium-term applications.** NQIT has work packages in areas such as sensors, sensor nets, and quantum simulation. While not quantum computing per se, these are already emerging as the first direct applications of quantum effects for computation and computing-related. They will provide early outputs for NQIT, making visible the real progress which has been made in this area; there are some high-profile applications such as "seeing" objects underground, detecting gases or pathogens, and very accurate navigation, which have clear implications and will help us to discuss emerging issues.

The overarching theme across these issues is the need to trust quantum computing and quantum technologies on a number of levels. Potential users of quantum computers are being asked to trust the results of computations which are difficult to verify, even "in principle". And beyond the immediate users of quantum technologies, there will be indirect users who will also be trusting in quantum results, affected by them but unable to verify results of which they may not even be aware.

Potential users of quantum computers are being asked to trust the results of computations which are difficult to verify, even "in principle"

Recommendations to address these issues on the horizon

1. **Consider the implications of the potency of quantum applications.** There is a risk of a "fall from grace" if expectations are set too high. Managing the combination of potency and fallibility will require a balance between enabling the advantages of powerful new capabilities, and ensuring that these capabilities lead to socially desirable outcomes
   - Make careful use of language to set the right expectations. The message is that there are potentially great benefits, but there remain challenges in implementation.
   - Recognise that there will be disappointments but also unexpected successes.
2. **Work to ensure trust in quantum computing by the users and by third parties who rely on these results.** Building this broader trust will depend on a continuation of the engagement work towards developing positive social narratives around quantum computing
   - Verification could be an area for misunderstanding; make sure that the verification work is not focussed solely on technical audiences, but reaches out to broader constituencies.
   - Work towards demystifying quantum computing, including demonstrations of quantum computing capabilities in ways that are meaningful for broader audiences.
3. **Take the time for learning and refining appropriate RRI processes.** Be honest about the risks as well as the benefits – which are two sides of the same coin. Early applications of NQIT technologies will be harbingers of universal quantum computing technology, presaging the social transformations that may follow.
   - Issues in quantum cryptography, sensors and simulation may feed existing controversies, intensify existing risks or upset the balance of interests within society. There is a risk of public resistance.
   - Conversely, successful management of these technologies will cement trust in quantum computing.

Future challenges

Section 3.3: these will emerge as applications of quantum computing precipitate societal transformations. While it is impossible to predict precisely what these will be, we can identify socio-technical trends that carry risk for society, as well as potential benefits, and explore what might happen if quantum computing were to be added to the equation. We cannot predict risks and controversies from entirely unexpected quarters, but we can continually anticipate and develop our sensitivities to the ways in which quantum computing may interact with existing trends.

Drawing on contemporary issues and sensitivities attached to innovation and technology, we have identified the following areas for likely challenges:

1. **Machine learning in the context of an emerging "algorithmic society"**. Machine learning underlies many of the algorithms that play increasingly powerful and important roles in governing our society. What greater influence might be yielded by quantum-enhanced machine learning?
2. **Defence and national security** are likely to be early adopters of quantum technologies, including computing. There are applications for secure communications, interception, and enhanced navigation and sensing. These uses add to fears of a "Surveillance Society". Changes in the capacities of different countries could have implications for geopolitics, and interact with current controversies around military technologies such as robots or drones.
3. **Ownership and access to quantum technologies.** Only large, state-level or very large corporations and research laboratories are likely to be able to sustain the resources required to operate quantum computers, at least for the foreseeable future. One effect may be to cement or increase imbalances of power between these powerful actors and ordinary consumers and citizens.

Each of these areas of social transformation raises implications for RRI. Algorithms and the "algorithmic society" have real-world implications (think about the power of a Google search); defence applications raise questions about the relative powers of the state and the citizen, and of one state relative to others; domination of "Big Technology" and "Big Science can create an aura of secrecy and may feed public disquiet.

As quantum computing intertwines with other technologies such as Internet, mobile communications, and the Internet of Things, it will become more complex and embedded in institutions of society

Recommendations for RRI to address social transformations

1. **Pro-actively encourage collaboration between scientists and engineers and RRI practitioners.** The strong potency of quantum technologies "on the horizon" will be made explicit in applications which will emerge as the technologies are appropriated by the market and by users. We cannot predict exactly what these applications will be, but we can foresee some areas of concern and possible tensions between diverse interests.
   - RRI should work formally and informally with others in the project to match social provisions to new technologies, identifying potential areas of conflict and developing "mid-stream modulation" to address these.
   - This should involve scientists themselves, RRI researchers, funding councils, the leadership of the project, and potential early adopters in open dialogue and interaction to resolve differences of interest between stakeholders.
2. **Start to develop forms of governance for these technologies.** As these applications emerge into regular use, it will be increasingly necessary to create the right institutions, practices, regulations, and fit with project structures.
   - At some stage, regulation may become necessary, to address specific risks, but such an approach is top-down and retrospective.
   - What is needed now is a mix of upstream, mid-stream, and, as the technologies move on, increasingly downstream adaptive governance.
   - Governance is not the same as government: it is more flexible, less top-down, and works with the existing structures rather than imposing new regulations.
3. **Use the RRI framework as a resource for acceptable outcomes.** Early applications of NQIT will be pointers to social transformations from quantum computing. As it intertwines with other technologies such as Internet, mobile communications, and the Internet of Things, quantum computing will become more complex and embedded in institutions of society.
   - This embedding should reflect the ethical acceptability, sustainability, and social desirability of the products of innovation
   - Towards this aim, RRI is a resource for creative thinking and adaptation as technology becomes more embedded in society.

Putting RRI into practice in NQIT

As well as the RRI implications of the outputs of NQIT, the Landscape looks at the features of NQIT as a project and how RRI fits into them. Section 3.4 draws on our interview data to explore some of the responses to RRI within NQIT, both positive and negative, and highlights the challenges of communicating the tenets of RRI and tailoring RRI to fit with the structure and objectives of NQIT.

NQIT and its relevance for RRI

- NQIT is end-to-end research, taking cutting-edge science and using it to build working systems. Driven by the overall goal of developing a working demonstration of a scalable quantum computer, together with medium-term spin-offs, NQIT is breaking down barriers between fundamental and applied research, and between Physics, Materials Science, Engineering, and Computer Science.
- Even fundamental research has social and ethical implications. **NQIT includes fundamental research but this is driven by impetus for a working system**. For researchers, this means that the project will lead to clear outcomes, and, they hope, will make the world a better place.
- There are many uncertainties about how these innovative technologies will make their impact in practice, which makes it especially hard to implement RRI. But other technologies are also difficult to understand and unpredictable in their impacts. RRI does not demand prediction, but anticipation – developing the capacity to think about what might happen, and to be prepared for future eventualities.

NQIT is breaking down barriers between fundamental and applied research, and between Physics, Materials Science, Engineering, and Computer Science.

Scientists and others participating in RRI in NQIT

- RRI should be thought of as a package of interconnected responsibilities that are appropriately distributed across various groups of stakeholders, not all on the shoulders of scientists. Scientists have a special responsibility from their expertise and deep understanding of what may and what may not be possible from their research; but decisions about areas of concern – such as privacy, security, or surveillance - are social and political, and need to be addressed at that level.
- RRI involves scientists at all levels. RRI needs a diversity of perspectives; RRI is trying to break down the "moral division of labour" between research strategy at a high level and the ground-level practices of research. Thus, the broad expertise of senior scientists is highly valued for high-level policy direction, while for more junior researchers developing the capacity and the skills of reflection and anticipation at an early career stage will reap benefits at all levels.
- Many scientists will engage strongly with RRI, while others may struggle to reconcile RRI with commitment to rigorous science. But a concern for responsibility is part of rigorous science, not inimical to it; responsible research should be part of good scientific practice.

RRI in current practices

- RRI offers tools for anticipation and reflection, but also embraces “de facto” responsibility activities which may not be labelled as such - such as informal discussions with colleagues over coffee, quantum science as a cultural activity, and outreach through science events and the press. Often, it will not be a question of “doing things differently”, but of maintaining an interested and aware attitude. Members of NQIT are already practicing RRI in this sense.
- RRI is not a wholly new set of practices; it builds on decades of research into the history of science and innovation. This research has shown that the path from science to innovation and impact on society is rarely simple and linear. Science moves mostly in incremental improvements, and works alongside technological and social innovations.
- RRI should not be framed as an “impossible” task such as predicting the future. Nor is it about removing all risk from innovation; failing to innovate introduces the risk that we will lose the benefits of innovation. Looking forwards, we can consider how quantum computing may play into existing trends and controversies, as well as opportunities. Indeed, the User Engagement programme in NQIT is already looking ahead and identifying application areas. This is an opportunity to enter into two-way dialogue to explore these implications.

Explaining our new technology development at the Quantum Showcase at the Royal Society / *Dan Tsantilis, EPSRC*

A framework for RRI in NQIT

Chapter 4 outlines our framework for continuing to conduct RRI in NQIT and describes the next steps. In preparation for the framework, we have identified a number of challenges and enablers for RRI in the specific context of NQIT, which summarise the findings in Chapter 3.

There are **challenges** in anticipating future social transformations when there is still much uncertainty around the form that quantum computing will take. The distance between active research and downstream applications, leads, in some cases, to doubts about the value of RRI. Looking ahead, it is important that social issues, such as a drift towards an algorithmic society, are not overshadowed by more high-profile issues, such as Shor's algorithm "breaking the internet".

But there are also **enablers** for RRI. There are positive portrayals around how the science is being channelled into societal benefits. In general, the quantum computing community, supported by professional organisations and associations, is aware and interested in the social implications of their work, and many researchers are motivated to see their work have positive impacts.

The framework for RRI in NQIT starts from the AREA framework but is attuned to NQIT. As a focussed project with a clear management structure, we are able to be quite specific in NQIT about the issue areas and the stakeholders; we can engage with stakeholders through events such as the project forum, and we can have a part in the outputs and make recommendations to the project leaders.

Although the framework is still a work in progress, we set out a series of practical actions to be developed further but already taking place as we practice RRI in the Hub. This forms a "toolkit" to identify areas of concern and suggest ways in which these might be addressed.

These activities include those we have already been doing for the Landscape study:

- interviews and workshops with project members
- review of quantum technology and related literature
- liaison with other Hubs

The framework should also open new spaces such as:

- workshops and case studies on key issues
- engagement with industrial partners and other potential early adopters,
- structured dialogue with the wider public
- engagement with civil society and other key communities
- risk assessment and technology assessment methods
- informal and formal structured foresight exercises

And in response to the framework, NQIT should be ready to act to:

- shape the direction of innovation, and to
- encourage responsibility as a part of professional research practice.

Final recommendations and next steps

In Section 4.3, we make overall recommendations for how to continue with RRI for the future of the NQIT programme.

4 **Continue to refine the framework for RRI in NQIT.** We have outlined a framework which will embed RRI and carry it forward. Further RRI work will refine the framework; we will do this as a form of action research, so that using the framework, to analyse RRI in NQIT, will at the same time strengthen it.

5 **Build capacity for RRI in NQIT.** This could involve self-directed study materials, workshops, posters, seminars, etc., supported by an RRI community – "RRI Champions" in NQIT, from researchers at all levels and people working in communications and technology transfer. This also includes forms of "de facto" RRI such as informal discussions, and activities, for example the current development of post-quantum cryptography, which act to address emerging issues.

6 **Include RRI in the structure of NQIT.** RRI will show the most benefits if it is supported by appropriate, well defined but not over-complex governance structures. In NQIT, RRI already has the effective status of a work package; RRI produces quarterly reports for the NQIT Management Board and contributes to the annual report. Increasingly, these reports will identify substantive issues, which may ask for a response from the project management.

7 **Disseminate RRI through channels to influence research policy.** We have identified a need fore a more defined route to enable the findings of RRI to reach those who are in a position to act on them. These channels – acting upwards as well as downwards - could take a number of forms: a repository of issues, co-ordinated across the Hubs, would be a useful technical resource.

8 **Engage in dialogue with the public and stakeholders.** This is not a new point but needs to be emphasised here. As well as interaction in public outreach, the kind of measured, interactive dialogue which will give a broader picture of the public perception of quantum technologies will require a wider pool of participants and a more structured approach.

The next steps for RRI, which we outline in Section 4.4, will bring stakeholders together for deeper investigation of the issues we have identified and presented in Chapter 3 and the challenges and enablers laid out in Section 4.1, This will be firstly through a set of case studies, centred on structured workshops, covering the dimensions from time distance to technology readiness, along one axis, and the contribution of quantum from "quantum-enhanced" to "could not be done without quantum", on the other axis. A workshop around the issues raised by the interest of defence in quantum computing is already in preparation.

Chapter 1: Introduction

This Overview of the Landscape of Responsible Research and Innovation (RRI) in NQIT is the outcome of a scoping exercise that forms part of the RRI strategy agreed by the NQIT Management Board in July 2015 (Inglesant et al., 2015). This activity has engaged the NQIT community in mapping the range of issues that may emerge in relation to the social acceptability of quantum information technologies. This report is based on a range of data sources, including thirty-one semi-structured qualitative interviews with NQIT researchers, directors, and Work Package (WP) leaders, funders, and innovators and commercial partners; a series of four scenario-based workshops and "RRI Roadshows"; a selected set of forty-three project-related documents; and over one hundred other articles on quantum technologies and computing.

Networked Quantum Information Technologies is one of four Hubs funded as part of the UK government's National Quantum Technologies Programme, a major public investment totalling £270 million. Quantum technologies are anticipated to have a profound impact on major industries worldwide. The Programme sets out an integrated approach across academic and industrial communities to consolidate the UK's strengths in quantum research.

Despite the best of intentions in terms of economic benefits and future wellbeing, it is also understood that new technologies may lead to questions, dilemmas, and unintended (and sometimes undesirable) consequences (Owen et al., 2013).

Recognising these challenges, RRI is an emerging approach aiming to ensure that research outputs and processes are not only ethically accessible and sustainable but also socially desirable (Von Schomberg, 2013). Going beyond unintended consequences, RRI is as much about what we do want science and innovation to do as about what we do not want them to do.

Quantum may be compared with other technologies that also work with "fundamental fabrics" – such as genetic modification, nanotechnology or synthetic biology, which can spark deep-seated fears about what these technologies may unleash. Such fears can delay or even derail research directions; for example, public concern contributed to the pausing of the Stratospheric Particle Injection for Climate Engineering (SPICE) project (Cressey, 2012, Macnaghten and Owen, 2011) – an experiment in geo-engineering which sparked protests from civil society organisations concerned that a technical fix would deflect political and scientific action away from greenhouse gas reduction. In an example widely cited in the RRI literature, public disquiet about genetically modified organisms (GMOs) and GM foods has significantly impeded the exploitation of GM

Responsible Research and Innovation (RRI) is an emerging approach aiming to ensure that research outputs and processes are not only ethically accessible and sustainable but also socially desirable

foods in Europe. One target of RRI in NQIT should be to recognise and avoid the mistakes of these earlier programmes of innovation to avoid any deep seated “quantum phobia” taking root. Lessons learned include the importance of building trust, being honest about scientific uncertainty, and taking seriously public concerns.

Quantum Information Technology also intersects with the context of our increasingly “digitally hybrid lives” (Jordan, 2009), where there are increasing concerns over how power and social control may be directed through digital networks. Thus, quantum computing, with its promise of powerful and potentially disruptive applications in cryptography, cryptanalysis, machine learning, simulation and sensing, emerges into an environment already full of fermenting controversies revolving around the use and ownership of big data, existential threats from Artificial Intelligence, military use of robots and drones, erosion of privacy and algorithmic control of society, amongst others. Quantum technologies offer an enormous range of potential applications and, in some examples, performance improvements of several orders of magnitude over conventional technologies (Pritchard and Till, 2014); many of these applications have clear military uses. Recently, the potential of quantum computing has been linked to the increasing fears of a Surveillance Society in a contested and highly polemic article (Cribb, 2016). Thus, a second target of RRI in NQIT should be to understand how the narrative of quantum technologies may unfold within the public sphere, and to provide resources to support scientists to engage in shaping that narrative in partnership with other societal actors.

The programme of RRI within NQIT reflects the EPSRC’s strong commitment to RRI as part of the UK National Strategy for Quantum Technologies (UK National Quantum Technologies Programme Strategic Advisory Board, 2015). As the strategy document makes clear, rather than being a barrier to innovation, RRI "potentially enriches the process and improves the chance of commercial success by stimulating creativity; informing standards, regulation and governance; and ultimately allowing products to be developed that are more likely to be embraced by the public".

Our strategy for RRI in NQIT focuses on the following:

- reflecting on the emerging social narrative of quantum computing and its applications to understand and anticipate where the issues may lie;
- moving forwards with the research in dialogue with publics, policy-makers, stakeholders in industry and the defence sector to build dialogue and trust;
- acting to identify and ameliorate potential hazards and to address concerns while it is still possible to do so.

1.1 Organisation of this report

The rest of this document is organised as follows. Chapter 2 discusses the historical background of RRI and why it has emerged as a pressing concern for the 21st century.

Chapter 3 presents the results from the analysis of interviews with participants in NQIT, with insights from quantum and RRI literature to identify properties of quantum technologies and the socio-technical issues which can be anticipated from them.

Finally, in Chapter 4 we make concrete recommendations for action and suggest areas for the focus of case studies which will deepen these insights during year 2 of NQIT. These actions include the development of a framework for RRI in NQIT based on the AREA framework, and suggestions for processes to maximise the impact of our RRI research on future governance and public policy for quantum technologies.

Chapter 2: Responsible Research and Innovation

Research and innovation is usually conceived with good intentions. Modern technology has transformed the quality and security of our lives. But there are unexpected and sometimes negative consequences, made all the more challenging by the ubiquity of technology in the modern world. Responsible Research and Innovation is emerging as a set of ideas and processes to address these challenges, and has been adopted by funding bodies and is an integral part of the UK National Quantum Technologies Programme, as set out in the document National Strategy for quantum technologies: A new era for the UK (UK National Quantum Technologies Programme Strategic Advisory Board, 2015).

Responsible Research and Innovation (RRI) has emerged over the last decade, initially to address risks associated with some novel areas of research including nanotechnology (Fisher and Rip, 2013), environmental sciences including geo-engineering (Stilgoe et al., 2013), and synthetic biology (Tucker and Zilinskas, 2006). The aim is to ensure that science and innovation are undertaken in the public interest by incorporating methods to involve affected stakeholder groups. More recently, RRI has expanded into fields such as Computer Science, Robotics, Informatics and ICT in general.

2.1 The challenge of responsible research and innovation with technology

Modern technology is inherently "big", affecting the lives of millions. We are often dependent on it for the infrastructure of our lives. At least in developed countries, it is no longer within our power to choose whether or not to engage with technology, underpinning our everyday activities. The very "long run" of technology creates a risk of "too much", not only in failure, but even in success (Jonas, 1982).

At the same time, the eventual outcomes of innovation are hard to predict, and may work in counterintuitive ways to produce greater risk via unwanted and unanticipated side effects. Examples from processes of industrialisation include pollution, nuclear waste and climate change. Even the current benefits of the Internet age such as much wider access to information also introduce problems of digital divides, erosion of worker rights in emerging digital labour markets, disruption to established business models and centralisation around powerful companies, and new forms of vulnerability to critical infrastructures via cyber-attacks.

As our sciences are becoming more powerful, seeking to harness the “fundamental fabrics” of nature (nano-, synthetic biology, quantum) for society’s ends, this inevitably raises fears of potentially hazardous consequences – some of which could undoubtedly be real, but others which may have little foundation or are exaggerated.

Against this backdrop, RRI has emerged to help to ensure that the products and outputs of research and innovation are socially acceptable and desirable. RRI aims to anchor science and technology in society’s values – at the same time, reflecting and deliberating on what those values should be. In this way, rather than closing down options, RRI can be a creative resource for thinking anew about the potential of innovation.

FIGURE 1: DIMENSIONS OF RRI: THE AREA FRAMEWORK DEVELOPED FOR EPSRC

(A)nticipation

We cannot know what the future holds. We **do not aim to predict the future**, but if we do not make plausible attempts to look forwards, the field will be left to the most powerful actors or to “linear extrapolators and hype-or-horror prognosticators” (Guston, 2013).

Anticipation asks questions such as: “what if?...”, “what else might it do? ...” (Owen et al., 2013), to open up the space of possibilities and to help us to consider what these might be.

(R)eflection

In contrast to anticipation, which is outward-looking and free-ranging (within the bounds of plausibility), reflection invites stakeholders to think about the **purposes and motivations** for research, and the associated uncertainties, assumptions, framings, and dilemmas, and the **social transformations** which may follow from innovation..

Deliberation and (E)engagement

Experiences have shown the importance of engaging with all groups of stakeholders. This is more than, communicating the results of research and innovation to a wider public. Deliberation is **inclusive**, opening up visions and impact to questioning, recognising that other insights are valuable.

Responsiveness and (A)ction

Responsiveness entails listening, being open to questions, readiness to **rethink goals** and strategies and so adapt the **trajectory and pace of innovation**; it is **dynamic** and iterative.

Figure 1: Dimensions of RRI: the AREA framework developed for EPSRC

2.2 RRI in practice

RRI coordinates a varied range of multi-level activities undertaken by multiple actors across the research and innovation lifecycles. These activities may include traditional ethics processes, risk assessments and foresight procedures, as well as activities tailored to specific domains or research streams that emphasize multi-stakeholder involvement (Stahl, 2013).

RRI has been incorporated into funding bodies' research agendas. The National Strategy (UK National Quantum Technologies Programme Strategic Advisory Board, 2015), of which NQIT and the other quantum hubs are a part, explicitly includes Responsible Research and Innovation. For the European Union it is an integral part of the Horizon 2020 programme, as targeted RRI actions are integrated within other research themes (European Commission, 2016a). In the UK, the Engineering and Physics Sciences Research Council (EPSRC) expects (but does not mandate) that RRI activities are built into research processes, offering the AREA framework[1] as guidance for researchers wishing to pursue an RRI approach (see Figure 1).

It is reasonable to consider at which stage in the research and innovation cycle is RRI most useful. If at the earliest stages of innovation, change in research trajectory may be relatively easy and have most impact; if later, it may be easier to identify the social implications of innovations, but the technology may have become so embedded in society that it is very hard to make any changes apart from retroactive regulations (Collingridge, 1982). Some research has suggested that during the middle stage between these ends there is a space for influencing the trajectory of outcome with more certainty than at the very earliest stages of development (Fisher et al., 2006).

RRI tools have been identified to be adopted at all of these stages:

Upstream

- Starting early in the research and innovation process.
- Anticipating the transformations and impacts of new products and processes.
- Giving a voice to multiple publics and stakeholders to explore the consequences of research and its desirability.

Midstream

- Being responsive to the dialogue.
- Adjusting the trajectory of research, to avoid negative consequences and take advantage of creative possibilities.

Downstream

- Creating the right policy and regulatory environment for the emergence of the technology.

RRI is not separate from the research and innovation process, but is integral to it; rather than adding more top-down governance, RRI links with existing principles and activities, research ethics and professional codes, but broadens these to engage a variety of stakeholders. RRI is also an active area of research in itself. The outcome of this research in NQIT is a framework for RRI specifically tailored for the Hub. Part of the work of RRI research in NQIT will be to refine the shape and contents of this framework, as we describe Section 4.2 and in the recommendations in Section 4.3. This is a set of tools and activities, using the structure of the AREA framework; in practical terms, this includes internal and outward-facing workshops, using outreach events

1 https://www.epsrc.ac.uk/research/framework/area/

to listen to public concerns, more structured dialogue with the public, and building links with external stakeholders such as potential early adopters and civil society. The process of shaping the framework will be on-going, as a form of "action research" – doing RRI, in practical terms, at the same time as refining the framework. As the National Strategy (UK National Quantum Technologies Programme Strategic Advisory Board, 2015) points out, the "the UK has an opportunity to produce the first comprehensive public perspective on quantum technology". In bringing RRI into NQIT, therefore, we are breaking new ground, not only for NQIT but also for RRI. In NQIT, RRI will address the specific properties of quantum computing where it will have to meet challenges including the following, which are drawn from our analysis which we present later in Chapter 3.

There is still much uncertainty about when quantum computing will be realised, and what form it will take. Quantum computing will raise new questions of verification, of computational results and trust in applications based on new computational principles.

Related to the above point, Quantum applications whose properties are underwritten by the laws of nature (for example, to provide "unbreakable" cryptography), will nevertheless be subject to errors in implementation; a responsible attitude recognises that technology in practice is always intertwined with human and social factors.

Some possible applications of quantum computing may raise the urgency of existing political debates in areas such as defence and dual use, machine learning and big data.

Recognising that quantum technologies, as a major area of innovation, carry correspondingly large uncertainties alongside opportunities and potentially profound impacts, our aim is to build capacity within NQIT to develop those technologies in line with principles of RRI.

NQIT is in a good position of having dedicated RRI research resources who can lead events such as this RRI workshop during a Project Forum
/ University of Oxford Department Engineering Science Media Unit

2.3 RRI as a collective responsibility

The prospect of implementing a programme of RRI may initially appear daunting because RRI appears to imply new and far reaching responsibilities that go beyond the capacity of any one individual to discharge. In reality, the process of RRI is less mysterious than it seems, and its responsibilities less onerous that one might fear. RRI often involves building upon existing informal practices, such as the reflective moments in team meetings where discussion touches upon the way quantum technologies may be used, and in this way "naturalising" RRI practices (Grimpe et al., 2014). And its responsibilities are not focussed entirely on any one individual, but rather distributed across networks actors within scientific establishments and beyond. RRI configures responsibilities between different actors as we outline below:

Specialist researchers in RRI: NQIT is in a good position of having dedicated RRI research resources who can act as an anchor point to identify how best to implement RRI activities across the Hub, and bring these together around events such as workshops and training

Scientists: Reflecting on their research, using their expertise to inform dialogue around quantum technologies, being prepared to adapt in response to the process. Some of these will be senior researchers who are also involved in policy formation and advisory boards.

- **Funders:** have a role to continue to support RRI throughout their funded projects, and specifically the Quantum Technology (QT) Hubs, identifying what is being done and addressing any gaps
- **QT hubs:** Taking a leadership role in the development and dissemination of an RRI framework and suitable QT-specific RRI tools for practical engagement.
- **Universities:** Acknowledging the importance of RRI and provide resources and guidance for their researchers.
- **Stakeholders**, including non-governmental organisations, to engage early and in an open manner.
- **Industry** to realise the commercial benefit of engagement with the RRI agenda.
- **Government** to include RRI in funding mechanisms and facilitate the production and adoption of suitably flexible governance mechanisms.
- **Media** to engage with researchers and provide a platform for open and balanced debate.

In our presentation of key findings around RRI in practice in NQIT in Section 3.4, while there is a concentration on scientists and research leaders in NQIT, it is important to bear in mind this interplay of many different kinds of actors.

RRI often involves building upon existing informal practices, such as the reflective moments in team meetings where discussion touches upon the way quantum technologies may be used

2.4 NQIT and Quantum Technologies

Devices which underpin many of the technological advances of the last century – including lasers, transistors, semi-conductors – rely on the effects of quantum mechanics, which by then were already well understood. Their development is sometimes called “the first quantum revolution” or “Quantum 1.0”. Quantum technologies, however, go beyond simply making use of quantum properties to actively harness quantum principles such as superposition and entanglement to produce usable results: “the second quantum revolution” or ”quantum 2.0”

Currently, the “quantum industry”, still far from fully mature, is dominated by specialist quantum components and parts manufacturers, largely supplying the needs of quantum researchers (Faullimmel and Stilgoe, 2015). A few companies are providing services making use of “Quantum 2.0”; an example is the Swiss company ID Quantique[2] which provides quantum key distribution and random number generation. The Canadian company D-Wave sells quantum computers of a kind, but these are short of the universal quantum computers which will be the end result of the long development process of which NQIT is just one part.

This is the background to Networked Quantum Information Technologies (NQIT) as one of four university-led quantum technology hubs. NQIT is led by the University of Oxford, and brings together another eight UK universities and many major companies. NQIT is highly inter-disciplinary, combining elements from advanced theoretical and experimental quantum physics including materials science, engineering, and computer science.

The flagship deliverable from the NQIT Hub, known as Q20:20, will represent a major step forwards towards realising a universal quantum computer – one which will be able to implement any quantum algorithm. Q20:20 will consist of an optically-linked network of cells each containing a number of matter (Ion trap) qubits. This architecture, based on a network approach, is designed to ensure that the NQIT approach will be inherently robust and scalable. Integrating and controlling the NQIT Q20:20 machine (Work Package 9) is an engineering as much as a physics challenge, bringing together hardware and software components to make a new solution. Q20:20 is designed to be sufficiently adaptable to cater for different implementations of ion traps and, potentially, other forms of qubits.

NQIT is also Networked Quantum Information Technologies, in the plural. Although Q20:20 is the flagship deliverable from the hub, to which all others are ultimately directed, there are other important areas of research, some of which are expected to produce some early-stage spin-offs, including quantum-based simulation, secure communications, and enhanced sensing.

To make this all work, NQIT is addressing fundamental problems in theoretical physics, materials, engineering and computer science. For the scientific researchers, this is a rare opportunity for theorists to see the results of their work in practical application; for RRI research, it is an opportunity to trace the path of responsible innovation from theory to lab to innovation (Section 3.4).

2 http://www.idquantique.com/

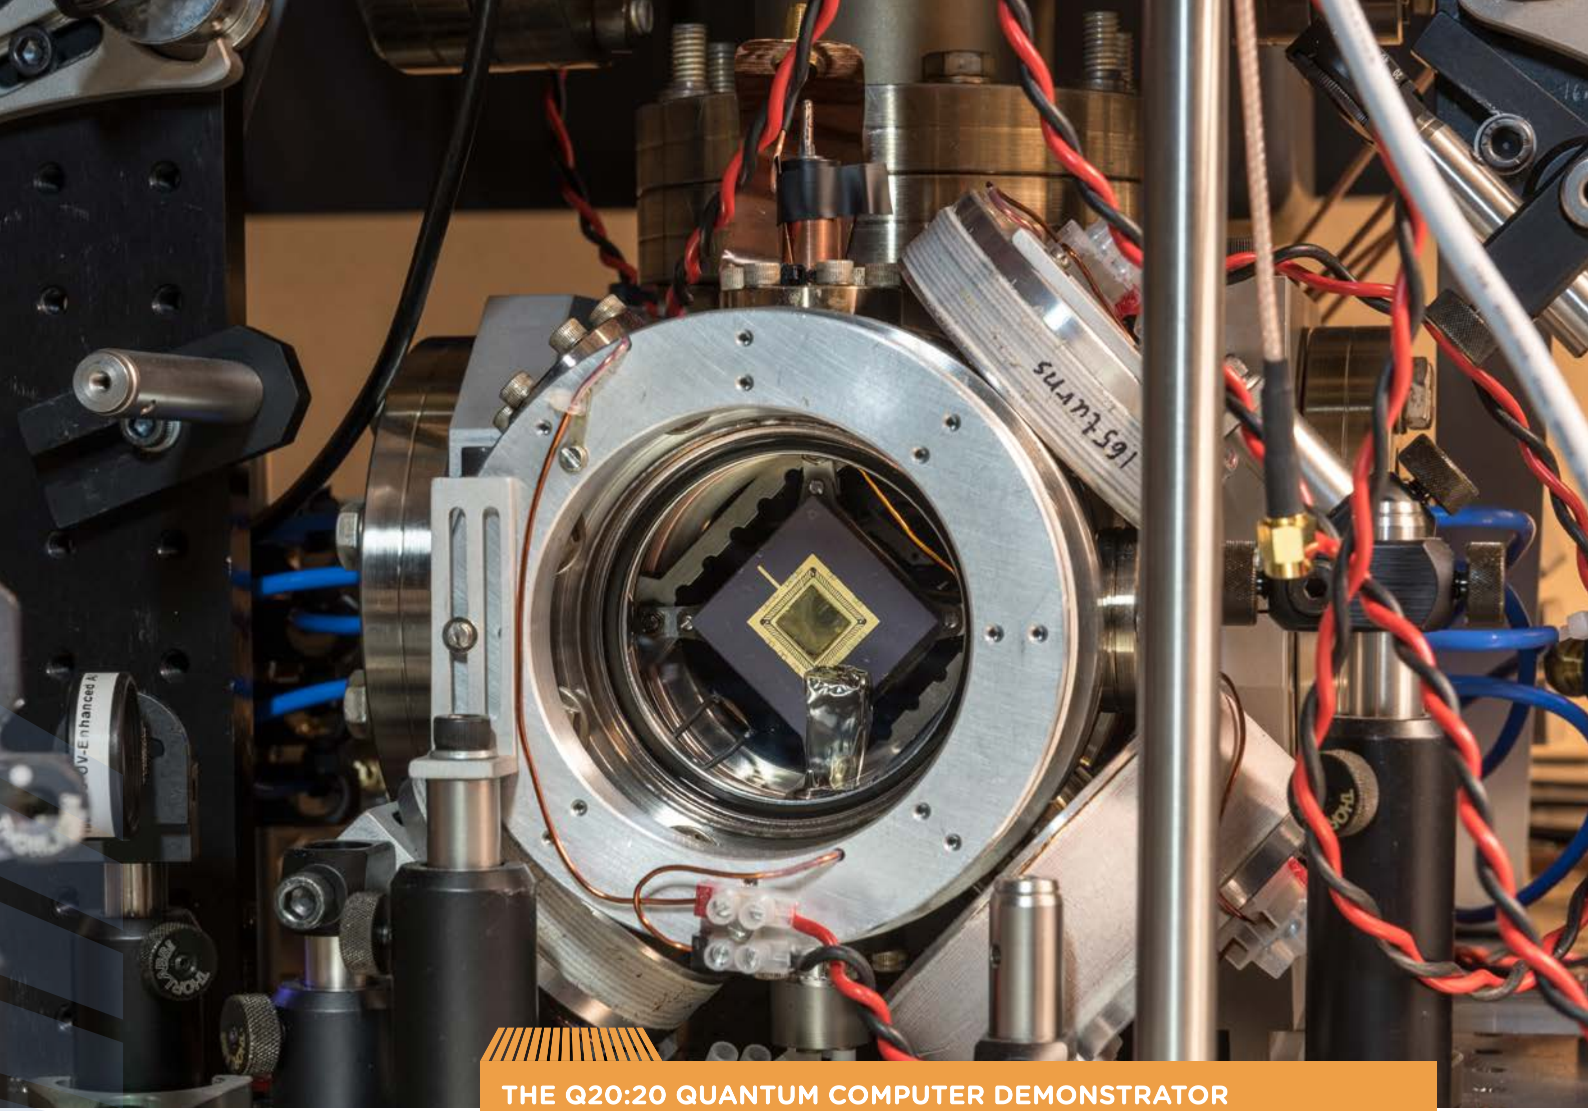

THE Q20:20 QUANTUM COMPUTER DEMONSTRATOR

The NQIT Hub's flagship goal is to realise the Q20:20 engine: a hybrid light-matter quantum computer involving twenty nodes, optically interlinked, where each node is a small quantum processor of twenty qubits.

Realising this ambitious target requires rapid development in multiple hardware efforts: these include engineering the nodal processors themselves, each of which requires integrated laser and microwave subsystems, as well as the optical links formed from fibres, photonic switches, splitters and detectors.

Each processing node will be an ion trap, a device within which a small number of charged atoms – ions – are held suspended in a vacuum and manipulated by laser and microwave systems. Each trap will be segmented, meaning that it can shuttle ions back and forth like beads on an abacus, and will be capable of holding two different atomic species simultaneously – we will be using calcium and strontium. A single unit of quantum information, one qubit, is embodied within the internal hyperfine states of each ion, and control of the qubits is achieved optically via integrated lasers and through microwave manipulation.

During the programme we will build a series of demonstrations that show the key steps in creating the Q20:20 engine. Initially we will connect two ion traps, then four using a switch. This will allow us to prove the concepts that we need to show work. Whilst we are doing this we will also be developing the miniature components that will allow us to shrink the system size to our target. We will use all of these components to build the final Q20:20 quantum computer demonstrator.

2.5 Prior RRI activities in Quantum Technology research

In comparison with RRI exercises in nanotechnology, synthetic biology and ICT, there has been far less attention given to Quantum. An analysis of government and research funding policy and a workshop in the EU project "Ethical Issues of Emerging ICT Applications" (ETICA)[3] (Stahl and Flick, 2010) identified quantum computing as one of eleven emerging ICTs, and argued that it is difficult at present to investigate the ethical and social implications because little is known and practical applications are hard to discern.(Stahl et al., 2009, Stahl, 2011). A study by Sciencewise (2014) could find no published examples of any systematic activities or deliberative dialogue to gather public attitudes towards quantum technologies, although there had been examples of engagement with the public to inform and to educate. In the absence of other data, the study made a survey of mainstream media reporting as a first attempt to capture public opinion around quantum technologies. As well as the well-known application examples around the potential of quantum computing for decryption, heightened by reports that US National Security Agency (NSA) is building a quantum computer, the media survey noted a perception that quantum computing and quantum-secured communication are "semi-permanently 'on the horizon'". Another recurring theme in media reports is that quantum technologies, like quantum mechanics, are extremely difficult to understand.

A round table held in March 2015, organised by the UCL Hub for RRI (Faullimmel and Stilgoe, 2015), started to lay out what they call the "social constitution" of quantum technologies, shaped by different interests and actors, and their "political economy": who can be seen to benefit from these developments, and in what ways? For example, the domination of research by public and private "giants" and the military interest in quantum technologies creates an aura of secrecy. If government or large corporations increase their capability to intercept communications or perform intrusive analysis of personal data, this could tilt the balance of power away from the citizen, eroding public trust.

Taking lessons from the failure of public acceptance of GM crops, the round table suggested that, "we need to be honest about what we know, but also what we don't know". At the same time, RRI needs to overcome the hype, which is also a challenge, that quantum technologies are set to be highly disruptive. The round table picked up the theme noted by Sciencewise (2014) of the perceived difficulty of understanding quantum technologies, and suggested a "new language" in order to promote as inclusive a debate as possible.

The RRI work package in NQIT is taking this forward as we orient RRI towards quantum computing and the other quantum technologies involved in NQIT on the road to Q20:20 and the other spin-offs.

3 www.etica-project.eu

Chapter 3: Key Findings

This section brings together emerging ideas in Responsible Research and Innovation and the nature of NQIT presented above with the findings from interviews and workshops to draw out implications for responsible research and innovation.

RRI research in NQIT is a work in progress; while we can already identify some specific areas of concern, this section starts to draw out the socio-technical landscape of NQIT as far as it can already be identified, in preparation for making recommendations for action and further research in the conclusion to this report.

It investigates the potential concerns, deeper and more transformative social implications, what the Round Table in RRI in Quantum Technologies (Faullimmel and Stilgoe, 2015) called the "social constitution" of quantum technologies: how different interests and groups of stakeholders will work together, who is seen as benefitting, the positive and negative consequences, and the determinants of public trust in the technologies.

Figure 2 below shows the layout of the landscape of RRI issues in NQIT, grouped according to the themes of narrative, trust and impact, and also according to a timeline which discerns:

- immediate issues that are live now
- those that are on visible on the horizon
- those that lie further downstream

The text on the righthand side of Figure 2, which covers these three (overlapping) time periods, indicates the main themes that we have identified: groupings of issue areas and overarching concepts: respectively, the narratives around quantum computing; the need for trust and how quantum computing will enhance trust; and social transformations which can be anticipated as applications come on-stream.

Cutting across these time periods, on the lefthand side of the figure, are active areas of research in quantum – we have allocated these roughly according to their maturity. The topics shown are indicative of the areas of research and application but are not intended to be exhaustive; the timeframe similarly is not intended to be to scale.

Immediate issues relate to the ways in which quantum computing and quantum as a whole is perceived by the public and others, rather than focussed on specific applications. Here we include narratives around quantum mechanics and quantum computing going back to their origins. Three key issues that we identify relate to the uncertainties around quantum computing, the key distinctions between quantum and classical computing, and the conceptual difficulties of quantum mechanics.

Issues on the horizon can already be discerned and are the subject of active research, but are not yet "live" in the sense that they require short-term actions by stakeholders. In this category are some strong claims made for quantum technologies "from the laws of nature", including well-known examples of quantum-guaranteed secure communications and, conversely, threats to existing secure communications. In contrast to this apparent certainty, some of the results from quantum computing can be hard to verify in terms of correctness and quantumness ("is it really using quantum effects?"). Some early implementations of computing (or computing-related) technologies in NQIT, sensing, simulation are also on the horizon.

Future challenges from quantum computing: Existing and historical technology trends and their risks provide some indication of the social transformations which might be precipitated when quantum computing becomes a reality. Use of quantum computing to give new powers to machine learning and big data analysis may be seen as threats to privacy and as leading to a changed relationship between citizen and state. Applications to defence may threaten the delicate balance of international relations. Quantum computers are currently "big technology", only available to governments, universities and large corporations. The ownership of this infrastructure, concentrated in the hands of these large organisations, or made widely available in a "cloud"-style architecture, will influence how they are used and perceived by the public.

Timeline

Whilst ordering issues according to those that are "live" today, those on the horizon, and the social transformations we can envisage yet further into the future, we do not wish to imply that action should be delayed on this basis. On the contrary, in anticipating issues in the future, we should take the opportunity, in the present moment, to lay down foundations so as to be able to manage those issues effectively as they arise. As we move forwards in time we will become more certain as to the several dimensions of quantum computing that are as yet unknown. We will understand better the form it will take, its power and the effects it will introduce into the world. But as far as possible we want to be in a position to anticipate and respond to these possible consequences before they emerge, as opposed to reactively and retrospectively managing difficult situations after the fact.

In anticipating issues in the future, we should take the opportunity, in the present moment, to lay down foundations so as to be able to manage those issues effectively as they arise

This involves a mix of preparatory dialogue, foresight work and careful attention to timing. Although we can be sure that unforeseen issues will arise, we can also see that many of the issues, such as surveillance and access or control of quantum technologies, are already concerns with classical computing; we can learn from experiences and public responses to these existing technologies.

In the following sections, we consider these issues and themes and their timing in greater detail.

Stakeholder events provide an opportunity for dialogue around the issues that may arise from quantum computing / *Kirsty Allen*

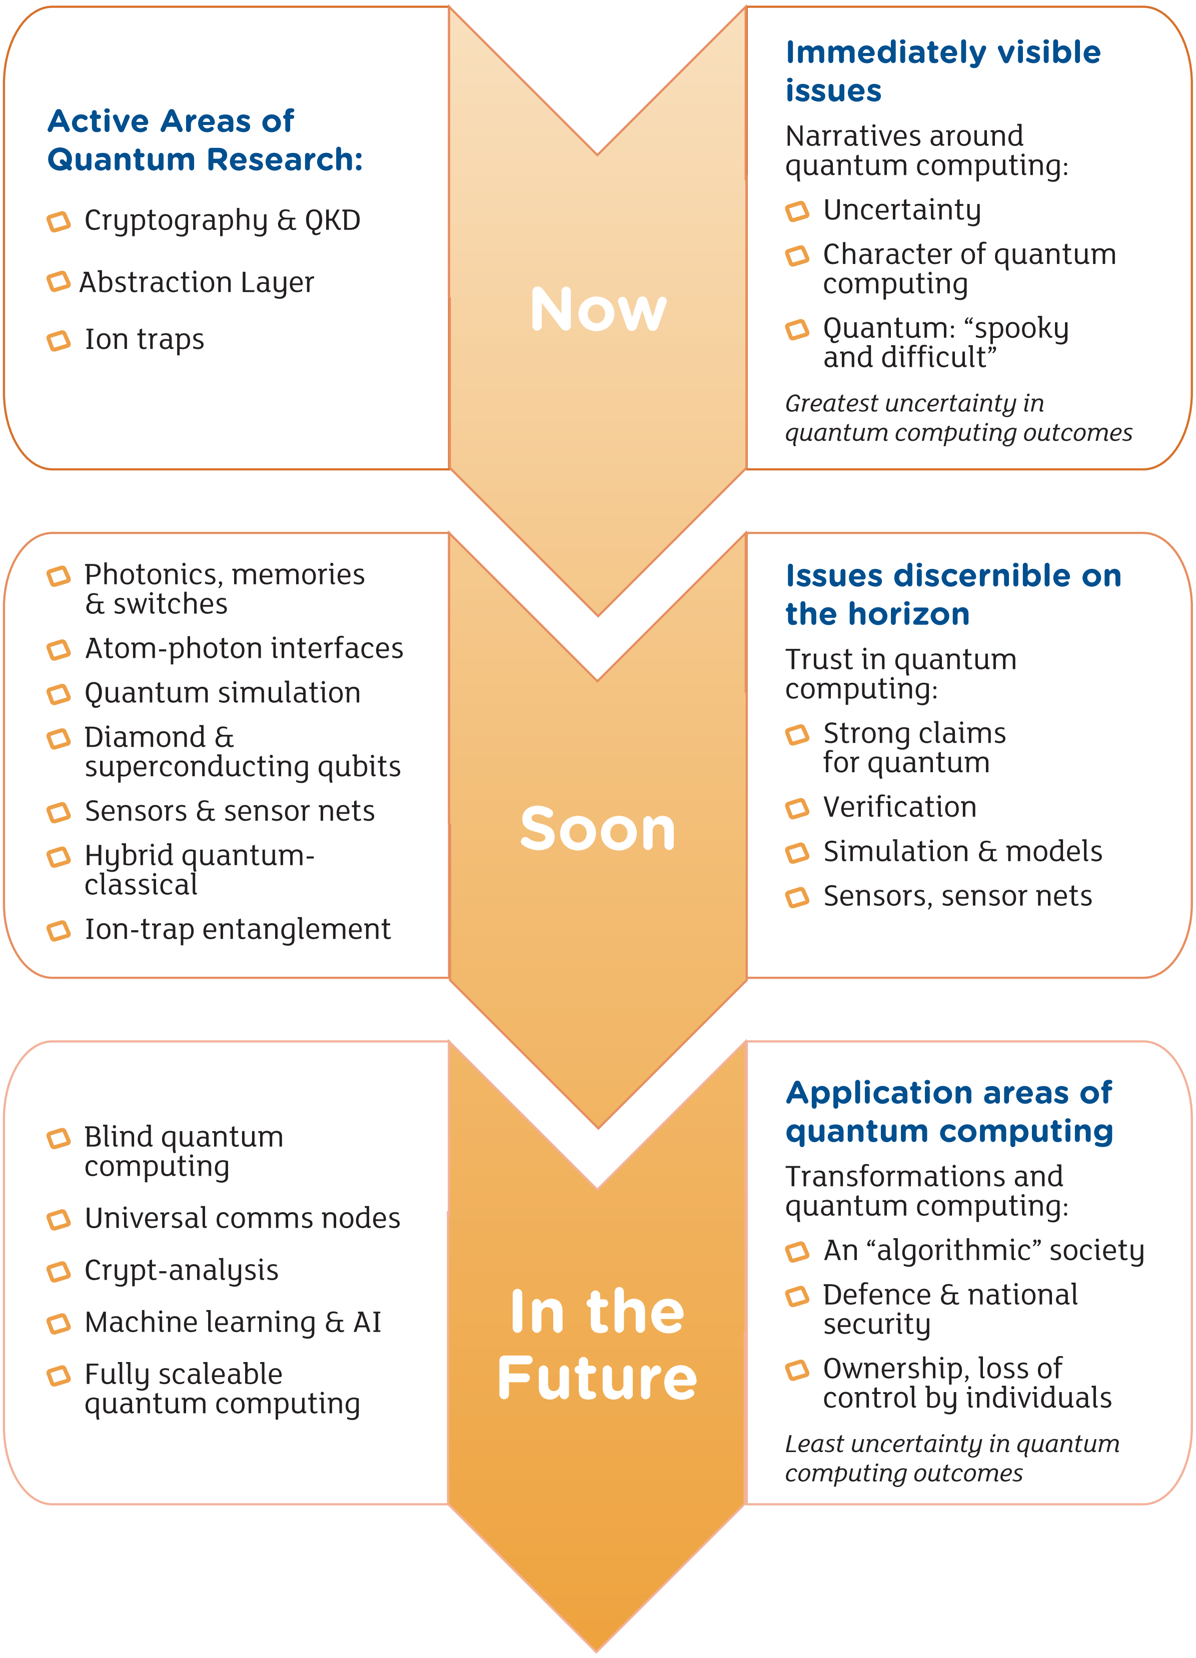

Figure 2: The Landscape of NQIT with relevance to RRI, showing time periods and issues with an indication of areas of quantum research

3.1 Immediate issues

3.1.1 Uncertainties around whether, what, and when

One of the most immediate concerns relating to quantum computers, in terms of Responsible Research and Innovation, is the high level of uncertainty, along several dimensions, which still surrounds these technologies. These uncertainties present several challenges: they **complicate** the timing of RRI activities; they make it difficult to ground dialogues with the public and with other stakeholders where the "if", let alone the "when" or the "what" of quantum computing, are hard to discern. They make it even harder to **forecast** how quantum computing will interact with some future configuration of society and technology. On the other hand, a significant risk of not tackling these challenges is that it will leave a **vacuum** where inaccurate and conflated narratives of quantum computing may flourish.

We have identified a hierarchy of uncertainty in the process of realising quantum computing; uncertainty at the higher level adds to and compounds the uncertainty at the level below:

Theoretical uncertainty – this perhaps provides the firmest ground – according to current accepted theory, quantum computing should be possible in principle, including achieving and maintaining coherence and fidelity.

- Issue: If quantum computing turns out not to be possible, this in itself would create new turmoil in our understanding of physics.

Engineering uncertainty – while quantum computers are theoretically possible, there are many ways in which they could be actually be implemented. Each of these brings engineering problems, and even if these can be solved, it is not certain that the resulting quantum computers would become economically viable.

- Issue: What is the most viable way to implement quantum computers? Can we be sure that it can be done?

Delivery uncertainty – uncertainties over when a viable device will reach market, and what capabilities it will have.

- Issue: When should we worry about quantum computing and when should we act?

Impact uncertainty – uncertainty over how quantum computers will be used and the effects that this will have on society.

- Issue: Makes it hard to calibrate our attitude to quantum computing – should we welcome it or be more cautious?

Thus, while there remains some uncertainty on the scientific level, there is more uncertainty around **what**, in the end, quantum computers will be able to do, **when** will they be able to do it, and **how** this will impact upon the economy and society.

These uncertainties make it difficult to open up honest conversations within the research communities and with other stakeholders about the future implications of quantum computing. Yet if these conversations are delayed until there is more certainty, it will be difficult to re-orient the trajectory of research and innovation towards desirable societal goals (Collingridge, 1982).

Theoretical Uncertainty

The concept behind quantum computing was first proposed in 1982 by Richard Feynman (Feynman, 1982). David Deutsch (Deutsch, 1985) asked whether quantum computers might be able to provide efficient solutions to certain problems that have no efficient classical solution. Peter Shor produced his famous algorithm in 1994 (Shor, 1999).

This is all rather a long time ago, and yet there are still no at-scale implementations. Even the milestone of the achievement of quantum supremacy – a real application for which quantum computers are significantly faster than classical (Preskill, 2012) - has still not been met, apart, possibly, for recent results by D-Wave (King et al., 2015). Indeed, an early "application" of non-universal quantum computation, what is known as Boson Sampling, if successful, would provide a convincing experimental demonstration of quantum computing supremacy (Aaronson and Arkhipov, 2011), but is not intended to solve any real-world problem.

The long time-span between the spark of the idea and implementation opens a space for narratives questioning the fundamental viability of quantum computers [4]. However, many of the theoretical requirements have already been met experimentally, and the path – albeit with many remaining engineering challenges - appears to be open towards scalable quantum computers. The corollary is that, if, for some reason we cannot foresee, it can be shown that quantum computers cannot be built, that in itself would represent a paradigm shift in physics:

> *"Even if the whole thing would turn out to fail, for some inconceivable reason, that would actually be one of the biggest scientific discoveries of the century"* – (Researcher in NQIT)

- a quote which echoes the thoughts of Scott Aaronson in 2008: *"In my opinion, the most exciting possible outcome of quantum computing research would be to discover a fundamental reason why quantum computers are not possible. Such a failure would overturn our current picture of the physical world, whereas success would merely confirm it"* (Aaronson, 2008).

Thus, while there are strong theoretical grounds to be optimistic that a Quantum Computer is possible, the translation from concept to reality will also be an important proving ground for theory.

Engineering Uncertainty

There is, however, another uncertainty around the actual form of quantum computing, how best the well-understood physics can be implemented practically. NQIT has chosen an implementation based on photonically linked ion traps, which currently are the most mature technology platforms, but, reflecting this uncertainty, NQIT is also keeping the pathway open for solid-state and cryogenic superconducting qubit implementations, with work packages in these areas parallel to the main Q20:20 pathway.

Like classical computers, quantum computing is subject to physically-induced errors; sources of errors in quantum states include decoherence and quantum noise. A simple example in classical computing may result in errors on a single bit, say, stored as a 1 but corrupted to a 0. Classical computing uses redundancy, making multiple copies of bits, and other techniques to reduce errors. However, it is not possible to simply make a copy of a qubit due to the no-cloning theorem[5]; instead, the information from a single qubit is spread across a number of entangled qubits. Theorists in Physics and Computer Science are continually developing techniques to

4 See Scott Aaronson's discussion at http://www.scottaaronson.com/democritus/lec14.html

5 https://en.wikipedia.org/wiki/No-cloning_theorem

Looking forward, it is not simply a question of whether these physics and engineering challenges can be overcome, but whether this can be done in an economically viable way.

overcome quantum errors; while there is a good understanding of the theory, making these methods work in experimental demonstrations is one of the greatest challenges to practical quantum computing. Moreover, because of error correction, many physical qubits (a factor of 100 or more, depending on the complexity (Gottesman, 2014, Bennett and DiVincenzo, 2000)) will be required for each error-corrected logical qubit available for calculation; scaling up from the relatively small number of physical qubits of existing experimental implementations will be essential before anything approaching a universal quantum computer can be achieved.

At the level of engineering challenges, as NQIT moves to Q20:20 and beyond, the development moves along what a director of NQIT has described in interview as a "pipeline" from being physics-led to more engineering-led. But looking forward, it is not simply a question of whether these physics and engineering challenges can be overcome, but whether this can be done in an economically viable way. The way quantum computing power will actually be delivered will depend to a large extent on the ways in which these engineering uncertainties are resolved: as large, centralised servers, or as small, affordable devices; as widely available computing power, or restricted to large corporations, governments, and wealthy individuals.

Delivery uncertainty

An obvious question to ask is, when will "quantum computers" be available in the market? But to consider this question, first it is necessary to clarify what is meant by "quantum computer", and, second, in what form quantum computing power, perhaps as a style of service rather than as consumer devices, will be available in the market.

Quantum computing in some form is already available in the market from D-Wave, but this, although sold as a commercial product, is highly experimental and is still not a universal quantum computer (UCQ). Other important quantum-based applications, around sensors and metrology and imaging for example, are already available or are close to market. Quantum simulation which can be used to model systems which are currently intractable (such as difficult quantum many-body problems) is expected to become available some years before universal quantum computers.

Quantum computers might beat classical simulations in terms of qubits around 2020 - looking ahead, the timeline becomes more uncertain

Given that different species of quantum devices will become viable within different time frames, and that these devices will sometimes exhibit computational properties (without being a universal quantum computer), sometimes combining different quantum technologies (sensors and storage, for example), there is a strong potential for conflation and misunderstanding as to the current status of quantum technologies and their potential. For example, an uninformed reading of the current availability of quantum computers could lead to the mistaken assumption that quantum cryptanalysis is just around the corner. As quantum research continues to deliver new technologies, it will be very important to continually underline distinctions between different classes of quantum applications. A useful exercise would be to establish a timeline for quantum innovation which would show the (regularly updated) timescales over which different types of applications may be expected to arrive. One source has considered the projected growth of qubit-based quantum computers as a parallel to Moore's Law for classical computing, and has estimated that the number of mutually entangled qubits implementable in a laboratory doubles approximately every six years (Weimer, 2011). On this basis, quantum computers might beat classical simulations in terms of qubits around 2020. Looking ahead, the timeline becomes even more uncertain, but if the trend continues at the same rate then 2048-bit RSA keys would come under attack somewhere between 2052 and 2059. Of course, by this time the technology may have developed beyond what we can see now, so this timeline is very speculative.

On the one hand, commodity products do not require the same level of adaptability as is needed in laboratory and experimental set-ups; this should make manufacture easier, cheaper, and smaller (consider the ubiquity of small laser devices in consumer products and industry compared with an optical table in a lab). On the other hand, quantum computing, however it is defined, is still far from the stage of being manufacturable in this sense, and not economically viable for all but the largest companies and governments. A director of NQIT described the current state of the art as the "valve stage", not yet even at the "transistor stage", by analogy with the history of classical computing.

A director of NQIT described the current state of the art as the "valve stage", not yet even at the "transistor stage", by analogy with the history of classical computing.

Quantum computing power would not necessarily be delivered in the form of consumer artefacts, however; there is uncertainty about the form factor which quantum computing will take, as a useful product, and as a reflection of the "character" of quantum computing (Section 3.1.2). It might turn out to be more viable to offer quantum-computing power as a cloud-style service; and, indeed, on a very limited scale, IBM has recently started to offer something akin to this.

It is important to recognise that quantum computers will not supplant classical; existing computers do many useful tasks very well - such as delivering email, providing applications for word processing and spreadsheets, and maintaining and processing records, and quantum computing will not change this. Besides, quantum computing, whatever form it takes, will require control mechanisms at the lowest level, and, at higher levels, operating systems and human interaction through classical devices.

These uncertainties around the form factor of quantum computing; the timeline; and the ways in which quantum computing will work alongside classical, are reflected in uncertainties in the economic benefits and social transformations from these technologies: which applications will emerge, when this will happen, and what new and currently unforeseen uses will be made of the technology.

Impact uncertainty

A number of application areas are posited for quantum computing as part of establishing the economic case for engaging in the research. Indeed, NQIT, and the other quantum hubs, are geared towards technologies that have an economic benefit to the UK. There is unevenness in the attention given to the various application areas possible for quantum computing. On the one hand, the potential effects of some applications, such as cryptography and cryptanalysis, have received considerable attention given the significant roles that secure communication plays in modern society. On the other hand, application areas such as those related to simulation, machine learning and big data have not received such high levels of scrutiny, and it remains somewhat sketchy as to precisely what role quantum computers may play in these domains (for a summary of quantum computing applications as we can foresee them, see Table 1).

We recognise that the eventual effects of technology are notoriously hard to pin down. Deterministic notions of technological progress, which imply inevitability associated with technological outcomes, often dominate technology policy, despite growing amounts of evidence across disciplines that the form and direction taken by technology as it interacts with society are open to collective ingenuity, economic priorities, cultural values, individual creativity, tensions between institutional interests, and many other factors (Stirling, 2008). What we do know is that the effects of technology are often unpredictable and surprising, perhaps presenting greater challenges in areas where we may least expect them to arise. This implies that we should keep all potential application areas on the radar, and be alert to forms of appropriation and use that we had not previously considered. It also suggests that we need to engage in continuous dialogue with commercial users of the technology, and maintain a watch on emerging technological trends and controversies to see how they may interact with upcoming quantum applications.

The effects of technology are often unpredictable and surprising.

Table 1: Application areas of quantum computing

Capability	Implications
Factorising large numbers, which are used for public-key cryptography (eg. RSA)	"Breaking the Internet" and destroying public confidence in e-commerce
Enabling very strong encryption, including by Quantum Key Distribution	for high-value or high-sensitivity transactions; perhaps at a lower level of security for ordinary users; risk that legal intercept will be impossible
Enhancing machine learning	disruptive potential for skills and professions; potential for surveillance; potential for financial applications
Searching through data (eg. Grover's algorithm)	Big data analysis; use by police and security services to find suspicious data (alongside machine learning)
Boson Sampling	Sampling (ie, deducing statistical properties) of the probability distribution of identical bosons; the "use" is to demonstrate quantum supremacy in a form of non-universal quantum computation
Highly accurate gravity sensors and accelerometers	"seeing through walls", even if the "vision" is not very clear, could have privacy implications; very accurate navigation could have useful applications but also, combined with tracking, surveillance implications
Quantum simulation	Discovery of new drugs and materials without "trial and error"; simulation of quantum-mechanical systems Greater understanding of physics; perhaps the most fundamental application
Optimisation	This is the focus of D-Wave. Many problems can be expressed as optimization: network optimization, optimal therapies, protein folding, finance, etc.

3.1.2 The still-emerging and uncertain character of quantum computing

An important feature of quantum computing is that is that it is not strictly analogous to classical computing, but rather has a very specific mode of operation that constrains what it may be able to do. In a sense, referring to it as "computing" is unhelpful in that the model conjured is that of classical computing. Quantum computers are programmed in a different way to classical computers. Quantum operations cannot have loops and must be reversible, and should be more thought of as a "pipeline" where framing of the problem at the outset leads to a series of entangled quantum states that resolve to the solution of the problem. However, quantum operations can be placed within the flow control, loops, tests, and conditionals of classical computing. Languages such as QCL[6] and Quipper[7] are being developed to enable programming at a much higher conceptual level in a hybrid classical/quantum computer.

The risk is that there is the opportunity for over-statements, or simply for misconception, about what quantum computers will and will not be able to do, or might be able to do.

While common explanations used to distinguish quantum from conventional computing relate how qubits can represent states between 0 and 1, and that a quantum computer can analyse an entire problem space instantaneously (for example, as described by Canadian Prime Minister Justin Trudeau[8]), these images also have the potential to be misleading. A crucial point further point is the special preparation needed to represent a problem in a way that is solvable by a quantum computer. Problems have to be expressed in a way such that incorrect solutions "cancel out" when the wave form collapses, leaving the "correct" result as the final state of the machine. Thus, although quantum computers will not have the generality of classical computing, their power arises because they can, in principle, solve some problems that classical computers cannot.

Quantum computing will enable solutions (in realistic, non-exponential timescales) for some problems which cannot be solved (in any reasonable time) by classical computers. There is a "zoo" of known quantum algorithms[9] which currently lists at total of 57 algebraic & number theory, oracular ("black box"- ask the computer a question and it gives a response, as part of a larger process), and approximation and simulation algorithms. These algorithms exist and are widely expected to work, but they are all theoretical, prior to the implementation of a working quantum computer.

Quantum computing will enable solutions (in realistic, non-exponential timescales) for some problems which cannot be solved (in any reasonable time) by classical computers.

However, as a generalisation, it is still an open question which classes of problems may be solvable by quantum computers. This can be expressed more accurately by stating that it is still not known which problems are actually in BQP, Bounded-Error Quantum Polynomial-time, generally considered to be *"the class of all problems feasible for a quantum computer"* (Aaronson, 2010). There are problems in NP that are not in BQP and, in particular, it is strongly suspected that quantum computers will not to be able to solve NP-complete problems in polynomial time, as illustrated in Figure 3[10].

However, misunderstandings over these crucial properties create a tendency to over-state the capabilities expected of quantum computers. An example is the common trope that they will be able to solve any NP-Complete problems (as propounded for example by The Economist in

6 http://tph.tuwien.ac.at/~oemer/qcl.html

7 http://www.mathstat.dal.ca/~selinger/quipper/

8 https://youtu.be/4ZBLSjF56S8

9 http://math.nist.gov/quantum/zoo/

10 Based on https://commons.wikimedia.org/wiki/File%3ABQP_complexity_class_diagram.svg. By User Mike1024 [Public domain], via Wikimedia Commons

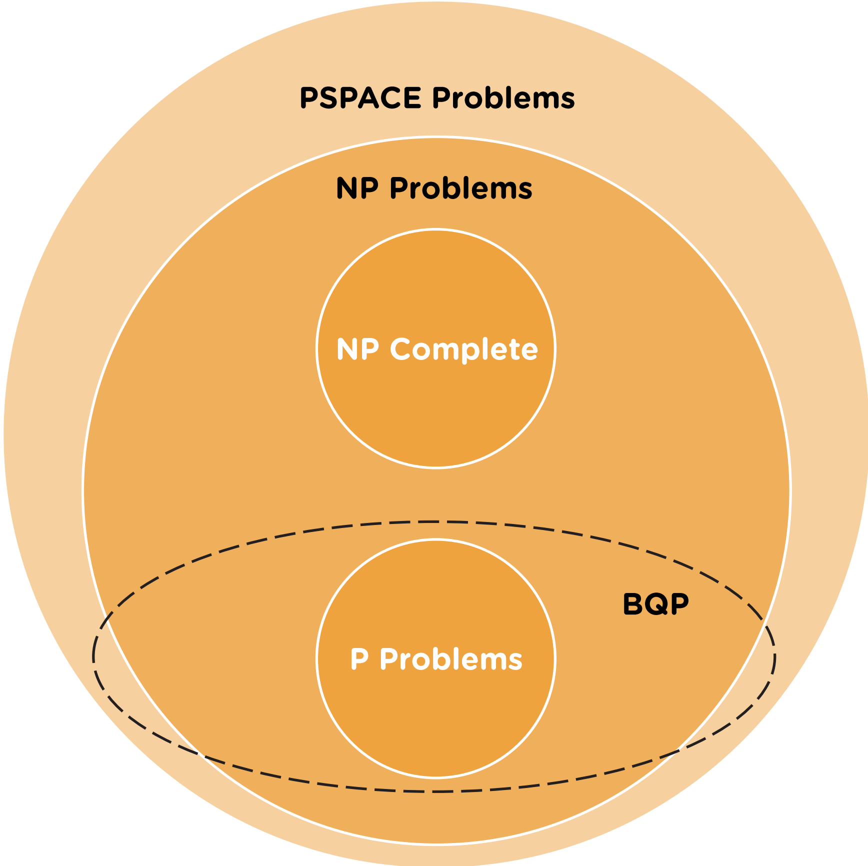

Figure 3: Suspected (but not proven) shape of the BQP of problem space, the range of problems easily solved by quantum computers.

February 2007[11]; and a long discussion following it in Scott Aaronson's "Shtetl-Optimised" blog[12]). Another example is an article in a newspaper claiming that *"superfast quantum computers … will spell the end of freedom as we know it"* (Cribb, 2016), on the basis of their supposed ability to analyse vast amounts of data.

So while the theory of quantum computing has been developing for over three decades, and now is well-defined – despite remaining challenges in implementation – the uncertainties around the timescale and the eventual form of quantum computing, combine with exaggerated claims for its eventual applications to create a mystique and potential fear on the part of the public. It is important to ensure clear and accurate articulation about what quantum computers are and what they will be able to do, as well as around their limitations – what they will not be able to do.

Form factors of quantum computing

Speaking of 'quantum computers' may mask the reality that in fact all quantum computing will be, to an extent, hybrid.

However, speaking of "quantum computers" may mask the reality that in fact all quantum computing will be to an extent hybrid, because it will always require classical computers for control and for input and output and interaction with the humans or other non-quantum devices.

The shape and form of quantum computing will have implications for RRI: will it be widely available, or restricted in availability? Will it be accessed as remote servers, or carried around in consumer devices, and does this matter from the point of view of the user? What impacts can be expected from non-universal quantum computing, and to what extent will this foreshadow the impacts of universal quantum computers, when they are eventually available? These and related considerations are discussed in terms of the ownership of quantum computing in section 3.3.3, but in this section we are concerned with the uncertainty as quantum computing is still at an early stage of realisation.

11 http://www.economist.com/node/8697464
12 http://www.scottaaronson.com/blog/

Quantum simulation and quantum-classical hybrids are the most likely early implementations of quantum computing, and specifically as demonstrators for the NQIT project. Beyond this stage, quantum computing will still take the form of co-processors, analogous to specialised graphics co-processors in classical computers, or perhaps specialised servers, rather than general-purpose computers comparable to a personal computer or laptop. Quantum digital simulation will probably be the first style of quantum co-processor; not a universal quantum computer, but powerful simulators enabling targeted design of new materials or drugs, for example, rather than trial and error discovery.

As physical devices, the actual implementation – using photonically-linked ion traps, in the case of Q20:20 – is contingent, rather than the necessarily only possible route, chosen as the most promising currently available implementation path but recognising that this will not necessarily remain the case. Just as with classical computing, the physical instantiation of quantum computing can be expected to develop; for example, prior to transistors, computers used short-term memory based on three-dimensional arrays of cores, which were very expensive and difficult to engineer.

The emergence of quantum computing as a practical reality is engaging computer science in the physics which underlies computing (DiVincenzo, 2000). There is every reason to believe that the physical underlay of quantum computing will evolve – indeed, different forms of qubit are actively in development. However, as physical devices, quantum computers, while having some new, very powerful capabilities, will also have physical limitations, as do all forms of technology. At the extreme, Lloyd (2000) has calculated theoretical limits to computation, which, in practice, are unlikely to be reached.

3.1.3 How can we talk about Quantum?

Quantum Physics has attracted the perception that it is difficult to understand and explain (Sciencewise, 2014). For responsible research and innovation, discourses that emphasise its supposed impenetrability may discourage wider engagement with the topic.

Discourses that emphasise the supposed impenetrability of quantum physics may discourage wider engagement with the topic.

The words which are used to describe fundamental quantum mechanical phenomena can unintentionally add to this perception. For example, Einstein's phrase "spooky action at a distance" - which has become arguably a *"bona-fide technical term"* (McEliece, 2006) has spread to become associated with quantum mechanics in general. Some have argued that "spooky" is a good candidate word to describe quantum technologies in order to capture the public imagination (for example, as *"spookytechnology" (*Tahan, 2007)), and for similar reasons "spookiness" has an appeal to journalistic presentations of quantum phenomena[13]. On the other hand, participants in more than one of our workshops have suggested that the word "spooky", used out of its technical context, risks sending a message that the science is difficult for non-specialists to grasp and perhaps even gives an impression that scientists themselves do not really understand it. A recent example of this sort of discourse (but there are many others) is in an article in The Guardian, in which we read that "the mysterious subatomic realm of quantum physics, … leads to some inconceivably large numbers, not to mention some mind-boggling working concepts"

Moreover, the word "quantum" has itself has become widely re-appropriated; used as an adjective it can be taken to mean a big step forward, or, in the context of science fiction stories,

13 For example, a report on recent experimental demonstrations of loophole-free quantum entanglement: http://www.livescience.com/52811-spooky-action-is-real.html?li_source=LI&li_medium=more-from-livescience

It is important to be aware of the metaphorical and informal uses of the word "quantum" and its connotations within popular culture.

to indicate a departure from usual physical rules[14]. There are companies that use the name "quantum" to suggest advanced technology of almost any kind. In an RRI dialogue that depends upon achieving a shared understanding of quantum phenomena it is important to also be aware of the metaphorical and informal uses of the word and its connotations within popular culture. Special care is needed in the selection of terms, given their potential or existing connotations. For example, describing quantum phenomena as "counter intuitive" may perhaps be a better approach than labelling it "spooky".

Yet, it is important not to deny that quantum phenomena do present a real challenge to our intuitive understanding of the world, and that they have opened important philosophical debates about the status of an "objective reality" in relation to how quantum phenomena respond to observation. It is, of course, the very "quantumness" that confers quantum computers with their advantage over classical computing (Nielsen, 2008), which implies that communicating about quantum computing needs to take on the challenge of conveying the unique properties of quantum phenomena.

One form of reassurance has been to add the caveat that these unsettling phenomena manifest only on very small scales and so are actually far removed from our everyday concerns. Such an approach, however, may lead to a risk for the acceptability of quantum 2.0 technology, whereby the new quantum technologies eliminate this distance between everyday concerns and quantum phenomena, transforming their counter-intuitive properties, and harnessing them to have consequences for society as part of people's direct experience.

In contrast to the discourses of quantum phenomena as spooky, unsettling or hard to understand are the solid, proven foundations enabling quantum science to underpin many existing technologies and applications. Quantum Mechanics, from this perspective, has been a profoundly successful theory capable of explaining many phenomena to the point of becoming a mundane and non-contentious part of the science curriculum; for example, first quantum concepts are currently taught in physics in England at key-stage 5 (years 12-13). Quantum phenomena are quite unremarkable in the sense that, like any physical phenomena, they consistently adhere to the laws that describe them. Ambainis (2014) has expressed this succinctly: "There are a few things that are different from classical physics and one has to accept those. But, once you accept them, quantum mechanics is simple and natural". Quantum laws may be surprisingly different to those that describe classical systems, yet the occurrence of such a difference is also less remarkable in the context of the landscape of physics which contains a number of overlapping but distinct paradigms, each with their own logics, including Newtonian mechanics, relativity and quantum mechanics. In adopting the perspective of any one of these paradigms, the logic of the others may appear to be "strange". Moreover, when phenomena register as "counter-intuitive" then this can be taken as simply revealing the limits to our intuition (which is always borne out of a localised perspective and localised sets of experiences) as opposed to evidence of mysterious or unreasonable phenomena.

14 https://en.wikipedia.org/wiki/Quantum_fiction

Quantum as a cultural phenomenon

Although quantum mechanics is well established among the scientific community, for quantum technologies to become acceptable there will need to be a wider societal process of coming to terms with the largely unfamiliar and counter-intuitive logics of quantum phenomena. Metaphor and symbolism play an important part in society, and cultural assimilation of technology is tied in with its portrayal in media, in fiction and in the visual arts, which establish connections between technological qualities and broader patterns of cultural meaning. Thus, as quantum technologies achieve a higher profile, one might anticipate a gradual increase in the amounts of drama, poetry and art that derive their themes from quantum phenomena, particularly as artists explore the relationship with between the unfamiliar quantum worlds and our familiar world of everyday experience[15]. Rather than distancing themselves from these developments, NQIT and the quantum hubs more generally could benefit from working with artists and designers to engage with the public, to explore quantum phenomena from new perspectives, and generate a better public understanding and cultural assimilation.

Overall we can recommend a strategy of discussing quantum phenomena in ways that at the very least avoid further mystification, and at best pursue a strong programme of demystification. This should be done in a way that acknowledges the counter-intuitive properties of quantum phenomena, but which also puts these into the context of our wider understanding of physics, drawing on the wide range of phenomena that lie outside our ordinary, everyday experiences.

15 For example, the Quantum Art Project: http://www.thequantumartproject.com/

The Quantum Art Project is an exhibition that aims to spread new research through art, and to create a platform for scientific inquiry

3.1.4 The unfolding narrative of quantum computing

I think the best you can do is come up with routes to try and predict the adverse reaction and, then, spread enough positive underlay, so that when that comes out, the public either won't believe it or will see it in proportion. - (Senior NQIT leader)

This is a quotation from a member of the NQIT leadership team in response to how we might handle the risk of negative, unforeseen consequences from some innovation. The interview up to this point had discussed some examples of negative consequences of innovations – tobacco, asbestos – not directly related to quantum, and now turned to considering how to deal with the risk of such an outcome in the case of quantum technologies. We did not at all discuss or try to speculate on specific issues but instead considered in quite general terms how to deal with such an eventuality.

The interviewee's comments highlight the importance of narratives which surround innovation in determining public acceptance. In this section we have considered:

1. The forms of uncertainty surrounding quantum technologies;
2. The distinguishing characteristics of quantum computing, and finally
3. The terminology and perceptions of quantum phenomena. Where these topics come together is in their implications for the unfolding narrative of quantum computing, and how that narrative is likely to be shaped by various social actors and their different perceptions.

Uncertainties surrounding quantum technologies leave space for many competing narratives to emerge; the distinctive character of quantum computing is open to false analogies with classical computing; and popular conceptions of quantum phenomena tend towards mystification as opposed to elucidation.

NQIT has to meet head on the challenges of engaging with dialogue under conditions of uncertainty, and of adopting a way of speaking about quantum technologies that avoids mystification.

The combination of uncertainty, misunderstanding and mystification provides a fertile ground for ill-informed and dystopian narratives to emerge and take hold (for example, an article published in the Sydney Morning Herald in January 2016 (Cribb, 2016)). If the quantum hubs do not themselves step up to provide a strong, informed narrative voice with integrity, then others will step in and claim this space. In providing this narrative voice, NQIT has to meet head on the challenges of engaging with dialogue under conditions of uncertainty, and of adopting a way of speaking about quantum technologies that avoids mystification, but equally does justice to their "counter-intuitive" properties.

The challenge for NQIT and other quantum hubs is to play an active role in shaping the narrative of quantum computing (and quantum technologies more generally) from their perspective of having a deep technical understanding of their capabilities, potential and operation. We are not recommending necessarily that theirs should be the only or always the dominant voice, but that that the hubs should aim to have a strong influence on the narratives that will inevitable compete to define quantum technologies in the coming years.

Crucially, the hubs' role in shaping these narratives should not be proclamatory, but rather taken as being part of an on-going process of dialogue that is responsive to multi-stakeholder concerns and perceptions. This will help to foster a shared narrative, one that many in society may find acceptable and easy to buy into.

Recommendations to address the uncertainty and emerging narratives of quantum computing

To summarise the practical recommendations from this section, we advise:

4 **Adopt a strategy of being honest about the uncertainties surrounding quantum technologies, but at the same time to find ways to articulate an informed understanding of what is likely to happen and when.** Creating a "timeline" in the form of a quantum application clock could be helpful, as could continually articulating the time between scientific advance and its application in the real world. This is common in news reports for medical advances, where the scientist articulates how long it will take to turn an advance in the lab into a new treatment (usually stated at about 10 years). Uncertainty should never be used an excuse to delay or avoid dialogue, especially since this creates a vacuum in which others can play a powerful role in shaping the quantum narrative. Take the opportunity for mutually informative dialogue with all stakeholders, throughout the lifecycle and not only as outputs become closer to market.

5 **Articulate the very specific character of quantum computing, both its powers and limitations, and the risks as well as benefits these can enable, as distinct from classical computing.** This is an area ripe for misunderstanding and where powerful mythologies can take root – such as the notion that quantum computing can solve any problem currently too hard for a classical computer. It may be useful to develop materials on very specific points of difference between classical and quantum computing (similar to a FAQ), particularly in relation to aspects most prone to misunderstanding. An example, which would also help to explain the processes of quantum computing, could be a poster or short pamphlet or a video describing in simple terms how a problem is represented in a special way before it can be "run" as a quantum program. It might be useful to explore ways in which an artist in residence can help to express complex scientific ideas in accessible and interactive ways. This has proven very effective in other projects[16].

6 **Develop and adopt a language with which to discuss quantum phenomena in a way which is faithful to quantum effects but which avoids introducing elements of mystification.** The example we have used is that "counter intuitive" is a much less loaded term than "spooky". At the same time, underline the effectiveness of quantum mechanics as a theory that has underpinned the engineering of everyday technologies, and also encourage artistic engagement with quantum phenomena.

16 For example, http://proboscis.org.uk/5302/digital-alchemy/

3.2 Issues on the horizon: emerging areas of interest and early applications of quantum technologies

This section considers a point in the near future when early quantum technologies are sufficiently mature to make the transition from the laboratory into real world applications. Earlier outputs from NQIT and other quantum hubs may include quantum cryptography, sensors and quantum simulation, which are likely to emerge much sooner than a large-scale universal quantum computer. The recommendations from Section 3.1 should ensure that there will already be a positive narrative surrounding quantum technologies to build upon. However, as technology moves from the laboratory towards working quantum prototypes, we anticipate that issues already on the horizon will become increasingly urgent. These issues, already in sight, are the focus of a number of research themes. These issues relate to:

- The social and technical embedding of quantum applications to realise the **strong claims** made for them, particularly in high-profile applications to secure communications and cryptanalysis;
- **Verification** of the results of quantum computing and of the "quantumness" of the implementation, by which we mean the extent to which the results are exhibiting quantum phenomena;
- Early application areas preceding full quantum computing, such as **quantum-enhanced sensors** and **quantum simulations**.

As a more fully Quantum 2.0 world emerges, narratives around quantum will become less speculative, but as these applications start to impinge directly or indirectly on people's lives, there will be a requirement for **trust**: there will be increasing reliance on the capabilities of quantum technologies, yet those who rely on the results and capabilities will not be able to certify them for themselves. This element of trust relates to the technology's quantum credentials, its correct implementation and proper operation, and its exploitation for the benefit of society. In turn, this level of trust will influence and be influenced by what the round table in RRI in Quantum Technologies (Faullimmel and Stilgoe, 2015) called the "social constitution" of quantum technologies: the industries surrounding these technologies, the large defence interests, the ownership of the technologies, and who is seen to benefit from them – which will emerge in the social transformations which are the focus of Section 3.3.

3.2.1 Strong claims

The transformative potential of Quantum 2.0 applications, directly harnessing and exploiting quantum phenomena, is the basis for **strong claims** about the potency of quantum-based applications. Thus, within grant applications, scientific publications, and press coverage there have been strong rhetorical statements claiming "unbreakable" cryptography or "perfect" sensors:

> *"quantum systems use the laws of quantum theory, which have been shown to be inherently unbreakable"* – BBC News, 9th October 2008

> *Networked quantum sensors can provide "an ultimate proximity or intrusion detection system". Light can be harnessed to act as "a near perfect information carrier".* – NQIT 12-page document, extracted from project proposal

These strong claims are justified because, in principle, quantum methods hold out the hope of unconditionally provable properties, for example in technologies such as perfectly secure communications[17], blind quantum computing (in which the input, output, and computation remain unknown to a quantum-based server) (Barz et al., 2012); and perfect randomness[18]. Conversely, known algorithms such as Shor's and Grover's algorithms make strong claims to be able to "break" existing standards of Internet security and privacy.

The common thread is that the properties of these applications are underwritten by quantum phenomena, and hence guaranteed by inviolable laws of physics. For example, classical cryptographic methods depend on certain mathematical problems being hard to solve, but do not claim that these problems could not be solved in principle. Quantum cryptography, on the other hand, would provide a form of cryptography that could never be violated by any future discovery or technological advance – not even in principle.

Claims of such strong guarantees for powerful applications with inviolable properties may store up a number of problems for the future.

The potency of such applications, should they be realised as anticipated, may engender very strong and polarizing forms of social change. The limitations and weaknesses of existing technologies mean that there is always room for forms of exploitation that prevent any one group of societal actors gaining a complete upper hand over another; there are always opportunities for circumvention and weaknesses to exploit. Much of this point is represented in the recent battle between the FBI and Apple to unlock an encrypted iPhone[19] to assist with their investigation of a terrorist suspect; weaknesses in the technology and a security exploit, in the end, allowed a rough balance to be struck between the interests of law enforcement, the commercial interests of Apple, and the interests of the law-abiding majority of iPhone users. In contrast, "perfect" or "inviolable" technologies would privilege the interests of those who are in a position to take advantage of them, to the complete exclusion of others.

Conversely, claims of inviolability may not be borne out in practice as completely as the claims that are made for them. The emergence of DNA fingerprinting, which was also accompanied by a similar, and subsequently contested, rhetoric of certitude, provide a model for the way that strong potential capabilities backed by sound science may play out in practice[20]. The claims of DNA fingerprinting to establish identity unequivocally do not only rely on the fidelity of the test but, in order for the test to speak its truths, also require the careful enactment of laboratory and intuitional practices that have evolved over time. In the reality of a criminal investigation, even when the best practices are followed, samples may be contaminated or incomplete; the DNA evidence may support or contradict other evidence, and is rarely able to provide irrefutable proof, despite claims made for it by defence or prosecution counsel. In other words, the strong epistemological guarantees provided by the technology itself need to be underwritten by equally powerful social guarantees - interlocking and intricate laboratory procedures, organisational practices and accountability structures - that the application of the technology is sound and trustworthy, and not subject to error, carelessness, or hidden interests and influence (Thompson, 2008). These procedures largely disappear into the background, but are nevertheless essential. Similarly, for quantum technologies, strong claims and powerful guarantees are bounded by the realities of implementation and use. Indeed, much of the theoretical work is addressing how these protocols can be made actually secure, up to a defined level, in situations of imperfect or unknowable implementation:

The strong epistemological guarantees provided by the technology itself need to be underwritten by equally powerful social guarantees.

17 For example, http://www.nature.com/news/quantum-communications-leap-out-of-the-lab-1.15093
18 For example, http://qrng.physik.hu-berlin.de/
19 https://www.washingtonpost.com/world/national-security/fbi-paid-professional-hackers-one-time-fee-to-crack-san-bernardino-iphone/2016/04/12/5397814a-00de-11e6-9d36-33d198ea26c5_story.html
20 For example, see http://www.councilforresponsiblegenetics.org/GeneWatch/GeneWatchPage.aspx?pageId=57

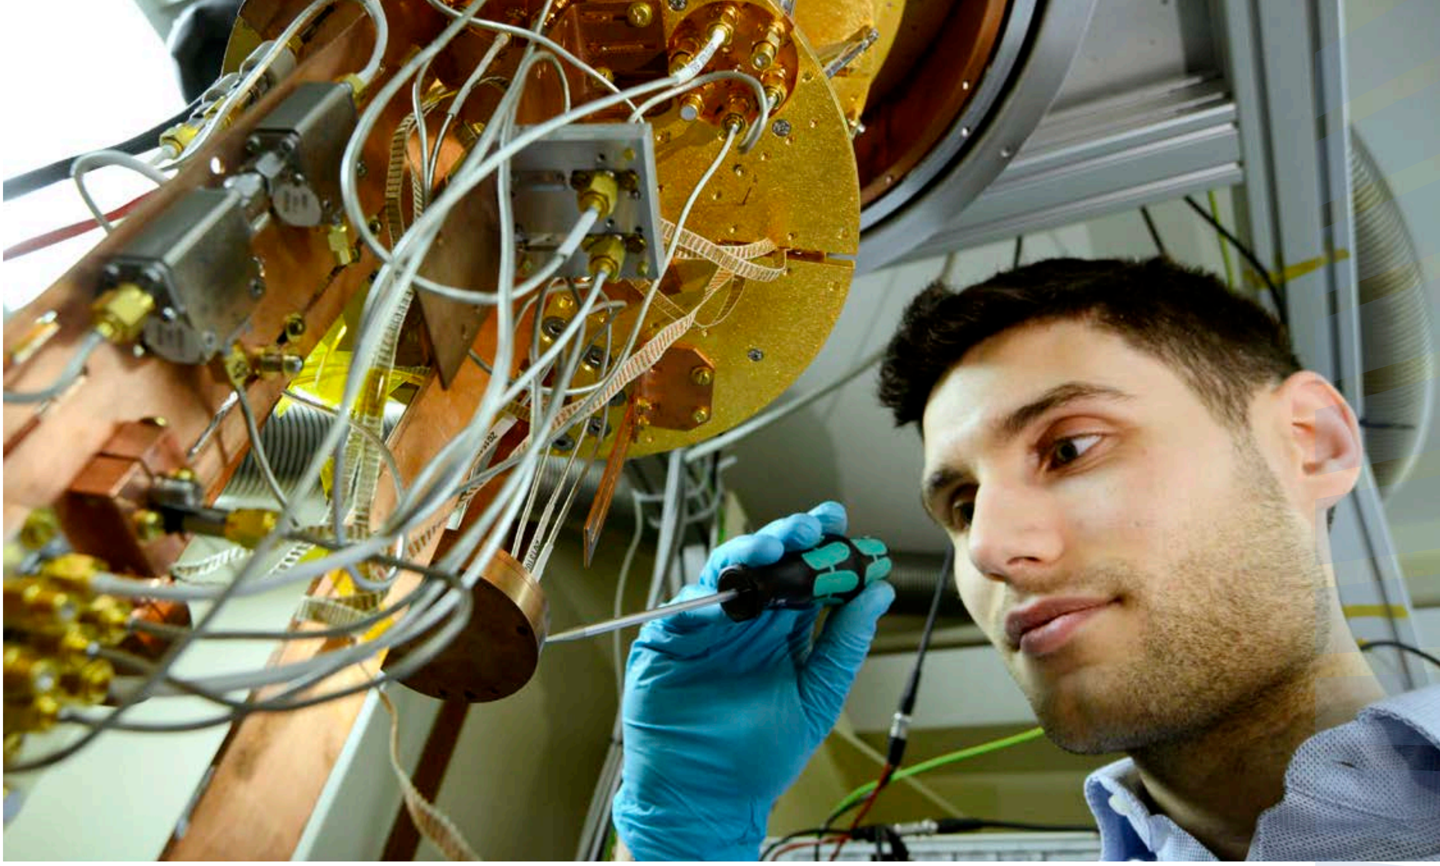

Superconducting qubit experiment using a customised dilution refrigerator for operation at 10 mK / *David Fleming*

"… there are protocols that are secure in theory, when there are no imperfections … but when you try to implement them in practice, then they become completely insecure, so I'm also interested …proving security against different types of adversaries …" – (NQIT researcher)

Concretely, in the context of quantum computing, on the one hand, there are the engineering difficulties of realising quantum computers – it turns out to be very hard, and the strong guarantees for quantum computers are circumscribed by quantum decoherence, problems establishing entanglement, controlling various forms of qubits, and other challenges in physics and engineering. On the other, even when the science and engineering has developed to enable this to be done, these theoretically "perfect" technologies will still be subject to errors of implementation, just as existing systems are, and to insecurities and failures "around the edges" from the social structures in which they are embedded and the assumptions of the wider society with which they are inscribed (Mager, 2012).

There is, thus, the challenge of navigating between claims for powerful properties of quantum applications, on the one side, and, on the other, the realities of what they may actually be capable of in the context of messy real world implementations that are removed from the purity of their theoretical underpinnings. Claiming too strongly risks undermining confidence in quantum technologies when examples of their fallibility are later revealed. Conversely, making claims that are too weak risks underplaying the significant transformations that quantum technology may bring about, and the work that we need to do to prepare for them.

But the very complexity of the interacting components leads to unintended and sometimes regrettable consequences, which in a complex real world situation are hard problems to solve. Information and communications technologies introduce new levels of interconnectedness which cannot be controlled by well-established principles of social, organizational or technological separation (Kallinikos, 2005). Attempts to overcome unexpected consequences by more technology can create a vicious circle of increasing complexity. And the more powerful the capabilities of a technology, especially in relation to communications, cryptography and cryptanalysis, the more likely it is that it will become entangled in complex regulation.

Quantum cryptanalysis

Quantum cryptanalysis and quantum cryptography have been the poster children of powerful and transformative implications of quantum technology, pointing out both the potential that a quantum computer will have for "breaking" existing internet encryption on the one hand, but simultaneously providing new powerful forms of cryptography which could in principle never be broken, on the other. These visions carry an implied paradoxical symmetry in the way that they are framed as being simultaneously devastating and beneficent, which may contribute unhelpfully to the mythology surrounding quantum technologies.

The Internet could be "broken" (Biever, 2013), in the sense that existing Internet security would no longer be trustworthy, if existing public-key cryptography algorithms could be cracked using a quantum method known as Shor's Algorithm and some other quantum algorithms. This would undermine the security of all e-commerce transactions and encrypted communications on the Internet – effectively destroying most applications that make the Internet useful. This is one of the first identified examples of an actual application, if not one of the first that will be implemented in the short term:

> *"what is a quantum computer going to do for, for anyone? And the usual things that people say are all good examples, whether it's code-breaking, which gets a little bit too much attention but it is a very strong example, but not, perhaps, very helpful to you unless you're the NSA,…"*
> – (NQIT senior manager)

The practical significance of this risk may, in the end, be limited, but this does not mean this it does not have important implications. There is currently debate as to whether the NSA has, or is developing, such a quantum computer (the consensus is that they do not, but have a budget and are trying) (Koblitz and Menezes, 2015). However, without the active involvement of cryptographic researchers, the threat of "breaking the internet" from a quantum computer capable of running Shor's or other quantum algorithms would sooner or later become a reality, and this has spurred the emergence of research into forms of "post quantum cryptography" - seeking cryptographic solutions that are resistant to quantum cryptanalysis.

There are a number of points to derive from this:

Firstly, "post-quantum" cryptography is a response to the mere potential of quantum cryptanalysis, even though quantum computers of a sufficient size are still many years from being realised. This reinforces the idea that a technology does not yet have to exist before eliciting a response; quantum technologies of many kinds always have the potential to elicit a powerful reaction, even when at their earliest, embryonic, stages.

Secondly, the response of "post quantum" cryptography represents a form of RRI in action (which we discuss later in section 3.4.2) in the way that consequences are anticipated for cryptanalysis and remedial action is taking place now, in advance and anticipation of such an eventuality.

A third point is that the omnipotence of quantum cryptanalysis, suggested by the phrase "breaking the Internet", is not actually borne out in practice, since, as we have said, there are existing and developing cryptographic approaches that are not (yet, at least) susceptible to cryptanalysis by quantum computers. The work of the post-quantum cryptography community is as much to explore how existing resistant cryptographic approaches can practically do the job as effective as current asymmetric key distribution solutions, as it is to find new solutions.

Nevertheless, post-quantum cryptography cannot be assumed to be a cure for the threat potentially posed by quantum computing. The analysis of which security protocols are or are not vulnerable is by no means trivial; for example, it was recently found that one particular lattice-

The response of "post quantum" cryptography represents a form of RRI in action in the way that consequences are anticipated for cryptanalysis and remedial action is taking place now, in advance and anticipation of such an eventuality.

based cryptosystem is vulnerable to a quantum-based algorithm in polynomial time (Bernstein, 2016b). Whether other similar schemes are vulnerable is still unknown. And, evidently, any such cryptography must also be equally strong against classical cryptanalytic attacks.

Even if a quantum-proof cryptosystem were immediately available, and if there was a great deal of confidence in its invulnerability, such a scheme might have performance and usability problems. And beyond this, deploying such a scheme to the millions of Internet servers and clients, many of which are never updated, would be a huge task. Furthermore, it is not certain that any new cryptographic scheme would be easily implementable in existing Internet browsers, apps, servers, and communication protocols.

Another reason that the potential to crack existing Internet security is already making its impact is that the capture and storage of extremely large quantities of data, even if it is encrypted, could allow malicious organisations to intercept and store data from encrypted conversations until such time as powerful quantum computers are able to decrypt it (Bernstein, 2014, Huttner, 2016)). Data which must be kept secret for a long time is therefore vulnerable for the future, and needs to be encrypted in ways which are believed to be quantum-secure, without reducing its resistance to classical decryption methods.

Quantum cryptography

Quantum cryptography holds out the promise of security which cannot be broken even by an adversary with unlimited computer power.

While quantum computers raise the threat of subverting existing forms of security, a parallel line of research – although not directly related – is the development of powerful forms of cryptography making use of quantum properties. The theoretical basis has been developed since the 1980's; there are already some practical implementations, and widespread commercial availability is likely to emerge significantly earlier than quantum cryptanalysis capabilities.

Quantum cryptography - of which quantum key distribution (QKD) is the most well known instance – holds out the promise of security which cannot be broken even by an adversary with unlimited computer power ("information theoretic security"[21]).

There is debate in the information security community as to how significant this is in practice, because – at least, until/if Shor's algorithm and other algorithms are able to break mathematical cryptography – existing symmetric and public-key algorithms are already the "strongest links" in the security chain, and therefore not the most likely to be a source of vulnerability; according to this argument, for the purpose of building more robust security, resources are better used elsewhere (Schneier, 2008).

Nevertheless, quantum cryptography is an active area not only of research but also of commercial interest and is already available as a product[22]. Even if existing cryptography is the "strongest link" in the chain (Schneier, 2008), and therefore not the most likely to be a source of vulnerability, there are several reasons why quantum-based cryptography may be preferred over classical.

Quantum cryptography would quickly become of great importance if quantum computers were to become able to undermine existing forms of Internet security, and if, by that time, post-quantum cryptography was not yet able to provide security by classical methods.

Moreover, quantum cryptography has properties which are not available to other forms of cryptography: unlike classical secure communications, which can be compromised without the

21 https://en.wikipedia.org/wiki/Information-theoretic_security

22 For example, ID Quantique http://www.idquantique.com/

legitimate users being aware of this, quantum cryptography would detect any eavesdropping, and, under certain conditions, it is possible to communicate securely even if the quantum devices are not trusted, or in imperfect conditions (Mayers and Yao, 1998).

A third advantage follows from the speed and security of key exchange. QKD, as the name implies, is involved at the key exchange part of the secure conversation; in a typical use scenario, the key exchange is quantum, and then the encryption, once established, is a classical symmetric scheme such as AES. Here QKD nevertheless has several advantages over existing key exchange using asymmetric cryptography (Alléaume et al., 2014). Firstly, QKD is not vulnerable to compromise of private keys; secondly, QKD could enable a more rapid rate of key renewal than classical schemes, an important consideration in increasing the security.

QKD, however, has a number of technical restrictions, limited by noise to around 150 km over optical fibre, with greatly reduced speed of key exchange over longer distances. The record so far is 307 km using a new form of photon detectors (Korzh et al., 2015). QKD over free space has been demonstrated over 144 km between two of the Canary Islands, and this leads to the promising approach of transmission via satellites, taking advantage of the lower atmospheric density. This is leading to a "Quantum Space Race" (Merali, 2012, Jennewein and Higgins, 2013, Vallone et al., 2015) - bringing China into global prominence, involving collaboration and competition between Europe, Canada, and China.

But, assuming for the sake of this argument that unbreakable security could actually be achieved, this would have a widely-discussed consequence that it would make it harder for legitimate national security and law enforcement agencies to intercept the communications of organised criminals and terrorists. This is serious concern, raised by an interviewee for example:

> *"GCHQ require the power of legal intercept. If you make it completely encrypted in a quantum way that nobody can break, then you deny them that. Now, in practice, of course, it's not as simple as that, because if you can go to the provider you can probably still get access, ok, so it's not necessarily true that it's secret … but it's the sort of thing that [researchers] don't even think about"* – (Research Council employee)

Assuming that unbreakable security could be achieved, this would make it harder for national security and law enforcement agencies to intercept communications.

A simple response is that quantum-secured communications will always be between "trusted nodes" which provide for legal intercept – this was the argument used in the Quantum Manifesto which successfully called for the European Commission to launch a quantum technology programme (European Commission, 2016b). But at the same time as enabling lawful intercept, this introduces a rather clear security weakness; and probably does not solve the initial problem, if even one part of a QKD network falls under the control of criminals and terrorists. How "trusted" are the trusted nodes, who defines the trust, and who does the trusting?

Interception is not only of interest for legal intercept by law enforcement agencies. There is a lesson from a recent example in the practical implementation of secure protocols; it has been announced that the forthcoming Allo messaging app from Google will not by default implement end-to-end security[23]. This announcement highlights interception practices by actors other than security services, in this case to support Google's Artificial Intelligence, which will be "super useful … [but] to help it help you you'll have to entrust it with your chat messages.". In Allo, whether to allow end-to-end encryption is an option for the user to choose, but Google are clear that they would really prefer that users do not choose to have too much security.

Regulation of quantum cryptography technologies is another point of intersect between technology and society concerns; we pointed out in the prior section that technologies with

23 Security and Privacy in Google Allo, blog entry by Google security Cyber Overlord Thai Duong https://vnhacker.blogspot.co.uk/2016/05

The lesson is that quantum cryptography will be underpinned by social, legislative, and regulatory systems as well as the technology, in which strategic decision of service providers will act to reinforce or undermine its power.

powerful properties may also lead to powerful effects within society, and also attract significant and complex regulation. Quantum cryptography (erroneously said to be synonymous with QKD) is explicitly listed in the list of dual-use goods whose transfer is controlled by the Wassenaar Arrangement[24]. ETSI, the standards organisation, has established an Industry Standards Group with the aim of transferring QKD from the laboratory to *"the real world, where business requirements, malevolent attackers, and societal and legal norms have to be respected"* (European Telecommunications Standards Institute, 2015).

The lesson is that quantum cryptography will be underpinned by social, legislative, and regulatory systems as well as the technology, in which strategic decision of service providers will act to reinforce or undermine its power.

Finally we return to the "strong claims" (Section 3.2.1) which surround quantum technologies in general and forms of quantum computing in particular. QKD claims to be not only strong; it claims to be unconditionally strong, because it is based in the laws of physics. But the laws of physics include Newton's law of gravitation (at least as a first approximation) and Maxwell's equations for electromagnetism; "theoretical", or actually unbreakable, QKD is not compatible with these laws (Bernstein critique (Bernstein, 2016a) although we might accept that the costs to Eve are higher than is realistically achievable – with current technologies). There is risk of generating a false impression of overwhelming potency, which, although justified as long as certain hypotheses are accepted (while others are set aside), depends on contingencies in context in order to be realised. The paragraphs above have shown that QKD is embedded in social processes, but in addition, all security schemes are circumscribed in practice, perhaps because of undiscovered security holes but more probably through failures and technical difficulties in implementation in use. Indeed, the practical challenges in extending QKD over long distances are a simple illustration of the challenges which separate theoretical perfectly cryptography from actually secure communications. There is weakness "around the edges", either technically in the form of side channels which leak information to an eavesdropper or figuratively in the sense of weak security practices. We cannot blame "the users" (Adams and Sasse, 1999). If users find it hard to comply with security policies requiring them to recall complex passwords, or to lock their screens every time they leave their desk (Sasse and Flechais, 2005), while they also find it hard to implement cryptographic schemes such as PGP (Whitten and Tygar, 1999), then even the strongest cryptography will not be invincible.

3.2.2 Verification

We believe that verification of quantum computing will be an important issue, particularly for early applications in the fields of simulation and modelling, for three reasons:

Firstly, because processes of verification and validation are already notoriously hard to apply in the context of scientific computing (i.e. in models and simulations).

Secondly, there is the challenge to of how certain results of quantum computations can be certified by entities that are unable to perform those computations themselves (Barz et al., 2013) - how can one check the results of a computation which can only performed on a quantum computer? We can express this as the question of whether a prover who has access to quantum-computational resources can convince a classical verifier that the prover can solve a given problem (Barz et al., 2013) (We note here that, clearly, some problems can be easily verified, even if it is not easy to find the solution; for example, if prime factors for a large number have been found, it is much easier to check that they are correct).

24 http://www.wassenaar.org/wp-content/uploads/2016/04/WA-LIST-15-1-CORR-1-2015-List-of-DU-Goods-and-Technologies-and-Munitions-List.pdf

A quantum simulator presents different kinds of verification challenges because in general there is not one result that can be verified (or falsified). It might be possible to compare results with a trusted classical simulation, or to compare several simulators. This will give us increasing confidence in the results. If the model, based on the simulation, is to be used to study a real system, then the results can eventually be compared against the real system. Rather than verifying the simulation, at each test we can state that the model is "not yet false". This, of course, is nothing more or less than the scientific method (Johnson et al., 2014).

The question of verification also extends to the question of demonstrating that a claimed random number sequence, for example from a quantum random number generator, really is random (see for example (Acín, 2014)).

Finally, how can we be sure that the results really are making use of quantum phenomena? These issues may contribute to a perception of a "mysterious" aspect to quantum computing, particularly relating to concepts of indeterminacy and chance, that may encourage scepticism about their capabilities.

Confidence appears to be an issue for early, specialised, quantum devices, for example, surrounding the D-Wave machine and its claim to exploit an underlying quantum phenomenon. On the other hand, incremental availability of smaller universal quantum computers[25] may help build and cement confidence over time.

NQIT has several strands of work aimed towards methods for verifying quantum computations. So called "blind" quantum computation (Barz et al., 2012) could be an important enabler of real-world quantum computing where a client with limited quantum capabilities delegates computation to a quantum server, somewhat similar to cloud computing, and does so in a secure way in which the server does not have knowledge of the input, computation or results.

The verification work in NQIT will allow trust and certification in results even from untrusted or uncharacterised devices. The user may also have privacy requirements which preclude sending data to a possibly untrusted third party. In parallel, there are forms quantum cryptography which allow for secure communication even if the quantum devices are not trusted, or even if the devices are under the control of an adversary (Vazirani and Vidick, 2014, Vazirani and Vidick, 2016).

The verification work in NQIT will allow trust and certification in results even from untrusted or uncharacterised devices.

But, to re-iterate the point that we made in the context of QKD above, the trustworthiness will only be as good as the social practices in which it is embedded. Ignoring these factors - placing unconsidered trust that the system must be invincible, that the results must be correct, and that there cannot have been any interception or deception - will make the systems less trustworthy, not more, and risks a high-profile failure or security breach, which could be catastrophic for public confidence in quantum technologies. Computer models are imperfect representations of reality, and so there will always be instances in which they fail to describe what comes to pass in actuality (Beck et al., 2009). A real example is the loss of trust, for some, in mathematical models to predict climate change, or much else (Saltelli and Funtowicz, 2014, Pilkey and Pilkey-Jarvis, 2007), largely driven not by a failure of the models themselves but by a selective use made of them by politicians and commentators. Models can be viewed as "truth-generating machines" (Beck et al., 2009) – and often are so viewed, when it happens to serve a purpose at the science-policy interface. In complex areas such as climate modelling – or any other system which we are likely to want to model – a sensitivity analysis would show very wide variation, and many of the underlying factors are uncertain (Saltelli and Funtowicz, 2014). Although crisp, apparently clear-cut results from a model might provide a convenient argument, in the long term they will lead to a loss of public trust in the modelling process.

25 For example, the recently announced IBM "quantum computing on IBM cloud" http://www-03.ibm.com/press/us/en/pressrelease/49661.wss

3.2.3 Sensors and sensor nets

Although quantum sensors and metrology are the focus of the Birmingham Hub in the UK National Quantum Technologies Programme, NQIT will make important contributions to quantum sensor research, particularly as part of the engineering and research in solid state nodes (diamond) and photonic networks of sensors. Within such networks, correlated states of two or more nodes, which can be widely spaced, provide enhanced precision per spin.

Sensors can have many applications[26], including

- sensing gravitational fields to detect sub-surface structures, such as underground pipes and so avoid unnecessary roadworks, or used in archaeology and oil exploration;
- accelerometers, enabling far more accurate inertial navigation without the use of GPS, useful indoors, below the sea or underground;
- magnetic sensors in medical settings, contributing to dementia research, easier MRI scans, detection of skin types and potentially enabling the brain to communicate directly with computers;
- In the remit of NQIT, microcavity-based sensors or "noses" could detect gases or pathogens. This could have applications in manufacturing, safety, security, and healthcare.

One way to consider the implications of new and powerful sensor technologies is to think of them in relation to the existing technological and social trends that are already being driven by innovation in sensor technologies. Cheap, ubiquitous sensors are partly responsible for the growth in personal or wearable devices, such as smart phones and Fitbits[27], and the growth of the Internet of Things.

Whilst inexpensive quantum-based sensors are further into the future, one can begin now to consider what the implications for higher fidelity sensors may be for these existing trends.

Whilst inexpensive quantum-based sensors are further into the future, one can begin now to consider what the implications for higher fidelity sensors may be for these existing trends. One possibility is that quantum accelerometers would allow more fine-grained location positioning. This could be in combination with GPS, but could also work indoors. Combined with other data, this would allow improved activity recognition, where computers not only know where you are, but also make a good guess at what you are doing. This could open up opportunities for innovative personalised services, but at the same time create new risks related to monitoring, profiling and social control. For example, insurance companies already strike bargains with young drivers who are encouraged to install a "black box" in their car to monitor their driving in exchange for reduced premiums. It is easy to extrapolate from this fairly benign application to more worrying situations that involve more significant erosion of personal freedoms, perhaps coupling high fidelity location and physiological measurements to tailored insurance premiums or rationed healthcare access.

The idea of brain-computer interfaces may seem extraordinary, but in fact research goes back to the 1970s and they have been used in cases of sensory-motor damage to enable communication or control of devices or, more invasively, to restore some level of vision. However, the interaction can be two-way; the brain could become part of an information technology system, including AI and connecting to the Internet. In time, the external systems could learn from the brain and pre-empt its needs (Wolpe, 2007). What if the computer became better at determining needs than the human who is supposed to be in control of it? In some disorders, brain interfaces might be able to replace psychoactive drugs for effective treatment; but would people with these disorders retain the right to these accept or withdraw from these treatments? On the other hand, maybe "enhanced" brains would become a desirable asset, with prosthetic enhancement giving advantages to those who have access to it, further entrenching inequalities (Heersmink et al., 2014, Human Brain Project, Undated).

26 The Birmingham-led Hub has a list http://quantumsensors.org/innovation/
27 http://www.fitbit.com/uk

Worrying uses of technology may not always arise from an extension of an existing trend, they may also emerge unexpectedly where new technologies are used in unanticipated ways. For example, a participant at one of our RRI workshops suggested that a gravity-based sensor in a stationary mobile phone is then no longer a locative device as an accelerometer, but instead one that can "see", or sense, what is happening in the room. Thus improved sensors, coupled with existing paradigms, such as the rapidly growing Internet of Things (IoT) and wearable computing, have implications by introducing new forms of accountability, creating new potentials for social control via activity aware computer technologies, and for surveillance in unexpected ways.

Improvement in sensor technology often plays into debates around privacy and how privacy is traded against other social values, as is visible in the debates around the use of full body scanners in airports[28]. Other applications of improved sensing capability may also have ambivalent outcomes. For example, increasing the sensitivity of medical tests may lead to benefits through earlier diagnosis of disease, but may also risk "overdiagnosis" where the increased sensitivity acts to alter the threshold for treatment, potentially leading to unnecessary medical interventions[29]. We explore some of the possible challenges posed by new sensor technology in defence applications later in this document in Section 3.3.2. As another example, a smartphone or smart watch could be able to give exact information to its user about location and match this with other sensor-based data such as temperature, motion, or heart-rate. Sent to a database or stored in the Cloud, this could give useful feedback to an individual, but also be mined for targeted marketing or, for example, used by medical insurers to tailor premiums against lifestyle choices. Note that the data analysis itself is not necessarily quantum, although here, too, quantum computing has been proposed for solving hard problems in large and complex sets of data (Lloyd et al., 2016).

The issues outlined above provide concrete starting-points for thinking about the ethical challenges attached to sensing and the sorts of care needed to achieve the best outcomes from sensor technologies whilst reducing the impact of potentially harmful side-effects. We believe, too, that these examples help illustrate the RRI process, since it is possible to see the chain of connections between the hazards of existing sensor technologies, how these may be exacerbated with the introduction of high fidelity quantum sensors, the sort of public debates and controversies that might be provoked, and anticipate the sorts of regulation that may help to ensure the technology delivers benefit and remains acceptable.

These examples help illustrate the RRI process, since it is possible to see the chain of connections between the hazards of existing sensor technologies and anticipate the sorts of regulation that may help to ensure the technology delivers benefit and remains acceptable.

3.2.4 Quantum simulation and modelling

Simulation and modelling are likely to be early applications of quantum to forms of computing; although not universal quantum computers, they will be able to solve some currently unsolvable problems. This is a significant and under-reported aspect of quantum computing, despite the attention given to the related D-Wave implementation. The round table on RRI in quantum (Faullimmel and Stilgoe, 2015) noted an unevenness in media coverage of quantum technologies generally, with the more futuristic, robotics and Human-Machine Symbiosis receiving more attention than potentially transformative but more mundane applications.

Simulation

Quantum digital simulation will probably be the first style of quantum processor; not a universal quantum computer but powerful simulators. Years before the very long-term and ambitious achievement of a universal, scalable quantum computer, and probably within the timescale of

28 For example, http://www.nytimes.com/2009/12/30/us/30privacy.html?_r=0
29 For example, http://www.nhs.uk/news/2012/05may/Pages/overdiagnosis-in-medicine.aspx

Quantum computing could speed up the process of scientific discovery itself.

NQIT, digital quantum simulation will have important practical applications to problems that are classically intractable - not least in modelling quantum processes themselves, the insight made by Feynman (1982) which is among the foundations of the concept of quantum computing. These hold the potential to accelerate discovery in areas such as drug design and the development of new materials (Johnson et al., 2014), rather than trial and error discovery.

There are grand challenges facing humanity, including climate change, and hard computations such as understanding protein folding with important medical implications. Computational chemistry is an established area of research since the development of high-powered computing. Potentially, these techniques, enhanced by quantum computing, could speed up the process of scientific discovery itself. But we know from other areas in which science moves faster than society is able to accept, for example synthetic biology and the case of genetically modified crops, that faster discovery must work alongside responsible innovation, identifying areas of genuine concern before they become blockages – especially when science is working with the fundamental fabrics of life and nature.

However, simulators are unlikely to replace other methods; there are many examples of scientific progress without simulators (Johnson et al., 2014). Simulation is only as good as the model, as we discussed above in Section 3.2.2; and indeed, one use of simulation is to test the model, rather than to reveal insights into the system of interest.

Simulation raises similar challenges of verification to those raised by computation as more generally understood; how to verify the results of analysis in situations in which a trusted classical simulation is not available.

Optimisation

Finding the minimum value of a function over a space of candidate solutions is equivalent to finding the optimal solution to a problem with a number of variables. Quantum methods for solving optimisation problems are already available, notably in the D-Wave machine, already being used to enhance consumer products such as Google Glass (Davies, 2013). This does not purport to be a universal quantum computer, but instead uses a technique known as quantum annealing, making use of what is called adiabatic quantum optimisation. A system is initialised to a simple state and the conditions are slowly ("adiabatically") changed to reach a complex final state, from which the solution can be "read off".

There are many important problems in optimisation which are hard, or in practice impossible, to solve using classical computing methods. For example, a supermarket wanting to maximise the efficiency of its supply chain, or a city authority responding to a civil emergency by closing some roads and opening others. In such cases, even if a classical solution is feasible, speed is essential.

3.2.5 Trust in Quantum Computing

We have been considering the phase of development where quantum computing and related technologies begin to roll off the conveyer belt and enter more fully into the public consciousness as they are incorporated into higher profile application areas: cryptography, cryptanalysis, sensing, metrology, and simulation; and, related to these applications, there are concerns such as verification of the results and the extent to which they really are quantum. Overall, we emphasise and the importance and limitations of the strong claims made on behalf of quantum technologies.

As these emerging applications begin to find their home in the world, we anticipate that they will also begin to face new challenges as well as success stories. Perhaps, like the D-wave machine, some applications will struggle to establish their quantum credentials. With others, it may be the very quantum nature of the underlying process may raise concerns as to whether the results they produce are indeed reliable. Yet this phase of increased proliferation of quantum applications will be crucial for building confidence for quantum technologies more widely, and paving the way for quantum computation which is still some years ahead.

Whilst verification and demonstrations will be important for generating trust within technical communities, building confidence within user communities and wider publics will depend upon a broader programme of engagement. Crucially, expectations will play a role in how quantum technologies are received. Heralding quantum technologies as possessing potent and inviolable capabilities underwritten by the laws of physics risks fear of powerful and perhaps uncontrolled transformation, on the one hand, and a loss of confidence should these be capabilities be perceived to be compromised, on the other. Thus there is tension between certainty provided by the laws of physics backed by a strong theoretical analysis, and uncertainty around how this will play out as they move into the world and encounter complex social realities.

Attention is required to the technical, social, organisational and regulatory frameworks that need to be in place both to effectively realise potent quantum applications, and also to ensure these potent capabilities are exercised under situations of control in so that they do not fundamentally distort existing social relations.

These issues on the horizon, already discernable but not yet live, have in common an element of uncertainty, and, consequently, a requirement for trust on the part of the wider public as well as researchers. Trust in quantum technologies will be key to gaining broad public acceptance of the technologies; and this trust, in turn, is impacted by the immediate uncertainties and narratives of quantum as being mysterious and hard to understand which we investigated in Section 3.1 as well as emerging tensions between certainty and uncertainty as discussed in this section.

This phase of increased proliferation of quantum applications will be crucial for building confidence for quantum technologies more widely, and paving the way for quantum computation which is still some years ahead.

Uncertainty in conditions of apparent certitude

We use the term "apparent certitude", not because there is doubt about the strength of the theory, but because this certitude is always contingent not only on implementation but also on the ways in which the technologies are used in practice.

In the case of quantum technologies, harnessing the laws of nature involves pushing the boundaries of scientific and engineering knowledge. No matter how perfect the laws themselves, implementation will always be less than perfect. There is always the potential for mundane mistakes and physical failures as the technology is developed and appropriated in use. There are established techniques to minimise and mitigate the risks implementations in classical computing and theoretical work in verification in quantum, but even if implementation could be made "provably perfect", actual use would not be flawless because *"what is done by computing"* (Kling, 1992) fits into a complex socio-technical system in which "the computer" itself, while a key enabler, is only one part.

Thus, highly-accurate sensors may provide a witness to activity, or may provide sufficient accuracy to guide a developer or detect gas or oil, but this is decontextualized data, which, even if perfectly accurate in itself, still does not provide certainty sufficient for unquestioning use as legal evidence. We can anticipate consequences if this data were to be combined with other datasets to empower or control individuals, for example by granting or denying access to goods or services.

The far greater accuracy in the case of sensors, or, to use another example, the potential for unbreakable secure communications, is in a tension with uncertainty introduced "around the edges" of the technology

Trust at this level, then, accepts uncertainty, but it is not "blind trust", but rather is "weak inductive knowledge". Giddens (1990) argues that this is not really trust at all, since it does not require "faith" on the part of the person doing the trusting, merely a conditional acceptance even though there is still a level of risk. Indeed, this subsection is called "uncertainty" rather than trust; trust, properly speaking, is discussed in the subsections which follow.

Trust in the outputs from Quantum Computing

The wider public and users of quantum technologies need good enough grounds to be confident of the results and to develop trust in quantum technologies.

Section 3.2.2 above shows how, quantum computing raises new challenges when computations are performed which cannot be verified other that by another quantum computer. These issues are revisited here to consider their social implications, specifically in terms of their potential to build or to weaken more widespread trust in the technologies. Scientists need to be convinced of the validity of the results, but the wider public and users of quantum technologies also need good enough grounds to be confident of the results and to develop trust in quantum technologies. Formal verification approaches, such as those outlined above, are hugely important – but it has to be borne in mind that technical verifications are one thing, but turning those verifications into trust is amongst wider group of stakeholders, is another.

Trust, then, even at the level of technical verification, will in part depend on acceptance, and be tied in with politics and culture as much as it will depend on objective demonstrations of veracity. So it will be important to consider an incremental strategy that keeps in mind wider questions of what will NQIT need to "prove", to whom, and when?

Second-order trust

The need to trust will extend to the first direct users of quantum technologies who may be, for example, pharmacologists or materials scientists using quantum simulation, or companies using Grover's algorithm to search for results in very large volumes of data, or using optimisation techniques to enhance a consumer product.

These immediate users are trusting in the outputs from the technology which they may not be able to verify directly for themselves. But beyond them are those who are impacted by the technology even though they have no direct contact with it, and maybe are not even aware of it. This is a **second-order level of trust;** trust in this case rests on the correctness of principles of which the person is perhaps only distantly aware and not in a position to check (Giddens, 1990). People are being asked to develop trust not only in the quantum technologies themselves, but also in infrastructure with has been engineered on the basis of quantum computing calculations, drugs and materials designed using quantum simulation, or, potentially in future public policies tested with the help of quantum-based models. And this may be in a situation in which the infrastructure is in the hands of a few large organisations, unavailable to ordinary citizens who will be required to trust the outcomes of systems over which they have no direct control; there is another kind of trust, not only in the technology, but in the organisations and the people in them.

Trust in Researchers

There is one another area of trust, which has to do with the trust which the funding bodies place in researchers, and, in turn, which the public (represented by the government, acting through departments and ministries) places in the funding bodies. This is an aspect of responsible research and innovation; this includes fundamental scientific principles of honesty and integrity, but also a responsibility to make the best use of public money. Rather than close control over details of research, funding councils use a light-touch approach, recognising the pressures on researchers to "self-police", and also realising that research is not always a straight line from theory to practice, but is often messy, contingent, and with results which are not the same as those expected at the outset.

Trust is an aspect of responsible research and innovation: this includes fundamental scientific principles of honesty and integrity, and a responsibility to make the best use of public money.

Of course, the responsibilities of researchers are not unique to quantum, but, for NQIT, as a major part of a technology-oriented programme, there is a stronger than usual link between science which is far from production and technology which may have direct applications. Although the funding is directed towards taking science towards the development of technology, doing so in the case of highly innovative and ambitious technology in NQIT requires more fundamental research in physics and materials, as well as meeting challenges in engineering which are themselves a form of research, pushing the boundaries of what is possible. The relationship between science and technology, with relevance to responsible innovation, is considered in more depth in Section 3.4.

Recommendations for RRI in emerging areas of research interest and issues on the horizon

There are high expectations for the potentially transformative implications of quantum technologies, for which the early applications identified in this section will be the first precursors. This, together with the need for adopters of these technologies to place their trust in them – and others who rely on their results, even though not directly implicated – creates both opportunities and risks. We make the following recommendations for responsible research processes:

1. **Consider the implications of the potency of quantum applications.** On the one hand these applications are very powerful, raising concerns that they may need to be tamed. On the other hand, quantum processes are themselves fragile, and also they must interact with non quantum components and systems subject to normal failures and implementation errors, and so there is a risk of a "fall from grace" if expectations are set too high. Managing this combination of potency and fallibility will require a balance between enabling the advantages of new potent capabilities, while at the same time ensuring that these capabilities lead to socially desirable outcomes. This recommendation is partly about careful use of language, and setting the right expectations. The message is that there are potentially great benefits, but that between current state of the art and achievement of this potential there remain challenges in implementation. There will be disappointments but also unexpected successes.

2. **Work to ensure trust in quantum computing by the users and by third parties who rely on these results.** Being able to demonstrate computation that makes use of quantum effects will be important for scientific acceptability - but there is a risk that others beyond this community may distrust applications or services based on quantum technologies. Building this broader trust will depend on a continuation of the engagement work towards developing positive social narratives around quantum computing. The realisation that results of quantum computations may not be easily verifiable could

be another area for misunderstanding; make sure that the verification work is not focussed solely on technical audiences, but reaches out to broader constituencies. As with the narrative work, there will be a need to identify myth and work towards demystifying quantum computing, including creating demonstrations of quantum computing capabilities in ways that are meaningful for broader audiences.

3 **Take the time for learning and refining appropriate RRI processes.** Acceptability and benefits from early applications will set the pathway for later applications. Be honest about the risks as well as the benefits – which are in many cases two sides of the same coin. Early applications of NQIT technologies will be ambassadors or harbingers of universal quantum computing technology, presaging the sorts of social transformations that may follow. They present an opportunity to show how quantum technologies can be safely and controllably brought into society. Successful management of these technologies presents a significant opportunity for cementing trust in quantum computing. We can identify a range of issues for quantum cryptography, sensors and simulation, which may feed existing controversies, intensify existing risks or threaten how interests are balanced within society. There is a risk of public resistance, but, conversely, managed well, these application areas will become use-cases for processes of successful social embedding.

Taking quantum computing to the Oxfordshire Science Festival - will the public trust and embrace quantum computing? / *Hannah Rowlands*

3.3 Future challenges from quantum computing in use

Further downstream, application areas of quantum computing will start to emerge and eventually become widespread in use. These are not yet specific applications, but areas of application which carry the potential for positive as well as less desirable implications arising as these applications are put into practice.

Since these technologies are not yet available and it may be years before they are widely adopted, and since not only the applications but even the form in which quantum computing will be viable are still uncertain, it is not possible to predict in any reliable sense the ways in which the technology will be appropriated and which new possibilities will arise from the combination of quantum computing with other quantum and non-quantum technologies; however, some technologies are already closer than universal quantum computing; for example, Section 3.2.3 suggests possible future uses of quantum sensors and accelerometers.

However, this does not imply that there is nothing we can or should do to address these as-yet unforeseeable implications. Firstly, we can **draw lessons from existing technologies**; for example, some likely application areas in Machine Learning or in defence and national security will open new potential avenues and enable users of these technologies to solve some currently intractable problems, but we can already see what these might be and how they might extend existing concerns and benefits from these application areas.

Secondly, building on the realisation that the AREA framework starts with anticipation, we can **start to imagine possible applications**, and to analyse what the implications of these might be. Even if our imagined futures are not realised in the way we expect ("flying cars"), anticipation is a useful exercise in itself to help us to be prepared for whatever the future does hold.

Thirdly, even though it is useful to imagine and be prepared for the future through these detailed scenarios, it is also important not to get too entangled in the technical intricacies; at the same time as anticipating particular technical capacities and their social implications, we also need to take a step back, and consider what responsible research and innovation involves and how we can enact it. (Owen et al., 2013). We need to reflect, in other words, on the **motivations which direct research**, and be prepared to change the direction or emphasis of aspects of the research and innovation pathways.

This wider vision is a reminder that many of the concerns around quantum technologies are also concerns with classical computing and other existing technologies. None of the areas of concern which we discuss below is specific to quantum computing, although quantum may heighten certain aspects of them. This does not lessen the significance of quantum computing in transforming existing trajectories; it will be in the intersection between quantum, classical computing, and other developments such as the Internet of Things and mobile computing that we can expect to find the most significant developments.

Social transformations

As these applications from quantum computing emerge, working alongside other existing and new technologies, we can expect that they will engender the kinds of social transformations that have followed other major technological developments The social transformations which we can start to discern can be considered in terms of their contributions to what we have called the algorithmic society; this includes applications to machine learning but more specifically, the increasing interaction between technologies and our real lives; applications to defence and national security, with implications around surveillance, automated decision-making, and the international balance of power; and the as-yet undecided ownership models through which quantum computing will become available outside research laboratories.

These categories are not distinct, but overlap and influence one another: the appropriate use of surveillance, for example, an important concern in discussions of national security, is also a major feature in uses of the intertwined use of algorithms and machine learning; ownership and control of quantum technologies may lead to imbalances of power between nations, but also within them, if large corporations and national agencies have use of these systems while smaller organisations and individuals do not. The overall shape of this section, showing the application areas we can anticipate and the issues that might arise from them, is illustrated in Figure 4.

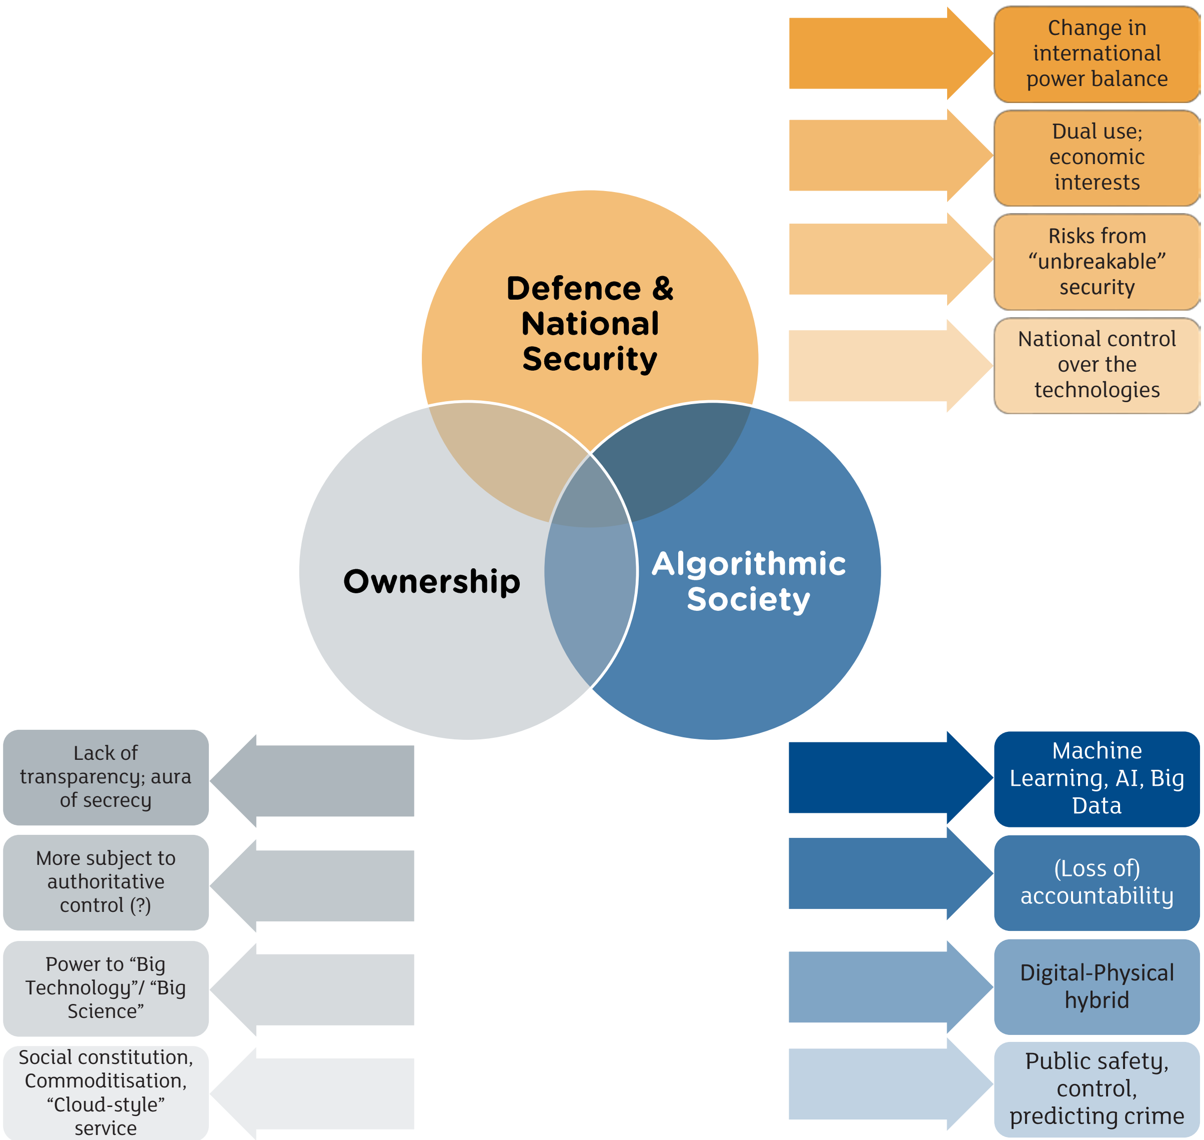

Figure 4: Areas of social transformation around quantum computing

3.3.1 Machine Learning, Big Data and an Algorithmic Society

In already-existing computing, our lives are increasingly shaped by unseen algorithms, in Google searches, Amazon recommendations, Facebook friend suggestions, locating Uber drivers, and inevitably for targeted marketing. It is not only at the online interface that algorithms impact our lives; for example, Section 3.2.3 outlined some emerging uses of sensors and accelerometers such as recording devices in cars which monitor driving patterns and adjust insurance premiums accordingly. Quantum computing can be expected to intensify emerging trends towards an algorithmic society, and we can already perceive, some of the trajectory along which quantum computing may travel by pushing the boundaries of the possible, and extending the reach of algorithms even further.

Quantum computing can be expected to intensify emerging trends towards an algorithmic society.

Simply put, an algorithm is a self-contained set of instructions to be performed. In principle, there is no room for doubt about the outcomes of the algorithm, given a set of data (Markov ref1954). However, the apparent neutrality of algorithms does not necessarily lead to neutrality of outcomes. Machine learning (ML) does not follow a static set of instructions, explicitly written by a programmer. Recently there have been findings that search engine advertisement results are significantly more likely to suggest a link to an arrest record for people with distinctively black names, (Sweeney, 2013) or to link to high-income jobs for men than for women based on Google profiles (Datta et al., 2015). Friedman and Nissenbaum (1996) have analysed forms of what they label pre-existing (societal), technical, and emergent bias in algorithms; for example, an airline reservation system which gives preference to same-carrier over third parties in multi-hop journeys will tend to favour local over international carriers.

Yet the results of algorithmic calculations are increasingly intertwined not only with our online lives, but also with the physical world, for example in the "black boxes" which monitor driving behaviour (Section 3.2.3). In this blending of the virtual and the physical into "digitally hybrid lives" (Jordan, 2009), algorithms are powerful force in regulating action and behaviour.

Who, then, is accountable for the results of these algorithms (Woolgar and Neyland, 2013)? Who is in control? The Google engineers? Datta et al (2015) say that they would like the "assign blame where it is due"; but this is not at all simple. Are search engine providers complicit with advertisers, or are advertisers circumventing the rules of the site? Maybe the outcomes are not the results of deliberate programming, but that advertisers, or the marketing services which they use, are matching previous results to a cache of searches with similar properties, effectively "amplifying" (Agre, 2002) underlying societal inequalities. However, the risk is real: *"There's a real threat that the negative effects of algorithmic decision-making will disproportionately burden the poorest and most marginalized among us."* (Gangadharan, 2015).

Who is accountable for the results of these algorithms? Who is in control?

Adding to the threat is the problem that the connection between input and output, mediated by algorithms, is opaque; partly, because it is an important part of the business model of the marketing companies, and partly because the results arise not from a simple procedure but from a complex array of human and machine configurations (Suchman, 2007), no doubt including contributions from machine learning.

Moreover, in the case of machine learning, it is not always the case that the machine is able to give an account of its reasoning. At a workshop in an NQIT Project Forum, participants discussed a scenario involving a doctor making a diagnosis with the aid of an ML assistant. One of the concerns raised was whether the computer would be able to justify its diagnosis, and whether human validation might still be preferable. For example, the ML assistant might provide a diagnosis which the doctor had not considered, but which she or he would subsequently validate. The doctor will know the patient; the doctor will be able to make judgements based on experience.

But the computer could also learn to know about the patient, with the advantage that it will not forget critical details, and could be kept up to date with medical advances.

Machine Learning is a high-profile topic in relation to quantum computing, even though functionally separate. We can begin to anticipate some implications if the power and capability of algorithms and ML were enhanced much beyond the current capacities. The QuOpal project[30], also hosted at Oxford, is exploring this interface. Quantum algorithms as subroutines, with most of the work being done by classical computation, could outperform purely classical methods for some ML tasks; eventually, purely quantum ML might emerge in which the data itself and the entirety of the ML algorithm is quantum – which would enable the use of quantum algorithms also in the "learning" task (Fernick and Gheorghiu, 2016). The non-deterministic, probabilistic nature of quantum computing also seems to lend itself well to ML, compared with early ML attempts which attempted to model human intelligence as deterministic, syntactical machines. Fernick and Gheorghiu (2016) give examples of ML processes where pure-quantum machine learning could offer exponential speed up over classical, although with some important caveats: there are many hurdles before ML will be a primary application of quantum computing (Adcock et al., 2015).

Given the influence of algorithms already in our lives, there is the risk that AI will be such a powerful tool that whoever controls it will have excessive and unaccountable power (the implications of who owns, controls or has access to quantum computing are discussed in Section 3.3.3), even if we don't expect machines to become so intelligent that humans lose control over them (as feared by Bostrom and his colleagues (Muehlhauser and Bostrom, 2014)). This power could become increasingly cemented in our lives by trends such as the Internet of Things and new forms of quantum-based sensors. Another scenario is that, in a parallel with classical computing today, we could become so reliant on AI that in practice it would be impossible to turn it off, even if negative effects become apparent.

In the related but different issues surrounding Big Data, it is likely that quantum computing algorithms will be one of the essential enabling technologies for the speed and rapidity required to analyse the vast datasets which are already being collected by state agencies and large companies. It may be the case that "*Grover's algorithm will not let Google search the Internet more quickly*" (Bernstein, 2016b) – but here, again, what actually will and will not be possible is not yet clear (as we have discussed in Section 3.1.1), since on the other hand quantum algorithms have been shown to be useful for some (topological) big data analysis (Lloyd et al., 2016).

New forms of data processing enabled by quantum computing could in turn change the nature of this data, and these relationships which are mediated and transformed by the data.

The fundamental political issue is the collection of this data, and what it means for the relationship between citizens and the state, and between individuals and corporations. But it is new forms of data *processing* enabled by quantum computing which could in turn change the nature of this data, and, with it, these relationships which are mediated and transformed by the data.

Combined with powerful advances in ML, the ability to analyse far greater *volumes* of data of many different *forms* very *rapidly* will raise ethical questions which, unlike the questions of privacy and freedom which overlap with existing political debates (governing data protection with a much more limited concept of what is meant by "processing" of data), are qualitatively new.

For example, advanced facial recognition could enable some new services, such as a shop assistant immediately "knowing" a customer who enters a store (the shop assistant might be a robot, of course). ID checks in public places, to respond to terrorist threats for example, would

30 http://www.quopal.com/

be obsolete, because individuals would be immediately recognised from images on security cameras (Schneier, 2015). Identification, however, is only the front-end, because behind the camera could be very large interlinked databases; companies or law enforcement agencies will be able to make inferences, from a variety of sources, of people's likes, political affiliation, sexual orientation, and more mundane preferences. The irony is that people are providing this valuable and often sensitive personal data, along with photos of themselves just as needed for facial recognition, absolutely free and without any coercion, on Facebook and other social media sites.

Rather than collecting huge databases of largely irrelevant meta-data, there is the possibility that agencies will be able to target their data far more accurately – no need to look for the needle in the haystack, because the needle will be immediately locatable. If really accurate database searches enabled near-universal surveillance, but with the assurance that targets of interest could be much more precisely identified, should ordinary people be less concerned, knowing that their data will be routinely screened but that no human would need to see details of the personal lives of innocent people? How would society cope with the inevitable false identifications? How would an innocent person be able to acquit himself or herself in the face of an all-knowing and apparently irrefutable algorithm?

One high-profile application of computer technologies is in predictive policing, marketed commercially as "PredPol"[31] – aiming to predict crimes, individuals who might commit a crime, victims of crime, or *profiles* of likely offenders (our italics). Is it acceptable to increase surveillance and perhaps intervene physically against individuals who may have been detected on the basis of their profile rather than on any actual intent to commit crime? The idea that there are "born criminals", who should be punished or at least confined before they commit any offence, goes back at least to the 19th century (Rafter, 1997).

Quantum computing is implicated in these issues because it will potentially enable more accurate facial recognition, and, alongside this, link identity to machine learning and more powerful database searches. For example, Wiebe Kapoor and Svore (2014) describe an approach to apply quantum computing to deep learning, the technique used by Facebook to recognise photographs.

This emerging technology, with or without quantum, could be subject to regulation – there are positive as well as more unwelcome applications, clearly, and these could be set out in a widespread engagement of stakeholders. Whether this will happen depends on the interests involved and whether there is public resistance. An existing example is the use of facial recognition at passport control, introduced with very little public debate. Clearing passport control more quickly is very convenient, but that is only one side of the coin; these developments are set against the background of a powerful discourse of public safety and the need to protect national borders, demanding much stricter controls of which the new technology is only the outward symbol.

The risk is that quantum computing could be linked in the public discourse with surveillance and control. Already there are signs that the public discourse is reflecting disquiet over these issues. The article in the Sydney Morning Herald (cited in Section 3.1.4) claims that quantum computers will lead to *"the end of freedom as we know it"* (Cribb, 2016). In reality it not likely to be quantum computing on its own which could have any of the implications for freedom, privacy and surveillance that this article claims, but rather quantum computing and other technologies working together which could enable powers which do raise important questions.

The risk is that quantum computing could be linked in the public discourse with surveillance and control.

31 http://www.predpol.com/

3.3.2 Defence and national security

This section considers not only defence but also the increasingly important overlapping area of national security, addressing risks from terrorism and organised crime, particularly concerning intelligence services. The defence industry is likely to be an early adopter of many quantum technologies ((UK National Quantum Technologies Programme Strategic Advisory Board, 2015) National Strategy) Many of the applications of quantum computing and communications, such as cryptography and cryptanalysis, have clear military, national security, and surveillance applications.

Issues around defence and national security inevitably raise questions about the relative powers of the state, of one state relative to others, of the citizen's relation to the state, and of the protection provided by the state to its citizens, tensions which an interviewee laid out very clearly:

> *"there are lots of people who are quite uncomfortable with the idea of snooping powers, but they're glad when a terrorist incident doesn't happen. So, … on the one hand we've got the costs of some individuals feeling that their privacy is being damaged, or interfered with, on the other hand we have some people who are not dead, so there's a trade-off, and I think that's, how I generally think about the public good, ok, in this sense. And I think this is where it does get tricky with quantum tech, because we're not entirely clear about what the public would see as being a good, and being beneficial and what they would like to happen"*
> – (Research council employee)

Quantum technologies also present challenges as well as opportunities for defence and national security, however. "Unbreakable" communications over long distances would be a huge advance over existing cryptography, but we have already noted risk that unbreakable communications could equally provide a hiding space for criminals and terrorists (Section 3.2.1). Conversely, much existing cryptography would be at risk; we could be back to the era of exchanges of one-time pads.

Or consider the less widely discussed implications of networks of sensors that might be able to detect bunkers and silos or submarines, while submarines will be able to navigate much more accurately in situations where GPS is not possible. A submarine equipped with quantum navigation would be able to know exactly where it was, while, informed by quantum sensors (Section 3.2.2), satellites or spy planes would also be able to discover the position of enemy submarines. This capacity could affect the delicate balance of international relations and the logic of deterrence, based on no one side possessing complete knowledge. The Reagan-era Strategic Defence Initiative (SDI) "Star Wars" was similarly intended to unbalance the doctrine of Mutually Assured Deterrence (Stares, 1985).

In these ways, and in consequences which we are not yet able to discern, quantum technologies will work alongside existing technologies, interwoven with the policies of national governments and international treaties, to enhance or weaken peace, stability, and economic development.

Quantum technologies will work alongside existing technologies, interwoven with the policies of national governments and international treaties, to enhance or weaken peace, stability, and economic development.

Another aspect of these technologies is that, while there is a great deal of interest by defence and national security agencies in the outputs of quantum research, at the same time quantum technologies have benefitted from defence-led research; the first demonstration of QKD over more than about 1km was made by a group in the British Telecom research laboratory and the British Defence Evaluation and Research Agency (DERA, now incorporated into DSTL) in 1993 (Townsend et al., 1993) and what is now known as public key cryptography was first conceived (although kept secret at the time) by J H Ellis at CESG (Ekert, 2001). Meanwhile, the UK Ministry of Defence is contributing around £30 million over five years to the UK National Quantum

Technologies Programme, in return gaining access to the research outputs from the programme (Freeman et al., 2015).

In this sense these technologies are "dual-use" (Parliamentary Office of Science and Technology, 2009) since they will also have important commercial applications - for example, in finance. CIA venture capital firm In-Q-Tel (the "Q" is a James Bond rather than a Quantum reference) is known to be investing in D-Wave, alongside Amazon (Simonite, 2012). At the same time, competing geopolitics suggest a convergence between national economic and military interests; maintaining defence control while also allowing restricted non-military uses is reflected in export regimes such as the multilateral Wassenaar Arrangement with 41 participating states (mentioned in Section 3.2.1) and the unilateral US International Traffic in Arms (ITAR) Regulations[32]. There is a tension, then – explicitly recognised in the National Strategy (UK National Quantum Technologies Programme Strategic Advisory Board, 2015) - between the benefits of international trade and co-operation and the progress of science from open sharing of discovery, and, on the other hand, protection of national economic and defence interests and maintenance of international stability. The "Quantum Space Race" – the race to teleport photons via satellites, with practical importance in the implementation of QKD (Section 3.2.1), and other applications – is a current illustration of these tensions; on the one hand, research groups in the University of Science and Technology of China are collaborating with the University of Vienna (Merali, 2012), while on the other hand China, Europe, and the USA are vying for leadership in these technologies.

These technologies are "dual-use" since they will also have important commercial applications.

3.3.3 Ownership of quantum computers

The supporting infrastructure for quantum technologies, such as fabrication and testing facilities and the equipment required for operation, are currently prohibitively expensive for all but very large corporations or governments. The round table on RRI in Quantum Technologies in March 2015 (Faullimmel and Stilgoe, 2015) noted that the research landscape is dominated by "giants" such as the US Universities Space Research Association, NASA, and very large enterprises including Google.

This "Big Technology", closely tied to "Big Science", has given rise to some dystopian visions (Ihde, 1991, Winner, 1989) and has implications for RRI. The round table found that the dominance of large public and private organisations, and lack of public dialogue around their interest in these technologies, creates an aura of secrecy – exemplified by the involvement of the CIA and In-Q-Tel which we mentioned above - which may feed public disquiet. The very large investments required exclude many actors - benevolent as well as ill-intentioned - from access to it. From one point of view, the need for significant investment and (potentially restricted) knowledge may keep quantum technologies out of the hands of terrorists or criminals, who might otherwise be able to use them not only for unbreakable communications or to break existing encryption (as we have noted several times already) but also, for example, for large data analysis to identify soft targets, in a mirror image of the use by state agencies for predictive law enforcement. Governments, at least if they are elected, and companies have a level of public accountability, matched by their tendency towards secrecy. Perhaps there could be internationally-agreed controls on their use. However, with a few exceptions, (very high-cost, high-risk technologies such as nuclear weapons), controls on access to technology meet resistance and are short-lived.

However, ownership by a closed group of powerful organisations would also represent a major change in the *balance of power*, in favour of those who have access to the technology. For ordinary users, this potentially oligopolistic ownership, combined with the technical opaqueness of the infrastructure, could be perceived as a major *loss in power, vis a vis* the owners or

Ownership by a closed group of powerful organisations would also represent a major change in the balance of power, in favour of those who have access to the technology.

32 https://www.pmddtc.state.gov/regulations_laws/itar.html

controllers quantum computing power. To return to the well-known example of "breaking the Internet", reflecting on the implications of Shor's algorithm is useful as a "what if?" question, even if we believe that post-quantum cryptography will be deployed well before sufficiently powerful quantum computers become a reality. Currently, anyone with a web browser can have "good enough" security, and, if they choose, can use encrypt emails and Internet connections to maintain their privacy except in the most serious cases justifying legal intercept; but the availability of quantum computers, if not matched by new forms of strong cryptography, would give power to governments with no corresponding power for ordinary users. A researcher in NQIT commented:

> *"That's a situation where you might have a small number of powerful people who actually have a quantum computer and no-one else does, and they're able to decrypt a lot of stuff which is right now still extremely difficult to decrypt, so, yeah, that would be hugely disruptive"* – NQIT Researcher

The technology may become commoditised and this may change the ownership models radically.

Of course, the technology may become commoditised and this may change the ownership models radically. It may be the case that access to powerful quantum computing becomes widely available within a few years of its development. But, in turn, the processes of diffusion and commodification will again raise new RRI questions, around the *"political economy"* (Faullimmel and Stilgoe, 2015) of the technologies: who will be relatively advantaged, who disadvantaged, in what ways? If commodification takes the form of on-line access in a Cloud-like model (for which there is, indeed, already an early instance[33]), then what will be the business model? If chargeable on a pay-per-use basis, will price create a new form of digital divide? If not paid-for directly, perhaps supported by advertising or the value of the data, then what secondary uses will be made of the users' data? Will cloud-based access to quantum computing be part of an "eco-system", with mobile front-ends supported by a mixture of free and premium services where the model relies on consumer tie-in?

Recommendations for RRI to address social transformations

1. **Pro-actively encourage collaboration between scientists and engineers and RRI practitioners.** The strong potency of quantum technologies, which we discussed "on the horizon" in Section 3.2, will be made explicit in applications which will emerge as the technologies are appropriated by the market and by users. We cannot predict exactly what these applications will be, but we can foresee some areas of concern and possible tensions between diverse interests (Fisher and Rip, 2013). There are two dangers: that RRI will attempt to direct this process, and risk losing the support of scientists; or that a breakdown in understanding will lead to a kind of RRI shaped in the interests of a subset of stakeholders. Rather, RRI should work more or less formally together with others in the project to gradually match social provisions to new technologies, identifying potential areas of conflict and developing "mid-stream modulation" (Fisher et al., 2006), to ensure that the technology continues to reflect strategic visions and, if necessary, adapts the trajectory of innovation (Fisher and Rip, 2013). This should involve scientists themselves, RRI researchers, funding councils, the leadership of the project, and potential early adopters in open dialogue and interaction to resolve differences of interest between stakeholders.
2. **Start to develop forms of governance for these technologies.** As these applications emerge into regular use, it will be increasingly necessary to create the right social provisions (institutions, practices, regulations, governance) in which quantum computing may be made comfortably at home. In a scenario which may see ownership of the technologies centralised

33 IBM Makes Quantum Computing Available on IBM Cloud to Accelerate Innovation, http://www-03.ibm.com/press/us/en/pressrelease/49661.wss

in large corporations and governments, other stakeholders, impacted by but not in control of the technologies, may lose trust in them; particularly if the applications have a large influence over life in an algorithmic society, and/or if many of the application areas are implemented behind a veil of secrecy. At some stage, regulation may become necessary, to address specific risks, but such an approach is top-down and retrospective (Lee and Petts, 2013). What is needed now is a mix of upstream, mid-stream, and, as the technologies move on, increasingly downstream adaptive governance, rather than government.

3 **Use the RRI framework as a resource to support the "(ethical) acceptability, sustainability and societal desirability" (Von Schomberg, 2013) of quantum technologies and innovation around them.** We suggested in the recommendations in section 3.2 that early applications of NQIT will be pointers to social transformations from quantum computing; in this section, we have anticipated the kinds of transformation that can be envisaged. These early applications will build into more complex institutionalisation of quantum as it intertwines with other technologies such as Internet, mobile communications, and the Internet of Things. To be acceptable, these transformations must reflect a "proper embedding of scientific and technological advances" (Von Schomberg, 2011, Von Schomberg, 2013) in society; but this begs the question of how this proper embedding is to be identified and put into practice. To meet this need, RRI is a resource for creative thinking and adaptation, rather than a set of additional activities added-on to the project.

Quantum computing may transform financial services

3.4 Perceptions, motivations, and practices of RRI in the processes of NQIT

Our interviews with NQIT researchers and others involved in the quantum programme express varying attitudes to RRI and ways of understanding what it might mean to them. This section presents and unpacks these concerns and explores how they may be addressed through a clearer articulation of the goals and processes of RRI.

To begin to articulate what RRI is, and what can be achieved in the context of NQIT, we first consider the nature of the NQIT research. NQIT makes an important contribution to understanding RRI as it **takes fundamental science and develops technology** and applications into a rich mixture of basic and applied research and engineering.

We then discuss perceptions of RRI within the NQIT Hub, which is based on analysis of **interviews with scientists at all levels in the NQIT project** to investigate their attitudes to RRI, as a perspective from which to understand the opportunities and challenges of implementing a programme of RRI in NQIT.

While in general supportive of the idea that researchers may consider the implications of their work, respondents **expressed difficulty** in seeing how individual researchers can contribute, as well as concerns about its implications for scientific progress. Others saw their work as being too far removed from the eventual uses of quantum computing to be able to make a difference to long-term outcomes. It is essential to understand and address researchers' reservations, because implementing a programme of RRI within NQIT needs to be consensual and tailored to researchers' needs.

In addition to these concerns about the concept of RRI in itself, there are challenges in actually **implementing RRI in practice**. RRI itself is not yet fully worked out, and there are no standardised practical templates that can simply be applied without regard to the particular discipline, context of innovation and area of research. Whilst several common, high level, guiding principles do exist, it can be difficult to see how these can translate into achievable, practical activities of the ground.

What we want to emphasise here is that, despite the inherent difficulty of addressing the issues of RRI, it is not impracticable. RRI is developing **a set of practical proposals**, but in addition, many of the actions which we suggest for RRI are currently being done as part of scientists' and researchers' work, but not explicitly. For RRI to work there needs to be an **overall framework** within which any individual activity can be seen and felt to be meaningful, so that any researcher's involvement in RRI adds to a collective effort that has an impact beyond that researcher's immediate sphere of influence. Finally, in the context of such a framework, there is an important requirement to assess what skills and understandings researchers need to acquire to be able to adopt its practices.

The overall aim of RRI in NQIT is to develop such a framework, and this Landscape report is a step on this road. In the following section, we set out some conclusions and recommendations in this direction.

3.4.1 End-to-end research to realise Quantum 2.0 technologies

NQIT is breaking new ground, not only in basic research and engineering of the Q20:20 machine, but also in the way it is bringing many developments together, towards the overall end of an important step towards quantum computing. Even before Q20:20, there will be spin-offs and nearer-term outputs - including quantum simulation, properties of diamond for qubits and sensors, advances in photonics and atom photon interfaces, and quantum mechanics in superconducting materials - along the way.

This end-to-end development brings a new perspective on distinctions between basic and applied research, with implications for how we think about and how we implement RRI. It is not simply a blurring of the distinction between fundamental and applied science, but that the "stack" of involvement is all directed, from the most fundamental to the most applied, towards a clearly defined end in the form of a demonstration of a scalable (if still relatively small) quantum computer. This direction gives a shape to the research, including research in RRI. This is RRI not as just a concept but as an activity, involving researchers across different disciplines, each bringing different perspectives.

End-to-end development brings a new perspective on distinctions between basic and applied research.

If RRI activities are not to be seen simply as a disingenuous way to smooth the introduction of technologies whose pathway is already set out, retro-fitted as a bolt-on, they need to be embedded from the outset (Owen et al., 2013). But in the early stages the specific areas of application are only dimly foreseen. RRI, therefore, is a continuing progress, developing alongside the applications which will become clearer, over the lifetime of NQIT and beyond in the development of quantum computing.

Mixing theory, experiment, fundamental and applied research

In our interviews respondents often spoke about the interplay between theory and experiment, and between fundamental and applied research, and both of these interactions seem to be important to NQIT approach:

> *"it's not just you cook up theories, but then you are very close to the experiments, so things move quite quickly, and then you push the theory you have to its limits, because you want to design something, then sometimes things don't add up,"* – (NQIT senior researcher)

> *"We're at this kind of tipping point, I guess, because it's the first time somebody's trying to build something, I think a lot of it is very fundamental, because you come up against problems that no-one foresaw before, but it's applied in the sense that we're trying to build something that has a technological application, so I think that the end goal is very applied, but the route to getting there is very kind of foundational"* – (Early career researcher)

Whilst neither of these articulations of the relationships between theory and experiment on the one hand and fundamental and applied research, on the other, are new, it seems that their juxtaposition and interactions are more compressed within a programme such as NQIT as it attempts to translate science into technology over an ambitious time frame.

To put this another way, the pressure of the outcome – a working quantum computer – drives the interaction between theory and experiment, and pure and applied research, much more strongly that might otherwise be the case. What is interesting too about this from an RRI perspective, is that those engaged with more theoretical or a foundational research have greater exposure to the practical matters of experiment and application, and all are having to dedicate some of their focus to the economical imperative that drives the programme – for example, by considering what choice of material may imply for the form factor or cost of an eventual quantum computer.

The pressure of the outcome drives the interaction between theory and experiment, and pure and applied research.

> *"Basic research is a slightly funny word, almost, these days, because quantum mechanics, the field that we work in, is largely figured out. So there aren't that many people who are really doing fundamental research in the field of quantum mechanics. There's a lot of discovery to do in terms of the way that materials behave and things like that, but well, I guess different*

people might call that fundamental science, or not, but I don't think it's any surprise or even necessarily a bad thing at all, that we are expected, as scientists, to justify why we're working on particular things and often, at times, that's a question of seeing a potential application or something. Maybe that's quite a bit down the road in some cases, but you should have your eye on what might this be useful for". – (NQIT senior leader)

"at the moment, I'm still what you would call basic research. I mean the point of NQIT is to kind of take, by and large quantum information is still mostly really a basic research field, it's, sort of, you know, feelers out into applications, and, I mean, there are some companies that have commercialised it, but, I mean, I think, yes, so, in some sense, at the moment, it's, basic, and the idea, I guess, of NQIT is to try and really make, a, by throwing a lot of people and a lot of money at it, make substantive step towards being applied, as well." – (Early career researcher)

This future-focussed and future-driven character to the research within NQIT provides a hook on which the forward looking aspects of RRI may be hung, in particular, those aspects that aim to anticipate what the social consequences of a new innovation may be. By framing RRI in the context of a future orientation that already exists and is already accepted and understood, we can start to demystify the elements of RRI that also call for these types of anticipations.

Motivations

Another theme emerging from the interviews was how researchers spoke what motivated them as scientists:

"I was always excited, and thrilled, … I want to understand the world, basically, so that's, my main motivation is to understand nature, and, it's very exciting when I see that, understanding nature can also lead to something, with the perspective of having an application" – (Early career researcher)

"Generally, scientists do research because it's fascinating to them, and because they think in some way it can make the world a better place, and they're right." – (Senior NQIT leader)

"I got into this, from day one, because I wanted to contribute towards knowledge… I want to advance knowledge, I want to do things which do good things, in a way, and, I hope that what we do here, we have an impact on the world, and I hope it's a positive impact." – (Early career researcher)

Respondents speak of a two-part motivation that is focussed on creating knowledge, driven by excitement and curiosity, but also of a desire to see that knowledge put to work.

Each of these respondents speak of a two-part motivation that is focussed on creating knowledge, driven by excitement and curiosity, but also of a desire to see that knowledge put to work, and especially ways that have a "positive impact" or which "make the world a better place". So RRI can find a place within these scientists' own sense of their scientific identity, which includes a desire for the positive outcomes of their research – and not only the pursuit of knowledge for its own sake. From this position, RRI can be considered an approach that can increase scientists' overall confidence that the result of research will be, to the larger part, beneficial.

3.4.2 Perceptions of RRI

In this key section we consider the different ways that researchers have responded to the idea of undertaking RRI within NQIT. This has enabled us to identify and counter some misperceptions, as well as to find ways of articulating RRI concepts in ways that build on existing practices and expectations.

Participants were asked whether researchers should think about responsibility issues - the interpretation of this was left to the interviewees. With a few exceptions, interviewees were in general supportive of the need to consider societal implications of research. This is of course a good basis for RRI, but, at the same time, some participants expressed reservations, particularly around the challenge of explaining the implications of quantum technologies. In this section we consider the responses of researchers, work package leaders, and senior investigators and directors of NQIT, to uncover sticking points and uncertainties about what RRI might involve. RRI recognises that predicting how technologies will work out is difficult and that there will always be unforeseeable contingencies. Nevertheless, we uncovered practices which we consider to be forms of RRI, whether or not they are identified with that label, which, taken in context with the involvement of stakeholders across the project and at all levels as we discuss in the following subsection, suggest ways forward towards clarifying how RRI may be put into practice in the context of NQIT.

Predicting the future is hard

Sometimes RRI is equated to a form of prediction where poor access to knowledge of what may happen in the future are seen to hamstring any precautionary activity in the present:

> *"It's not a bad thing for scientists to think about. So many scientists are working on things whose, potential consequences are so ill-defined that it's not clear that they can, get much out of it, that they can get much traction."* – (Early career researcher)

> *"It's actually harder to research the consequences of research than it is to do the actual research, and probably in many cases impossible."* – (Senior NQIT leader)

These quotes highlight some of the challenges which RRI is designed to address; it may be "impossible" to foresee the consequences of research, but that does not mean that there is nothing that can be done to forestall them. The outline AREA framework (Figure 1) includes the dimension of anticipation; this is not prediction, but developing the capacity to be prepared for future eventualities.

On respondent suggested a pragmatic approach of iteratively responding as issues become apparent:

> *"All we can say is that this is likely to have the same kind of impacts as the IT industry had, although a bit more narrowly focussed 'cos it's not as general as overall IT, and that there will be positive and negatives, and that when those start to become apparent, then we should be intelligent about how things get, about the rules that come into play, I just, I don't know how to be more productive than that on that question."* – (Senior NQIT leader)

Clearly this is very sensible, and reflects much of the approach suggested within this landscape report. But by *anticipating* outcomes, even if we do not have any pretensions that we can make predictions of the future, we can start to *take care* and uncover hidden assumptions about what the effects will be. While it seems reasonable that impacts can be extrapolated from IT generally, bearing in mind the potentially powerful but limited ("deep and narrow") applications of quantum computing, on the other hand the scope of transformation can be unpredictable; it will almost certainly be more in some directions, and less in others, than we can currently expect. This mitigates towards keeping an open mind, and taking care in how we "sell" quantum technology.

Our approach can explore, in the first instance, how quantum computing may play into existing trends and controversies, as a way of exploring proactively what those positives and negatives

By anticipating outcomes, even if we do not have any pretensions that we can make predictions of the future, we can start to take care and uncover hidden assumptions about what the effects will be.

Informal discussions at conferences are a part of “de facto” RRI activities - UK Quantum Conference 2015

may be. Certainly creating the connection between these concerns and regulation is important, and perhaps this is the most crucial area to have a proactive approach since it is common for any existing regulation to “run behind” technology. The rhythm of the regulatory process needs to be taken into account. Creating new regulation and governance structures takes time, and requires a preliminary effort of raising awareness and understanding the issues.

RRI in action, or “de facto” RRI

Even in the absence of RRI activities recognised as such, researchers and innovators frequently do undertake responsibility activities. Researchers are naturally inquisitive people, and are interested in the downstream implications of their work, sometimes as an extension of their fascination with their research and sometimes as a major motivating factor for them personally.

“De facto” RRI includes simple activities such as informal discussions with colleagues over coffee and reading of the popular scientific and general press.

We include in this *“de facto”* RRI simple activities such as informal discussions with colleagues over coffee and reading of the popular scientific and general press (Randles et al., 2016). We have found that researchers are more likely to talk informally about their work than about possible down-stream implications, but, on a personal level, do sometimes discuss this amongst themselves:

> *“In terms of the potential applications, the potential impact on the real world, I suppose, I do think about this, and I wonder whether or not there might be any sort of side effects that I can’t really foresee. Lacking any specific suspicions for applications such as that, I suppose, my baseline has been naïve technologism, the idea that it will allow people to do things that would be interesting and not terrible, that weren’t possible to do before”* – (Early career researcher)

"And certainly it's good, right, every so often, it's good to try and, try and imagine really far forward, and see if you can come up with some, and it, just. There's no harm in speculating, as long as you know you're speculating. So, yeah, I think that, and I think people do, quite a bit" – (Early career researcher)

3.4.3 Sharing responsibility for RRI

There is an on-going debate as to the ways in which the responsibility for Responsible Research and Innovation is shared between scientists and researchers and others groups within society; we want to avoid the sense that somehow RRI is something that should be shouldered only by scientists (Lee and Petts, 2013). On the one hand, decisions about large areas of concern – such as privacy, security, surveillance - are social and political, and need to be addressed at that level; the outcomes, good and bad, from NQIT will be determined by structural factors such as governance, regulation, and adoption by industry, as much as by individual actors. On the other hand, scientists do have a special responsibility, as they are often best-placed to understand the upcoming implications, but without subscribing to the myth of the "all-powerful" scientist: *"Modern "Frankensteins" are not intentionally created by a single actor, but (if they arise) are more likely to result from the unforeseen side effects of collective action"* (Von Schomberg, 2013).

Limits to responsibility

Some of the respondents suggested that researchers, junior researchers in particular, do not have time to get engaged in RRI, and there seemed to be an implied concern that RRI may compete or interfere with the core scientific mission:

"Yeah, and I think, for researchers, there's only so much we can reasonably expect them to do in terms of speculating about what they public want and don't want, but I think, and we don't want them spending 99% of their time worrying about that rather than doing their research, so there is an issue of balance, but I think that there are other dimensions to making sure something's a public good which are more tractable to them." – (Research council staff member)

"I think that young researchers, in particular, whilst it's always good to think about consequences of your research, there's never enough time to just do your research, and so I think it's probably not a good use of a starting postdoc's time to require them to put a significant percentage of their time into thinking about the consequences of their research, cos it's just not a very well-defined question, it's very important, but that doesn't mean that you can actually help that area by taking people's time and putting it into it" – (Senior NQIT leader)

"I think we're busy enough already" – (Early career researcher)

For researchers in science, there is an implication that responsibility in research only goes so far. Several participants suggested that, while they would naturally prefer their research to have positive implications, this is not under their control: downstream, the path to innovation is not direct and responsibility can usually not be, or should not be, traced to an individual:

"it depends, also, on, on, the bigger picture, is, has quantum computing revolutionised the, the world we live in, as many people say, and there's a bad example, where, where it's been used to, er, to do bad, yes, I mean, it is regrettable, and, really not good, but, I would probably then still not feel that everything I've done was, was, bad, or I shouldn't have done it" – (Early career researcher)

A senior leader in NQIT made the argument that RRI should be handled by the senior scientists as part of a political process:

> *"So, I think, what tends to happen is that as scientists reach a more senior stage in their career they're more and more likely to be involved in committees and panels and things that look at this, and I think that's a pretty reasonable way to do it, and those committees and panels feed into the decision-making processes, along with political views and so on, and that seems to be healthy."* – (Senior NQIT leader)

RRI does not present a radically different vision to what is already in place, but rather aims to make these activities more connected, systematic and purposeful.

In one sense, this viewpoint constrains RRI to be mainly an activity for particular, senior scientists acting in particular, special situations on committees and panels. However, it also points to a way forward for RRI. RRI can build on these existing expectations and roles, especially processes that are already in place to make the important connections between science and politics to guide policy formation and regulation.

This is an example of an existing activity which is already consonant with RRI; RRI does not present a radically different vision to what is already in place, but rather aims to make these activities more connected, systematic and purposeful.

RRI for everyone

However, it is important that all scientists at all levels are encouraged to contribute to RRI in various ways, which could be as simple as considering the outcomes from their work from time to time. RRI needs a diversity of perspectives – and for another reason that we outline below.

The practice of scientists at a high level of seniority and acting in that capacity advising policy-makers or setting a strategy for research is well established. In general, it works, but, in response to a number of failures in these processes (Grove-White et al., 2000, Von Schomberg, 2013), RRI is trying to bring science and society much closer together, to break down the *"moral division of labour"* (Stilgoe et al., 2013, Fisher and Rip, 2013). In this way, responsibility is expressed not only in formal policy structures but embedded both in the daily practices of scientists and in choices made by those – funders, policy-makers, horizon-scanners – who set the trajectory "upstream"; and by Innovators – people working in knowledge transfer, research impact, small specialist companies, large corporations, governments, early adopters – who influence the paths which new technologies will follow "downstream".

For junior researchers, familiarity with these topics will be part of their apprenticeship, and will enable them to more effectively to contribute to committees and panels as their career progresses.

Developing the capacity for RRI

Responsible Research and Innovation requires the capacity to continually reflect, and readiness to anticipate on the part of researchers, innovators, funders and policy-makers. These capabilities and habits of thought require preparation; engagement and responsiveness are structured processes. RRI demands openness to reconsideration of underlying rationales, if it is not to simply be reactive and therefore closed, self-referential, and narrowly focussed (Pellizzoni, 2004).

While we do not deny that high-level policy direction requires the broad expertise of senior scientists, we suggest that, for junior researchers, familiarity with these topics will be part of their apprenticeship, and will enable them to more effectively to contribute to committees and panels as their career progresses. We have found that many junior researchers (although not all) have been interested and enthusiastic to contribute to RRI workshops and made valuable contributions.

This style of contribution has many advantages and few disadvantages. Junior researchers in this way can develop the capacity for sensitivities and skills that serve them well in as their career develops; and they can make a contribution that is carried forward by the RRI process to wider influence, which makes their contribution a meaningful one, without being an excessive use of their time. In Section 4.3, one of our recommendations gives more details about how we suggest that the capacity for RRI can be developed.

3.4.4 RRI in science and society

Finally, we can draw on researchers' perceptions of RRI, the activities which they may already be doing, and the shared responsibilities of different stakeholders, at different levels, and at different stages in the emergence of the technologies, to express, in more concrete terms, what RRI may mean in NQIT.

Researchers have recognised that responsibilities for the impacts of innovation are shared through different functions within society. Perhaps it is not the science that is the problem, so much as how the science is utilised when it enters society:

> *"I'm, I've told you, like, everything around you can be used in a bad way and in a good way, so, as you can, you know, save lives, by having an army, you can also, like, kill people, it's the same thing"* – (Early career researcher)

> *"Whether it has detrimental properties or not will be nothing to do with scientists, I mean, we made the gun and gave them the gun and they chose to shoot each other"* – (Senior NQIT leader)

RRI would see this separation between "science" on the one hand, and "society" on the other, as being less clear cut as these quoted views may suggest; researchers are after all part of their society. As one respondent commented:

> *"I think as a human being then you have a responsibility to do that and, yeah, scientists are human beings and I think we have that responsibility"* – (Senior NQIT researcher)

Whilst essentially it is true that technology can be pressed into the service of a variety of agendas, sometimes unpalatable, this does not also imply that there are never any opportunities to forestall some of these "mis-uses". For example, the "honour system" underpinning the SMTP protocol allows the issue of spam to proliferate[34]. Clearly, there are alternative models for the protocol that would be much more resistant to spam, and that would have avoided negative uses of e-mail; but the early specification of the protocol, designed for sending mail internal to an organisation or around academic institutions, lacked any form of anti-spam or even the most basic security.

34 Fixing the SPAM problem once and for all: Ou, 2003 http://www.lanarchitect.net/Articles/SPAM/FixingSPAM/. This, and other suggestions for major changes to the Internet mail relay protocol, have not been widely implemented. Instead, mail servers rely on various ad-hoc fixes such as SMTP server authentication and spam filtering.

RRI as part of "scientific citizenship"[35]

Yet on the other hand, some lines of research are seen as entangling scientists in wider responsibilities:

> *"So I think, obviously, if you're a scientist being asked to create an Ebola variant which is far more violent and deadly, then you should think about the consequences of your research, which are almost certainly that whoever's paying you for that wants to drop Ebola on someone, but that's an extreme case."* – (Senior NQIT leader)

> *"I think you have a responsibility to not blind yourself to those questions, and there are examples historically where scientists have been working on stuff which, in hindsight, you just think, come on, guys, what did you think was going to happen?"* – (Senior NQIT leader)

RRI exists in a more moderate space where there are roles for scientists, and a variety of other actors, to contribute part of the policy making undertaken by politicians, hold a dialogue that moves out into public spaces; and to engage in activities that promote reflection and anticipation.

RRI, as we conceive it, does not focus on these extreme and polarising situations. It is not about the implications of researching pathogens (a research activity that is likely to be heavily regulated anyway), nor is it about solving every dilemma that technology creates for society. It exists in a more moderate space where there are roles for scientists, and a variety of other actors, to contribute part of the policy making undertaken by politicians, hold a dialogue that moves out into public spaces; and to engage in activities that promote reflection and anticipation. To reiterate, this is not meant to imply that scientists have some ultimate responsibility for how technologies are used, or what their effects, positive or negative, may be, but to suggest that there are practical ways that scientists can be involved in and contribute to the wider societal processes that shape the relationships between technology and society.

Scientists spoke of a desire to reflectively consider the wider ramifications of their research from a professional perspective, involvement of public money, and broader notion of "scientific citizenship":

> *"… something like NQIT really prompts [speculating about implications], because, basically, people have given you a lot of money and you're really, you know, it's not an academic exercise, you're really trying to build and do some things so, but that's actually, that's quite a good, there's nothing like, sort of, the stimulus of, you know, presumably the same thing happened when people were, you know, thinking about going into space, and someone said, we're actually going to do it, yeah, so, so, then you do have much more of a responsibility to try and really think out the real-world implications of academic, kind of, [issues]"* – (Early career researcher)

Demystifying quantum technologies

There is a perception, which we explored in Section 3.1.3, that quantum physics is intrinsically hard to understand, "spooky", and "mysterious". We have suggested that these rather deterring discourses could be replaced by expressions such as "counter-intuitive". Nevertheless, it is hard to avoid the conclusion that quantum mechanics is hard to understand, and that most members of the public will not have any background in understanding it even at an elementary level. There are some popular books that make a reasonable pass at explaining the basics of quantum mechanics to non-specialists (for example, *"Alice in Quantumland"* (Gilmore, 1995)) but even these are not something which could be grasped in, say, a half-day workshop, and they still require the concentrated attention of motivated readers.

35 For an interesting discussion of the term "scientific citizenship", see http://blogs.nottingham.ac.uk/makingsciencepublic/2014/12/14/scientific-citizenship/

The conceptual difficulties of quantum mechanics pose a challenge for RRI, once we start to engage beyond the scientific and technological communities. According to this narrative, engaging with the public will not produce useful results, because the science is so hard that it is not realistic to expect non-expert members of the public to reflect in any meaningful way on the research and innovation process.

> *"It would be a mistake to imagine that then, the person, even a very alert and engaged individual, is in a position to contribute back some thoughts about the detailed direction a researcher can go in, because you are just ultra-simplifying things, to a level at which you're not informing the person so that they can make decisions"* – (Senior NQIT leader)

This reflects the "deficit model" (Wynne, 1993) – an assumption that the public are ignorant of science and hence excessively risk-averse and unreflective. Whether a non-expert can make useful decisions about science depends, of course, on the kinds of decisions they are asked to make: probably their suggestions for experimental design or refinement of a theory would not produce any great advance in knowledge, but this does not imply that their insights are not useful in other ways. The challenges of understanding the science may seem to be particularly great in the case of quantum technologies, but other areas of science and technology also have aspects which are hard to understand. It is not necessary to understand the underlying science in depth in order to grasp the implications. Public dialogues have produced useful insights from the public around other complex technologies such as nanotechnology (for example, (Royal Society, 2004)); and brain science (Goldschmidt and Renn, 2006). A UK-government sponsored organisation, Sciencewise[36], whose survey of media coverage of quantum has already been mentioned (2014), was highly experienced in conducting these dialogues, but is currently without funding; a dialogue commissioned by EPSRC is in the process of seeking tenders from suitable organisations or companies. Expert participants have specialised knowledge which members of the public do not have, but the purpose of engagement in RRI is to open up to a broad range of perspectives, identify areas of potential contestation, in dialogue with multiple publics and diverse stakeholders (Owen et al., 2013).

Moreover, lay people, although not expert in a particular area of science, bring their own areas of expertise, and different viewpoints, not least that they are acutely aware that the path from lab to technology to innovation is rarely linear (Stirling, 2008). Experience in the UK with social scientists and lay members of the public on expert advisory groups has repeatedly shown that people of varied backgrounds are able to engage in meaningful and intelligent dialogue with experts and to produce useful results (Goldschmidt and Renn, 2006, Macnaghten and Owen, 2011). But, just as scientists do not shoulder all of the responsibility, neither does the "general public", or policy makers.

Lay people, although not expert in a particular area of science, bring their own areas of expertise, and different viewpoints.

Demystifying RRI

But just as there is an aura of mystery around quantum technologies, which represents a barrier between scientists on the inside and non-experts on the outside, so also the concepts and practices of RRI can seem "impossible" and can be thought to imply an excessive burden on scientists.

We have identified on the part of researchers and scientists a willingness to act, a recognition of responsibility as part of good research, which is given a particular imperative when a project is charged with a major – perhaps even dramatic – visible outcome. But what, concretely, can researchers, and other stakeholders, do to enact RRI? The issue is not so much that RRI is perceived as being unimportant, but that it is hard to see clearly how it might be implemented,

36 http://www.sciencewise-erc.org.uk/

and hard to steer clear of the apparent implications that it "asks the impossible" (as expressed in the quotes near the start of this section) – such as predicting the future, or controlling societies' use of technology.

As one respondent said:

> *"So you can't sort of pretend that you're planning ahead on the knowledge of how everything will develop decades ahead. … So, you guys are experts in RRI, so, what should I be doing, what should we be doing, in your opinion?"* – (Senior NQIT researcher)

On one level, this is not an easy question to answer. But this statement is useful, because, while no doubt honestly asking and showing a willingness and interest to follow the ideas of RRI, it reveals a couple of assumptions. Firstly, that RRI requires some sort of knowledge of the future; if this was the case, we would certainly never be able to innovate responsibly. On the other hand, if we somehow did have perfect foreknowledge, then we would not need RRI; who would do research or develop new products in the certain knowledge that they would lead to disastrous outcomes? Who would fail to do research which is known for certain to lead to great benefits? But in the real world, we never do know, with certainty, how things will turn out. In the real world, we are rarely dealing with unambiguously disastrous outcomes, or unalloyed benefits: some people will gain, some will lose; we might invent a new technique which is somewhat useful but at the same time "uninvent" (MacKenzie and Spinardi, 1995) older techniques which might at some future date have proven useful once more.

The aim of Responsible Research and Innovation is to provide supportive frameworks for researchers and other stakeholders. The shape of these frameworks will depend on the situation; it might not involve "doing anything differently", but rather it might suggest a variety of ways in which researchers can contribute to responsible innovation: examples could include making their expertise available for the public benefit, public engagement at events and accessible publications, using social media, or intervening in policy debates. In many cases, simple awareness of the potential downstream implications, and readiness to change the trajectory and pace of innovation, may be the most appropriate response in the immediate term.

As we have already suggested in the recommendations in Section 3.3, where appropriate RRI demands a willingness to change trajectory in response to social needs and changing circumstances (Stilgoe et al., 2013). This might involve a laboratory adopting new training materials, or taking steps to remain aware of wider implications of their research, as a form of reflexive *"mid-stream modulation"* (Fisher et al., 2006, Fisher and Rip, 2013).

RRI is not a check-list which can guarantee socially desirable outcomes, but is a developing framework and set of practices which can be integrated into the work of NQIT at all levels.

This is the "Action" dimension of the AREA framework (Figure 1), but it is perhaps better expressed as "responsiveness", which is the word used in the original working out of this framework by Owen, Stilgoe et al. in the Responsible Innovation book (Owen et al., 2013). The verb "respond" conveys the idea that action is not necessarily "doing something", but is a way to cope with and influence uncertainties and areas of concern. These areas of concern, in turn, are identified by processes of anticipation, reflection and engagement in dialogue with the public, policy-makers, early adopters, and scientists in other disciplines.

This leads to the idea that RRI does not have all the answers, but instead can help us to ask the right questions. RRI is not a check-list which can guarantee socially desirable outcomes, but is a developing framework and set of practices which can be integrated into the work of NQIT at all levels. The next actions in RRI in NQIT, described in Chapter 4, will be steps towards such a framework.

3.5 Drawing together the key findings for RRI in NQIT

This section has started to uncover the ways in which the science and technology of NQIT works to produce socio-technical outcomes which have social implications.

Predicting the social consequences of research is difficult. The FRRIICT project (Eden et al., 2013, Stahl et al., 2013) suggested that this is particularly difficult for ICT research. In (classical) ICT, it is not that research does not have hoped-for societal impacts or potential for negative consequences, but rather that the range of applications is so wide. In contrast, in quantum computing, at least up to the current date, the range of identified applications is rather sparse, but the potential applications are very wide; understanding their social implications requires anticipation.

At the same time, in addition to uncertainties which parallel to those of classical ICT, there are uncertainties over the basic theory, the engineering to actually make it work, when and in what form quantum computing can be delivered, and not least, in the ways it will work alongside other ICTs and other technologies as it is adopted by society. These uncertainties about the trajectory towards quantum computing are exacerbated in public perception by a number of myths and "spookiness". Currently, the public is generally supportive but this could change.

Nevertheless, some early applications can be discerned; some, relatively close up, using forms of quantum simulations; others, still far off but well-discussed and for which algorithms have already been defined.

Quantum computers requiring expensive infrastructure could change the balance of power even further in favour of governments and large companies, for example if algorithms such as Shor's algorithm or Big Data analysis could disadvantage ordinary computer users. In a similar way, quantum computing could exacerbate trends towards what has been called the "algorithmic society". By this we do not simply mean that algorithms can circumscribe our actions, but that these algorithms themselves reflect, and perhaps amplify, assumptions of the wider society. Often, the algorithms are opaque; in quantum computing, the need for verification intensifies the need to trust results, particularly in situations where it not possible to trace back through the logic in traditional computing style.

There are specific challenges of quantum technologies for RRI, but also enablers. Embedding RRI throughout the NQIT Hub is also an opportunity to expand our understanding of RRI. Technology develops from results discovered in research, but they ways in which it develops are determined by they ways in which it is adopted by society as much as by the inherent capabilities of the science. In the long span of NQIT, taking fundamental science and developing it into a usable quantum computer, we will be able to see this social development as it plays out in the early stages of a technology. We will be able to reflect on the motivations of the different stakeholders, and understand how they work together: for some, scientific discovery in itself is a motivation; for others, doing some good, and the overriding imperative to develop a working demonstration, as part of what the funders hope will be a long-term investment in UK competitiveness.

All of these issues working together have implications for Responsible Research and Innovation: will the uncertainty of the direction that the technology might take allow a space for very negative public perceptions? How can we deal with the potential applications of quantum technologies, such as the risk that existing forms of information security will become obsolete? Further downstream, will quantum computing lead to major social changes, and will these be positive or negative for society as a whole?

In the following final section of this Landscape, we draw out recommendations, and suggest ways forward in developing the framework which will enable RRI as a practical set of actions in NQIT.

Chapter 4: Recommendations and Future Plans

We conclude this Landscape report with recommendations for how to continue with Responsible Research and Innovation in NQIT, lifting RRI out of the NQIT internal space which has been the focus so far and developing practical suggestions for process which can actualise RRI for the future of the Hub.

We begin with a summary of the challenges and enablers for RRI, as we see them, drawn from the findings presented in section 3. Then we set out some concrete recommendations for process, starting with the continuing development of a practicable framework for RRI in NQIT. Finally, we show our plans for developing this framework, building on a series of case studies and other research.

4.1 Challenges and enablers for embedding RRI in NQIT

The key findings presented in Chapter 3 can be summarised as a set of challenges and enablers for RRI in NQIT, following from:

The "live" issues around:

- Uncertainty - can it be done theoretically, how can it be engineered, what will quantum computing be able to do?
- The emerging and still-uncertain character of quantum computing
- The counter-intuitive properties of quantum and how we can talk about these

    - **We bring these together under the umbrella of "Narratives"**

Issues "on the horizon" – discernable, and the focus of active research, but not yet immediate:

- Strong claims for the potential of quantum technologies
- How to verify quantum results and, how can we be sure?
- Sensors and sensor nets; quantum simulation and modelling

➡ **We bring these together under the umbrella “Trust” in quantum technologies**

Future applications of quantum computing in use:

- An “algorithmic society”, particularly concerning Machine Learning and AI
- Applications for defence and national security
- Ownership and control of quantum computers

➡ **We bring these together under the umbrella “Social Transformations”**

4.1.1 Challenges

Challenges from the uncertainty surrounding quantum computing

1. There is still a high degree of uncertainty about what form quantum computing will take and when it will produce useful results.
2. Whereas classical ICT artefacts are perceived as “logically malleable” (Eden et al., 2013), in reality there are limits to what they can do. Quantum computing will widen the scope of what ICT will be able to do, particularly in the form of hybrid computing, but it is still not clear what the scope and limitations will be.
3. Other narratives that quantum mechanics is “spooky” and hard to understand may increase public concern. There is a risk that the uncertainty combined with the “spookiness” will give rise to unhelpful narratives around the threat of quantum computing that may be ill informed or dystopian.

Challenges from research issues on the horizon

1. Areas of research interest – on the horizon, already being developed but not yet accomplished – each emphasise the potency of aspects of quantum technologies, and this will increase the trust required from users, not only directly but also as second-order trust in infrastructure, materials, and processes which have been developed with the aid of quantum technologies.
2. We have identified some scepticism amongst some researchers as to the benefits of RRI and their ability to contribute.

Challenges which can be anticipated from the social transformations around quantum

1. Important and real social issues, such as a drift towards and algorithmic society and changes in power towards already-powerful actors, need to be considered. There is a risk that these may be overshadowed by more high-profile issues (such as “breaking the internet”) and other unhelpful narratives.
2. There are challenges in engaging wider range of stakeholders in science which is perceived as hard to understand.
3. Some sensitive and sometimes classified applications of quantum computing present difficulties in gaining access to key informants, for example in companies and government organisations that operate in the defence industries.

4.1.2 Enablers

Enablers in the emerging narratives surrounding quantum computing

1. The narratives around quantum technologies are still emerging. Alongside a strong narrative of "spooky, mysterious" science, there are more positive portrayals around how the science for is being appropriated for societal benefits. We can draw on these narratives to assess perception among a wide group of stakeholders.
2. Quantum computing can draw on a tradition of socio-technical methods in ICTs and, more recently, on progress in bringing Responsible Research and Innovation into ICTs (Eden et al., 2013).
3. The quantum computing community is already aware and interested in the social implications of their work; although some question whether they can contribute usefully to RRI, many researchers are motivated to see their work have positive benefits, whether through applications that can be already be discerned or further downstream.
4. Professional organisations such as the Institute of Physics encourage their members to take a responsible attitude. Many of the researchers in NQIT are active members of professional organisations. Other associations such as the Royal Society also provide fora to raise social issues around science and innovation.
5. NQIT is highly multi-disciplinary and as such can draw on traditions of ethics and responsibility from different areas of research. RRI in NQIT is strongly supported by the senior management of the project and, throughout the UK National Quantum Programme, is supported by the funders and included in the National Strategy (UK National Quantum Technologies Programme Strategic Advisory Board, 2015).

Enablers to embed RRI in emerging social transformations

1. Although we do not know with any certainty the applications for which quantum computing will eventually be adopted, we can discern types of application, and to envisage the societal transformations which may follow from them. So we can already anticipate and start to be prepared for these societal changes, and, where necessary, reorient the trajectory of development.
2. The User Engagement programme within NQIT is looking ahead and working with potential future users as well as existing industry to identify application areas. This is an opportunity to introduce responsible innovation into this process, and at the same time enter into two-way dialogue as we explore the implications of these technologies.

4.2 A framework for RRI in NQIT

4.2.1 Reframing RRI

The interdisciplinary, end-to-end research in NQIT, which we have discussed in section 3.4, with theorists interacting with experimentalists and with the overall aim of developing the Q20:20 machine and spin-offs, combine with existing *"de facto"* practices and the general openness to responsibility as good scientific citizens to lay a good foundation for RRI.

The need, now, is to address researchers' uncertainties about how to put this into practice with a framework for RRI, tailored to the NQIT project. The shape and form of such a framework will be refined *through the process of working with RRI.* Here, we set out the basic points and principles of this framework.

Building a framework for RRI will not be a one-way process, and so one answer to what scientists and researchers might do differently is that, rather than changing anything in their own research, they could make their expertise available as a resource for developing the framework. We, in turn, will work across the Hub in dialogue with stakeholders.

The perspectives of NQIT researchers point some reservations, which need to be clarified as we proceed with developing the framework. The following points may help:

1. RRI is not a wholly new set of practices. It builds on existing activities such as senior scientists' involvement in policy formation, formal and informal discussions about long-term implications from even the most fundamental research, a common interest in what science might lead to, outreach to the public, and decades of research into the history of science and innovation.
   - The aim of RRI is to build on existing practices to make them more connected, structured and purposeful.
   - Often, it will not be a question of "doing things differently", but of maintaining an interested and aware attitude.
2. RRI is a shared responsibility. Scientists have a responsibility which comes from being in the forefront of research leaders and from their special expertise, but governments, funders, industrial partners, and early adopters of technology have at least as large a role to play.
   - Instead of an individual responsibility, RRI should be thought of as a package of interconnected responsibilities that are appropriately distributed across various groups of stakeholders.
   - This can be put into practice through workshops and other venues which engage a wider group of stakeholders, and by working with these other stakeholders as we have been doing with NQIT participants.
3. RRI should be proportionate to the capacity, interest, and experience of the researcher. This involvement need not be lengthy nor a regular time commitment for any given researcher, although a programme of RRI activities as a whole should be as comprehensive and systematic as it can be.
   - Many researchers will engage strongly with RRI, while others are less motivated. While responsible research should be part of good scientific practice, we can build most effectively on those who are most interested, particularly as researchers grow in experience and seniority.
4. A commitment to RRI raises complex issues about the responsibilities of scientists. But a concern for responsibility is part of rigorous science, not inimical to it.
   - It is rarely the case that fundamental research will not have social or ethical implications; it is good for all stakeholders to maintain awareness of these potentials, even if they appear to be very downstream.
5. The path from science to innovation and impact on society is rarely simple and linear. Science moves mostly in incremental improvements, and working alongside technological and social innovations. The overall impact might be dramatic or small, but will depend on many factors, not only on one scientific discovery.

    - It would be a mistake to attribute research outcomes, good or bad, only to scientists, and especially not to any one individual scientist.

6. RRI should not be framed as an "impossible" task such as predicting the future, with many uncertainties between lab and product. Nor is it about removing all risk from innovation; failing to innovate introduces the risk that we will lose the benefits of innovation.
    - Instead RRI should always be considered as a very practical activity that is concretely specified, properly resourced, achievable and appropriate to the role of the person undertaking it.

4.2.2 Basic principles for the framework for RRI in NQIT

The framework will benefit from being based on clear and sound foundations. We suggest the following principles, which cut across the dimensions of the Anticipation, Reflection, Engagement, and Action (AREA) which will structure our thinking as we solidify the framework:

- All researchers should have the opportunity to engage in RRI related activities (although they may choose not to take advantage of these opportunities) and the support to do so. Supporting the broadest possible engagement will embed RRI more deeply throughout the Hub.
- There should be RRI activities occurring on different levels to enable a broad range of participation. This will have two benefits: it will bring in a variety of viewpoints – experienced and senior scientists as well as fresh insights from early career researchers; and it will develop the capacity for doing RRI. As Owen et al. (2013) have written, "responsibility is a learned behaviour: our children are not born with the capacity to be responsible, or to take responsibility. It is a social ascription that is learned … Capacity for responsible innovation must be nurtured, across and within our institutions of science and innovation."
- All RRI activities need to sit within a wider framework so that even small contributions are contributing to a larger whole.
- Elements of this framework should be connected up and down through any hierarchy. Thus a senior scientist giving advice to a government committee should have access to the digests of more local RRI activities. Conversely, the work of scientists should be informed by the overall aims, including public policy aims, to which it is directed.
- RRI should involve partnerships. Some key partnerships will be with social scientists, who can help with understandings of how science and society interact. Others will be with RRI specialists who will play a facilitator role. They will have some technical knowledge and can help initiate, coordinate and join up RRI activities.
- An RRI framework has to span between research activities, communication, commercialisation pathways, policy formation and funding directions. This means that it should not be a burden solely falling on the shoulders' of scientists, and that these other actors in the research and innovation ecosystem have a role to play.
- Dialogue with the public, with early adopters, with civil society, and other stakeholders is crucial to ensure that the research delivers results which will be widely recognised and welcomed. Public communication exercises should aim to encourage dialogue with the public, not only in specifically tailored activities such as a workshop to systematise and amplify informal reflections on research processes and outcomes, but also in outreach events of all kinds. The public is interested to give their opinions at events such as science

fairs, so, for example, questions at the end of public presentations should be encouraged, and noted down as useful pointers to public concerns.

- RRI activities should be connected into project structures. This means that those responsible for coordinating RRI report to the project executive. An RRI coordinator should produce an RRI annual report that the project executive are bound to respond to.

This framework should be supportive, rather than prescriptive, avoiding unnecessary engagements and onerous use of time and resources. The RRI framework in its own turn must follow the precepts of responsible research, continually evolving, reflective and responsive to changes in technology and society and incorporating new insights.

4.2.3 Outline of the initial framework

The framework for RRI in NQIT starts from the AREA framework (Figure 1), but whereas the AREA framework is quite general, adaptable to any area of research and innovation, our framework is attuned to NQIT and much more specific and detailed. Moreover, as a focussed project with a clear management structure, we are able to be quite specific in NQIT about the issue areas and the stakeholders; we can engage with stakeholders through project events such as the project forum, and we can make recommendations to the project leaders.

Although the framework is still a work in progress, we can set out a series of practical actions, to be developed further but already taking place as we practice RRI in the Hub, as a "toolkit" to identify areas of concern and suggest ways in which these might be addressed. In this report we have unpacked the ethical and social implications of the uncertainty and "mysteriousness" which surrounds quantum computing, the narratives and the strong claims which demand a new level of trust from the public, and likely early application areas with their potential for public enthusiasm or disquiet.

These activities include those we have already been doing for the Landscape study:

- interviews and workshops with project members
- review of quantum technology and related literature
- liaison with other Hubs

The framework should also open new spaces such as:

- workshops and case studies on key issues. opening these up to external stakeholders;
- engagement with industrial partners and other potential early adopters,
- engagement with civil society and other key communities
- structured dialogue with the wider public
- risk assessment, technology assessment
- informal and formal structured foresight exercises
- public engagement, or more exactly, public dialogue

And in response to the framework, NQIT should be ready to act to:

- shape the direction of innovation and to
- encourage responsibility as a part of professional research practice.

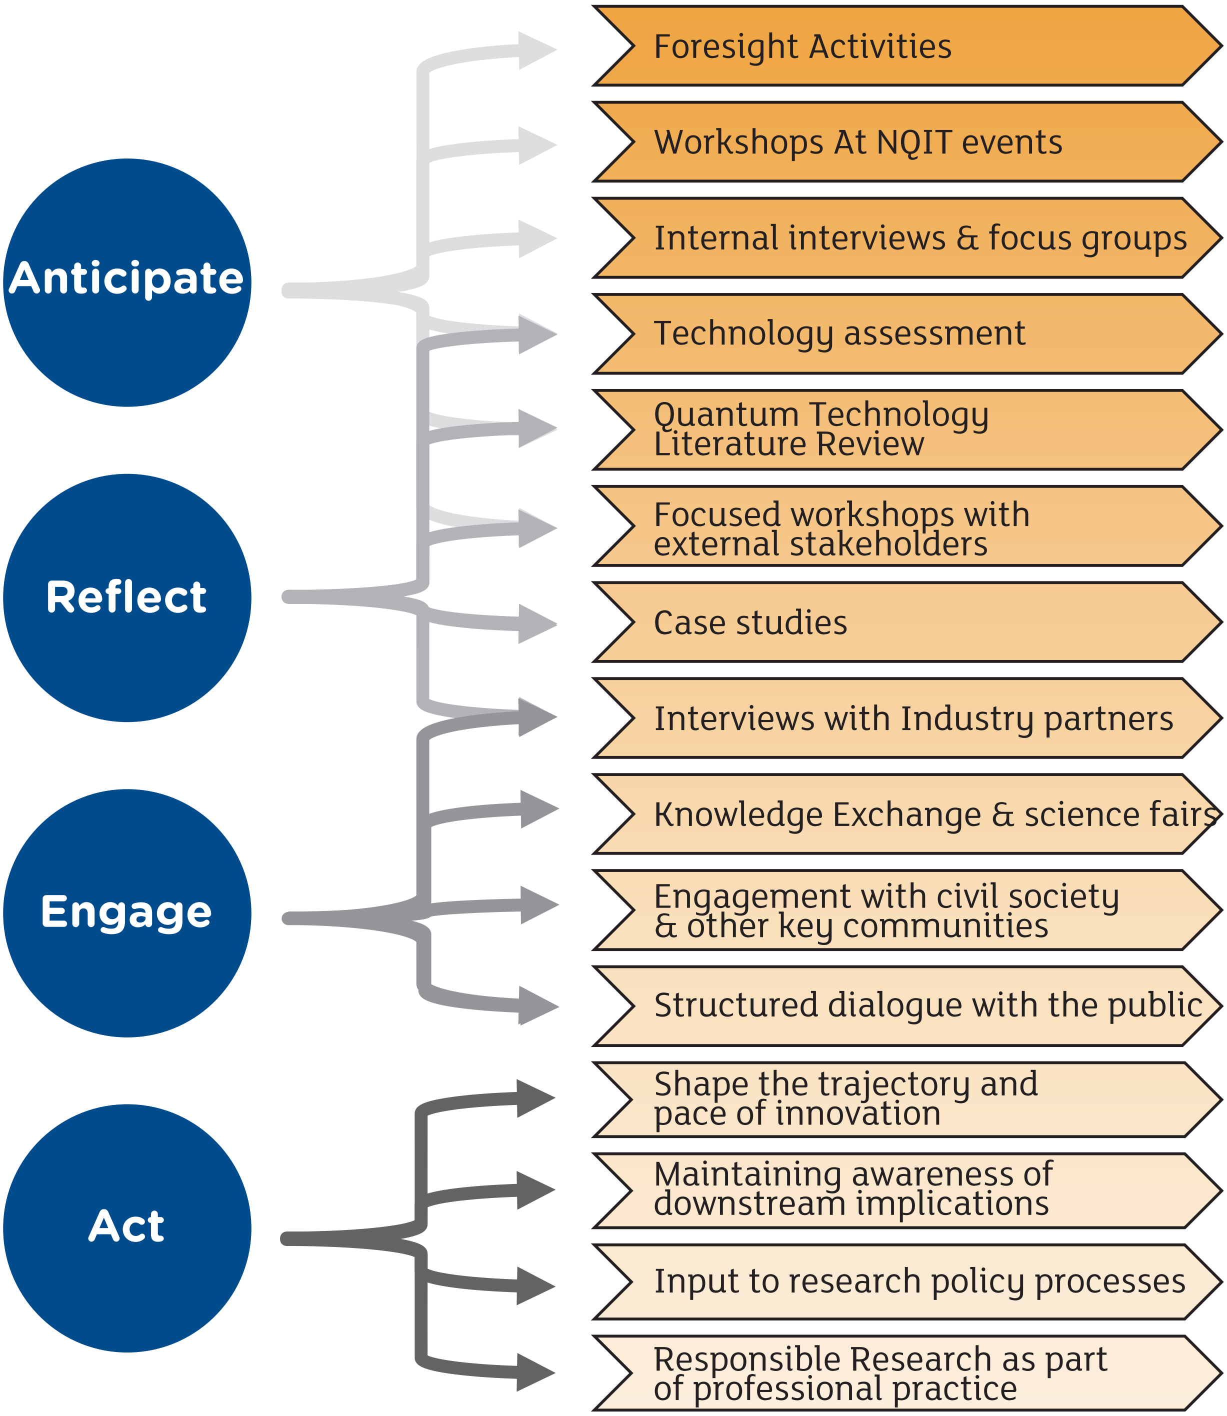

Figure 5: Framework for RRI in NQIT, related to the AREA framework

These activities are shaped by the AREA framework (Figure 1), but tailor it to the specifics of NQIT. Our further recommendations at this stage relate to activities to strengthen the framework by building capacity for RRI in NQIT, include RRI in the structure of NQIT, co-ordinating RRI across the Hubs, and disseminating RRI issues upwards as well as within Hubs to influence research policy.

4.3 Recommendations

We have seen, in the findings of Section 3.4, that scientists have a special, but not unique responsibility for RRI; that these responsibilities are shared with other stakeholders, guided by public policy and by the ways in which technologies are adopted by society; and that researchers, even though they are in principle concerned to research responsibly, struggle to know how best they can put this concern into practice. There is often a great distance between the work of researchers and application in the real world, so that social implications are seen as too downstream to be meaningfully considered; this is a challenge, but also an opportunity, because it is at this stage that changes will have greatest long-term effect, and the longer timespan allows for preparedness and anticipation even if we cannot yet be clear about the details of the impacts.

Here we make some recommendations for the process of RRI, which will support all stakeholders in NQIT to work together in implementing RRI throughout the Hub: developing a framework, building capacity, co-ordinating with other Hubs for maximum effect and to be able to influence upwards to public policy, and engaging to maintain public trust.

4.3.1 Co-Develop a framework of practical actions for RRI in NQIT

From the key issues identified in interviews, workshops, and document analysis in section and from application of RRI principles to the NQIT architecture, we are now in a position to start to suggest activities to put RRI into action in NQIT.

We have clarified and reframed RRI as the basis for practical activities (Section 4.2.1), and drawn these activities together into a framework (Figure 5) which is a major outcome of the RRI research in NQIT. With further refinement as we continue to work with other in NQIT and outside, this framework will embed RRI practices within the future NQIT work, at the laboratory/office, institutional, and public policy levels (Stilgoe et al., 2013).

In fact, the activities which have gone into this landscape report – workshops at project forums and the UK Quantum Conference; RRI Roadshows at each of the hubs; many interactions formally and informally with researchers and research leaders; upcoming activities including an RRI-oriented video and a knowledge exchange event – are already “RRI in action”; it is not to be understood that we will prepare for RRI, produce a framework, and then only later actually perform RRI activities; we are already doing them.

This is a form of “action research” – promoting science and innovation that will have desirable outcomes, in the public interest, and uncovering issues that this raises – while at the same time, these activities serve to refine and solidify the RRI framework, will in turn lead to new activities, and join them together. Thus this framework should be seen both as an already useful toolkit, and as a work in progress which will build a well-founded and well-tried basis to carry RRI forward throughout the lifecycle of NQIT.

4.3.2 Build RRI capacity in NQIT

Alongside a workable and practical framework, training and advocacy within the Hub will build capacity so that the practices developed in the framework are embedded and become part of the regular activities across the hub. This training echoes the need identified in the National Strategy to develop high-calibre expertise in quantum technologies alongside multidisciplinary business skills.

This training would aim to develop broad critical thinking beyond the confines of focus only on the particular scientific results. This would be a cultural change, to encourage people to consider the consequences of their actions, to become part of professional responsibility.

This training could be part of post-doctoral, DPhil and MSc training, and of continuing staff development. For example, as well as the Hubs the QT programme also includes doctoral training, knowledge transfer under the aegis of Innovate UK, and other activities. RRI could be especially appropriate as an element of doctoral training, educating the new generation of Quantum researchers with the habits of anticipation and reflection. Within NQIT, a particular group of researchers and researcher leaders could become "RRI Champions", ensuring that the RRI work is sustained throughout the lifespan of the programme.

Training methods could include workshops, group exercises, and self-directed study. It would draw on the findings of the RRI research in NQIT and on the now considerable body of experience in other research areas in actively working to promote good practices to bring together scientific and societal values.

RRI Champions

The development of an RRI community in NQIT can revolve around the role of "RRI Champions" within a project, across a funding programme, and in the project's dealings with the wider innovation ecosystem. These can be recruited from those researchers, at all levels, who have put themselves forward to volunteer for RRI related activities; it should extend to senior managers of the project, subject to their time constraints, and to communications and technology transfer specialists. The aim is not to have responsibility set aside as another, external activity, but rather that there is a transformation of research and innovation practices towards *"responsibilisation"* (Lindner et al., 2014), which is internalised by participants, enabled by appropriate organisational conditions and governance mechanisms, and carried out in daily practices and in processes of research and innovation.

4.3.3 Include RRI in the structure of NQIT

Our interviews, together with what we have learned of the practices of NQIT and other hubs, have shown that, even if not formally labelled as such, researchers and research leaders do tackle the social issues arising from their research, and, in public engagement, receive feedback even if not formally sought as structured dialogue. RRI emphasises the importance of engaging with all groups of stakeholders, including the public who, after all, will eventually accept or reject the products of research, however far downstream these may be. In this way, much engagement can take the form of *"de facto RRI"* or *"RRI in-the-making"* (Randles et al., 2016).

However, these informal activities, welcome as they are in themselves and as a pointer to future good practice, are not systematic and may not penetrate into the fabric of NQIT; their effects might be dissipated if there are not channels to record and act on them. The capacity for RRI embedded throughout the Hub should be supported at a more managerial level by appropriate, not over-complex but realistic *governance processes.*

The work “governance” emphasises that this is more flexible, and more of a partnership between regulator and regulated, and involves many groups of people working together (Lee and Petts, 2013). The European Commission-funded Governance of Responsible Innovation (GREAT)[37] project had a remit to develop Governance for Responsible Innovation and, in particular, produced guidelines for research and innovation processes. These guidelines follow a key RRI principle – reflexivity, re-assessing the guidelines themselves throughout the life cycle, including by regular review.

This is governance not only for outcome – anticipating and reacting to unintended consequences – but also governance of intent (Owen et al., 2013) – setting, at the level of direction of the project, the long-term outcomes, including societal outcomes, that the Hub is aiming to achieve. The GREAT project (Pellé and Reber, 2013) makes the distinction between engagement aiming to justify decisions so as to maintain acceptance of technology and the more fundamental question of the social acceptability of technology as a whole – is the technology not merely tolerated (because it is very useful, because there is no alternative), or is it positively desirable and beneficial?

As will be clear by now, RRI is not a single topic but encompasses a range of thought, and so the GREAT guidelines are purposefully wide-ranging, bringing in links to different conceptions of RRI and existing assessment frameworks. This gives us the freedom, but also the responsibility, to built structures to support RRI fit for the special position of NQIT, as a major project in its own right and as part of a very large public research and technology investment.

In thinking out how this is applied in NQIT, one of the features of the Hub as a project is its partition into a number of work packages, from Work Package 0 which is the overall architecture, through to Work Package 9 developing capabilities and control systems for the overall Q20:20 machine. These first and last work packages are cross-cutting, and between these ends are specialised work packages which each make important contributions to knowledge and potential spin-offs, which also building towards the Q20:20 machine. At the same time, there are cross-Hub partnership initiatives to work with industry partners and other early adopters, thinking towards the future about what will be possible and what technical solutions will be needed to achieve these ends.

This interdisciplinary cross-work package structure allows a space for RRI, with effectively the status of a work package, but able to act across the Hub. RRI research works closely with the communications manager to engage with the public and others, and with the technology associates who link promising outputs from NQIT with industrial partners, enabling us to get a view on likely early applications.

With the status of a work package, RRI produces quarterly reports for the NQIT Management Board and contributes a section to the annual report. But, so far, these are reports of RRI activities done, rather than of results from RRI research. At NQIT progresses, and especially as spin-offs and applications emerge, we would expect that these reports will increasingly suggest upcoming societal and ethical concerns.

Our recommendation is that these concerns should have a formal status, to which Management Board will be asked to respond. Such a response could include, for example, passing a recommendation up the research policy chain, raising the issues with funders or others for action or amended policy, or leaving an issue “on the table” so that the senior management are made aware of future developments. In some cases, recommendations for future regulation may be appropriate.

37 http://www.great-project.eu/

NQIT RRI workshop where researchers discussed a future scenario involving an artificial intelligence doctor / *University of Oxford Department of Engineering, Science Media Unit*

4.3.4 Co-ordinate responsibility in the hubs, across the hubs, and for the QT programme as a whole

RRI is still developing as a set of processes, with no simple solutions. Every technology area presents specific challenges for RRI, as well as commonalities with other technologies. In quantum technologies, there are differences but also commonalities between the Hubs, and so RRI will have greater impact if there is strong liaison between the Hubs.

NQIT is a major project, large enough in itself to justify having dedicated RRI resources. By this we mean resources for RRI as a research area and Work Package within NQIT, as well as resources to support the project as a whole to develop a responsible attitude and practices.

We in NQIT have been in liaison with each of the other UK Hubs, and we have also been keen to support RRI activities by taking part in data-gathering "Roadshows" around each of the Hubs (including Oxford) under the leadership of EPSRC. Cross-hub liaison provides a forum so that each Hub can benefit from the perspectives of the other Hubs and, as the Roadshows have demonstrated, can start to provide a base of training and awareness of RRI at each Hub.

Many of the issues we have identified in this Landscape Report are common to other Hubs, such as issues relating to quantum security and in some cases to quantum sensors. Others are unique, and the comparison between the Hubs provides useful insights. A cross-hub RRI strategy would avoid duplication of effort, and enable a combined dialogue with stakeholders including NGOs, policy-makers, regulators, industry, and defence; a piecemeal approach is likely to meet some resistance from these stakeholders.

4.3.5 Disseminate RRI through channels to impact research policy

The UK Quantum Technologies Programme is a major high-profile public investment, of interest to science and technology policy at the highest levels of government.

Working together and separately, the Hubs can take the findings from RRI research to influence the trajectory not only of their own research and technology development, but also of the UK technology programme as an instrument of public policy. As Guston (2004) has asserted, research into Ethical, Legal, and Social Impacts (ELSI) has a contribution not only to the outcomes but also to the ways in which research is done, right up to the policy which guides research in the first place: *"ELSI research must be more directly plugged back into the policy process. ELSI programs should include more technology assessment and 'research on research'"* (Guston, 2004). Indeed, a finding in one of the RRI Roadshows was that, while individual hubs may identify RRI issues which have implications for research policy, there is no clear channel for these issues to feed upwards towards policy-makers.

So a further recommendation is to build on the links already established between the hubs and, in conjunction with EPSRC and Innovate UK, to develop new channels to *disseminate* RRI ideas and practices throughout the QT Programme and into public policy around quantum technologies, to provide a route for these findings to reach those who are able to act on them. This would build on the existing roles that senior scientists already play in liaising with government; but in addition, the workshops that we are initiating as part of the next steps in RRI will connect together relevant stakeholders to start a "policy conversation" around key areas.

These channels – acting upwards as well as downwards from research and user engagement to research policy and beyond – could take a number of forms. One simple way to implement this would be by building a repository of issues and methods to identify commonalities across the Hubs. RRI co-ordinators could pass findings up to Hub management boards (as we proposed in the recommendations in Section 4.3.3), and from there, if appropriate, they could be taken to the UK National Quantum Technologies Strategic Advisory Board (QT SAB). In this way, results from RRI research can reach representatives of industry and academia, and, through funders and knowledge exchange organisations such as EPSRC and Innovate UK, can feed to policy-makers in government departments.

4.3.6 Engage in dialogue with the public and interested stakeholders

The public places a significant amount of trust in researchers to develop appropriate technologies on their behalf, but this trust should not be taken for granted. Researchers need to be prepared to engage in discussion with the people who support their work and who will ultimately be beneficiaries of it.

All of the Hubs take part in public outreach activities, and receive audience feedback which is generally supportive. These events offer opportunities for dialogue – that is, two-way communication, listening as well as informing. However, this is usually limited in scope, because this is not a representative sample (most of the people at a science fair are interested in science), and because questions from the audience are only a small part of the event. In fact, from comments in the Roadshows, it seems that attendees at these kinds of events rarely express concern for downstream implications of the science. But insights from RRI show that even apparently benign technologies will have consequences which are unintended or more mixed in their impact – in quantum computing, we have indicated challenges in verification of results and movements towards an algorithmic society as examples of issue areas that we can perceive. Meanwhile, there are parallel negative discourses – such as the article in Sydney Morning Herald

(Cribb, 2016) – which may emerge unexpectedly, in contrast to the optimistic vision presented at public events. We suggest that in public engagement activities, scientists should be prepared to acknowledge areas of potential concern and uncertainty. Taking "ownership" of the downsides as well as the upsides of quantum technologies would leave science much less exposed when negative reactions do appear.

At the same time, the kind of measured, interactive dialogue which can give a broader picture of the public perspective on technology, and identify potentially difficult issues, will require a wider pool of participants and a more structured approach to capture and act on the results. This is public dialogue of the kind which has been led by Sciencewise in other application areas of research and technology; without such dialogue, public attitudes can only be glimpsed from publications aimed at the public (Sciencewise, 2014). To reinforce a point made earlier, this dialogue is not viable at the level of an individual hub, except on a small scale, but would be much more effective across all of the hubs

The two forms of dialogue, one reaching out to an interested audience, the other more systematically bringing in a population sample, can support one another. A suggestion from a Roadshow was that it would be useful to have advice on how to respond to potentially difficult questions at public events; but our view is that such questions, rather than to be avoided, would a useful resource and a pointer to public concerns. RRI can help those who are engaging with the public to prepare to receive this feedback and to give prior anticipation and reflection to the issues which might arise. Conversely, the NQIT community can feed back into RRI by building a record of difficult questions, and collectively we can develop a strategy for answering them in an honest, understandable and objective way.

Jo Johnson MP, UK Minister for Universities and Science, visiting the NQIT labs in Oxford and discussing the UK National Quantum Technology Programme with Professor Ian Walmsley, NQIT Director, / *John Cairns*

4.4 Next steps towards for the RRI framework for the NQIT Hub

The next steps in the Strategy in RRI in NQIT will bring key stakeholders together for deeper investigation of the issues we have identified. This will work across the range of NQIT research, from basic science to application areas, emphasising the overall contributions of NQIT to the UK's capacity in quantum technologies.

The work in this stage will ensure that NQIT is matched to meet its policy and innovation objectives and that those objectives follow a trajectory which includes RRI considerations. At the same time, this will build the network of RRI supporters within NQIT and beyond, seeding the acquisition of expertise and building capacity for RRI into the future.

4.4.1 Case studies

A series of case studies will investigate the issues identified in this Landscape report in greater depth.

The case studies will produce a summary report, but the final aim of the process, of which the case studies are a part, is to start to refine the framework which we presented in Section 4.3.1. This will be "building RRI (the RRI framework) by doing RRI". The case studies will also contribute to a process of education by involving project participants. .

The studies will give insights into social concerns and wider implications from quantum computing, expanding on issues uncovered in this Landscape, including:

- Economic impacts and ownership models
- Surveillance and privacy impact
- Trust in quantum computing
- Uncertainty about the shape of future quantum computing
- Emerging narratives around quantum technologies, including public fears and hopes

The cases will be chosen to cover the identified areas of concern, across two dimensions (Figure 6):

- distance in time from technology readiness or distance from commodification as product or service, and
- and the extent to which the application is "essentially quantum" (could not be done without quantum) or "quantum-enhanced" (quantum is expected to add a step change in performance).

At least one case study will focus on an application area as a lens through which to understand these social implications, but other case studies will be cross-application, for example looking at defence and national security, or pathways to market, rather than at a small range of specific applications

The "essentially quantum" – "quantum enhanced" distinction will allow us to dig deeper into the rhetoric of quantum as enabling powerful new unbreakable, invincible, or irrefutable properties. Some application areas simply could not be done without the use of quantum methods; this includes quantum secure communication and "blind" quantum computing. Many other applications are so much speeded up by quantum methods that operations that are unfeasible

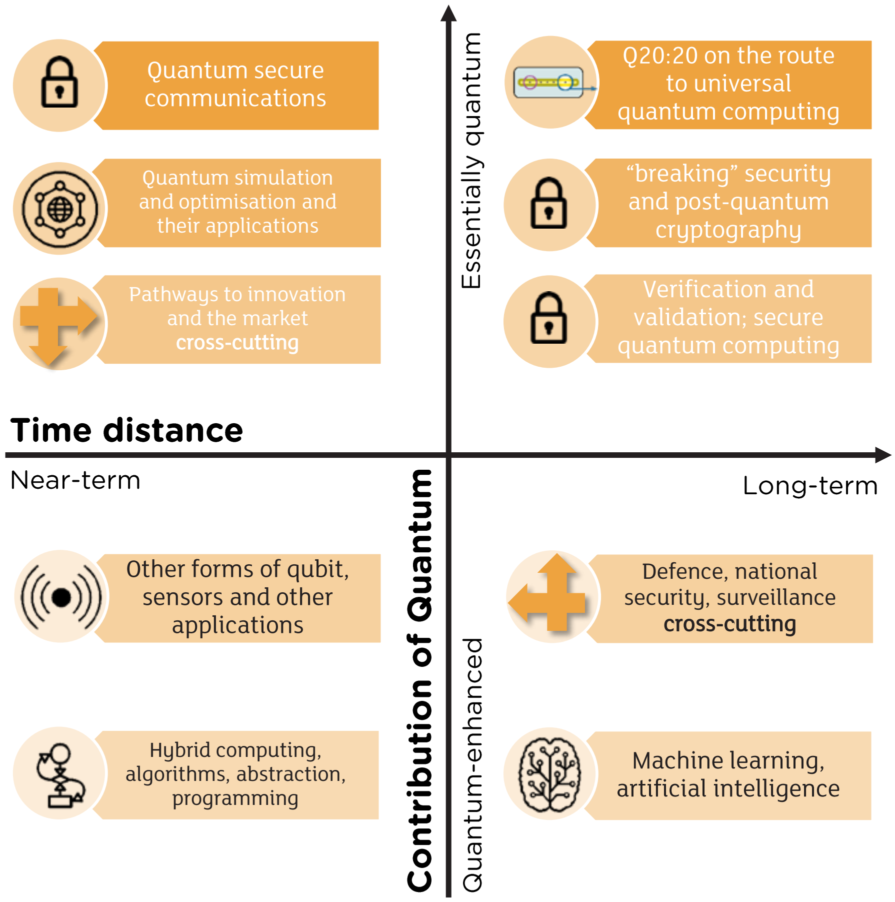

Figure 6: Potential case studies across dimensions of time from commoditisation and "quantum enhanced" to "essentially quantum"

by classical methods become practically possible; examples include most of the known quantum algorithms and methods of optimisation such as with D-Wave. Other application areas are already established without the need for quantum computing; these will gain some advantage but the exact contribution of quantum to the existing methods is not yet clear. A case study would enable us to understand in more depth the implications of these distinctions and the limitations and opportunities for quantum computers in practice.

The exact set of case studies has yet to be finalised, but a discussion has been held with some work package leaders and others in NQIT. Covering all of the possible topics would be beyond the scope of this programme. From feedback from this discussion and from our own analysis during the preparation of this Landscape report, we have drawn up a list from which the three studies will be chosen (some could be combined into a wider study):

High profile cross-cutting areas of interest which will throw light on a range of issues:

- Defence and national security, and dual use applications
- Pathways to innovation and the market; the form and business models of quantum computing

Topics reflecting the work packages of NQIT, including:

- Early applications of quantum computing, including quantum simulation and optimisation
- Hybrid computing, algorithms and abstraction and the path to programmable quantum computers
- Q20:20 on the route to universal quantum computing; what might a real quantum computer look like, how will it evolve, what would it enable and how can we know this?
- Other forms of solid state qubit, with implications as sensors and other applications
- Verification and validation, with implications for secure quantum computing

Other possible application areas, not so much the focus of NQIT

- Secure communications, "breaking" security and post-quantum cryptography
- Machine Learning, AI, Big Data and quantum implications, and how quantum would enable new features

Defence and national security will be a likely area for a case study, to reflect the important contribution of these application areas to the UK National Programme. These application areas will touch on several of the concerns we have identified, in security, surveillance, and "big technology" with a lack of transparency. Conversely, we expect that many of the applications of interest for defence could have non-military uses, but one of the purposes of the case study will be to uncover unexpected uses that we have not yet identified. The wide and quite general topic of defence and security will be further refined in preparation for the workshop to focus the discussion and to ensure that it remains focussed on the social implications of quantum technologies, in a defence or national security context.

Pathways to innovation will consider the "social constitution" of quantum computing: what the market might look like: for example, SMEs competing to offer Internet Service-style quantum computer services and added-value applications; big corporations using quantum computing for their own uses; an array of specialised SMEs with hi-tech products; cloud computing style online, on-demand access. A case study would consider these in high-level terms as far as is possible at present, and the RRI implications for each of these.

The suggested topics in applications in NQIT reflect the work packages and main themes of NQIT. Some of these could be combined; for example, "what might a real quantum computer look like" could include hybrid computers and quantum simulation. Other forms of qubit could cover some of the concerns around quantum sensors as well as some of the potential of diamond qubits and what it might mean if much smaller, more robust quantum computers could be made. The overall aim is to explore the major contributions which NQIT is making to the development of quantum computing, not only by its scientific contributions but also in building a level of skill in the technology and awareness of what quantum may be useful for, so that the UK will be prepared for the emergence of quantum technologies.

The case studies will be a part of the overall strategy to embed RRI in the NQIT Hub. But they will also be a contribution in themselves, producing a report as a resource across the UK Quantum

Technologies Strategy, with a focus less on the science and technology than on the social implications of NQIT.

As far as possible, their focus will not be abstract but will be on concrete applications. For example, modelling quantum processes using simulation would go to the heart of quantum computing and raise interesting theoretical questions, but considering how this might contribute to the discovery of new drugs and new materials and new possibilities for medicine and new products would be a more relevant case study. A participant in the discussion around topics for case studies gave a real-life example of planning for food security, to avert the real risk of famine if climate change and population growth continue at current rates. This is an example of an issue which resonates with a wider group of stakeholders, beyond those who are already interested in quantum.

4.4.2 Exploratory workshops

At the heart of each case study will be one or more exploratory workshops, each of around one half-day, in Oxford, London (for a more high-profile event), or alongside the Quantum UK conference. These will bring together domain experts in the topic of the case study along with industrialists, technology transfer specialists, policy makers, representatives of interested stakeholders in civil society, social scientists, science and technology studies, and philosophers and ethicists, alongside RRI specialists. Allowance has been made in the project budget for participants to travel and, if necessary, be accommodated to take part in workshops.

By bringing together a diversity of viewpoints, these workshops will introduce new insights to the issues we have identified in this Landscape study and introduce and promulgate new ideas. The workshops will deepen our analysis of key issue areas around the narratives, uncertainties, identified concerns and eventual social transformations of quantum computing.

As well as their contribution to refining the framework and to weaving RRI into the fabric of NQIT, these workshops will also initiate and promulgate dialogue between the groups of stakeholders, creating a link between our findings and wider debates. Recall that we emphasised earlier the need for informed, authoritative narrative, where there is still much uncertainty and the potential for the “mysterious” nature of quantum mechanics to lead to misunderstanding (Section 3.1.4).

This will be one way in which we can start to identify competing narratives and distinguish real issues, and take a view of the relative level of concern. Will post-quantum cryptography be implemented before quantum computer are a threat to Internet security? Is QKD a major contribution to security or a specialised interest? Will QKD provide a hiding space for criminals and terrorists? These are some of the well-publicised social issues in quantum computing for which the workshops will enable an informed view. In addition, we have identified concerns around the need for trust if quantum technologies are to be accepted, imbalances in the ownership of these technologies, and increasing disquiet as well as the capacity for improving our lives in what we have called the algorithmic society.

Each workshop will include presentations of the area under discussion and of the RRI issues which have been identified, as an introduction to a free flowing and wide ranging discussion. Participants would be invited to discuss, among other things, how they might introduce RRI into their own areas of work, as a way to carry back the findings to other areas.

The workshops, however, will be only one part of each case study. In preparation, there will be desk research and familiarisation with the topic, and refining the questions to maximise the usefulness of the workshop. There will also be one-to-one interviews, following on from those which form the core empirical work in this Landscape study. Laboratory visits and observations may also be a part of the data for the case studies.

Appendix

5.1 Methodology

The major source of empirical research underlying this report comes from 31 semi-structured one-to-one interviews, conducted between July 2015 and February 2016, with NQIT senior investigators and Work Package leaders, researchers, funders, and potential industrial innovators. This does not pretend to cover all stakeholder groups – that will be a task for the more in-depth case studies later in the project – but the dual aim was to investigate attitudes towards responsible innovation within the project and those close to it, and at the same time to gain a better understanding of the structure of the project.

The interviews were recorded, with the participants' permission, and either fully transcribed or noted in detail. The notes and transcripts were analysed thematically using NVivo[38], a qualitative analysis tool which simplifies mark-up of the raw textual data according to themes of interest.

Interviewees were selected firstly from the investigators (PI and Co-Investigators) listed in the initial EPSRC funding application, and, from there, to researchers recruited into each of the teams of the investigators, funder and finally ad-hoc additions to include some potential early adopters (this included a representative from a standards organisation).

In addition, two RRI scenario-based workshops have been held alongside project events. These workshops had the dual aim of introducing participants to the concepts of RRI and capturing the discussion as a part of the empirical data for this study.

We have also liaised with RRI activities in the other UK QT hubs, and more recently, a series of RRI "Roadshows" is enabling us to identify cross-cutting issues at other hubs.

In addition to interviews, this study drew on a collection of 43 documents around the project and over one hundred pieces of literature relating to quantum technologies and quantum computers, and web pages.

5.1.1 Research questions

The semi-structured interviews followed a schedule of 15 questions, which was developed and amended slightly during the course of the data collection. Being semi-structured means that the interviewer was free to take the questions in any order, depending on the responses, keeping to an informal, conversational style, avoiding a "stimulus-response" model of interaction.

The schedule was oriented around the two research questions which underlie this landscape study:

> *RQ1: What are the main thematic areas in NQIT at technical, engineering, application and innovation levels, how can they be classified, and how are they inter-connected?*
>
> *RQ2: How do the stakeholders within the project and commercial partners understand the economic, environmental, and social impacts from their research and innovation?*

---

38 http://www.qsrinternational.com/product

The first question is clearly relatively technical and the second relatively social, so the aim of this study is to show how these two dimensions are inter-related.

And, the broader question:

> *What are the challenges and opportunities of NQIT, and what are the implications for Responsible Research and Innovation?*

5.1.2 Interview schedule

Outline questions:

1. About the participant and their relation to NQIT.

   It will be important at this stage to understand the perspective of different participants, but to take care not to introduce unnecessary details which might reveal individual identities

   a. Could you describe the organisation you work for, your position in organisation, and the place of that organisation in NQIT?

   b. Are you funded from the NQIT project, or some other source?

   c. How long have you been working on this?

   d. What are your particular areas of expertise?

   e. Please describe your area of responsibilities (day to day…, this week, ongoing. …)

   f. What other areas, outside NQIT? (for example, admin responsibilities, future funding, supporting others in the group, work on other projects)

   g. A typical day, or today

   h. Are you a member of any professional organisations (eg. British Science Association)

   i. Of any relevant campaigns or NGOs or civil society organisations?

2. How would you describe the contribution of your work to NQIT: (open question)

   a. [Prompt] Is it basic research, applied, a mixture …?

   b. How does it fit with the work packages and the work plan as a whole?

   c. What is your place in this? leading - doing – administrating …

3. How would you describe the expected results from your work:

   a. In the short term - say, towards the end of NQIT?

   b. ultimately from the project and beyond?

4. In what ways do you see your research influencing the real world in:

   a. By the end of NQIT?

   b. 10-15 years?

5. If they say it's too early to say:

   a. Where good or bad outcomes may or may not be anticipated, and may even be impossible to foresee?

b. Try to think of 2-3 things, even very speculative, not-yet-possible things, which might come out of your research

6. In what ways do you think your research will contribute to the "public good"?
    a. Are there ethical considerations in your research?
    b. Legal, social, implications …
7. What do you feel about the way your area of Quantum Technology expertise is communicated to, or understood by (for example, in specialist, popular science, or other media):
    a. QT researchers in other areas?
    b. Computer scientists and people working in IT?
    c. Other disciplines?
    d. Potential adopters?
    e. The general public?
    f. How would you involve the general public and other stakeholders in dialogue?
8. What motivates you in your research?
9. How would you feel if there was some really bad outcome, perhaps years in the future, from your research?
    a. They can define "bad" in any way they like
10. Are there any areas where you can see (perhaps potentially) conflicting perspectives between different groups of stakeholders? Within NQIT? Other university departments? Beyond?
    a. Examples (prompt) — privacy and data linking — surveillance — changed relations with law enforcement — environmental or energy issues
    b. How would you think if the intention – by you or other researchers – is good but the outcome is bad? [prompt: tie into their expressed motivations]
    c. From Stilgoe round table: secrecy, surveillance, technological ownership and technological access
11. Do you think the outcomes from your research could possibly change people's behaviour?
    a. In what ways?
12. How do you justify spending public money on this research?
13. Do you think it is part of a researcher's role to ask themselves questions about the implications of their research, or should they concentrate on their core expertise?
    a. How do you deal with this? Examples (prompt) are there structures or guidelines to identify and manage potential impacts, risks, and uncertainties in your research area?
14. Are you familiar with the concept of Responsible Research and Innovation? This is a dynamic, iterative process aiming to enable all stakeholders involved in the research and innovation process— including academics, funders, commercial partners, but also

professional organisations, civil society, and the wider public — to engage with one another and to share responsibility to ensure that the outcomes are socially desirable and that the research and innovation process is responsive, transparent, reflective, and inclusive. At the same time, this process should enable researchers to think creatively about social impacts at all stages.

a. Does your university [institution], department or research group have best practice guidelines, as far as you are aware?

b. How do you see this as different from self-regulation by peer review and PIs?

c. What do you think about this?

d. Could you suggest ways of introducing RRI into research projects?

5.1.3 Consent form for interview participants

Department of Computer Science,
University of Oxford
Wolfson Building
Parks Road
OX1 3QD
Oxford

Participant Consent Form

Responsible Research & Innovation in Networked Quantum Information Technology

Please read and initial those points below that you agree with.

	Initials
1. I confirm that the purpose of the study has been explained to me and I have had the opportunity to ask questions about the research and have had these answered satisfactorily.	--------
2. I understand that my participation is voluntary, and that I am free to withdraw at any time without giving any reason and without any implications for my legal rights.	--------
3. I am allowing the researcher to audio record me and take handwritten notes as part of the study. The recording will be transcribed. I understand that anonymised quotes may be used in presentations or publications arising from the research but not in any way that might allow for identification of individual participants.	--------
4. [If applicable]I am allowing the researcher to video record or photograph me as part of this study. Recordings will be transcribed. I understand that still images may be used in publications arising from the research but that faces and other identifying features will be pixelized or disguised.	--------
5. I agree that any audio or video recordings may also be viewed and discussed with, but not made directly available to, close colleagues of the RRI in NQIT research team.	Yes/No
6. I understand the data will be kept confidential at all times.	--------
7. I understand that transcripts and other data will be archived and made available for further research, but not in any way that might allow for identification of individual participants	--------
8. I agree to take part in this research.	

Name of Researcher:

Signature: Signature:

Date: Date:

If you have questions or concerns about any aspect of this project you may contact the principle investigator: Marina Jirotka, Department of Computer Science, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK, +44 (0) 1865 610613, or by e-mail at marina.jirotka@cs.ox.ac.uk who will do her best to answer your query. If you remain unhappy and wish to make a formal complaint, please contact the Research Ethics Committee at the University of Oxford at ethics@socsci.ox.ac.uk; +44 (0)1865 614871; Social Sciences & Humanities Inter-Divisional Research Ethics Committee, Oxford University, Hayes House, 75 George Street, Oxford, OX1 2BQ, UK.

5.1.4 Documents

The following documents, relating to NQIT, were specifically analysed as part of the research, in addition to a thorough review of relevant QT and RRI literature

- FRRIICT:
  - Framework for Responsible Research and Innovation in ICT.pdf
  - RRIinICT-draft-AB.docx
- NQIT documents:
  - 2015_2_11_QT_Showcase_Kai Bongs.pptx
  - 130304_DM_final.pdf
  - 0505151.pdf
  - Brassard et al 1999 9901035v1.pdf
  - Control System PB.docx
  - Costs of LOQC - 1504.02457v2.pdf
  - DCOB NQIT Project Forum integration overview.pptx
  - developer job ad - for finance issues.pdf
  - Governing_at_the_Nanoscale.pdf
  - Horsman et al - models and computers Arxiv 1309.7979.pdf
  - J O'Brien et al - nphoton.2009.229.pdf
  - Minutes of Ion Trap Meeting.docx
  - NQIT WP report form WP8.doc
  - NQIT WP report form WP9.docx
  - NQIT WP TAB report form WP9.docx
  - NQIT_WP6_NetworkedQuantumSensors_20150720.pdf
  - Nunn et al 2013 oe-21-13-15959.pdf
  - p52-valiron - programming quantum.pdf
  - Q2020 Systems Vision.pptm
  - TAB report 2015 pictures.pptx
  - TAB report 2015 V0.1.docx
  - TAB report 2015 V1.0.docx
  - Wallden et al 2015 - PhysRevA.91.042304.pdf
  - WP leaders reports:
    - July 2015 MB:
      - NQITMB5WP0.pdf

NQITMB5WP1.pdf

NQITMB5WP2.pdf

NQITMB5WP3.pdf

NQITMB5WP4.pdf

NQITMB5WP5.pdf

NQITMB5WP6.pdf

NQITMB5WP7.pdf

NQITMB5WP8.pdf

NQITMB5WP9.pdf

WP6 TAB report_1.docx

WP9 meeting 1 minutes 2015-01-08.docx

Press and other public discourse:

Quantum computing and the dawn of the quantum tyranny.pdf

5.1.5 Interviews

Date	Anonymous Code	Interviewee Type
24/07/2015	LR1	PDRA
29/07/2015	LR2	PDRA
29/07/2015	LR3	PDRA
30/07/2015	LR4	PDRA
11/08/2015	LR5	PDRA
12/08/2015	LR6	PDRA
14/08/2015	LR7	PDRA
19/08/2015	LR8	PDRA
02/09/2015	LR9	PDRA
10/09/2015	LR10	PDRA
16/09/2015	LR11	PDRA
17/09/2015	LS1	Senior researcher or WP leader
24/09/2015	LS2	Senior researcher or WP leader
02/10/2015	LD1	NQIT Director or co-director
16/10/2015	LS3	Senior researcher or WP leader
19/10/2015	LS4	Senior researcher or WP leader
20/10/2015	LD3	NQIT Director or co-director
21/10/2015	LD2	NQIT Director or co-director
28/10/2015	LS5	Senior researcher or WP leader
03/11/2015	LI1	Industrial partner, innovator, or funder
04/11/2015	LS6	Senior researcher or WP leader
05/11/2015	LS7	Senior researcher or WP leader
12/11/2015	LI2	Industrial partner, innovator, or funder
16/11/2015	LS8	Senior researcher or WP leader
01/12/2015	LI3	Industrial partner, innovator, or funder
02/12/2015	LI4	Industrial partner, innovator, or funder
10/12/2015	LS9	Senior researcher or WP leader
17/12/2015	LI5	Industrial partner, innovator, or funder
05/01/2016	LI6	Industrial partner, innovator, or funder
03/02/2016	LD4	NQIT Director or co-director

References

Aaronson, S. 2008. The Limits of Quantum. Scientific American, March 2008 p 8.

Aaronson, S. BQP and the polynomial hierarchy. Proceedings of the forty-second ACM symposium on Theory of computing, 2010. ACM, pp 141-50.

Aaronson, S. & Arkhipov, A. 2011. The computational complexity of linear optics. Proceedings of the forty-third annual ACM symposium on Theory of computing. San Jose, California, USA: ACM.

Acín, A. 2014. Randomness in quantum physics. Euresis, 6 p 35.

Adams, A. & Sasse, M. A. 1999. Users are not the enemy. Communications of the ACM, 42, 12 pp 40-6.

Adcock, J., Allen, E., Day, M., Frick, S., Hinchliff, J., Johnson, M., . . . Stanisic, S. 2015. Advances in quantum machine learning. Quantum, Forthcoming arXiv preprint arXiv:1512.02900v1.

Agre, P. E. 2002. Real-Time Politics: The Internet and the Political Process. The Information Society, 18, 5 pp 311-31, 10.1080/01972240290075174: 10.1080/01972240290075174.

Alléaume, R., Branciard, C., Bouda, J., Debuisschert, T., Dianati, M., Gisin, N., . . . Zeilinger, A. 2014. Using quantum key distribution for cryptographic purposes: A survey. Theoretical Computer Science, 560 pp 62-81.

Ambainis, A. 2014. What Can We Do with a Quantum Computer? IAS Ideas[Online: 04/07/2016]. Available: https://www.ias.edu/ideas/2014/ambainis-quantum-computing [Accessed 22/07/2016].

Barz, S., Fitzsimons, J. F., Kashefi, E. & Walther, P. 2013. Experimental verification of quantum computation. Nature physics, 9, 11 pp 727-31.

Barz, S., Kashefi, E., Broadbent, A., Fitzsimons, J. F., Zeilinger, A. & Walther, P. 2012. Demonstration of blind quantum computing. Science, 335, 6066 pp 303-8.

Beck, M., Gupta, H., Rastetter, E., Shoemaker, C., Tarboton, D., Butler, R., . . . Harmon, T. 2009. Grand challenges of the future for environmental modeling.

Bennett, C. H. & DiVincenzo, D. P. 2000. Quantum information and computation. Nature, 404, 6775 pp 247-55.

Bernstein, D. J. 2014. A subfield-logarithm attack against ideal lattices. cr.yp.to [Blog] 13/02/2014. https://blog.cr.yp.to/20140213-ideal.html [Accessed 21/07/2016].

Bernstein, D. J. 2016a. Is the security of quantum cryptography guaranteed by the laws of physics? sidechannels.cr.yp.to [Blog]. https://sidechannels.cr.yp.to/qkd/holographic-20160326.pdf [Accessed 21/07/2016].

Bernstein, D. J. 2016b. Security fraud in Europe's "Quantum Manifesto". cr.yp.to [Blog]

Aaronson, S. 2008. The Limits of Quantum. Scientific American, March 2008 p 8.

Aaronson, S. BQP and the polynomial hierarchy. Proceedings of the forty-second ACM symposium on Theory of computing, 2010. ACM, pp 141-50.

Aaronson, S. & Arkhipov, A. 2011. The computational complexity of linear optics. Proceedings of the forty-third annual ACM symposium on Theory of computing. San Jose, California, USA: ACM.

Acín, A. 2014. Randomness in quantum physics. Euresis, 6 p 35.

Adams, A. & Sasse, M. A. 1999. Users are not the enemy. Communications of the ACM, 42, 12 pp 40-6.

Adcock, J., Allen, E., Day, M., Frick, S., Hinchliff, J., Johnson, M., . . . Stanisic, S. 2015. Advances in quantum machine learning. Quantum, Forthcoming arXiv preprint arXiv:1512.02900v1.

Agre, P. E. 2002. Real-Time Politics: The Internet and the Political Process. The Information Society, 18, 5 pp 311-31, 10.1080/01972240290075174: 10.1080/01972240290075174.

Alléaume, R., Branciard, C., Bouda, J., Debuisschert, T., Dianati, M., Gisin, N., . . . Zeilinger, A. 2014. Using quantum key distribution for cryptographic purposes: A survey. Theoretical Computer Science, 560 pp 62-81.

Ambainis, A. 2014. What Can We Do with a Quantum Computer? IAS Ideas[Online: 04/07/2016]. Available: https://www.ias.edu/ideas/2014/ambainis-quantum-computing [Accessed 22/07/2016].

Barz, S., Fitzsimons, J. F., Kashefi, E. & Walther, P. 2013. Experimental verification of quantum computation. Nature physics, 9, 11 pp 727-31.

Barz, S., Kashefi, E., Broadbent, A., Fitzsimons, J. F., Zeilinger, A. & Walther, P. 2012. Demonstration of blind quantum computing. Science, 335, 6066 pp 303-8.

Beck, M., Gupta, H., Rastetter, E., Shoemaker, C., Tarboton, D., Butler, R., . . . Harmon, T. 2009. Grand challenges of the future for environmental modeling.

Bennett, C. H. & DiVincenzo, D. P. 2000. Quantum information and computation. Nature, 404, 6775 pp 247-55.

Bernstein, D. J. 2014. A subfield-logarithm attack against ideal lattices. cr.yp.to [Blog] 13/02/2014. https://blog.cr.yp.to/20140213-ideal.html [Accessed 21/07/2016].

Bernstein, D. J. 2016a. Is the security of quantum cryptography guaranteed by the laws of physics? sidechannels.cr.yp.to [Blog]. https://sidechannels.cr.yp.to/qkd/holographic-20160326.pdf [Accessed 21/07/2016].

Bernstein, D. J. 2016b. Security fraud in Europe's "Quantum Manifesto". cr.yp.to [Blog] 16/05/2016. https://blog.cr.yp.to/20160516-quantum.html [Accessed 21/07/2016].

Biever, C. 2013. Don't worry, my algorithm won't break the internet… yet (Interview with Peter Shor). New Scientist, 220, 2944 p 29.

Collingridge, D. 1982. The social control of technology. London: Pinter.

Cressey, D. 2012. Cancelled project spurs debate over geoengineering patents. Nature, 485, 7399 p 429.

Cribb, J. 2016. Quantum computing and the dawn of the quantum tyranny. Sydney Morning

Herald, January 10, 2016.

Datta, A., Tschantz, M. C. & Datta, A. 2015. Automated experiments on ad privacy settings. Proceedings on Privacy Enhancing Technologies, 2015, 1 pp 92-112.

Davies, C. 2013. Google using $10m quantum computer to understand Glass future. SlashGear[Online: 10/10/2013]. Available: http://www.slashgear.com/google-using-10m-quantum-computer-to-understand-glass-future-10300936/ [Accessed 22/07/2016].

Deutsch, D. 1985. Quantum theory, the Church-Turing principle and the universal quantum computer. Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 400, 1818 pp 97-117.

DiVincenzo, D. P. 2000. The physical implementation of quantum computation. arXiv preprint quant-ph/0002077v3.

Eden, G., Jirotka, M. & Stahl, B. C. 2013. Responsible Research and Innovation in ICT: Summary of key issues, recommendations, challenges and enablers.

Ekert, A. 2001. Quantum entanglement and secrecy. International Journal of Modern Physics A, 16, 19 pp 3191-202.

European Commission 2016a. Horizon 2020 Work Programme 2016-2017: 16. Science with and for Society.

European Commission. 2016b. Quantum Manifesto: A New Era of Technology. [Online: 17/05/2016]. Available: http://qurope.eu/system/files/u7/93056_Quantum%20Manifesto_WEB.pdf [Accessed 21/07/2016].

European Telecommunications Standards Institute. 2015. Quantum Key Distribution. Available: http://www.etsi.org/images/files/ETSITechnologyLeaflets/QuantumKeyDistribution.pdf [Accessed 21/07/2016].

Faullimmel, N. & Stilgoe, J. 2015. Responsible Research and Innovation in Quantum Technologies: A briefing note (Report of interviews and round table workshop, March 2015).

Fernick, J. & Gheorghiu, V. 2016. Quantum Computational Intelligence. Institute for Quantum Computing Blog [Blog] 29/03/2016. https://uwaterloo.ca/institute-for-quantum-computing/blog/post/quantum-computational-intelligence [Accessed 12/07/2016].

Feynman, R. P. 1982. Simulating physics with computers. International journal of theoretical physics, 21, 6 pp 467-88.

Fisher, E., Mahajan, R. L. & Mitcham, C. 2006. Midstream modulation of technology: governance from within. Bulletin of Science, Technology & Society, 26, 6 pp 485-96.

Fisher, E. & Rip, A. 2013. Responsible Innovation: Multi-Level Dynamics and Soft Intervention Practices. In: Owen, R., Bessant, J. R. & Heintz, M. (eds.) Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society. Chichester, UK: John Wiley & Sons Ltd.

Freeman, J., Hellgren, T., Mastroeni, M., Paoli, G. P., Robertson, K. & Black, J. 2015. Innovation Models: Enabling new defence solutions and enhanced benefits from science and technology. Available: http://www.rand.org/pubs/research_reports/RR840.html [Accessed 12/07/2016].

Friedman, B. & Nissenbaum, H. 1996. Bias in computer systems. ACM Transactions on Information

Systems (TOIS), 14, 3 pp 330-47.

Gangadharan, S. P. 2015. Introduction. In: Gangadharan, S. P., Eubanks, V. & Barocas, S. (eds.) Data and discrimination: Collected essays. Open Technology Institute and New America Foundation.

Giddens, A. 1990. The Consequences of Modernity. Stanford: Stanford University Press.

Gilmore, R. 1995. Alice in Quantumland. Springer Science & Business Media.

Goldschmidt, R. & Renn, O. 2006. Meeting of minds-European citizens' deliberation on brain sciences: final report of the external evaluation. Available: http://elib.uni-stuttgart.de/bitstream/11682/5500/1/AB05_GoldschmidtRenn_EvaluationReport.pdf [Accessed 12/07/2016].

Gottesman, D. 2014. Fault-tolerant quantum computation with constant overhead. Quantum Information & Computation, 14, 15-16 pp 1338-72.

Grimpe, B., Hartswood, M. & Jirotka, M. Towards a closer dialogue between policy and practice: responsible design in HCI. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2014. ACM, pp 2965-74.

Grove-White, R., Macnaghten, P. & Wynne, B. 2000. Wising Up: The public and new technologies. Available: http://csec.lancs.ac.uk/docs/wising_upmacnaghten.pdf [Accessed 12/07/2016].

Guston, D. 2004. Forget Politicizing Science. Let's Democratize Science! Issues in Science and Technology, 21, 1.

Guston, D. 2013. "Daddy, Can I Have a Puddle Gator?": Creativity, Anticipation, and Responsible Innovation. In: Owen, R., Bessant, J. & Heintz, M. (eds.) Responsible Innovation: Managing the responsible emergence of science and innovation in society. Chichester, UK: John Wiley & Sons, Ltd.

Heersmink, R., van den Hoven, J. & Timmermans, J. 2014. ETICA Deliverable 2.2 Normative Issues - Final Report.

Human Brain Project. Undated. What are the ethical issues involved in simulating a human brain and in technology derived from human brain simulation? Available: https://www.humanbrainproject.eu/faq/ethics [Accessed 12/07/2016].

Huttner, B. 2016. Quantum Technologies and Real World Information Securit Challenges. Cloud Security Alliance Industry Blog [Blog] 25/02/2016. https://blog.cloudsecurityalliance.org/2016/02/25/quantum-technologies-and-real-world-information-security-challenges/ [Accessed 21/07/2016].

Ihde, D. 1991. Instrumental realism: The interface between philosophy of science and philosophy of technology. Indiana University Press.

Inglesant, P., Hartswood, M. & Jirotka, M. 2015. NQIT and Quantum Technologies Hubs Responsible Research & Innovation: Strategy. Internal NQIT report.

Jennewein, T. & Higgins, B. 2013. The quantum space race. Physics World, 26, 03 p 52.

Johnson, T. H., Clark, S. R. & Jaksch, D. 2014. What is a quantum simulator? EPJ Quantum Technology, 1, 1 p 1 arXiv preprint arXiv:1405.2831.

Jonas, H. 1982. Technology as a Subject for Ethics. Social Research, 49, 4 pp 891-8.

Jordan, B. 2009. Blurring boundaries: The" real" and the "virtual" in hybrid spaces. Human

Organization, 68, 2 pp 181-93.

Kallinikos, J. 2005. The order of technology: Complexity and control in a connected world. Information and organization, 15, 3 pp 185-202.

King, J., Yarkoni, S., Nevisi, M. M., Hilton, J. P. & McGeoch, C. C. 2015. Benchmarking a quantum annealing processor with the time-to-target metric. arXiv preprint arXiv:1508.05087.

Kling, R. 1992. Audiences, narratives, and human values in social studies of technology. Science, Technology, & Human Values, 17, 3 pp 349-65.

Koblitz, N. & Menezes, A. 2015. A riddle wrapped in an enigma. [Online: 03/12/2015]. Available: https://eprint.iacr.org/2015/1018.pdf [Accessed 22/07/2016].

Korzh, B., Lim, C. C. W., Houlmann, R., Gisin, N., Li, M. J., Nolan, D., . . . Zbinden, H. 2015. Provably secure and practical quantum key distribution over 307 km of optical fibre. Nature Photonics, 9, 3 pp 163-8.

Lee, R. G. & Petts, J. 2013. Adaptive governance for responsible innovation. In: Owen, R., Bessant, J. R. & Heintz, M. (eds.) Responsible innovation: Managing the responsible emergence of science and innovation in society. Chichester, UK: John Wiley & Sons, Ltd.

Lindner, R., Kuhlmann, S., Randles, S., Bedsted, B., Gorgoni, G., Griessler, E., . . . Mejlgaard, N. 2014. Navigating Towards Shared Responsibility in Research and Innovation Approach, Process and Results of the Res-AGorA Project. Available: http://res-agora.eu/news/navigating-towards-shared-responsibility/ [Accessed 12/07/2016].

Lloyd, S. 2000. Ultimate physical limits to computation. Nature, 406, 6799 pp 1047-54.

Lloyd, S., Garnerone, S. & Zanardi, P. 2016. Quantum algorithms for topological and geometric analysis of data. Nature Communications, 7:10138 10.1038/ncomms10138: 10.1038/ncomms10138 [Published online 25/01/2016].

MacKenzie, D. & Spinardi, G. 1995. Tacit knowledge, weapons design, and the uninvention of nuclear weapons. American journal of sociology, 101, 1 pp 44-99.

Macnaghten, P. & Owen, R. 2011. Environmental science: good governance for geoengineering. Nature, 479, 7373 p 293.

Mager, A. 2012. Algorithmic ideology: How capitalist society shapes search engines. Information, Communication & Society, 15, 5 pp 769-87.

Mayers, D. & Yao, A. Quantum cryptography with imperfect apparatus. Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on, 08/11/1998-11/11/1998 1998 Palo Alto, CA, USA. IEEE, pp 503-9.

McEliece, R. J. 2006. Life, the Universe and entropy. Review of: Seife, C.: Decoding the Universe: How the New Science of Information is Explaining Everything in the Cosmos, From our Brains to Black Holes. Nature 441, Issue p 409.

Merali, Z. 2012. The quantum space race. Nature, 492, 7427 pp 22-5.

Muehlhauser, L. & Bostrom, N. 2014. Why we need friendly AI. Think, 13, 36 pp 41-7.

Nielsen, M. 2008. Quantum computing for everyone. Michaelnielsen [Blog] 28/08/2008. http://michaelnielsen.org/blog/quantum-computing-for-everyone/ [Accessed 04/07/2016].

Owen, R., Stilgoe, J., Mcnaghten, P., Gorman, M., Fisher, E. & Guston, D. 2013. A Framework for

Responsible Innovation. In: Owen, R., Bessant, J. & Heintz, M. (eds.) Responsible Innovation: Managing the responsible emergence of science and innovation in society. Chichester, UK: John Wiley & Sons, Ltd.

Parliamentary Office of Science and Technology 2009. The Dual-Use Dilemma: POST Note Number 340. Available: http://researchbriefings.parliament.uk/ResearchBriefing/Summary/POST-PN-340 [Accessed 12/07/2016].

Pellé, S. & Reber, B. 2013. Governance of Responsible Innovation: The Theoretical Landscape. Deliverable 2.2. Available: http://www.great-project.eu/deliverables_files/deliverables03 [Accessed 12/07/2016].

Pellizzoni, L. 2004. Responsibility and environmental governance. Environmental politics, 13, 3 pp 541-65.

Pilkey, O. H. & Pilkey-Jarvis, L. 2007. Useless arithmetic: why environmental scientists can't predict the future. Columbia University Press.

Preskill, J. 2012. Quantum computing and the entanglement frontier. arXiv preprint arXiv:1203.5813.

Pritchard, J. & Till, S. 2014. UK Quantum Technology Landscape. Available: https://www.epsrc.ac.uk/newsevents/pubs/dstl-uk-quantum-technology-landscape-2014/ [Accessed 18/07/2016].

Rafter, N. H. 1997. Creating born criminals. University of Illinois Press.

Randles, S., Laredo, P., Loconto, A. M., Walhout, B. & Lindner, R. 2016. Framings and frameworks: six grand narratives of de facto RRI. In: Navigating Towards Shared Responsibility in Research and Innovation. Approach, Process and Results of the Res-AGorA Project. Fraunhofer Institute for Systems and Innovation Research (ISI).

Royal Society. 2004. Nanoscience and nanotechnologies: opportunities and uncertainties. Available: https://royalsociety.org/topics-policy/publications/2004/nanoscience-nanotechnologies/ [Accessed 21/07/2016].

Saltelli, A. & Funtowicz, S. 2014. When all models are wrong. Issues in Science and Technology, 30, 2 pp 79-85.

Sasse, M. A. & Flechais, I. 2005. Usable security: Why do we need it? How do we get it? In: Cranor, L. F. & Garfinkel, S. (eds.) Security and Usability: Designing secure systems that people can use. Sebastopol, CA, USA: O'Reilly.

Schneier, B. 2008. Quantum Cryptography. Schneier on Security [Blog] 21/10/2008. https://www.schneier.com/blog/archives/2008/10/quantum_cryptog.html [Accessed 22/07/2016].

Schneier, B. 2015. Automatic Face Recognition and Surveillance. Schneier on Security [Blog] 05/10/2015. https://www.schneier.com/blog/archives/2015/10/automatic_face_.html [Accessed 12/07/2016].

Sciencewise 2014. Public attitudes to quantum technology. Available: http://sciencewise-erc.org.uk/cms/assets/Uploads/Quantum-Technology-Social-IntelligenceFINAL.pdf [Accessed 22/07/2016].

Shor, P. W. 1999. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM review, 41, 2 pp 303-32.

Simonite, T. 2012. The CIA and Jeff Bezos Bet on Quantum Computing. MIT Technology

Review[Online: 04/10/2012]. Available: https://www.technologyreview.com/s/429429/the-cia-and-jeff-bezos-bet-on-quantum-computing/ [Accessed 12/07/2016].

Stahl, B. C. 2011. What does the future hold? A critical view of emerging information and communication technologies and their social consequences. In: Researching the Future in Information Systems. Springer.

Stahl, B. C. 2013. Responsible research and innovation: The role of privacy in an emerging framework. Science and Public Policy, 40 pp 708-16.

Stahl, B. C., Eden, G. & Jirotka, M. 2013. Responsible research and innovation in Information and Communication Technology. In: Owen, R., Bessant, J. R. & Heintz, M. (eds.) Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society. Chichester, UK: John Wiley & Sons Ltd.

Stahl, B. C. & Flick, C. ETICA Workshop on Computer Ethics: Exploring Normative Issues. IFIP PrimeLife International Summer School on Privacy and Identity Management for Life, 2010. Springer, pp 64-77.

Stahl, B. C., Rogerson, S. & Wakunuma, K. J. Future technologies: The matter of emergent ethical issues in their development. Future Computing, Service Computation, Cognitive, Adaptive, Content, Patterns, 2009. COMPUTATIONWORLD'09. Computation World:, 2009. IEEE, pp 603-7.

Stares, P. B. 1985. The militarization of space: US policy, 1945-1984. Ithaca, NY, USA: Cornell University Press.

Stilgoe, J., Owen, R. & Macnaghten, P. 2013. Developing a framework for responsible innovation. Research Policy, 42, 9 pp 1568-80.

Stirling, A. 2008. "Opening up" and "closing down": Power, Participation, and Pluralism in the Social Appraisal of Technology. Science, technology & human values, 33, 2 pp 262-94.

Suchman, L. 2007. Human-machine reconfigurations: Plans and situated actions. Cambridge University Press.

Sweeney, L. 2013. Discrimination in online ad delivery. Queue, 11, 3 p 10.

Tahan, C. 2007. Spookytechnology and Society. arXiv preprint arXiv:0710.2537.

Thompson, W. C. The potential for error in forensic DNA testing (and how that complicates the use of DNA databases for criminal identification). Council for Responsible Genetics (CRG) National Conference: Forensic DNA databases and race: Issues, Abuses and Actions, 2008.

Townsend, P., Rarity, J. & Tapster, P. 1993. Enhanced single photon fringe visibility in a 10-km-long prototype quantum cryptography channel. Electronics Letters, 29, 14 pp 1291-3.

Tucker, J. B. & Zilinskas, R. A. 2006. The promise and perils of synthetic biology. The New Atlantis, 12 pp 25-45.

UK National Quantum Technologies Programme Strategic Advisory Board 2015. National strategy for quantum technologies: A new era for the UK. Engineering and Physical Sciences Research Council.

Vallone, G., Bacco, D., Dequal, D., Gaiarin, S., Luceri, V., Bianco, G. & Villoresi, P. 2015. Experimental satellite quantum communications. Physical Review Letters, 115, 4 p 040502.

Vazirani, U. & Vidick, T. 2014. Fully Device-Independent Quantum Key Distribution. Physical Review Letters, 113, 14 p 140501 arXiv preprint arXiv:1210:1810v2.

Vazirani, U. & Vidick, T. 2016. Erratum: Fully Device-Independent Quantum Key Distribution [Phys. Rev. Lett. 113, 140501, (2014)]. Physical Review Letters, 116, 8 p 089901.

Von Schomberg, R. 2011. Prospects for Technology Assessment in a Framework of Responsible Research and Innovation. In: Dusseldorp, M. & Beecroft, R. (eds.) Technikfolgen abschätzen lehren: Bildungspotenziale transdisziplinärer Methoden. Wiesbaden, Germany: Vs Verlag.

Von Schomberg, R. 2013. A vision of responsible research and innovation. In: Owen, R., Bessant, J. R. & Heintz, M. (eds.) Responsible innovation: Managing the responsible emergence of science and innovation in society. Chichester, UK: John Wiley & Sons Ltd.

Weimer, H. 2011. Moore's Law for Quantum Computers. Quantenblog: Having fun with science and technology [Blog]. http://www.quantenblog.net/physics/moores-law-quantum-computer [Accessed 23/6/2016].

Whitten, A. & Tygar, J. D. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. 8th Usenix Security Symposium, 23/08/1999-26/08/1999 1999 Washington, D.C., USA. pp 169-84.

Wiebe, N., Kapoor, A. & Svore, K. M. 2014. Quantum deep learning. arXiv preprint arXiv:1412.3489.

Winner, L. 1989. The Whale and the Reactor. A Search for the Limits in the Age of High Technology. Chicago and London: The University of Chicago press.

Wolpe, P. R. 2007. Ethical and social challenges of brain-computer interfaces. American Medical Association Journal of Ethics: Virtual Mentor, 9, 2 pp 128-31.

Woolgar, S. & Neyland, D. 2013. Mundane governance: Ontology and accountability. Oxford University Press.

Wynne, B. 1993. Public uptake of science: a case for institutional reflexivity. Public Understanding of Science, 2, 4 pp 321-37.

For more information about NQIT, please visit our website:

http://www.nqit.ox.ac.uk

Or send us an email:

Or get in touch

email: contact@nqit.org

@NQIT_QTHub

NQIT Networked Quantum Information Technologies

UK NATIONAL QUANTUM TECHNOLOGIES PROGRAMME