

Quantum Information Processing Landscape 2020: Prospects for UK Defence and Security

Andrew Middleton and Stephen Till

Lt Cdr Matt Steele RN

DSTL/TR121783
June 2020

**Cyber & Information
Systems Division**
Dstl Porton Down
Salisbury SP4 0JQ
Strategic Command
Jt User Intelligence &
Cyber

© Crown Copyright 2020

UK OFFICIAL

Release Conditions

This document has been prepared for MOD and, unless indicated, may be used and circulated in accordance with the conditions of the Order under which it was supplied.

It may not be used or copied for any non-Governmental or commercial purpose without the written agreement of Dstl.

© Crown Copyright 2020

Intellectual Property Department
Defence Science and Technology Laboratory
Porton Down, Salisbury, Wiltshire SP4 0JQ

Authorisation		
Role	Name	Date
Project Manager	Fiona Beeching	
Technical Reviewer	Dr Andrew Bennett	
Authors	Andrew Middleton / Stephen Till / Lt Cdr Matt Steele RN	

UK OFFICIAL
Draft for comment

Executive summary

Powered by the transistor and, later, the integrated circuit, technologies belonging to the 'First Quantum Revolution' came from understanding the nature of quantum mechanics. These devices are ubiquitous and critical to daily life: we wear them, we access and process information through them and they enable us to communicate. The 'Second Quantum Revolution' will see the introduction of technologies that exploit the subtler quantum effects and it is expected to have even greater impact.

Many regard the rush to develop quantum computing as a new 'space race'. Through counterintuitive phenomena such as quantum entanglement and quantum superposition, a large-scale quantum computer is expected to surpass the best digital super computer by orders of magnitude and allow us to simulate the complexities of the natural world as never before. Globally, progress over the past 5 years has been remarkable, and in October 2019 Google claimed to have achieved 'quantum supremacy' – that is, a programmable quantum computer able to solve a problem that classical computers practically cannot.

Within the UK during 2014, a coherent National Programme was planned and a first phase of five years Government support for translational R&D (exploiting decades of previous publicly-funded fundamental research in universities) to build a sovereign quantum technology industry sector was announced in the Autumn Statement. A further five years Government support for the National Quantum Technology Programme (NQTP), which spans quantum timing, sensing, communications and computing and simulation, was announced in the 2019 Autumn Statement including the creation of a National Quantum Computing Centre. Concurrently, the Defence Science and Technology Laboratory (Dstl) launched a DARPA-like project in 2014¹ to develop quantum technologies at pace which principally focussed on quantum sensors for position, navigation and timing (PNT) and comprised clocks, gravity sensors and inertial force sensors. The project complements work in the NQTP and Dstl and the UK Ministry of Defence (MOD) have been NQTP Partners since the National Programme's inception.

In 2014, Quantum Information Processing (QIP) was judged too immature for near term Defence and Security benefit and only technology watch activities (with an emphasis on quantum algorithms) have been carried on. However, the progress achieved both nationally and globally has exceeded early expectations and a technology watch stance is no longer appropriate for MOD. Accordingly, Dstl and MOD have surveyed the national and global positions in QIP and this landscape report is one of the products produced. The intended readership comprises senior policy and decision makers in Defence and Security.

This document critically surveys UK R&D in QIP hardware and software to inform what MOD investment should be considered to give Defence and Security benefit from QIP during the near term (Era 1, 2020 – 2025), medium term (Era 2, 2025 – 2030) and further into the future (Era 3, 2030 and beyond). Special attention is directed towards any disruptive benefits which are likely to result from early adoption of this emerging technology during Era 1.

¹ During Financial Year 2017-18, the project was renamed Quantum Sensing and currently has funding until 2022

As was done with the UK Quantum Technology Landscape document prepared by Dstl in 2013 (DSTL/PUB75620), this version is being released to UK colleagues and Stakeholders for comment, correction (if errors or misunderstandings are found in the document), or for additional input.

Section 5 - A strategy for UK Defence and Security capability in QIP is self contained; it may be read independently of the rest of the document although references to other parts of the text are included where more detail would help understanding.

After review, the document will be finalised and issued.

The list of **Key Points** below gives a synopsis of the sections and sub-sections throughout the text.

List of Key Points

	Section	Page	
KP1	1.1	1	This document critically surveys emerging Quantum Information Processing technologies in the UK and the rest of the world. In the approaching second ('Quantum 2.0') Information Age, QIP is expected to transform business functions and wider society, and could give disruptive early adopter benefits. The document identifies technologies and applications expected to appear during Era 1 (2020 – 2025) and recommends MOD re-consider its current technology watch-only stance on QIP
KP2	1.2	2	The UK Budget of 2018 announced government funding to develop a large-scale quantum computer, including new quantum algorithms, programming languages, techniques for error detection, applications and a QIP-ready workforce
KP3	2	6	Quantum computers are analogue machines which operate in fundamentally different ways to conventional analogue and digital computers. At least initially they will be hybrid quantum-digital machines
KP4	2.1	7	Quantum computers represent information using qubits which differ from bits used by classical computers by being able to represent not just the values 0 and 1 but also all possible intermediate numbers, including complex numbers, at the same time
KP5	2.2	8	There are different types of quantum computer which operate in very different ways and at different levels of commercial maturity. Defence and Security should also be aware that these compete with each other and with digital silicon solutions to offer the 'best' solution to specific problems
KP6	2.2.1	9	Circuit model quantum computers have similarities to conventional digital computers but have recently been demonstrated (by Google's Sycamore device) to have superior performance for some tasks
KP7	2.2.2	10	Adiabatic quantum computers (or quantum annealers used to solve optimisation problems) are large machines which have been in development for over a decade and are available commercially from D-Wave Systems
KP8	2.3	10	Benchmarking is essential to determine whether digital technologies may be superior to early quantum computers in the near-term

KP9	3	12	Quantum computers work by carefully manipulating qubit states using lasers and/or electromagnetic fields. Configured qubits form quantum gates and a set of quantum gates collectively implement a quantum algorithm
KP10	3.1	12	Future computer-based decision support systems (DSSs) will exploit 'Big Data' and QIP will help to manage the data deluge
KP11	3.2	13	Although QIP systems are crucial to meeting the computational demands of future decision support systems, existing (classical) algorithms must be recast to run on quantum computers
KP12	3.2.1	14	Out of the quantum algorithms found in the quantum algorithm zoo; many are of academic interest only and just a few are expected to be of value to Defence and Security or the wider economy
KP13	3.2.2	15	Five Noisy Intermediate Scale Quantum (NISQ) algorithms are expected to provide value to Defence and Security in Era 1
KP14	3.2.3	15	A further five NISQ algorithms may have value for Defence and Security during Era 1
KP15	3.2.4	16	Two quantum algorithms requiring large, fault-tolerant quantum computers may benefit Defence and Security during Era 3
KP16	3.3	16	Artificial intelligence is increasingly being used to deliver business functions
KP17	3.3.1	17	Automated data analysis using digital machines is becoming mature
KP18	3.3.2	18	Neural nets on digital machines have been developed for control systems, sensor processing, AI / machine learning, situational understanding and other applications
KP19	3.3.3	20	There is the potential for overwhelming quantum speedup by running neural nets on quantum computers ('quantum neural nets')
KP20	3.3.4	21	Training neural nets is not trivial and requires careful selection of training sets. However, this training data also can be used for quantum neural net calculations provided it is suitably modified to be input to the quantum computer
KP21	4	23	QIP can be applied across all Enterprise activities where IT is already in use
KP22	4.1	23	Over the next 15 years QIP will impact the sifting, compiling, extraction and presentation of information to human decision makers through faster AI execution
KP23	4.1.2	24	For Enterprise Management, the value of QIP is expected to be accurate, rapid pattern matching

			within 'Big Data' allowing high quality information to be extracted and used to facilitate timely, accurate decisions. Benchmarking, as QIP and digital platforms develop, will confirm or refute this
KP24	4.2	25	Quantum Computing is expected to impact secure data transmission over networks, authentication of both traffic and participant identity and to improve network immunity to attack
KP25	4.2.1	26	Civilian enterprises have the same requirements as Defence and Security, but on larger scale and early adoption for commercial applications is essential to attract the investment and R&D activity needed to mature the technology. Modern network enabled capabilities depend on network-orientated methods of communication and control to maximise the success of Defence and Security missions. Unimpeded, secure flow of information at high data rates is critical and can be achieved by a quantum network
KP26	4.2.2	27	'Quantum signatures' use quantum encryption methods to verify the identity of a message sender and are immune to attack by quantum computers
KP27	4.2.3	28	Recently invented biomimetic cyber defence systems protect networks in a way which is analogous to the immune systems of biological organisms. Quantum processing may help such systems identify threats
KP28	4.3	28	Quantum computers are only expected to provide quantum advantage in areas for which quantum algorithms are superior to conventional methods. One such area is image processing and is the subject of this Section
KP29	4.3.1	29	Quantum enabled automated searches for features in all the world's images would allow the recognition and tracking of events
KP30	4.3.2	29	Quantum image processing (QulmP) has seen much more research and development than has quantum signal processing
KP31	4.3.2.1	30	QulmP algorithms for circuit model machines are at low Technology Readiness Levels because of the immaturity of the necessary computing platforms
KP32	4.3.2.2	30	Artificial neural nets are a mature information processing architecture and are well suited to Noisy Intermediate-Scale Quantum (NISQ)

			computers, especially D-Wave machines. Programming (termed ‘training’) neural nets can be challenging and requires careful selection of training data. Neural nets running on NISQ machines are expected to have many applications in Defence and Security
KP33	4.3.2.3	31	US industry has successfully used D-Wave quantum annealers for machine learning and image analysis for over a decade
KP34	4.4	32	In the context of Information processing, management problems, in general, comprise the effective presentation of complex data in a comprehensible way. Machine intelligence is a promising solution to this problem
KP35	4.4.1	32	QIP will reduce data deluge and enable better understanding to be extracted from large data sets through identifying correlations which could not be found using classical tools. QIP’s adoption is not justified simply by massive speed-up; the impact of the speed-up will be the key driver
KP36	4.4.2	33	QIP best extract the fullest information from future quantum sensors
KP37	4.4.3	33	QIP using neural nets is likely to offer solutions in situational understanding in Era 1 via image analysis and pattern detection
KP38	4.4.4	34	Verification is critical for accurate situational understanding and may adopt methods similar to those pioneered by commercial organisations analysing social media
KP39	4.5	35	By rapidly and accurately recognising the component parts of its environment a quantum computer running neural nets should be able to navigate, calculate orientation, avoid obstructions and ‘understand’ a robot’s environment through machine vision. Compact, low power quantum computers will be needed and possible chips are the subject of R&D programmes
KP40	4.5.1	36	Quantum neural nets are expected to have transformational impact for autonomous vehicles by facilitating a step change in machine vision
KP41	4.5.2	36	Quantum computers are likely to accelerate the use of ‘intelligent’ systems controlling mechanical handling, storage and transport systems within individual machines and not just the Enterprise management network
KP42	4.5.3	37	Quantum neural nets will revolutionise future delivery of medical care for all communities

			around the globe, including in hazardous situations
KP43	4.5.4	38	Price is expected to be the principal constraint inhibiting the adoption of QIP in domestic systems. If the technical and ethical challenges can be overcome, self-driving vehicles would transform society
KP44	4.6	38	QIP could contribute to future combat systems through Network Quantum Enabled Capability (NQEC). There are challenges and issues which must be considered and resolved before the technology is available so that adoption will be as rapid as possible. The authors believe the principal technical challenges are machines' understanding of their environments, planning, and navigation. Other challenges include compatibility with military doctrine, health and safety concerns and regulations
KP45	4.7	40	Computer based education and learning has been increasingly utilised since the 1950s for reasons of effectiveness and cost. Virtual Reality and AI technologies have added realism to training simulators and have been enabled by developments in neural nets running on CPUs and GPUs. Quantum neural nets will empower improved Training and Simulation technologies
KP46	5	43	In the UK since 2014, government and other investment totalling about £1B has ensured the UK is world leading in the development of quantum technologies and aims to build a future sovereign quantum manufacturing sector. In QIP, the National Quantum Computing Centre (NQCC) will accelerate the development of low TRL R&D and produce prototype quantum hardware and software. Although the UK has a strong (conventional) computer software sector, which is expected to diversify into quantum software, it lacks a large computer systems integrator which may inhibit growing a QIP industry with full-stack capability. The recent Industrial Strategy Challenge Fund (ISCF) Wave 3 Quantum Technology Challenge, in part, seeks to rectify this situation but gaps remain in the NQTP QIP technology portfolio. Modest investment by MOD would address these gaps benefiting many of its business functions and providing disruptive advantage in some areas
KP 47	5.1	43	At the fundamental level, all information is quantum in nature and very different to the

			classical information processed by digital computers. Quantum physics clearly identifies the advantages of processing quantum information using a quantum processor including the ability to solve some problems much faster than digital computers. For many years, building such a quantum processor has been an elusive prize but functioning prototypes are evolving at increasing rates. Era 1 (2020 – 2025) offers the potential to identify early applications and will be a stepping-stone to fully scalable machines. Era 1 is a critical time for business entities to carefully consider QIP investment strategies
KP 48	5.1.1	44	QIP capabilities represent significant opportunities and threats, especially for Defence and Security, and these are sufficiently significant and novel that organisations need to explore applications now to be ‘quantum-ready’ for the future. It is expected to take years to build capabilities and identify useful applications and it will be difficult for organisations that have not engaged early on to catch-up
KP 49	5.1.2	45	QIP is in the early stages of development and the dominant hardware platform is still not clear. State actors and companies are investing heavily to attempt to ensure early advantage. As with current digital technology, algorithms critically important and some can be executed on the NISQ machines expected to be available during Era providing a window of opportunity to accelerate progress and shape developing markets
KP 50	5.2	46	The UK NQTP is currently a diverse ecosystem of funded R&D, supported technology development in industry and other initiatives including the development of a National Quantum Computing Centre. This has created world class capabilities in QIP and determined efforts are being made to establish a sovereign, full-stack capability
KP 51	5.3	48	Mid TRL, translational QIP research is supported by the Oxford-led Quantum Computing and Simulation (QCS) Hub
KP 52	5.4	48	IUK is supporting commercialisation of QIP through Wave 3 of the Industrial Strategy Challenge Fund and the Department for Business, Energy and Industrial Strategy is leading a programme to establish a National Quantum Computing Centre which will accelerate translation of QCS Hub R&D into commercialisable technology

KP 53	5.4.1	49	In 2020 IUK has made 10 grant awards worth £25.7M for QIP projects under the first Wave 3 ISCF Quantum Technologies Challenge call which address technologies across the full quantum computing stack
KP 54	5.4.2	50	Hardware projects span the leading platforms and address key challenges including systems engineering and scalability. Software projects address qubit control, operating systems (including for hybrid digital / quantum machines) and application software
KP 55	5.5	51	ISCF Wave 3 projects funded during 2020 span a broad range of QIP technologies not including Quantum Neural Nets (QNNs). These have been intensively studied and could benefit all business enterprises, especially Defence and Security. If action is not taken now, it is possible that the UK may be left behind in this important area. The second tranche of ISCF Wave 3 funding, expected in 2021, could support the development of QNNs for machine learning and managing complex systems
KP 56	5.5.1	53	An exemplar ISCF Wave 3 Challenge in QNNs could be the development of a NISQ algorithm to identify, localise and track arbitrary features in imagery data
KP 57	6	54	NISQ computers are available now and businesses should begin assessing the opportunities and threats expected from large scale machines expected to appear within the decade. Broadly the NQTP spans all QIP technologies but has no work on quantum neural nets (QNNs) which, on existing commercial machines, could create near-term 'early wins'
KP 58	7	55	MOD should work with NQTP Partners to formulate and propose a QNN Challenge to be supported by the tranche of ISCF Wave 3 funds expected to be released during 2021. MOD should also ensure it has adequate SQEP to derive full early-adopter advantage from the technologies developed through the QNN Challenge

Table of Contents

Release Conditions

Executive summary	i
1 Introduction	1
1.1 Purpose of the document	1
1.2 Background.....	2
1.3 Structure of the document	5
2 Classical and quantum computers	6
2.1 Quantum information.....	7
2.2 Types of quantum computer.....	8
2.2.1 Circuit model quantum computers	9
2.2.2 Adiabatic quantum computers	10
2.2.3 Competitive alternatives to near term quantum technologies	10
3 Quantum Information processing in Defence and Security	12
3.1 Information Processing in Defence and Security	12
3.2 Quantum algorithms.....	13
3.2.1 The Quantum Algorithm Zoo	14
3.2.2 NISQ algorithms of primary importance in Era 1	15
3.2.3 NISQ algorithms of secondary importance in Era 1	15
3.2.4 Quantum algorithms for Era 3	16
3.3 QIP for Automation.....	16
3.3.1 Automated Data Analysis - The Automated Statistician	17
3.3.2 Neural nets	18
3.3.3 Quantum neural nets	20
3.3.4 Training neural nets	21
4 Distributed Enterprise Management IT	23
4.1 Civil and Military Enterprise Management	23
4.1.1 System Characteristics	23
4.1.2 Military Enterprise Management (C4ISR)	24
4.2 Information Security	25
4.2.1 Data Networks	26
4.2.2 Authentication	27
4.2.3 'Auto Immune' network defence	28
4.3 Data and Information Processing	28
4.3.1 Automated image searching	29
4.3.2 Image data processing	29
4.3.2.1 QulmP using circuit model quantum computers	30
4.3.2.2 QulmP using quantum neural nets	31
4.3.2.3 QulmP algorithms for D-Wave, an adiabatic quantum computer	32

4.4	Management Aids (AI)	32
4.4.1	Quantum enabled understanding	33
4.4.2	Situational Awareness	33
4.4.3	Situational understanding from situational awareness	34
4.4.4	Verifying accuracy	35
4.5	Robotics and The Enterprise	35
4.5.1	Transport Automation	37
4.5.2	Logistical Systems	37
4.5.3	Medical Systems	38
4.5.4	Domestic Systems	38
4.6	Future combat systems	38
4.7	Training & Simulation.....	40
5	A strategy for UK Defence and Security capability in QIP	43
5.1	Quantum Information Processing	43
5.1.1	Opportunities and threats – the case for Government investment	44
5.1.2	The global technology race	45
5.2	A commercial sovereign computing capability	46
5.3	Targeted UK translational quantum computing research	48
5.4	Quantum computing sovereign capability development	48
5.4.1	ISCF Wave 3 QIP projects	49
5.4.2	The current ISCF Wave 3 quantum computing portfolio	50
5.5	A QIP Strategy for Defence and Security	51
5.5.1	An example NISQ-value challenge:	53
6	Summary and conclusions	54
7	Recommendations	56
8	Acknowledgements	57
9	List of abbreviations	58
APPENDIX A	An introduction to quantum computers	61
A.1	Data representation within classical and quantum computers.....	61
A.2	Quantum computing paradigms.....	62
A.2.1	Figures of merit: quantum volume	63
A.2.2	An overview of circuit model quantum computers	63
A.2.3	An overview of adiabatic quantum computers	64
A.2.4	An overview of measurement based quantum computers (MBQCs)	65
A.2.5	An overview of topological quantum computers	66
A.2.6	An overview of emulators of quantum processors	67
APPENDIX B	Quantum computing hardware	68
B.1	Introduction.....	68
B.2	IBM.....	68
B.3	Google.....	70

B.4	Intel	70
B.5	Microsoft	71
B.6	Honeywell	71
B.7	D-Wave	72
B.8	Competing digital silicon technologies	77
B.8.1	Hardware and Software	77
B.8.2	Quantum and Digital Hardware Competition	78
B.8.3	Low-cost and Highly Parallel non-Quantum Computers	79
B.8.3.1	Cloud Computers	79
B.8.3.2	GPU (Graphical Processing Unit) based Solutions	79
B.8.3.3	Bespoke neural net chips	80
B.8.3.4	Niche optimisation solutions	81
B.8.3.5	Competition through Software	82
B.9	Photonic quantum computers	82
APPENDIX C	Quantum algorithms of primary importance in Era 1	83
C.1	Introduction	83
C.2	Shor's Algorithm:	83
C.3	Grover's algorithm:	85
C.4	Quantum Fourier transform:	86
C.5	Quantum annealing	87
C.6	Machine intelligence and learning:	89
APPENDIX D	Quantum algorithms of secondary importance in Era 1	92
D.1	Introduction	92
D.2	Variational quantum eigensolver, VQE	92
D.3	Quantum approximate optimisation algorithm, QAOA	93
D.4	Data-driven quantum circuit learning, DDQCL	93
D.5	Quantum Auto-Encoder, QAE	94
D.6	Population transfer, PT	95
APPENDIX E	Quantum algorithms for Era 3	97
E.1	Introduction	97
E.2	Deutsch's algorithm	97
E.3	Quantum Simulation:	97
E.4	Solution of linear equations:	99
APPENDIX F	Boltzmann, Hopfield, Hamming and Neural Networks	100
APPENDIX G	QulmP algorithms for circuit model quantum computers	102
APPENDIX H	TensorFlow Quantum (TFQ)	104
APPENDIX I	UK and other National quantum technology programmes	105
I.1	Australia	105

I.2	Canada.....	106
I.3	China.....	107
I.4	European Union.....	109
I.5	Japan	110
I.6	Russia	112
I.7	United Kingdom.....	113
I.8	United States.....	115
I.9	Summary: UK prospects.....	115
APPENDIX J Industry, Supply Chains, UK Capability		117
J.1	Introduction.....	117
J.2	End-to-end capability.....	117
J.2.1	D-Wave	117
J.2.2	Google	117
J.2.3	IBM	118
J.2.4	Intel	119
J.2.5	Microsoft	120
J.2.6	Oxford Quantum Circuits	121
J.3	Quantum software for the systems and application layers	122
J.3.1	ATOS	122
J.3.2	Cambridge Quantum Computing	122
J.3.3	Riverlane	123
J.4	Commercial providers of quantum hardware	123
Initial distribution		124
Report documentation page v5.0		125

1 Introduction

1.1 Purpose of the document

Key Point 1: This document critically surveys emerging Quantum Information Processing technologies in the UK and the rest of the world. In the approaching second ('Quantum 2.0') Information Age, QIP is expected to transform business functions and wider society, and could give disruptive early adopter benefits. The document identifies technologies and applications expected to appear during Era 1 (2020 – 2025) and recommends MOD re-consider its current technology watch-only stance on QIP

In 2014, with funding from the UK Ministry of Defence (MOD) Chief Scientific Adviser's R&D programme, Dstl began a DARPA like programme to develop quantum sensor demonstrators including quantum clocks and inertial force sensors. The programme was based on a prior study of UK quantum capabilities summarised in the UK National Quantum Technology Landscape.² The Landscape identified clocks, communications, sensing and computing as broad quantum technology areas which would begin to come to market in roughly that order and also formed the evidence base for the National Quantum Technology programme (NQTP); the MOD and NQTP programmes were conceived to be mutually complementary. (The Landscape document was refreshed in 2016,³ to serve as an evidence base for the rebid of the NQTP as a Phase 2 programme, which followed seamlessly from Phase 1 on 1st December 2019.)

An essential part of the Dstl strategy was to focus narrowly on sensor demonstrators while keeping a close watch on developments in quantum communications and computing both in the UK and overseas. This was because the underpinning hardware, software, and algorithm maturity of quantum computing and communications meant quantum information processing (QIP) or quantum information science⁴ systems were too immature to justify MOD investment. However, over the past few years, significant advances across all three of these disciplines have been achieved and it is now sensible to review MOD's position regarding investment in QIP technologies.

The principal driver of the NQTP was to position the UK at the forefront of new, emerging, 'Quantum 2.0' technologies⁵ and create a sovereign UK quantum industry from which ongoing economic and national security benefit could be gained. Of the four identified technology areas (clocks, communications, sensing and computing), QIP will have by far the greatest impact, both economically and on national security. This conclusion follows by analogy with the information revolution which has followed the invention, development and application of very large-scale integrated silicon circuits. The International Data Corporation⁶ estimated that the global information technology industry is expected to reach or exceed \$5 trillion in 2019 (growing at about 4% according to CompTIA⁷) with the US contributing about 31% of this figure. Within this global spend, devices contribute ~22% and emerging technology 17% with software

² See <https://www.epsrc.ac.uk/newsevents/pubs/Dstl-uk-quantum-technology-landscape-2014/>

³ See <http://uknqt.epsrc.ac.uk/files/ukquantumtechnologylandscape2016/>

⁴ https://en.wikipedia.org/wiki/Quantum_information_science

⁵ The 'Quantum 2.0' label implies that the subtler aspects of quantum physics, such as superposition and entanglement are exploited unlike 'Quantum 1.0' technologies which largely exploit the quantisation (of energies, momenta, etc.)

⁶ <https://www.idc.com/>

⁷ <https://www.comptia.org/resources/it-industry-trends-analysis>

and services comprising the remainder. Thus, even a national QIP industry focusing only on software development and services will be addressing a global market worth billions of dollars annually. Particularly noteworthy in the coming second ('Quantum 2.0') Information Age will be cybersecurity which critically underpins modern society.

1.2 Background

Key Point 2: The UK Budget of 2018 announced government funding to develop a large-scale quantum computer, including new quantum algorithms, programming languages, techniques for error detection, applications and a QIP-ready workforce

Continuous improvements in information technology have driven economic and societal changes since the emergence of "high speed" digital computers in the 1960s which made use of integrated electronic circuits ("chips") manufactured 'en masse'. Early chips used bipolar transistors but from the 1970s Metal-Oxide-Semiconductor (MOS) technology began to replace bipolar designs, culminating in Complementary-Metal-Oxide-Semiconductor (CMOS) technology which currently dominates large scale digital chip design. Ever higher component densities on the chips led to an exponential growth in computing power (first noted by Gordon Moore in 1965 and, inevitably, called Moore's Law) which began to slow in about 2015.

The anticipated slowdown gave impetus to research in quantum computing which had begun as a research area after Richard Feynman, in 1981, had questioned whether a classical, universal computer could simulate any physical system, especially a quantum system. He concluded the answer was 'No' because the memory requirements rise exponentially quickly as the system size increases. However, the Hilbert spaces⁸ spanned by quantum systems have the same size dependency leading Feynman to propose that 'quantum computers' could simulate quantum systems.

When realised, quantum computing will be a fundamentally different approach to information processing compared to classical computing. This is because of the very different ways information is represented in quantum and classical systems. A simple overview may be found in **A.1** in **Appendix A**. For solving certain problems, it is expected to be superior, and identifying the problems for which quantum computers will possess a 'quantum advantage' is the subject of intense research. Some cases where quantum advantage exists are well known and algorithms have been developed. They include searching large datasets (Grover's algorithm where there is a square root quantum speed-up) and factoring large numbers (Shor's algorithm which has an exponential quantum speed-up). It is widely believed that algorithms to simulate chemical and biological systems will be identified which will also show quantum advantage raising hopes for step change improvements in [big] data analytics, Artificial Intelligence, automation, drug development, materials design and the design of novel biological processes (e.g. artificial photosynthesis).

⁸ A Hilbert space, named after David Hilbert who was a mathematician active during the late 19th and early 20th centuries, generalizes the notion of ordinary (Euclidean) space. It extends the methods of vector algebra and calculus from the two-dimensional Euclidean plane (in which unique points are identified by two co-ordinates, x and y, say) and three-dimensional (x, y, z) space to spaces with any number of dimensions

The NQTP, and other work globally, has created rapid progress since 2015 in the realisation of various hardware platforms for producing and manipulating quantum information (qubits). The NQTP has allowed the UK to build on the many previous decades of investment in the underlying quantum physics and this is reflected in the UK's leading position developing ion trap and, emergently, superconducting platforms. Aspirations within the second phase of the NQTP are to demonstrate Noisy Intermediate Scale Quantum (NISQ) computers comprising 50 – 100 qubits within the next five years and large-scale quantum computers following some 5 – 10 years later. However, in addition to the Phase 2 Hub investment, ISCF Wave 3⁹ calls from 2019 onwards include quantum computing and the National Quantum Computing Centre (NQCC, announced in the Budget of 2018), will focus on engineering a large-scale quantum computer. In parallel, work to understand the capabilities of quantum computers will accelerate and include research to develop new quantum algorithms, programming languages, techniques for error detection and correction (or error avoidance if topological quantum computers are realised), applications and a QIP-ready workforce.

Close-coupled hardware and software development will benefit from the rapidly increasing interest from and involvement of industry. Real world problems and user needs will stimulate progress and it is not unreasonable to expect a similar trajectory of the development of the second Information Revolution to that seen during the first. The UK economy, society and national security will all benefit from this work.

Early expectations were that applications including machine learning, artificial intelligence (AI) and big data would be the first to show quantum advantage. This is proving to be correct¹⁰ and commercial providers from large players (such as Rigetti) to SMEs, start-ups¹¹ and incubators (such as Creative Destruction Lab) are beginning to provide access to quantum machine learning (QML). However, a revolution in AI is expected to be some time away. Nonetheless, there may now be genuine opportunities to introduce QIP into the toolset used by the military to augment aspects of defence procurement, training and operations and making possible a sea change in situational awareness:

⁹ Industrial Strategy Challenge Funded technology development was launched by the White Paper 'Industrial Strategy: Building a Britain fit for the future' first published in November 2017 and subsequently revised, see <https://www.gov.uk/government/publications/industrial-strategy-building-a-britain-fit-for-the-future>

¹⁰ "Quantum Computing: A Research Project or a Practical Computing Architecture for Machine Learning?" M Brisse and C E Sapp, Gartner Report 3791363 (2017), <https://www.gartner.com/en/documents/3791363>; "Top 10 Strategic Technology Trends for 2019: Quantum Computing" D Cearley, B Burke and M Horvath, Gartner Report 3904279 (2019), <https://www.gartner.com/en/documents/3904279>

¹¹ Spin-outs and start-ups are names used commonly in the media, sometimes interchangeably. They do have crucial differences, however. A start-up is a business, usually carried on through a limited company, that has recently begun operating. They are owned by their founders who provided the business ideas, know-how and intellectual property required to start the business. Spin-outs, in contrast, are not solely owned by their founders but also have minority shareholders which are often the "parent" universities or other higher educational institution employing at least some of the spin-out's staff. Spin-outs involve the parent moving some of its assets (often intellectual property) into the new spin-out company which is then run as a separate trading entity.

Whatever the application, future quantum information processing ‘could lead to an understanding of what’s actually happening, as opposed to an approximation of what might be happening.’¹²

The underlying physics of QIP hardware is unfamiliar to many and, for brevity, this document assumes the reader has some familiarity with the basic ideas. (The 2016 UK Quantum technology Landscape² provides background to underpinning theory and recent developments in quantum technology.)

The technologies under development for the practical realisation of quantum information processing can be imagined in a layer structure, see figure 1. The hardware layer comprises the physical devices, equivalent to transistors, resistors, capacitors etc. in a classical computer, which manipulate light or electric charge during a computation in order to operate on the quantum information encoded in the light or charge. However, this information is fragile so an error correction layer is essential to ensure ‘correct’ answers are obtained.

The highest layer, the application layer, comprises the algorithms and software through which a problem is entered into the machine, solved and the solution output to the user.

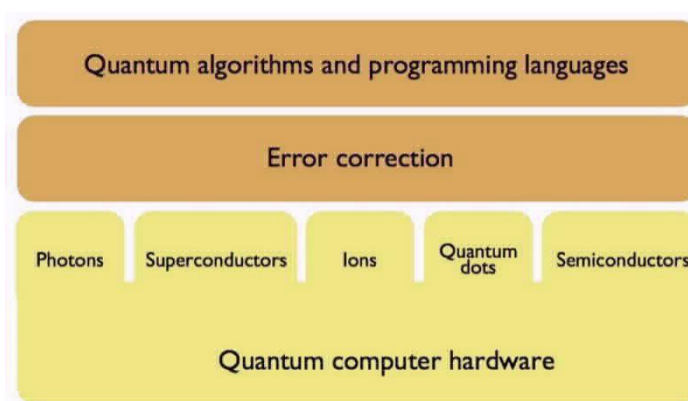


Figure 1: Technology layers for a quantum information processing machine. The application layer (algorithms and software) is the quantum digital interface

In the early days of classical computing, computer scientists were routinely concerned with the physical characteristics of the electronic circuitry but the huge improvements in hardware and software reliability and sophistication have been such that these fundamental computer science considerations are almost never a concern to today’s programmers.

This is not the case for quantum computing; algorithms cannot yet be considered in isolation from the hardware although compilers are beginning to appear and abstract some aspects of QIP away from the hardware. This process will continue and accelerate, but for the near to medium term, using a quantum information processor will remain an activity for which a considerable degree of expert knowledge is required. Thus, adoption of QIP requires a suitably skilled workforce; such individuals will be in

¹² Vijay Swarup, ExxonMobil VP for R&D, p.12, 2018 IBM Annual Report, IBM.com

high demand initially and early preparation, through targeted recruitment, will be necessary if the full benefits of quantum techniques are to be enjoyed.

Just as computer science involves the creation of algorithms which can be modelled as Turing machines¹³ and run on a classical computer architecture, so quantum computer science involves the creation of quantum algorithms which can be modelled as quantum Turing machines and run on one of the quantum computer architectures. A quantum computer manipulates quantum representations of data to solve problems and uses quantum physics to do so in a way which is far superior to the manipulation of bits according to the laws of classical physics, allowing a quantum machine to solve problems beyond the capabilities of classical machines. The problem is that building such machines is a huge engineering challenge.

1.3 Structure of the document

This document will briefly summarise the principal areas of Defence and Security businesses in which the authors judge that the complex, emerging discipline of QIP has the potential, within 10 years, to have transformative effects. In some niche areas of business, disruptive benefit may be possible as soon as 2025 by the early adoption of current commercial QIP systems.

Section 2.1 introduces the concept of quantum information while **Sections 2.2.1** and **2.2.2** give a brief description of circuit model and adiabatic quantum computers which are different – and currently the best - realisations of quantum Turing machines able to manipulate this information by exploiting quantum physics to achieve enhanced performance compared to classical computing. **(Appendix B** briefly summarises the quantum computers being developed for commercial purposes by leading IT companies well as competing digital silicon technologies.)

Section 3 addresses Quantum Information Processing (QIP) and introduces the concept of a quantum algorithm. A simple introduction to the implementation of quantum algorithms on a quantum computer is given in **Section 2.2** (more detail is available in **Appendices C, D and E**).

Section 4 considers how QIP may benefit Defence and Security in the future in the areas of Situational Awareness and Survivability, Communications, Command and Control systems, Logistical, Medical and Combat Robotics, and Training and Simulation.

¹³ Turing machines were invented by Turing in 1936 and, although conceptually simple, can simulate computer algorithms of arbitrary complexity. At its simplest, a Turing machine comprises an infinitely-long paper tape which stores information at fixed positions. If the storage positions hold the symbols 0, 1 and ' ' (blank) the device is called a 3-symbol Turing machine. The machine has a read/write head which passes over each storage position and which can (i) read the symbol under the head, (ii) edit the symbol (by writing a new symbol) and (iii) move the tape left or right by one storage position.

2 Classical and quantum computers

Key Point 3: Quantum computers are analogue machines which operate in fundamentally different ways to conventional analogue and digital computers. At least initially they will be hybrid quantum-digital machines

Quantum computers, or Quantum Information Processors, are a new form of analogue computer, similar in concept to the early computers adopted successfully by the military in the first half of the 20th Century. The massive battleship gun platforms which then dominated naval warfare had something which for many decades was vastly faster and more effective than the early electronic digital computers developed by Turing and others at Bletchley Park¹⁴ during World War II. They used analogue computers comprising complex mechanisms of gears and cams that computed shell trajectories (given wind, platform and target movement) in order to set bearing and elevation for the guns.

Analogue computers used many different physical mechanisms including hydraulics, pneumatics and later analogue electronics. They dominated computing, even addressing areas such as national economic modelling. Analogue models mirrored the complex mathematics of a ‘problem’ and ‘solved’ it to give answers on dials, meters or even directly to gun turret servos.

Analogue computers were superseded from about 1960 onwards as transistor-based digital computers were developed.¹⁵ Like quantum computers, these semiconductor devices exploit quantum constraints but the physics is simpler; in transistors the allowed energies of the electrons which carry charge are limited to fixed discrete values (‘Quantum 1.0’ devices) while in quantum computers (‘Quantum 2.0’ devices) more complicated phenomena (such as superposition¹⁶ and entanglement¹⁷) are exploited. The astonishingly rapid evolution of digital computers was driven initially by military investment but later commercially. Digital electronic computers quickly displaced analogue computers not because of superior computational speed (although that happened later) but because of greatly reduced programming complexity and cost.

¹⁴ The UK code-breaking Establishment during the Second World War.

¹⁵ The Turing-Welchman ‘Bombe’ machine (later called Agnes) which broke the German ‘Enigma’ codes from 1940 was electromechanical. At the time, the only electronic technology available was based on thermionic valves. Owing to their short operating lives valves were poorly suited to making digital computers but worked well enough to establish concepts that remain the foundations of modern digital computing, see www.tnmoc.org.

¹⁶ Superposition arises from the ability of quantum particles to exist in different states or places at the same time and explains the ‘double slit’ experiment in which a beam of identical particles incident on adjacent, narrow, slits create an interference pattern demonstrating that light and matter have both particle- and wave-like properties. In practice, superpositions cannot be observed, only the consequences of their existence. Superposition, from wave-particle duality, is essential to imaging technologies such as electron microscopy and underpins the operation of many sensors including atomic clocks. Energy absorbed from sunlight at “antenna” sites in chlorophyll is rapidly transported via a superposition of energy-transmitting pathways to different sites (reaction centres) to initiate photosynthesis.

¹⁷ Entanglement is a property of a composite system of quantum particles which forbids its state being accurately described in terms only of the states of the component particles; the quantum state of the system as a whole is in a definite state although the parts of the system are not. Entanglement is key to potentially unbreakable quantum key distribution (QKD), extended networks of clocks that would measure time independently of location, and imaging to resolutions below the diffraction limit.

Fundamentally, analogue computers are extremely slow and costly to program; digital computers transformed that situation through software which flexibly translated ‘problems’ into the native ‘language’ of the digital computer.

Quantum computers will be used in conjunction with digital computers; ‘real’ systems will be quantum plus classical hybrids, similar to the architectures which appeared during the 1980s in which floating point operations were carried out on a dedicated co-processor chip. Quantum computers will be co-processors, dealing with particular mathematical problems better-suited to quantum solution, tasked and returning answers via digital computers.

2.1 Quantum information

Key Point 4: Quantum computers represent information using qubits which differ from bits used by classical computers by being able to represent not just the values 0 and 1 but also all possible intermediate numbers, including complex numbers, at the same time

In a classical digital computer, a Turing Machine, data is stored in memory using ‘bits’ - binary digits - taking the values 0 or 1. A state of the memory is represented as a fixed-length string of bits (current hardware commonly uses 64 bits). Physically, bits correspond to the voltages of transistors in their ‘on’ and ‘off’ states. All information – numbers, text, sound or images – is stored by a collection of bits of suitably large size. The choice for the length of strings determines the precision of arithmetic, resolution of an image etc. In a similar way, a quantum computer represents data using ‘qubits’ – quantum bits – but these differ from bits by being able to represent not just the values 0 and 1 but also all possible intermediate numbers, including complex numbers, *at the same time*. Physically, qubits could be represented by the spin-up and spin-down states of an electron in a magnetic field or the orthogonal polarisations of a single photon (horizontal- and vertical- or left- and right- polarised). Information – again, numbers, text, sound or images – is represented by a collection of qubits. **Section A.1 in Appendix A** gives more detail.

The laws of quantum physics, which allow qubits to be simultaneously in two states called a superposition, is the property which makes quantum computers more powerful than their classical equivalents. However, if a qubit is measured (‘read out’) only a single value (‘0’ or ‘1’) is obtained¹⁸ and the quantum-nature of the data is lost. Quantum mechanically, a measurement is made when the qubit interacts with the measuring device. This is a very general concept; even the influence of the qubit environment (for instance electro-magnetic fields or the jostling of atoms resulting from their thermal motion) constitutes a measurement. When the measurement occurs, the quantum state collapses and the qubit is said to decohere. Thus, qubits are very fragile and must be isolated from all outside influences if they are to exist for sufficiently long to manipulate them successfully during a computation. Qubits with long coherence times preserve quantum information well and are said to have high fidelity.

The qubits comprising a quantum computer must also be entangled (see footnote 16) as well as in superposition states. This allows multiple quantum states to be acted on simultaneously during operation of the quantum machine. Classical bits in a classical machine, however, can only have one value at a time. Even in a ‘parallel processing’

¹⁸ The superposition is said to collapse

computer, where a computation is split into parts which are executed simultaneously on different processors attached to the same computer, the computation is not truly 'parallel'. Thus, entanglement is an essential resource in a quantum computer and is the origin of the superior performance of a quantum compared to a classical machine. Additionally, entanglement allows 'superdense coding' which enables a single qubit to represent two bits of classical information.

2.2 Types of quantum computer

Key Point 5: There are different types of quantum computer which operate in very different ways and at different levels of commercial maturity. Defence and Security should also be aware that these compete with each other and with digital silicon solutions to offer the 'best' solution to specific problems

There are numerous paradigms¹⁹ of quantum computing, some of which are more easily understood and implemented than others. For brevity, only the two best developed types of quantum computer will be considered here.

The operating principles of **circuit model quantum computers** derive from their classical analogues in which switches and relays implement Boolean logic and arithmetic. As discussed above, a classical machine stores information as bits and these are subjected to a sequence of operations that may be reduced to sets of interconnected one- and two-bit operations ('gates'). In an analogous way, the qubits comprising a quantum machine are subjected to a series of precisely timed interactions producing what are called "quantum gates" which may be broken down into primitives of one or two qubit operations. The machine must have a structure that permits the required interactions and measurements, which must be made at the correct points in the calculation. The performance of a circuit model quantum computer is quantified by its quantum volume (introduced by IBM²⁰) which reflects the number of qubits comprising the machine and their connectivities as well as other things such as errors which occur during individual gate operations and the gate-level implementation of the algorithm being run (software).

Adiabatic quantum computers require that the system of qubits is prepared in a state that can be represented by a function²¹ and is commonly thought of as an energy surface. The surface is then slowly, and adiabatically,²² distorted into the shape that represents the problem; the lowest point on the final surface corresponds to the state of the system equivalent to the 'answer'. The best-known examples of this type of quantum computer are the series of machines built by D-Wave Systems.

Although D-Waves are slowly becoming accepted, the engineering goal is the fabrication of scalable, fault tolerant, circuit-model quantum computers able to run any quantum algorithm. Ideal, perfect qubits are unlikely ever to be achieved and so practical architectures will comprise large numbers of physical qubits working synergistically as a smaller number of near-perfect (logical) qubits. Theoretical

¹⁹ The Oxford English Dictionary variously defines a paradigm as a pattern or model; an exemplar; a typical instance; an example. A given paradigm can be built using different qubit types

²⁰ <https://arxiv.org/abs/1811.12926>

²¹ A mapping from a defined space into the real numbers

²² A process which happens without transfer of heat or mass of substances between a thermodynamic system and its surroundings.

estimates of the numbers of physical qubits required for these error correction schemes range upwards from 100s or 1000s. Within the circuit model paradigm, academic R&D teams broadly favour trapped-ion qubits while industry prefers superconducting qubits.

Over the next 5 years, engineering complexities are expected to restrict trapped-ion based machines to small numbers (~50) of fully connected, very high fidelity qubits which operate at ambient temperatures but need ultra-high vacuums; gate operations take typically ~10 μ seconds. Superconducting qubit chips are CMOS compatible (and hence intrinsically scalable) but architectures with extensive connectivities (beyond nearest neighbour connections) are difficult to engineer; complications arise from the need for cryogenic cooling but gate operations are fast (typically ~10 nanoseconds).

Prototype circuit model machines have begun to appear. D-Wave adiabatic machines with regularly increasing qubit numbers, and recently greater qubit connectivity, have been available for over a decade. Experimentation will establish which, if either, architecture is preferable for a general purpose machine or whether the technologies are better suited to specific applications. These early machines have been called Noisy Intermediate Scale Quantum (NISQ) computers reflecting the limited numbers of partially connected qubits which are relatively poorly isolated from background noise which causes decoherence. Low Noise Intermediate Scale Quantum (LNISQ) machines are a near term (0 – 5 year) technology target.

2.2.1 Circuit model quantum computers

Key Point 6: Circuit model quantum computers have similarities to conventional digital computers but have recently been demonstrated (by Google’s Sycamore device) to have superior performance for some tasks

After many years when little progress (measured in terms of number of connected, controllable qubits) was slow, the past five years has seen rapid progress with IBM and Google the frontrunners in a race to demonstrate a quantum computer able to outperform a conventional machine (‘quantum supremacy’). Although quickly challenged by IBM, Google claimed in October 2019 to have demonstrated quantum supremacy with a 53 qubit device²³ (and is optimising a 72 qubit machine).

Circuit model quantum computers are the most often used to explore the potential of quantum computing. Their operating principles may or may not ultimately be seen as the most appropriate paradigm as the ideas are derived from considering the operation of classical digital computers. These in turn are derived from the paradigm of Boolean logic, arithmetic, and switches and relays. Those are familiar, but do not necessarily sit well with the behaviour of quantum objects.

In a classical machine, the information is stored as bits of information that are subject to a sequence of operations that may be reduced (ultimately) to a complicated set of interconnected one- and two-bit operations. A quantum circuit model takes an array of qubits and subjects them to a series of precisely timed interactions via an arrangement of what are known as ‘quantum gates’. These may be broken down into primitives of

²³ <https://www.nature.com/articles/d41586-019-03213-z>. Google’s ‘Sycamore’ quantum computer took 200 seconds to solve a problem which Google estimated would take an IBM Summit (~1 million core supercomputer) 10, 000 years but IBM subsequently challenged this and claimed a different algorithm would require only 2.5 days to find a solution. See <https://www.sciencemag.org/news/2019/10/ibm-casts-doubt-googles-claims-quantum-supremacy>.

one or two qubit operations.²⁴ The machine needs to be set up with a structure that provides the appropriate interactions and measurements at the right point in the calculation. That arrangement can be very complicated. The operation is typically represented by a diagram such as that in figure 2 where the progress of the qubits $|0\rangle$ and $|1\rangle$ is from left to right and the operations defined by the boxes.




Figure 2: Representation of a typical quantum 'circuit' for three qubits

2.2.2 Adiabatic quantum computers

Key Point 7: Adiabatic quantum computers (or quantum annealers used to solve optimisation problems) are large machines which have been in development for over a decade and are available commercially from D-Wave Systems

The best-known examples of this class of machine are produced commercially by D-Wave Systems and use superconducting qubits. Over more than a decade, machines and software have been developed capable of solving problems from many areas including logistics, artificial intelligence/machine learning, materials science, drug discovery, cyber security, fault detection and financial modelling.

In October 2018, the D-Wave 2000Q was made available to the public through D-Wave's cloud service and in February 2019 the Pegasus chip was announced and claimed to be 'the world's most connected commercial quantum system' with a maximum of 15 connections per qubit, > 5, 000 low-noise qubits and will be available commercially in mid-2020. D-Wave are believed to be on-track to engineering a 10, 000 qubit machine by 2023.

2.2.3 Competitive alternatives to near term quantum technologies

Key Point 8: Benchmarking is essential to determine whether digital technologies may be superior to early quantum computers in the near-term

Emerging digital silicon technologies may provide competitive, or better, solutions to some computing problems than quantum approaches can in the NISQ era. Relevant technologies which should be investigated and bench-marked against NISQ machines ranging from software running on classical supercomputers, which mimics perfect quantum computers, to custom designed digital silicon hardware, inspired by quantum principles, which solve very narrow classes of problem in computationally efficient ways.

Software approaches comprise:

²⁴ It is beyond the scope of this paper to provide a deeper explanation, however see http://en.wikipedia.org/wiki/Quantum_computer for more information

- ‘Ad hoc’ methods, for instance for feature recognition in images (examples being automatic number plate recognition and facial recognition);
- Quantum computer emulators,²⁵ however, practical considerations severely limit the circuit sizes which can be emulated.²⁶

Competition from conventional digital silicon machines comprises a number of different approaches:

- Large numbers of processors arranged in highly parallel architectures including cloud computers;
- Reduced Instruction Set Computers (RISC machines) including Graphical Processing Units (GPUs);
- Bespoke neural net chips.
- ‘Niche’ devices such as the Digital Annealer recently announced by Fujitsu.

Appendix B gives more detail about types of quantum computer and reviews their near-term, digital silicon competition.

²⁵ These are valuable in their own right since they provide a way to use ‘perfect’ quantum computers ‘today’ (and hence develop skills) and also develop understanding of practical issues (such as qubit connectivity or ‘noise’-induced decoherence of qubit states) which affect the performance of real machines. This will give an understanding of the expected impact of near term NISQ machines and, potentially, how different engineering approaches might improve their performance.

²⁶The memory and execution time requirements increase exponentially with the problem size; thus, the memory required to simulate circuits comprising 26, 29 and 45 qubits is 2 GB, 16 GB and 1, 048, 576 GB.

3 Quantum Information processing in Defence and Security

Key Point 9: Quantum computers work by carefully manipulating qubit states using lasers and/or electromagnetic fields. Configured qubits form quantum gates and a set of quantum gates collectively implement a quantum algorithm

From the 1970s onward, systems of quantum particles (electrons, nucleons, atoms, molecules, photons, etc.) began to be viewed not as phenomena found in nature requiring explanation but as systems which could be designed, engineered and exploited. These systems can only properly be described using quantum theory and the (configuration-, spin-, momentum-, ...-space) parameters needed to fully specify the system is called 'quantum information'. A very simple example is a system comprising an isolated electron in a magnetic field for which knowledge of the magnetic spin quantum number, $m_s = \pm 1/2$ (i.e., 'spin-up' or 'spin-down'), is sufficient to specify the state of the system for most purposes. This quantum information can be manipulated (in this simple example, by reversing the magnetic field direction) and machines built to do so systematically are called 'Quantum Information Processors' (QIPs) or 'Quantum Computers'. More generally, a quantum computer implements a quantum algorithm by manipulating the states of quantum particles using lasers and/or electric and magnetic fields.

3.1 Information Processing in Defence and Security

Key Point 10: Future computer-based decision support systems (DSSs) will exploit 'Big Data' and QIP will help manage the data deluge

Whether at a local, tactical or an international, strategic level, Defence and Security requirements for information processing comprise the abilities to accurately gather, identify, process, understand and respond to mission-critical information so that decisions may be taken and actions completed in a timely and sustainable way.

Evidence-based decisions were facilitated by the first Information Age which allowed the use of more information than can be usefully and opportunely processed manually. In operational scenarios, faster analysis of more, richer, sensor data was possible; in logistical scenarios, the design, development, acquisition, storage, distribution, maintenance and disposition of materiel and personnel can be optimised more efficiently and carefully; in intelligence scenarios, denser and more diverse data can be cross-correlated more quickly giving greater confidence in predictions of future events. Computer-based (algorithmic) decision support systems (DSSs) have become ubiquitous across all of Defence and Security's businesses.

The business scope of Defence and Security is very broad and it is difficult to write a concise list of algorithms with potential value but, broadly, there are four areas:

- Situational Awareness and Survivability;
- Communications, Command and Control systems;
- Logistical, Medical and Operational Robotics;
- Training & Simulation.

These areas will all benefit from information processing developments expected during the second Information Age. In particular, QIP will allow both mitigation and exploitation

of the ‘data deluge’ associated with ‘Big Data’²⁷ and it is critical that UK Defence and Security plays a leading role in the development and adoption of QIP.

Quantum computers demand quantum algorithms to process data and although many have been developed only a small number have practical value for Defence and Security (see **Section 3.2**).

Today, data analysis has become an essential part of many decision-making processes and machine learning-based tools have been developed to perform the analysis in a semi-autonomous way. The authors believe the general class of machine learning algorithms will be of huge importance for Defence and Security and, in a number of important areas, quantum algorithms suitable for existing commercial NISQ machines already exist. Intensive efforts are continuing to develop fully automated tools which can analyse data, extract patterns and compose reports summarising the results using text and charts (see **Section 3.3.1**). The research has received considerable funding over two to three decades and deep learning (DL) techniques based on artificial neural networks²⁸ (see **Sections 3.3.2** and **3.3.3**) have been developed for many applications including speech and image recognition and processing, automatic financial trading, drug discovery, bioinformatics and autonomy.

Having summarised quantum algorithms of value, **Section 4** will consider applications of these algorithms in more detail.

3.2 Quantum algorithms

Key Point 11: Although QIP systems are crucial to meeting the computational demands of future decision support systems, existing (classical) algorithms must be recast to run on quantum computers

An algorithm is defined here as a logically-constructed set of instructions that takes an input, A (which may be null), and produces an output, B. If the instructions can be ‘understood’ by a computer (or ‘compiled’ into such a set), the set comprises a computer algorithm. More than one algorithm can usually be combined to execute a specific task and some algorithms achieve their objective more efficiently (using fewer computational resources) than others. Algorithms have a wide variety of applications including sorting by some criterion, searching or evaluating a function from given input data.

As anticipated earlier in the document (page 4), algorithms which can be run on quantum computers cannot yet be considered in isolation from the hardware although it is essentially true that algorithms constructed for conventional computers may be regarded as hardware-independent. It is desirable to be able to run existing algorithms on quantum computers but the technology to do that is still many years away and so it will be necessary for the foreseeable future to construct new quantum algorithms to exploit QIP hardware. Furthermore, because quantum and classical information is very

²⁷ ‘Mitigation’ because QIP intrinsically scales exponentially with problem (data) size. ‘Exploitation’ because current methods are severely limited in the amount of data which can be analysed and the subtlest detail will require analysis of vast amounts of different types of data

²⁸ There are many texts which discuss neural nets in differing levels of detail. <https://towardsdatascience.com/a-gentle-introduction-to-neural-networks-series-part-1-2b90b87795bc?gi=a47ec4e90e08> is a simple online introduction.

different, the forms of quantum and classical algorithms which solve the same problem are also very different when their purposes are the same. There is significant work ongoing to produce quantum compilers but until these become widespread, the implementation of quantum algorithms is still a challenging task.

Quantum algorithms have been a growing area of research (and an area of UK strength) since Deutsch described the first. Deutsch's algorithm was important, not because it solved a useful problem but because it was the first algorithm to exploit quantum physics and in doing so, achieved an exponential speed-up compared to a classical algorithm (see **Section E.1, Appendix E**). However, quantum algorithms are not magical, do not allow problems to be solved which cannot, in principal, be solved on a classical computer and do not always give exponential speed-up compared to classical counterparts.

The literature describing quantum software, theory and error correction is extensive and has resulted in a multitude of algorithms (see **Section 3.2.1**) but comparatively few have any practical applications and the mathematics required to understand them is challenging. Consequently, the potential of QIP is poorly understood and the general opinion seems to be that quantum software will lag behind developments in QIP hardware unless greater resources are committed. Currently, it is thought that quantum computing will significantly outperform classical computing only for a few algorithms, although research, and associated breakthroughs, continue to be made. One of the latest developments is concerned with evaluating certain characteristics of systems of simultaneous linear equations²⁹, possibly using a hybrid classical / quantum approach.

3.2.1 The Quantum Algorithm Zoo

Key Point 12: Out of the quantum algorithms found in the quantum algorithm zoo; many are of academic interest only and just a few are expected to be of value to Defence and Security or the wider economy

A list of published quantum algorithms, the Quantum Algorithm Zoo,³⁰ is maintained by Stephen Jordan and, at the time of writing (May 2020), comprises 63 algorithms categorised as: Algebraic & Number Theoretic, Oracular and Approximation & Simulation. Quantum algorithms can be categorized:

- by the methods used (which include phase kick-back, phase estimation, the quantum Fourier transform, quantum walks and amplitude amplification³¹);
- by the type of problem solved;

²⁹ A quantum computer yields an expectation value associated with the solution rather than the solution itself and so calculations must be repeated many times to establish the probability that a particular solution is correct. See <https://arxiv.org/pdf/1302.1210v1.pdf>

³⁰ <https://quantumalgorithmzoo.org/>

³¹ Terms such as phase kick-back, phase estimation and amplitude amplification are impenetrable for non-experts but are part of the language of quantum algorithms and reflect the fact that quantum states, in general, are described with complex numbers, i.e. a number which has a real and an imaginary part or, equivalently, can be written as an amplitude (A) multiplied by a phase, $e^{i\varphi}$. As an example, consider a quantum state which is acted on by quantum gates; the result is $f(\varphi)Ae^{i\varphi}$ where $f(\varphi)$ is a function of the form $e^{Mi\varphi}$. The phase kick-back is the term $e^{Mi\varphi}$ and is a useful concept in quantum algorithm design providing a common framework to understand many quantum algorithms.

- by the speed up (over classical algorithms) achieved.

Most of the quantum algorithms which inhabit the Quantum Algorithm Zoo are arcane with few, if any, practical applications. An alternative to the grouping used by the Quantum Algorithm Zoo is categorisation into those requiring many qubits, and with theoretically proven exponential speed-ups, and those which are more practical, heuristic, requiring fewer qubits and tolerant of some noise.

3.2.2 NISQ algorithms of primary importance in Era 1

Key Point 13: Five Noisy Intermediate Scale Quantum (NISQ) algorithms are expected to provide value to Defence and Security in Era 1

Major effort is being devoted commercially to developing NISQ hardware but the uncertainties about how much resource is required for a given fidelity of operation mean that it is impossible to predict accurately algorithm speed-up compared to classical analogues. However, broadly, only lesser (polynomial) speed-ups are expected. **Appendix C** presents a critical overview of five quantum algorithms executable on the NISQ machines available now or expected to be realised before 2025. These algorithms are

- Shor (**C.1**): factorisation of integers into two prime numbers;
- Grover (**C.2**): searching an unsorted database;
- quantum Fourier transform (QFT) (**C.3**): a key component of many quantum algorithms;
- quantum machine learning (QML) (**C.4**): a quantum algorithm for the analysis of classical data on a quantum computer;
- quantum annealing (**C.5**): finding the global minimum of a function using quantum fluctuations.

The authors believe these could provide value to Defence and Security if adopted.

A crucial nugget lies within generic group ‘quantum machine learning’. That nugget may well eclipse all others as regards value, at least over the next 10 years and so receives extended mention (see **Section 3.3.2**). Quantum computers are able to run a family of algorithms termed ‘neural nets’ in a fashion which side-steps the speed limitations that have inhibited the widespread take-up of neural nets to date.

Neural nets are a mature class of pattern recognition algorithms which are usually ‘run’ on digital computers. They are widely used and the UK already has skilled neural net programmers. Neural net algorithms are one of very few examples where very similar ‘code’ can be made to run on either digital or quantum computers.

Towards the end of Era 1, as quantum computer volumes increase, pattern recognition in large data sets using QFTs will become of increasing importance. Shor’s and Grover’s algorithms will increase in importance also as more network traffic, including sensor data – is encrypted.

3.2.3 NISQ algorithms of secondary importance in Era 1

Key Point 14: A further five NISQ algorithms may have value for Defence and Security during Era 1

Appendix D describes a further five NISQ algorithms which will also be executable on the NISQ machines available now or expected to be realised before 2025 but potentially will have niche value to Defence and Security.

As a group, these algorithms have wide ranging applications of interest to Defence and Security including:

- Variational Quantum Eigensolver: VQE (**D.2**): general optimization problems, quantum simulations, quantum chemistry;
- Quantum Approximate Optimisation Algorithm: QAOA (**D.3**): many applications but of particular importance are machine scheduling, image recognition and the layout of electronic circuits;
- Data-Driven Quantum Circuit learning: DDQCL (**D.4**): computer vision, speech synthesis, image and text analysis and molecular design for drug discovery;
- Quantum Auto-Encoder: QAE (**D.5**): quantum simulation, data distribution across nodes in a quantum network, reducing quantum memory requirements in quantum communication channels and simplifying quantum circuits;
- Population Transfer: PT (**D.6**): protein folding.

3.2.4 Quantum algorithms for Era 3

Key Point 15: Two quantum algorithms requiring large, fault-tolerant quantum computers may benefit Defence and Security during Era 3

Some quantum algorithms require many qubits and take significant times to execute. This algorithmic class requires zero error rates over these (relatively) long execution times and so will only be practical on the more mature quantum hardware not expected to be available until Era 3. All of the algorithms from the Quantum Algorithm Zoo will run on these machines and show exponential speed-ups compared to implementations on classical computers but, as commented in **Section 3.2.1**, few are known to have any practical applications. **Appendix E** describes the first invented quantum algorithm plus two others which are expected to have value for Defence and Security:

- Deutsch's algorithm: historical interest;
- Quantum Simulation: prediction of quantum properties of large numbers of entangled quantum objects;
- Linear equations: determination of simple relationships between an outcome and one or more variables that drive that outcome.

3.3 QIP for Automation

Key Point 16: Artificial intelligence is increasingly being used to deliver business functions

Increasingly, business functions are being replaced by automation. This implies the need for artificial intelligence (AI) which can be realised by two different approaches.

One is to create a general purpose 'brain' that would be recognised by people as human-like. An approach would be required that delivers all that a human brain does (and maybe more). This is termed Strong AI. Little progress has been made since digital computers were invented and it is unclear how it would be possible to decide if such an AI was intelligent or not.

A far more fruitful approach has been that of Weak AI. This considers several topic areas which collectively might be combined to make something that delivers complex adaptive behaviour within a real environment such that it emulates intelligence. Sometimes these techniques have been combined to make software interact as if it were a person, searching databases or controlling systems in response to commands

from a human. These systems are often called Intelligent Agents (or ‘bots’) and the technology Agent technology. It combines one or more of the following technologies:

- Ontologies (the symbology of situation representation and reasoning);
- Situational Understanding;
- Reasoning;
- Planning;
- Learning;
- Natural Language Processing (to communicate with people);
- Social Intelligence (to help communicate with people);
- Agent Technology Human-Computer Interfaces to convert people into ‘commanders’ of an Agent workforce.

Effective realisation requires massive information handling architectures, special AI related chipsets and software toolkits. It can involve robotics technologies e.g. for an information gathering asset.

3.3.1 Automated Data Analysis - The Automated Statistician

Key Point 17: Automated data analysis using digital machines is becoming mature

In an increasingly digitally-empowered world, data analysis has become an essential part of many decision-making processes. Numerous tools have been developed and some are freely available (such as Microsoft’s Power BI Desktop³²) but training and experience are required to exploit these tools effectively and the need to develop automated machine learning tools has been recognised increasingly over the past decade and many different methods have been developed. Microsoft’s AutoML³³ is a well-known state of the art package, but its use requires significant human interaction and interpretation. However, work is ongoing in Ghahramani’s group at the University of Cambridge, funded in part by Google, to develop software (the Automatic Statistician³⁴) which can analyse data, extract patterns and compose reports summarising the results using text and charts. The project was reviewed in 2015 by the MIT Technology Review.³⁵

The Automatic Statistician is addressing two of the more significant shortcomings of Machine Learning which inhibit the widespread adoption of ML methodology. The first is that substantial human input into the process is still required to identify features in the data and develop models. The second is that the results of the automatic analysis, while accurate, are obtained by processes which often are very difficult to understand and, therefore, trust (for technical, legal or ethical reasons). The Automatic Statistician uses Bayesian model selection to address these limitations and has been successfully applied to a diverse range of problems³⁶ including non-stationary temporal variations in airline passenger numbers, sun spot activity and smoke produced by wood burning stoves. With efficient use of computer resources, the software package aims to:

³² <https://docs.microsoft.com/en-us/power-bi/desktop-what-is-desktop>

³³ <https://azure.microsoft.com/en-us/services/machine-learning/>

³⁴ <https://www.automaticstatistician.com/about/>

³⁵ <https://www.technologyreview.com/s/535041/automating-the-data-scientists/>

³⁶ <https://www.automaticstatistician.com/examples/>

- automate the process of feature selection and data analysis assuming Gaussian processes (or other models such as regression or classification and for multiple data types);
- automatically account for missing data values, outliers and different types of data;
- search over a large space of models to identify the best model that reliably describes patterns in the data;
- produce reports explaining the patterns found to the user.

A particular success of the project to date is claimed to be the explanations, in plain English, of the results found. The flexibility is achieved by systematically constructing data representations as combinations of a set of functions (constants, and linear, squared exponential, periodic and white noise) subject to ‘+’ (addition), ‘x’ (multiplication) and ‘CP’ (Change Point) operators.³⁷

A search over all the models generated is performed to identify the optimal model in the search space. This is not guaranteed to be the ‘best’ model since not all possible models are generated. Finding the globally ‘best’ model is not usually essential, if a good-enough model is found in an acceptable time. If needed, other ways of searching and evaluating models can be used to find the global best model.

At the end of the model selection process, the one with the highest ‘figure of merit’ is used to generate a natural language description of the model by converting to a standard form, picking the dominant function type using a pre-specified preference and expressing the other functions in the model in predetermined natural language expressions to produce a plain text report; the report on sun spot behaviour has text of the form³⁵

‘This component is approximately periodic with a period of 10.8 years. Across periods the shape of this function varies smoothly with a typical length scale of 36.9 years. The shape of this function within each period is very smooth and resembles a sinusoid. This component applies until 1643 and from 1716 onwards.

This component explains 71.5% of the residual variance; this increases the total variance explained from 72.8% to 92.3%. The addition of this component reduces the cross validated MAE by 16.82% from 0.18 to 0.15.’

Graphical representations of the models are also included (for this example the full report runs to 15.5 pages and includes 33 figures).

3.3.2 Neural nets

Key Point 18: Neural nets on digital machines have been developed for control systems, sensor processing, AI / machine learning, situational understanding and other applications

Neural nets allow a more general approach than that described in **Section 3.3.1**, and one which does not require human intervention. Neural nets have been developed³⁸,

³⁷ Steinruecken *et al* in F. Hutter *et al.* (eds.), Automated Machine Learning, The Springer Series on Challenges in Machine Learning, https://doi.org/10.1007/978-3-030-05318-5_9

³⁸ Some technical communities might refer to Boltzmann Machines or ‘RBM’s. See **Appendix F**

essentially since digital computers came into use, and applied to many problems including:

- **Financial trading / sociology** – neural nets are able to recognise complex patterns and react accordingly, at vast speed. Recently, similarity has been observed between the behaviour of these AI systems and that of humans³⁹ suggesting a useful tool for Defence and Security;
- **Pattern identification** – detecting features in images but also in data that makes no intrinsic sense to Humans. Neural nets are agnostic as to where patterns come from – they are as effective with electronic emission data or multi-spectral imagery as anything else. They are not limited to being trained to patterns that humans can detect;
- **Control systems** in aircraft, missiles, fire control systems and defensive aid systems. Classical control systems usually use a simple mathematical relationship between the control signal and what is intended to happen. But real systems have edge effects, need to deal with instabilities and the ‘ideal’ control laws are often too complex to determine. Neural nets ‘learn’ and can create exceptionally effective control loops able to create complex optimised responses to events;
- **Sensor data processing** e.g. data fusion, navigation, resolving signals in noise, interference and jamming;
- **Machine Learning** – neural nets have proved very effective in the implementation of a very broad range of techniques designed to allow machines to ‘understand’ and react to their environments. This impacts Autonomy, robotics, automated logistical handling and human-machine communication;
- **AI Situational Understanding** – to date only applied to simple situations because very large neural nets are ‘un-runnable’ on current computers. But the limitation appears to be only one of computational speed; quantum computers will dramatically lift that limit. Wargaming, threat detection, response generation at tactical and national levels appear possible. Conventional computers have to be ‘told’ what to recognise and highly complex systems escalate in programming man-hour cost. Neural nets ‘learn’ by being presented training data and are potentially capable of providing highly complex analysis and response in an affordable way;
- **Identifying warning markers** – in conventional computing a programmer knows what pattern has to be ‘spotted’ and programs his knowledge into the computer. So only known ‘markers’ can be automatically recognised. Neural nets derive their own indicators by training from mass data. This is a huge strength but can be a problem, as the net is not generally able to explain its reasoning and the method may, for some reason, only be valid within the training data.

Classical computers break down problems into a sequence of many small steps which are then individually executed at very high speeds. Neural nets are conceptually parallel (all calculations performed concurrently) and with suitable algorithms can be ‘run’ on a classical machine comprising many interconnected processing units so that the ensemble of sequential processors approximate a parallel processor. Experience

³⁹ <https://techfinancials.co.za/2020/01/18/financial-trading-bots-have-fascinating-similarities-to-people/>

has shown that neural nets work extremely well but they apply huge computing loads and the difficulties of providing the necessary computer resources limit their effectiveness and accuracy. Despite the compute-intensive requirements of neural nets, they are proving to be invaluable; for instance, the BBC recently (January 2020) reported⁴⁰ Google Health's success using AI to diagnose breast cancer from analysis of mammograms and neural nets were shown to outperform six radiologists.⁴¹

3.3.3 Quantum neural nets

Key Point 19: There is the potential for overwhelming quantum speedup by running neural nets on quantum computers ('quantum neural nets')

Quantum computers are intrinsically parallel and can 'execute' a neural net in one machine cycle instead of many thousands or millions of sequential steps. In addition, they have the potential for high levels of truly parallel node connectivity by exploiting superposition and/or entanglement in addition to options for conventional 'wiring' as used by digital machines. In contrast, a digital computer running a neural net has to go through all the links in sequence doing the calculation, summing and weighting results, then feeding onto the next node. The potential for overwhelming quantum speedup is clear.

Although circuit model machines are not yet sufficiently advanced to challenge current conventional supercomputers, reports in the literature indicate that the D-Wave machine can run neural nets of some complexity. This capability has attracted very large private investment by Google, IBM, Intel, Microsoft and others. There has also been investment by US government entities including Los Alamos, and NASA, and Lockheed-Martin. All of these organisations are exploring the use of quantum neural nets (i.e., neural nets running on quantum computers). Google, Deep Mind and Intel, in particular, are investing heavily in neural net research as well as quantum computing.

Neural nets are often used in machine learning, where the weightings between nodes represent what is learned. This would map well onto a QIP / Digital hybrid where the digital computer captured and stored the best weightings, and the QIP performed the weighted multi-layer neural net calculations. Neural nets only have local connectivity to neighbours in the next layer and not all-to-all connectivity. A neural net with many layers is called a Deep Neural Net and Deep Machine Learning usually centres on Deep Neural Nets. Three layers would be a 'normal' net depth, as opposed to a 'deep' net.

Neural net calculations are not programmed in the usual sense of the word. Instead, a process of 'training' is used.

⁴⁰ <https://www.bbc.co.uk/news/health-50857759>

⁴¹ <https://www.nature.com/articles/s41586-019-1799-6.pdf>

3.3.4 Training neural nets

Key Point 20: Training neural nets is not trivial and requires careful selection of training data sets. However, this training data also can be used for quantum neural net calculations provided it is suitably modified to be input to the quantum computer

Training (and validation) is not a trivial process⁴² and assistive tools have been developed for classical neural nets.⁴³ Without care, poor quality solutions, long training times and possible failure to find any acceptable solution at all are likely.

At least two datasets are required. The training dataset referred to above is used to evaluate the parameters (weights of connections between the nodes) of the neural network using a ‘supervised learning method’ (a fitting process). The set of parameters is refined iteratively to give a set of model weights which allow the training set data to be reproduced to an acceptable accuracy. Subsequently, a validation dataset, containing data previously ‘unseen’ by the neural net, is analysed using the model weights. The validation step can highlight shortcomings in the training set and the neural net must be retrained. For instance, overfitting gives model weights which reproduce the training data to high accuracy but can only give poorer quality fits to new data.

Essentially the training process is the optimisation of the model weights to give the best fit to the training data however the dimension of the search space is exponential in the number of inter-node connections. Even moderate sized neural nets can generate search spaces with millions of dimensions.

Conceptually, the optimisation process is a search for the solution (set of model weights) which finds the lowest point on a landscape (surface) of hills and valleys. Overall, the best solution is that with the smallest error between the actual and fitted surface (mathematically this is called a non-convex optimisation problem). In addition to the high dimensionality of the landscape, the problem is numerically challenging because algorithms have problems with regions where the surface has zero slope or if there are shallow valleys (minima) which correspond to a sub-optimal solutions and the algorithm is unable to converge onto the deepest valley.

Other considerations during neural net training include the amount and type of training data both of which depend on the complexity of the problem (and on the training algorithm used). It is not possible to generalise but sufficient data is required so that it fully spans the complexities of the feature space to be described and allows accurate evaluation of the model weights. In practice, an empirical approach is required; data sets are augmented until stable solutions are obtained. (Note: if linear optimisation methodology is used, at least hundreds of data items are likely to be needed while non-linear methods will require thousands, or more, data items.) This can result in the need for training sets comprising hundreds of thousands (for ‘average’ problems) to tens of millions (for ‘hard’ problems)⁴⁴ of data items and it is because of this ‘big data’ nature of neural net training that truly parallel quantum computers, when available with

⁴² See https://ml4a.github.io/ml4a/how_neural_networks_are_trained/ for an introduction

⁴³ <https://playground.tensorflow.org/>

⁴⁴ <https://machinelearningmastery.com/much-training-data-required-machine-learning/>

sufficiently large quantum volumes, are expected to outperform conventional supercomputers.

Training data selected for (classical) neural nets can be used to run quantum neural net calculations. The only additional effort required is the conversion of the training set data into a form readable by the quantum computer.

4 Distributed Enterprise Management IT

Key Point 21: QIP can be applied across all Enterprise activities where IT is already in use

4.1 Civil and Military Enterprise Management

The problems of coordinating a distributed activity with many people apply equally to the military and to the majority of medium to large size government and commercial enterprises. To a significant extent this has been driven by available computing and communication structures. While civil standards drove data transmission the military drove 'network architectures' such that the formerly military system, Transmission Control protocol/Internet Protocol (TCP/IP), is now near-universal. Conceptually it assumes a network of nodes which collectively perform a joint activity, with communication between nodes and the ability to access locally held data in the network from any node, subject to permissions.

Military enterprise management is described using acronyms, one of the most common being C4ISR (Command, Control, Communication, Computers, Intelligence, Surveillance and Reconnaissance⁴⁵). The acronym mixes facilities (such as computing and communications) with activities (such as Command, Surveillance and Reconnaissance) and misses out crucial components such as Logistics and Personnel Management / Training; it remains a convenient shorthand.

Commercial enterprises contain very similar components, even if military shorthand conceals the fact. For example, Intelligence is the business of taking in information on 'what is happening', derived from internal reports and external surveillance, and converting it into management information against which decisions can be taken (called Command but meaning Management). It might be argued that Reconnaissance is 'different' in that it is a very resource intensive managed process, consuming a bigger fraction of enterprise resource than is common in commercial enterprise. But even that is more a question of scale than nature.

Which leads to the conclusion that quantum computing impact on Enterprise Management will be just as important as on the Military.

4.1.1 System Characteristics

Key Point 22: Over the next 15 years QIP will impact the sifting, compiling, extraction and presentation of information to human decision makers through faster AI execution

Both civil and military Enterprise management systems share common characteristics

- A TCP/IP based system is used such that people (or groups of people) are treated as 'nodes' in a network, with ability to hold data in one node yet access it from others. Nodes can be individual people at their desk, or a department or function which may also have its own network. The network is designed to be reliable and resilient;
- It is possible to transfer data and messages quickly between nodes;
- Data is partly automatically entered, reflecting stock and movement information, where assets are and their readiness levels / activity, and partly entered by

⁴⁵ <https://en.wikipedia.org/wiki/C4ISR>

people. The military might refer to outward looking data gathering as surveillance and reconnaissance;

- Reports / messages are compiled from interpreted data (usually by people) and distributed along with information on actions taken. Sometimes there is a degree of automation in that compilation. The military term for the part of this activity focussed outside the organisation is 'Intelligence';
- There is time pressure - activity has to take place quickly;
- There is management of authorities / permissions and often systems to authenticate, verify or audit that activity is as claimed and permissions adhered to;
- The parts of the Enterprise share a common goal and work within an operating plan designed to align local activities so as to achieve that outcome. Military activities are usually highly planned and designed to achieve particular outcomes e.g. use of effects-based operations⁴⁶.

Quantum information systems are expected to impact the security and resilience of the network, the ability to authenticate activity and the ability, through AI, to automate the conversion of mass data into situation reports matched to the roles of the nodes. This will include searching data to find events and patterns.

QIP is not expected to impact the mechanisms by which mass raw digital data is stored, retrieved and transferred. The TCP/IP mechanisms are aligned to digital computer implementation. Emerging Quantum Communications technology could, however, affect these aspects of the Enterprise.

In principle 'nodes' could be automatic or partly so. Civil examples might be remotely controlled pumping stations or process plant; military examples might include highly automated reconnaissance drones or swarms.

4.1.2 Military Enterprise Management (C4ISR)

Key Point 23: For Enterprise Management, the value of QIP is expected to be accurate, rapid pattern matching within 'Big Data' allowing high quality information to be extracted and used to facilitate timely, accurate decisions. Benchmarking, as QIP and digital platforms develop, will confirm or refute this

Computing in C4ISR systems has been exclusively digital to date and quantum computers do not offer superior mechanisms for the fast mass storage and retrieval of data. However, emerging hybrid quantum / digital computers will provide extra capability where digital computers do not excel. The quantum computing strength lies currently in pattern matching within large data sets and the identification of events and possibilities.

Surveillance and Reconnaissance are managed activities designed to deliver the data required to take decisions. Interpretation of that data lies as the heart of Intelligence. Their civil counterparts are the data flows describing progress in manufacture, stock, distribution or whatever, depending on the nature of the Enterprise. Information is collated, interpreted and fed to the human decision maker. The level of automation in collation and analysis may be high, including use of AI.

Information needs to be converted into 'Intelligence', that is to say interpreted such that it is easily used for decision making. There are usually many decision makers, each

⁴⁶ https://en.wikipedia.org/wiki/Effects-based_operations

focussed on a different aspect of delivering a combined plan. Each requires a 'picture' he can readily understand so it must be compiled and presented. Failure to sift, extract, compile and interpret will result in data overload and obscuration of key facts / events leading to poor decisions. This is the aspect of C4ISR and enterprise management systems that quantum computing is most likely to address.

One of the best algorithms for pattern recognition is the neural net, which is able to process many types of data including pictures. Neural nets are also able to adapt and 'learn' from data, hence they are commonly used by AI researchers focussed on Machine Learning. An early application of neural nets was financial trading where they proved able to recognise patterns and automatically generate responses, far faster than a human could.

In a C4ISR system, neural nets, executed at extreme speed on a quantum computer, could be used to extract relevant features from images, spot patterns in data and generally move from a deluge of raw data to collated 'Intelligence' at very high speed. In this a quantum computer does not offer a new capability but the transformation of an old method, one un-used because too slow. Neural nets present extreme loads to a digital computer because they present a truly 'parallel' algorithm that has to be mapped onto many serial processors and the results accumulated. Unfortunately, that process scales as the square of the number of data items and is further slowed due to the problem of coordinating many processors together as they work on parts of the data. A quantum computer with enough qubits to cover all the data points at once can compute the correlation in one machine cycle. Thus, quantum neural nets are expected to enable superior AI systems but benchmarking on quantum and digital machines, as they develop, is necessary to confirm or refute this.

4.2 Information Security

Key Point 24: Quantum Computing is expected to impact secure data transmission over networks, authentication of both traffic and participant identity and to improve network immunity to attack

Basic network immunity to subversion is usually accomplished by:

- Firewalls (internal and external) that limit the use of network TCP/IP protocols to only those required to conduct the Enterprise's business, which constrains subversion options;
- Checking programs and data for the presence of code associated with historical attacks (virus checkers);
- Intruder detection systems, which monitor internal network traffic and look for anomalies or other 'foreign activity';
- Use of password protection and 'permission hierarchy' to inhibit system and data changes that are unauthorised.

Quantum computers are able to identify patterns in both data and activity and manage very high processing throughput. This combination is likely to offer extremely strong performance in 'intruder detection' by real-time monitoring of traffic.

4.2.1 Data Networks

Key Point 25: Civilian enterprises have the same requirements as Defence and Security, but on larger scale and early adoption for commercial applications is essential to attract the investment and R&D activity needed to mature the technology. Modern network enabled capabilities depend on network-orientated methods of communication and control to maximise the success of Defence and Security missions. Unimpeded, secure flow of information at high data rates is critical and can be achieved by a quantum network

A local quantum network sends quantum information (entangled qubits) from one quantum processor to another. Multiple quantum networks can be connected together to form a quantum internet and, for most anticipated applications, only modest quantum processors are expected to be needed. Protocols such as quantum key distribution (QKD) in quantum cryptography systems requires processors that have to prepare and measure only single qubits, although this needs to be done at high speed to achieve the desired data rates.⁴⁷ The introduction and adoption of quantum networks should proceed with urgency since this lessens the future risk that data intercepted 'now' can be decrypted when quantum computers have the maturity to do so.

The architecture of a quantum network is analogous to a classical network.

Applications run on end nodes which comprise quantum processors which can manipulate at least one qubit. As well as at the processors terminating the network ends, manipulation of quantum information including error correction is needed at any quantum repeaters in the network. The end nodes are connected by communication lines, usually standard telecoms fibres, which transport the qubits (called 'flying qubits') encoded in polarisation states of laser light. To make maximum use of existing infrastructure, optical switches are needed to deliver qubits to the intended quantum processor. The switches must preserve quantum coherence, which makes them more challenging to realize than standard optical switches.

Light loss within the fibres requires regular signal amplification which is not problematic if this can be done without risk of attack. In a quantum network, because quantum states cannot be copied⁴⁸, this cannot be done. In QKD systems, trusted repeaters can be used to build long distance networks but true quantum repeaters are required for end to end generation of quantum entanglement, and - by using quantum teleportation - end to end qubit transmission.

Transmission over longer distances, without the use of repeaters, is possible using free space communication (laser beam propagation freely through the atmosphere or space). Although there are still problems which are active research areas, free space systems are the only option for military operations which preserve freedom of manoeuvre. Systems can use either satellites in low earth orbit for long distance networks or, more affordably, constellations of unmanned air vehicles for local area networks over a battlefield. Jian-Wei Pan, who leads China's massive quantum technologies programme and was the first to demonstrate a long distance, fibre-based

⁴⁷ Note this is in contrast to quantum computing where useful applications can only be realised if the quantum processor can manipulate many qubits so that it outperforms a classical computer

⁴⁸ The no-cloning theorem, see https://en.wikipedia.org/wiki/No-cloning_theorem#_blank

QKD system, led the Micius team which demonstrated a satellite QKD system in 2017⁴⁹ and has likened the two principal challenges of free space networks as (single photon detection) an Earth based observer seeing the light from a single burning match on the Moon and (signal resolution) an Earth based observer reading the licence plates of a vehicle on one of Jupiter's moons.

Quantum teleportation⁵⁰ is the name used for the communication process in which quantum information is transmitted from one location to another rather than the qubits themselves. Quantum teleportation was first realised using single photons but later demonstrated using atoms, ions, electrons and superconducting circuits. The current record for long-distance quantum teleportation is 1,400 km held by Jian-Wei Pan's group using the Micius space satellite.⁵¹

4.2.2 Authentication

Key Point 26: 'Quantum signatures' use quantum encryption methods to verify the identity of a message sender and are immune to attack by quantum computers

Despite the claims of 'complete end-to-end security' made by some commercial suppliers of quantum communications systems, the protocols used do not verify the identity of the message sender. The need to certify the origin of information is essential in any secure communication system whether it is used by civilian or government organisations and is an area of active research and has come to be known as 'quantum signatures' by analogy to handwritten signatures.

Quantum signatures make use of asymmetric (public) encryption keys⁵². Thus, the sender signs the message by creating a pair of keys; one is the signature and one the corresponding public key. The public quantum signature key can be created easily from either a classical bit string or a string of quantum qubits and the process makes use of the Heisenberg Uncertainty Principle so it is impossible even for a quantum computer to compute the inverse. A practical scheme must provide security against tampering by the sender, receiver or a third party and when the signature validity is tested, the result (true or false) must be the same for any recipient.

⁴⁹ <https://spectrum.ieee.org/tech-talk/telecom/security/china-successfully-demonstrates-quantum-encryption-by-hosting-a-video-call>

⁵⁰ 'Teleporting an Unknown Quantum State via Dual Classical and Einstein–Podolsky–Rosen Channels', C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, W. K. Wootters, Phys. Rev. Lett. **70** (13), 1895–1899 (1993); also see https://en.wikipedia.org/wiki/Quantum_teleportation#cite_note-0-8 for more details

⁵¹ 'Ground-to-satellite quantum teleportation', J-G. Ren, P. Xu, H-L. Yong, L. Zhang, S-K. Liao, J. Yin, W-Y. Yu, W-Q. Cai and M. Yang, Nature, **549** (7670), 70–73 (2017)

⁵² There are two approaches to encrypting information. Symmetric encryption is the oldest and best-known method; a private key, known only to the message sender and recipient, is used to encode and decode the message but this technique is vulnerable to third parties acquiring the key. Asymmetric encryption uses two keys, one which is freely available (public key) and a second which is kept private. Data encrypted by using the public key can only be decrypted by applying the same algorithm but using the corresponding private key. (The inverse is true; any message encrypted using the private key can only be decrypted by using the corresponding public key.)

Nielsen and Chuang's well-known book gives more detail.⁵³

4.2.3 'Auto Immune' network defence

Key Point 27: Recently invented bio-mimetic cyber defence systems protect networks in a way which is analogous to the immune systems of biological organisms. Quantum processing may help such systems identify threats

Methods which address the problem of identifying malicious traffic in vast amounts of proper traffic but introduce bottlenecks are usually unacceptable. Consequently, sophisticated systems for the detection of intruders in networks can successfully be applied to small networks but not large ones. To protect a network the 'detection system' has to be both dispersed (to minimise bottlenecks) yet centralised (to see the 'big picture' of events across the network) and quantum neural networks may provide both superior pattern recognition and extreme speed. There is also a response time issue. Attacks by an automated attack tool can be extremely rapid.

Recent ideas ('Qiangwang' or cyberspace power) have been proposed by Wu Jiangxing⁵⁴ which mimic biological immune systems. These bio-mimetic 'cyber mimic defence' (CMD) systems are being developed rapidly and in 2018 a Chinese cyber defence system, demonstrated in Nanjing, China, successfully withstood more than 2.9M attacks of various types in a first international challenge by 22 hackers from China, US, Russia, Germany, Japan, Iran and Poland.

The bio-inspired system uses dynamic, redundant, software to change the network's external appearance allowing it to adapt to a hostile environment. This ever-changing software environment makes it difficult for conventional hacker-attacks to locate a target. Theoretical analysis and computer simulations have shown that CMD can significantly increase the difficulties for attackers, enhancing security by at least a factor of ten.

4.3 Data and Information Processing

Key Point 28: Quantum computers are only expected to provide quantum advantage in areas for which quantum algorithms are superior to conventional methods. One such area is image processing and is the subject of this Section

In general quantum computers have completely different strengths and weaknesses compared with digital computers. It is therefore unlikely that they will challenge digital computers except in areas where there is a special 'quantum advantage', something derived from the Quantum Algorithms set out in **Appendices C, D and E**.

Section 4.3 concentrates on image processing and image searching because this is expected to deliver the most rapid and far-reaching change from using existing quantum computer approaches.

However other special quantum computer abilities, such as solving highly parallel vector manipulation problems (digital maps are stored as vector data) will become increasingly important as Circuit Model quantum computers mature and further

⁵³ M Nielsen and I Chuang, 'Quantum Computation and Quantum Information', Cambridge University Press, ISBN-10: 9781107002173 (2010)

⁵⁴ 'Cyberspace Mimic Defense: Generalized Robust Control and Endogenous Security' Wu (Springer) 2020

algorithms are ‘invented’. This area of the QIP Landscape will be extended as research and understanding develops.

4.3.1 Automated image searching

Key Point 29: Quantum enabled automated searches for features in all the world’s images would allow the recognition and tracking of events

It has become customary to search for events and information on the internet, or through large databases, using search engines to carry out ‘keyword searches’ (ie, searching for specified words or phrases).

Search engines operate by building and maintaining tables of which documents contain which words, and often some kind of connectivity data between words to allow phrase searches. So, to find a document containing a phrase using the engine it suffices to search an index which identifies the documents containing the search term. More words exclude more documents, finally producing a shortlist of items to retrieve. This is enormously faster than searching every document in response to every query.

A similar approach to search for image content has not been achieved but is the subject of massive research e.g. by Google. There are 4 broad approaches to automatically searching for images, which are

- a) Searching for the identical (including size). This might be direct byte by byte comparison of two files, or by decoding the file and performing a pixel by pixel comparison.
- b) Searching by image statistics. There are very many different approaches that have been trialled but, while partly effective, they tend to achieve a ‘shortlist of candidates’ that might surprise a human observer, who would not regard them as ‘similar’ at all.
- c) Identification of features in the image that have meaning to humans. Examples would include people, faces, vehicles, bridges. This imposes enormous computing load but is otherwise attractive.
- d) Identification of features not meaningful to a human, but from which features as described in c) can be identified. Examples would be short lines and their angle, corners, uniformly textured or coloured areas.

Approaches c) and d) lead to identification of features within images that humans can relate to and are therefore attractive, but both c) and d) demand massive computing power. This appears to be a key area that quantum computers can effectively address and where they have a large inherent speed advantage compared to digital processing.

Given the relentless expansion of available data, quantum enabled automated searches for features in all the world’s images (internet, databases, real time imagery from satellite and aircraft reconnaissance etc.) would allow the recognition and tracking of events.

4.3.2 Image data processing

Key Point 30: Quantum image processing (QulmP) has seen much more research and development than has quantum signal processing

Surprisingly, the field of quantum image processing (QulmP) has seen much more research activity than has quantum signal processing, although it is only recently that studies progressed beyond simulations on classical computers. Classical image data is converted to quantum data; in general, the classical signal (e.g., photon counts) is sampled and converted to a quantum signal to give an internal representation of the data. Fidelity of this classical → quantum step is critical; even noise as small as ~5 % in each pixel of the quantum image essentially destroys the appearance of the image because of the loss of edges and textures.

Not only must the image data quantisation be carried out with very low error rates, its transfer to the quantum processor from the imager (via a classical to quantum converter, if needed) must be achieved with similarly low errors. Fortunately, ways to transfer quantum information reliably have been developed as part of quantum communications R&D. Long-distance quantum teleportation protocols to transfer quantum information carried by states of quantum light were first investigated in the 1990s but suffered from huge photon losses. Recent work has developed protocols for reliable quantum communication in the presence of noise by introducing additional sub-systems comprising superconducting ‘quantum oscillators’ at both ends of the quantum channel. Transfer of a quantum state between superconducting qubits is achievable with microwave photons, already used for classical signal transfer, operating according to noise-tolerant protocols. Because the thermal noise affects both oscillators equally, the noise signal can be eliminated by precise coupling to the waveguides in the quantum channel. Long-distance noise-tolerant quantum channels seem feasible but have been demonstrated only over a few hundred metres; however, this is perfectly adequate for coherent image input to a quantum processor.

Quantum image processing can be carried out using any of the quantum computer types discussed in **Section 2.2**, and **Section 4.3.2.1** presents a very brief overview, otherwise adiabatic (D-Wave) computers will be discussed in this section.

Several implementation issues are common to all paradigms of quantum computing and comprise active and difficult areas of research because of the very different nature of quantum ‘information’:

- What are classical-to-quantum and quantum-to-classical interfaces and how are they implemented physically?
- What is the best internal (quantum) representation of the image?
- What is the output of a QulmP algorithm?
- How is the output from the QulmP algorithm used?
- How much noise in the end-to-end QulmP process can be tolerated?
- How does the computational cost scale with image size?

4.3.2.1 QulmP using circuit model quantum computers

Key Point 31: QulmP algorithms for circuit model machines are at low Technology Readiness Levels because of the immaturity of the necessary computing platforms

There are many quantum image processing algorithms for circuit model machines but they are difficult to understand and currently at low TRLs (because of the immaturity of the computing platforms). The principal algorithms are:

- Quantum Boolean Image Processing (QBIP);

- Flexible Representation of Quantum Images (FRQI);
- Novel Enhanced Quantum Representation (NEQR);
- Quantum State Tomography (QST);

and a very brief overview can be found in **Appendix G**.

4.3.2.2 QulmP using quantum neural nets

Key Point 32: Artificial neural nets are a mature information processing architecture and are well suited to Noisy Intermediate-Scale Quantum (NISQ) computers, especially D-Wave machines. Programming (termed ‘training’) neural nets can be challenging and requires careful selection of training data. Neural nets running on NISQ machines are expected to have many applications in Defence and Security

Conceptually a neural net, introduced in **Section 3.3.2**, is ‘computed’ in one cycle, with all the input weighted sums being calculated and summed concurrently. One could envisage a conventional electronic architecture using, for example, summing amplifiers at each node to achieve this. However, the vast number of node interconnections required by realistically sized neural nets defeats fully parallel digital chip implementation. Wiring on silicon chips is essentially 2D unlike the 3D interconnection of living neurons⁵⁵ and more than a few ‘crossing wires’ defeats the chip designer.

Digital implementation of neural nets, currently the norm, is essentially sequential. One or more processors are applied to each node in sequence to calculate a weighted sum. Instead of one cycle, the neural net is computed in more than N^2 cycles, where N is the number of nodes in the neural net hidden layers. The wiring limit is overcome with telephone exchange type message switching, which also requires additional compute cycles but this allows some degree of parallel computing. Ultimately the ‘telephone exchange’ limits the end result.

Quantum computers, where each node is a qubit, use quantum entanglement in addition to ‘wiring’ and so are able to achieve higher levels of truly parallel connectivity. They can execute neural nets of moderate complexity in a single cycle. Or if they have plenty of qubits, several smaller neural nets at the same time.

Hence if a digital computer had the same clock speed as a quantum computer running the same neural net then the quantum computer would be N^2 times faster. For a 100 node image processing net this would mean 10,000 times faster. Even with 10 digital processors running a single neural net, the speed-up would be more than 1,000 times and very much greater for a larger neural net or if computing several neural nets concurrently. The principal limit is the number of qubits; D-Wave plans to release a platform offering in excess of 5,000 qubits during 2020 and an upgrade to a machine with more than 10,000 qubits is expected within two years.

An array of qubits can be constructed relatively easily to provide connectivity between local qubits but not more distant ones. When searching for a feature in an image that feature will be identified from the inter-relation of the pixels that form part of the feature,

⁵⁵ Biological neurons are highly interconnected but do not have miniscule ‘error rates’ and appear to use low resolution serial communication messaging. This encourages the authors to believe that current error rates in quantum computers (which are good but far from perfect) are adequate for neural net use

and not the whole image. As a result, the connectivity and qubit count already achieved in current quantum computers appears adequate to task.

An image recognition neural net only needs enough nodes to cover all the pixels in the feature being sought, not the whole image size. A likely approach would be to 'scan' large images with a much smaller neural net of, say, 100 nodes (10x10 pixels). If the quantum computer had 5,000 qubits it could run 50 such neural nets concurrently, all within the same cycle.

Features can be areas of uniform colour or brightness, edges, corners, shapes. A common approach in image analysis by machine is to form a database of constituent features and then pass on that list for further analysis. Biological systems are believed to use exactly this approach. 'Features', as the term is used here, are therefore akin to the words or phrases that make up a text document and are amenable to indexed searches such as Google offers for the written internet but now extended to images. Note that an everyday example of this is OCR (optical character recognition) which is routinely offered on office scanners.

4.3.2.3 **QulmP algorithms for D-Wave, an adiabatic quantum computer**

Key Point 33: US industry has successfully used D-Wave quantum annealers for machine learning and image analysis for over a decade

As early as 2009, Hartmut Neven described work by his Google research team which used an early D-Wave processor based on the C4 Chimera architecture for binary image classification (see **Appendix B.7** for a summary of notation used to describe D-wave machine architectures). The algorithm was trained to identify cars within images and better performance was claimed than classical solvers running on classical computers.

After the 512-qubit D-Wave Two was released, NASA's Quantum Artificial Intelligence Laboratory (QuAIL), Google and the Universities Space Research Association began a multi-year collaboration investigating the use of D-Wave Two (updated to a 2000Q in 2017) for various applications including machine learning, pattern recognition and anomaly detection. In late 2015, Los Alamos National Laboratory (LANL) procured and commissioned a D-Wave 2X and during 2016 and 2017 ran seminars presenting LANL funded work including D-Wave algorithm development for machine learning, feature detection within images and image classification. In September 2019, LANL announced it had signed a contract to upgrade its D-Wave installation to the new 5000 qubit Advantage during 2020.

Today, D-Wave Systems highlight⁵⁶ a number of applications including machine learning, formulated as an optimisation problem, for feature recognition within images. By sampling from a probability distribution similar to a Boltzmann distribution, probabilistic algorithms are being developed for image analysis.

4.4 **Management Aids (AI)**

Key Point 34: In the context of information processing, management problems, in general, comprise the effective presentation of complex data in a comprehensible way. Machine intelligence is a promising solution to this problem

⁵⁶ <https://dwavefederal.com/applications/>

Section 3.3.1 described the Automatic Statistician. This is an interesting solution to the particular ‘management problem’ of presenting complex statistical data in a comprehensible fashion. The more general problem, presenting complex data matched to roles within an enterprise, is crucial to organisational efficiency.

In simple form Management Aids require that all data is somehow ‘tagged’ such that humans performing roles can do their own selection, and it requires that all data is made available to all possible users. This scales very badly in large enterprises with a lot of data and is often referred to as ‘Data Deluge’.

The alternative, on which this section is focussed, makes use of ‘machine intelligence’ such that computers can filter, extract and re-represent information.

4.4.1 Quantum enabled understanding

Key Point 35: QIP will reduce data deluge and enable better understanding to be extracted from large data sets through identifying correlations which could not be found using classical tools. QIP’s adoption is not justified simply by massive speed-up; the impact of the speed-up will be the key driver

Much understanding can be derived from the correlating of information derived from multiple sources. For a pan-optic view, all data sources must be fused together, and this can be expected to result in large data sets and data deluge. Topological data analysis is increasingly being applied to identify correlations and, with large data sets, Lloyd *et. al.*⁵⁷ have described a quantum algorithm for circuit model machines allowing the analysis of ‘big data’ with an exponential speedup compared to classical machines. Early circuit model machines with perhaps only a few hundred qubits are expected to be available during Era 2 and could be sufficient to give solutions which are impossible to find classically. If, therefore, solutions have been identified allowing quantum processing of large amounts of data, the remaining bottle neck is reading-in to the quantum machine the data in either a raw or pre-processed state.

4.4.2 Situational Awareness

Key Point 36: QIP will best extract the fullest information from future quantum sensors

An integral part of survivability is the awareness of **what** is happening **where** and **when** (i.e., situational awareness) and the use of this knowledge to achieve situational understanding in order to determine the future actions which will have the highest probability of achieving desired objectives. This definition can be used at the tactical level by a unit commander or the strategic level by a Supreme Commander.

Situational awareness is obtained from multiple sources including:

- Sensors (cameras, microphones, radio antennae, environmental sensors, etc.);
- Telecommunications;
- Media (newspapers, radio, television, etc.) and Social media;
- Professional and academic publications (papers, conferences, professional associations, etc.) and public data (government reports, official records, etc.).

⁵⁷ <https://arxiv.org/abs/0811.3171>

Existing sensors output **classical** information⁵⁸ but quantum sensors which output **quantum** information are in development and QIP will best extract the fullest information from that output. Imaging data has been most studied.

4.4.3 Situational understanding from situational awareness

Key Point 37: QIP using neural nets is likely to offer solutions in situational understanding in Era 1 via image analysis and pattern detection

In considering the list of AI technologies listed in **Section 3.3**, in the near term QIP using neural nets seems likely to offer solutions in ‘Situational understanding’ in Era 1 via two mechanisms:

- Image decomposition and processing;
- Pattern detection.

There are particular reasons why QIP will slowly impact other areas, concerned with complexity. Quantum computing elements, qubits, are general purpose problem solvers if all qubits are ‘entangled’, that is all qubits affect all other qubits. In practice, it is difficult to achieve more than very local entanglement without errors, so planned quantum computers usually have ‘clusters’ of entangled qubits in some kind of much larger array. A fully entangled 40 qubit QIP would be considered extremely ambitious at present⁵⁹ (> 10 years to realise) while partially entangled QIP of 1000+ elements can be built now. Consequently, problems involving particularly large patterns, for instance, will take longest to solve by quantum methods. Important problems difficult for current quantum computers include:

- AI Ontologies (languages for machine reasoning) since they are complex and contain huge numbers of symbols. Ontologies are difficult to process with small numbers of qubits.
- Planning often brings a similar ‘size’ problem, in that each entity to be taken into account will need at least one (and probably several) qubits.
- Learning poses a particular problem to QIP in that they essentially ‘solve’ problems in a single step and are then unable to store the answer. Answers, or intermediates, have to be read out and moved into a digital computer, quickly. Basically, pure QIP systems have poor memories, but a QIP + digital hybrid is very much more capable.

Neural nets have been found to be powerful in many forms of pattern matching. They have the useful property that the programmer does not have to work out a mechanism by which he can detect what he is looking for. Instead he feeds in sets of data, each comprising many patterns which either have or do not have the feature whose identification is to be learned. As each data set is presented to the neural network the presence or absence of the feature is declared and the network adapts its weightings and computation to best resolve the two situations (feature present, or not) and the resultant weightings are ‘the program’. This process is called ‘Training’.

⁵⁸ For example, a quantum imager responds to individual photons but outputs pulses of electric charge

⁵⁹ The key phrase is ‘fully connected’; circuit model machines being developed by IBM, Google and others typically have only nearest neighbour connections. D-Wave machines have slightly superior connectivities but significantly less than all-to-all.

4.4.4 Verifying accuracy

Key Point 38: Verification is critical for accurate situational understanding and may adopt methods similar to those pioneered by commercial organisations analysing social media

Defence and Security use Situational Understanding derived from data which in its rawest form is a partial capture of ‘ground truth’ recorded at a particular time. Analysis generates Information by bringing together, and interpreting, multiple data sources in order to ‘tell a story’ and use the resulting Information to drive decisions. Clearly the accuracy of the ‘story’ which is synthesised, and of the decisions derived from it, is critically dependent on the accuracy of the raw data and its interpretation.

Verification is the process of correlating and corroborating Data, and derived inferences, and should be independent of the personnel and processes used for the gathering and analysis.

A possible approach to verification is the use of machine learning and this has seen intense research activity over two to three decades. Techniques can be broadly classified as deep learning (DL) or semantic learning (SL). DL is based on artificial neural networks and can be supervised, semi-supervised or unsupervised and has been widely applied to speech and image recognition and processing, drug discovery, bioinformatics and autonomy. In contrast, SL seeks to understand structure in data by reference to larger, often massive, data sets.

Many SL techniques were developed in academia in the 1970s and 1980s but more recently research has invested principally in DL techniques for image analysis, especially by large commercial organisations including Microsoft, Google, Facebook and Amazon. These commercial players, for social media purposes, have also focused on understanding behavioural patterns. Currently, much research for commercial applications is developing data-centric DL methods, which do not require storing or processing significant amounts of factual information, knowledge. These DL approaches attempt to identify structures and anomalies correctly in datasets which is sufficient to allow individually targeted advertising. For non-commercial applications, processing and storing knowledge extracted and refined from multiple data sources is very useful for complex tasks and allows resources of learned ideas consistent with hypotheses to be established.

4.5 Robotics and The Enterprise

Key Point 39: By rapidly and accurately recognising the component parts of its environment a quantum computer running neural nets should be able to navigate, calculate orientation, avoid obstructions and ‘understand’ a robot’s environment through machine vision. Compact, low power quantum computers will be needed and possible chips are the subject of R&D programmes

At present machines to assist humans in the transportation, lifting, stacking, storing, processing and even situation monitoring are essentially human operated. They will be automated to a limited degree, to increase the capacity of the human directing them.

The level of autonomy built into powerful moving machinery working in close proximity to humans is currently very low. Autonomous road vehicles promise to be a notable exception but also illustrate generic concerns about humans and powerful robots moving in close physical proximity. A ‘standard’ solution to this is to segment the areas

in which people and machines operate. This is true for human operated machines and, for the same reasons, will be true for autonomous machines. It would be optimistic to assume that 'better' AI will fully remove the risks of close inter-working since that is not fully achieved by human intelligence.

However, there are major advantages to operating concepts where mobile robotic assistants work among people, without segregation. In this situation Health and Safety considerations encourages the use of smaller and less powerful robots.

An established example is the compact drone such as a quadcopter which has some ability to 'lift and shift'. 'Small' has the major advantage that mistakes (such as collisions) can be much less serious.

But the engineering challenges in building small automata are immense. Two of the most difficult are powering the robot (batteries remain a major inhibitor) and provision of compact and lightweight AI processing.

A mobile automaton must be able to precisely navigate, orient itself, avoid obstructions and plan activity in the context of its environment. Until it can do this, any special capability it has above humans (such as being able carry a heavier load) will be hard to exploit. In a factory or indoor setting there are technical solutions that may allow low cost implementation (such as fast communication to a remote computer providing the 'intelligence') but in a military context there is minimal pre-existing infrastructure and both Global Navigation Satellite Systems (GNSSs) and high-speed communications can be denied.

This is the technical 'blocker' that quantum computers may be able to remove. By rapidly and accurately recognising the component parts of its environment a quantum computer running neural nets should be able to navigate, calculate orientation, avoid obstructions and 'understand' the robot's environment through Machine Vision (MV). This has long been a major objective of MV research but demands huge computing resources. Quantum computers offer 'the same but much faster' in that they can 'run' neural nets, potentially at vastly greater speeds than a digital computer. This arises because neural nets demand parallel computation and a quantum computer can 'compute' a neural net in one machine cycle where a digital computer would need thousands of cycles.

This is believed to be the reason for very large quantum computing investment by the computer chip company Intel. Not only are Intel investing heavily to create single chip quantum computers with 1000+ qubits, they are also developing special neural net processing chips to replace conventional digital computers 'running' neural nets. These chips can already be purchased by research teams and 'built into' commercially available image frame stores able to interface directly with cameras. Eventually quantum computers will be 'chips', until then Intel are creating a digital stop-gap. The market for small automata is potentially enormous. Other possible compact quantum computers could be chip-based photonic quantum devices.

It is unclear at present how rapidly quantum computers may 'open up' this route. At present they are all too large to be useful, though this is unlikely to persist. In the same way that digital computers were once immobile 'mainframes' but progressed to become compact (single chips) and so are now routinely embedded within portable equipment.

4.5.1 Transport Automation

Key Point 40: Quantum neural nets are expected to have transformational impact for autonomous vehicles by facilitating a step change in machine vision

Autopilots / autolandings / auto-dockers in ships and aircraft, self-driving lorries and cars are very different in part because of the importance of environmental predictability to the operation of autonomous vehicles. The transformational impact of quantum computers is expected to lie in their potential ability to reason about a changing environment, plan actions and do it very quickly using quantum neural nets. The 'same but faster' immediate quantum computer 'enabler' is machine vision but vector computation is potent for planning and route-finding.

4.5.2 Logistical Systems

Key Point 41: Quantum computers are likely to accelerate the use of 'intelligent' systems controlling mechanical handling, storage and transport systems within individual machines and not just the Enterprise management network

While Logistical IT has tended to focus on management, the essence of a logistical service is that it manages physical objects that must be handled, stored and transported. Quantum computers, especially compact systems, are likely to accelerate the use of 'intelligent' systems controlling mechanical handling, storage and transport systems within individual machines and not just the Enterprise management network.

An enabler for impact by QIP will be the availability of compact, affordable systems (which might be realised by photonic quantum computers) and this could be transformative, breaking away from 'logistics and computers = data management'.

Small quantum computers will underpin autonomous machines that can safely work alongside humans, even in complex environments, without posing any physical threats to the human co-workers. The spaces, such as warehouses and factories, in which humans and machines work no-longer need be segmented, simplifying working practices, increasing efficiency and reducing costs.

For the military, and military/civilian groups responding in disaster relief scenarios, battlefield engineering will benefit, for instance buildings and utility / roads / bridges reconstruction will proceed more quickly with man and machine working side by side.

For the military, 'opposed environments' present special challenges. Logistical systems are expected to be deliberately attacked and there is significant merit during peacekeeping operations in exposing machines and not people to harm. This both reduces casualties and protects military tactical capability. Civilian situations can also be very dangerous; apart from disaster relief search and rescue in adverse conditions, responding to accidents where there are chemical, biological or fire hazards or fire would benefit greatly from autonomous logistical support.

The wider world, and the military, have to manage backwards compatibility and new systems have to work alongside old, safely and reliably. This can limit the ambition of systems as they evolve or require increasingly complex control mechanisms to handle multi-generational systems.

New ideas in research and development include the use of swarms / flocks / shoals of autonomous assistants. As the swarm size increases, the capacity of conventional control systems will be exhausted and quantum control systems, powered by QIP, will

be required. These will likely integrate complex sensor suites, in much the same way that 'smart' phones utilise multiple sensor modalities, to deliver their functionality and require QIP systems for control.

4.5.3 Medical Systems

Key Point 42: Quantum neural nets will revolutionise future delivery of medical care for all communities around the globe, including in hazardous situations

Revolutionary future end-to-end medical care system will begin with a quantum neural net based expert system diagnosing patients' ailments and subsequently 'smart' medical systems will autonomously, or semi-autonomously, move patients between care stations, monitor and operate machines delivering diagnostic and treatment functions and control logistical supply chains of materiel, ensuring that medical supplies are at hand when needed. Such innovations will significantly reduce costs making healthcare affordable for many for the first time.

Hazardous situations which endanger medical staff, for instance where patients have highly contagious diseases or in disaster relief areas, will benefit from autonomous systems and telemedicine will allow the unique skills of gifted physicians and surgeons to reach around the world. Secure, high-speed communications links will be a key enabler as will robot vision. Some complex but straightforward medical procedures, such as cataract surgery, could be delivered to communities which otherwise have very limited access to western medicine.

4.5.4 Domestic Systems

Key Point 43: Price is expected to be the principal constraint inhibiting the adoption of QIP in domestic systems. If the technical and ethical challenges can be overcome, self-driving vehicles would transform society

Unlike industrial automation where there is a capital multiplier effect, the authors believe that price will be the principal driver for uptake of QIP in domestic systems. Unless compact, low cost photonic quantum computers can be realised, QIP is unlikely to be encountered in the home environment. Currently, there are only a few robotic systems which have been widely adopted in domestic settings. Robotic lawnmowers, vacuum cleaners and domestic security systems (which automatically call the emergency services in the event of intruders or fire) are beginning to be adopted.

Long heralded, but technically and ethically challenging, autonomous cars are potentially a major market for compact QIP systems and would have a disruptive impact on society.

4.6 Future combat systems

Key Point 44: QIP could contribute to future combat systems through Network Quantum Enabled Capability (NQEC). There are challenges and issues which must be considered and resolved before the technology is available so that adoption will be as rapid as possible. The authors believe the principal technical challenges are machines' understanding of their environments, planning, and navigation. Other challenges include compatibility with military doctrine, health and safety concerns and regulations

Future combat systems will be ‘systems of systems’⁶⁰ which military planners have perceived will:

- Improve strategic capabilities;
- Increase battlefield effectiveness and survivability;
- Reduce logistics demands;
- Reduce through life equipment costs.

Real-time, network enabled capability (NEC) is central to future combat systems because it will allow individual military units to share information across the network and the commander to respond rapidly to changing battlefield conditions and co-ordinate the actions of unit under his command. NEC is most effective when orders are issued on the basis of well informed decisions and the network can seamlessly handle high data rates across the network without loss or corruption of data. In the limit, NEC can span a nation’s entire military and the information systems required are extremely complicated and require huge data storage and information processing resource.

Sections 4.1 – 4.5 have discussed how QIP can contribute to individual aspects of future combat systems which might be termed Network Quantum Enabled Capability (NQUEC). This section will consider briefly the challenges and issues which will need to be resolved for QIP to enhance successfully the performance of future combat systems.

The authors believe the principal technical problems are machines’ understanding of their environments, planning and navigation.

Information can be extracted from sensor data (**Section 4.3** outlined how this can be done for image data using current QIP) and the fusing of emerging quantum sensors into current sensor suites will be critically important. Electronic and image data will be augmented with data from gravitational, magnetic and other sensors and processed using quantum machine learning in the ways briefly described in **Section 4.4** to give situational awareness labelled with the degree of confidence allowed by the correlated data sets. Planning will use generalisations of quantum artificial intelligence (QAI) techniques being developed for image analysis (**Section 4.3.2**) and QAI may become an essential decision support tool at the strategic level while managing military resources down to the individual tactical unit. The movement of units will require accurate navigation systems, which for resilience, must operate without the need for GNSS. Although quantum enabled inertial navigation systems are being developed actively, the challenges for deployable systems are many and breakthrough developments will be needed if they are to be viable. Information security techniques described in **Section 4.2** will be essential.

Broadly, power will not be the inhibitor to the wide adoption of quantum enabled technologies but rather the overall system size and fragility arising from extreme sensitivity to electric, magnetic and gravitational fields as well as noise from platform vibrations. With only few exceptions, solid-state, chip-scale systems will be the technology of choice, benefitting from the high surface densities of micro- and nano-

⁶⁰ A ‘system of systems’ is a collection of dedicated sub-systems each of which has a specific role and which working together create a unique functionality and performance that is greater than the sum of the constituent sub-systems. Systems of Systems Engineering has been practiced, if not in name, for a long time but the advent of sophisticated control systems has created the possibility of more ambitious systems than were possible previously

scale components which can be achieved. The technology development will very likely emulate integrated electronic circuits through the 20th Century. Photonic quantum devices are, by their nature, integrated solid-state systems and photonic QIP is an R&D field to which, arguably, more resources should be directed.

Other than technology challenges, future quantum enabled combat systems must be compatible with military doctrine, health and safety concerns and regulations.

Current UK military doctrine requires 'man-in-the-loop' (low autonomy) systems wherever lethal force might be employed. Similar systems are mandated wherever there are health and safety (perceived or actual) concerns, for example, arising from the manipulation of dangerous materials or operating in dangerous environments.

With current UK doctrine, offensive systems would need to be low autonomy and the impact of QIP probably would be limited to individual parts of NQEC such as enhanced situation awareness, mission planning but not execution etc. Other activities which could be acceptable would include deployed logistics - resupply under fire, possibly medevac, surveillance and reconnaissance, self-protection e.g. of ships, bases etc. There is a need for organisations such as MOD's Development, Concepts and Doctrine Centre (DCDC) and Science and Technology (S&T) such as Dstl's Autonomy Programme, to extend work on autonomy to address any quantum specific issues so that there are few if any blockers to adoption when the NQEC becomes available.

4.7 Training & Simulation

Key Point 45: Computer based education and learning has been increasingly utilised since the 1950s for reasons of effectiveness and cost. Virtual Reality and AI technologies have added realism to training simulators and have been enabled by developments in neural nets running on CPUs and GPUs. Quantum neural nets will empower improved Training and Simulation technologies

For Defence and Security, efficiency and effectiveness are critical and, as with almost all its activities, the military has looked to technology to augment and enrich its training and education programmes. Just as some of the earliest computers were military computers, some of the earliest computer aided learning was developed by the military in the 1950s and were credited in 1988 with a key role in developing this technology.⁶¹

One of the earliest examples was PLATO (Programmed Logic for Automated Teaching Operations) which was designed for the presentation of instructional material and was revolutionary for its time in the use of digitised graphics and primitive animation displayed on a plasma screen alongside text. Developments in computer-based learning were closely linked to both the development of the hardware (through Moore's law and regularly increasing computer speeds) and of artificial intelligence (which allowed computers autonomously to create instructional material on demand and in a near-conversational manner). Pioneered by Uttal and Carbonnell in the 1960s and 1970s, the technology came to be known as 'intelligent tutoring systems' and, with later attention to portable systems, has hugely reduced the costs and time for training delivered anywhere from classrooms to battlefields. The current 'state of the art' is a

⁶¹ After a review, the US Congressional Office of Technology Assessment stated that "The military has been a major, and occasionally, the major player in advancing the state-of-the-art ... without [military research and development] ... it is unlikely that the electronic revolution in education would have progressed as far and as fast as it has" 'Power On! New Tools for Teaching and Learning' (OTA-SET-379, 1988) p. 158

collection of standards and specifications for internet-delivered eLearning (Sharable Content Object Reference Model⁶²) which has been almost ubiquitously adopted.

In the same way, the military has been developing computer-based simulation as an instructional technique to represent the visual, auditory, haptic and olfactory sensations of the operational world. Computer-based learning concentrates on teaching whereas computer-based simulation aims to enable learning through interaction with 'real world' experiences. In addition to cost and time benefits, such training can be delivered in any weather and in complete safety; for instance, a novice pilot 'crashing' a fast jet simulator walks away uninjured.

Training for tasks, especially those dubbed 'incredibly complex', must compress years of on-the-job experience into very short periods of time. Realistic simulation of surroundings and events is essential and has proved to be especially effective for those less comfortable with a traditional academic approach to learning. Examples include training sonar operators, avionics engineers and medical personnel and the technology is being widely adopted.⁶³ It has proved valuable in learning how to use operational procedures and tactics to make command decisions in confused and time-pressured environments. An early example of teaching success was the training of military jet pilots in combat situations where many multimodal stimuli must be interpreted and prioritised to create a plan of action in real time while performing continuing to fly the aircraft, attack targets and execute complex avoidance manoeuvres to avoid missiles arriving from anywhere around the aircraft.

Simulation has always been an important tool, but the added urgency and realism from virtual reality rapidly accelerates learning, reducing the time from novice to ace. The technology has benefited greatly from commercial gaming technologies which have become highly sophisticated beginning in the 1980s. As well as virtual reality, capabilities today include face and voice recognition, control by gesture, '4K' display technology, wearable technology and augmented reality; the augmentation can include visual, auditory, haptic, somatosensory and olfactory modalities. For instance, combining air temperature control with background sounds and mixtures of volatile odourants released into the air, the virtual reality experience of being beside the sea can be significantly enhanced. In pace with civilian gaming technology development, the UK MOD has adapted and adopted the technology and, most recently, has been trialling a new virtual reality training platform based on the same gaming engine as Fortnite.⁶⁴

In part, these advances have been made possible through special purpose chips, such as GPUs and cloud computing (see **Appendix B.8.3**), providing faster processing speeds but the principal enabler has been artificial intelligence and there are many commercial tools, some of which have been available for some time including Microsoft's cloud based Azure which was first released (as Azure after earlier products) in 2014.

⁶² https://en.wikipedia.org/wiki/Sharable_Content_Object_Reference_Model

⁶³ The technology is particularly successful for teaching anatomy in medical schools removing the need for hazardous and expensive cadaver sourcing, preparation, care and disposal

⁶⁴ <https://www.gov.uk/government/news/gaming-technology-trialled-in-training-uk-armed-forces>

Google released 'Tensor Flow Quantum'⁶⁵ (TFQ) in March 2020 which is an open-source library for the rapid prototyping of quantum machine learning (QML) models, analogous to 'Tensor Flow' released in 2017 which runs on CPUs and GPUs.⁶⁶ The authors expect TFQ will provide a step change in Training and Simulation technologies.

TFQ integrates an open-source framework for NISQ algorithms (Circ⁶⁷) with Tensor Flow and can represent and manipulate quantum data (which exhibits superposition and entanglement and is described by joint probability distributions that potentially needing exponential classical computational resources to process or store). Such data is noisy and typically entangled before measurement but QML can maximise the useful information which is extracted. TFQ provides primitives for processing the data and identifying correlations.

TFQ also uses the concept of hybrid quantum / classical models; these are mandatory because the limitations of near term NISQ processors (qubit numbers, connectivities, qubit coherence lifetimes etc.) require they work in conjunction with classical computers. Google claim TFQ is a natural platform because TensorFlow already supports working across multiple computer platforms such as CPUs and GPUs.

TFQ contains the basic elements required for quantum computations and user-defined calculations and can be executed on simulators or real hardware. It has been used for quantum-classical convolutional neural networks, machine learning for quantum control, quantum dynamics, generative modelling of mixed quantum states and 'learning to learn' with quantum neural networks via classical recurrent neural networks.⁶⁸ More information is given in **Appendix H**.

⁶⁵ <https://ai.googleblog.com/2020/03/announcing-tensorflow-quantum-open.html>

⁶⁶ In 2016, Google announced its Tensor processing unit (TPU), a custom chip, targeted at machine learning and tailored for TensorFlow. The TPU is a programmable AI accelerator for high throughput of low-precision (8-bit) arithmetic and intended for using neural net models rather than training them. Google have declared an order of magnitude better-optimised performance per watt for machine learning applications

⁶⁷ <https://ai.googleblog.com/2018/07/announcing-cirq-open-source-framework.html>

⁶⁸ <https://arxiv.org/abs/2003.02989>

5 A strategy for UK Defence and Security capability in QIP

Key Point 46: In the UK since 2014, government and other investment totalling about £1B has ensured the UK is world leading in the development of quantum technologies and aims to build a future sovereign quantum manufacturing sector. In QIP, the National Quantum Computing Centre (NQCC) will accelerate the development of low TRL R&D and produce prototype quantum hardware and software. Although the UK has a strong (conventional) computer software sector, which is expected to diversify into quantum software, it lacks a large computer systems integrator which may inhibit growing a QIP industry with full-stack capability. The recent Industrial Strategy Challenge Fund (ISCF) Wave 3 Quantum Technology Challenge, in part, seeks to rectify this situation but gaps remain in the NQTP QIP technology portfolio. Modest investment by MOD (about £5M for an initial 5-year programme) would address these gaps benefiting many of its business functions and providing disruptive advantage in some areas

5.1 Quantum Information Processing

Key Point 47: At the fundamental level, all information is quantum in nature and very different to the classical information processed by digital computers. Quantum physics clearly identifies the advantages of processing quantum information using a quantum processor including the ability to solve some problems much faster than digital computers. For many years, building such a quantum processor has been an elusive prize but functioning prototypes are evolving at increasing rates. Era 1 (2020 – 2025) offers the potential to identify early applications and will be a stepping-stone to fully scalable machines. Era 1 is a critical time for business entities to carefully consider QIP investment strategies

Classical information (such as text or speech or video which can be digitised) may be represented in a digital form and manipulated by digital computers. Quantum information⁶⁹ is fundamentally different to classical information. Information is represented by qubits which may be manipulated by quantum information processors (new types of analogue computers) which can solve some problems much faster than a digital computer by exploiting the quantum nature of qubits.

However, quantum information is fragile and current qubit technologies rapidly lose quantum information held in them (a process called decoherence which is caused by external electromagnetic fields and mechanical vibrations) and it is proving to be an enormous scientific and engineering challenge to build a quantum computer of sufficient size to have real value. Different types ('platforms' the most promising of which use superconducting circuits or ions held in space by electromagnetic fields⁷⁰ as physical qubits) of QIP are the focus of intense R&D to realise engineered systems of large numbers of qubit which can carry quantum information for long times.

This document refers to the periods 2020 – 2025 as Era 1, 2025 – 2030 as Era 2 and after 2030 as Era 3. It is not yet clear which platform technology is superior and small prototype systems of all types are appearing during Era 1⁷¹ and technology down

⁶⁹ See **Section 2.2** for a more detailed discussion

⁷⁰ An ion is a neutral atom which has lost or gained an electron. Promising platforms include Ca⁺ and Yb⁺

⁷¹ Systems comprising no more than 72 qubits have yet been demonstrated with the exception of the D-Wave machine; this is a quantum annealer in development since 2009 and roughly

selection is expected to happen in late Era 1 or early Era 2. These early machines – known as Noisy Intermediate Scale⁷² Quantum (NISQ) computers - offer the potential to identify early applications and will be a stepping-stone to fully scalable machines. Era 1 is a critical time for business entities to carefully consider QIP investment strategies.

The UK National Quantum Technology Programme (UKNQTP), which has been running in the UK since 2014, has ensured the UK is a world leader in many aspects of QIP. (The total UK investment in all quantum technologies since 2014 is now about £1B; for more details about the UK and other leading programmes see **Appendix I**). The objective of the NQTP is to exploit decades of Research Council investment in basic quantum physics to develop a world class quantum industry from which the UK derives economic, societal and National Security benefit.

5.1.1 Opportunities and threats – the case for Government investment

Key Point 48: QIP capabilities represent significant opportunities and threats, especially for Defence and Security, and these are sufficiently significant and novel that organisations need to explore applications now to be ‘quantum-ready’ for the future. It is expected to take years to build capabilities and identify useful applications and it will be difficult for organisations that have not engaged early on to catch-up

Quantum computing systems (hardware and software plus services) promise significant economic benefit - of the order of \$100Bs globally in the next few decades, comparable in magnitude to artificial intelligence (AI).⁷³ This value will be reaped by those developing these technologies, their components, and sectors benefitting from their application (pharma, health, logistics, IT, energy, chemicals, finance as well as defence and security and others).

Estimates suggest it will take at least 10 years before the full scope of the value chain becomes clear during Era 2 but significant benefit is still expected within 10 years (in financial terms, probably £10Bs).⁷⁴ The likelihood of finding useful applications increases as Era 3 approaches and fully scalable, fault-tolerant machines, mature enough to run a range of applications, become available. At any point during Eras 1 to 3, the identification of valuable applications is likely to dramatically accelerate the demand for and, subsequently, availability of quantum information processing systems.

QIP capabilities represent significant opportunities and threats especially for Defence and Security. Early stage systems (principally quantum annealers – see **Appendices B.7** and **C.5**) are already being used by government and industry laboratories and work is underway⁷⁵ to apply QIP to defence-relevant optimisation problems (e.g. of communication networks) and develop machine learning solutions (e.g. to analyse images). Promising results are being obtained even with the small circuit model NISQ

following a ‘Moore’s Law’-like technology development. A 5640 (superconducting) qubit machine is expected to be released in 2020. See **Appendices B.7** and **C.5**

⁷² ‘Noisy’ because the qubits comprising these machines lose information due to the influence of external electromagnetic fields or vibrations (‘noise’); ‘Intermediate Scale’ because the machines comprise only relatively small numbers of qubits

⁷³ <https://www.bcg.com/publications/2018/next-decade-quantum-computing-how-play.aspx>

⁷⁴ The estimates here and elsewhere in the text were made before the Covid-19 emergency which began in early 2020

⁷⁵ https://www.dwavesys.com/sites/default/files/D-Wave_Webinar_280519.pdf

platforms available today⁷⁶ and when fully scalable, fault-tolerant machines become available, applications are widely expected to include materials modelling to improve equipment design and rapid analysis of big data through a step change in the capabilities of artificial intelligence systems.

Large-scale quantum computers are also anticipated to threaten the integrity of many of the current encryption techniques, putting the UK economy's secure data and communications at risk.⁷⁷ This means Governments will need to maintain some level of sovereign or assured capabilities in order to understand and mitigate against the threat.

Taken together, these opportunities and threats are significant enough and novel enough that organisations, including Government, need to explore applications now to be 'quantum-ready' for the future. It is expected to take years to build capabilities and identify useful applications and it will be difficult for organisations that have not engaged early on to catch-up.

5.1.2 The global technology race

Key Point 49: QIP is in the early stages of development and the dominant hardware platform is still not clear. State actors and companies are investing heavily to attempt to ensure early advantage. As with current digital technology, algorithms critically important and some can be executed on the NISQ machines expected to be available during Era providing a window of opportunity to accelerate progress and shape developing markets

Remembering that the current state of digital computing has been reached only after eighty years of continuous R&D,⁷⁸ QIP is still in the early stages of technological development and the dominant hardware platform is still not clear.⁷⁹ The cutting-edge science and engineering required is attracting some of the brightest and best scientists and engineers but rapid progress will require significant scientific discovery to be closely coupled to the solution of engineering problems (of a 'Grand Challenge' type) in many different areas.

State actors and companies are investing heavily to attempt to ensure early technological and economic advantage. Global private and state investment is estimated to be of the order of \$100Ms pa and rising in countries such as the US,

⁷⁶ Proof of principle results for the solution of non-linear partial differential equations using the IBM 20 qubit Poughkeepsie NISQ platform have been described, see <https://arxiv.org/pdf/1907.09032> (<https://journals.aps.org/pr/abstract/10.1103/PhysRevA.101.010301>)

⁷⁷ Activities are already underway across Government to mitigate against this threat. These are not considered within the scope of this document.

⁷⁸ In a paper titled 'On Computable Numbers' published in 1936, Alan Turing wrote down the principles for a programmable 'Universal Machine' which could be programmed to solve any problem and, with Gordon Welchman, built some of the world's first (electromechanical) digital computers ('Bombes') at Bletchley Park during World War II although the work was classified. At the same time, in Germany Konrad Zuse was building similar machines and is credited with the demonstrating the world's first programmable, fully automatic digital computer, the Z3, which had 2000 relays and 22-bit words, operated at a clock frequency of about 5–10 Hz and could carry out arithmetic using floating point numbers (see Konrad Zuse, 'Der Computer. Mein Lebenswerk' ('The computer. My Life's Work'), 3rd edition, Springer-Verlag, 1993)

⁷⁹ See **Appendix B** for a discussion of the principal quantum processing hardware currently being developed as well as emerging digital technologies which may challenge the capabilities of early quantum hardware platforms

Canada and Germany (who invested €650M over 2020-22).⁸⁰ China is reported to have invested in the order of £1B pa to catch-up on its counterparts.⁸¹

The experts' view is that the race to produce QIP hardware from which significant benefit will be derived is far from won and technology and systems readiness levels remain mid-scale.⁸² None have yet achieved unassailable technological superiority or demonstrated essential applications.⁸³

However, as with current digital technology, much benefit lies with algorithms. Some algorithms can be executed on NISQ machines while others require more sophisticated hardware; the algorithms described in **Appendix C** are believed to have value for Defence and Security (and other business entities) which could be realised during Era 1 (up to 2025). This presents a window where intervention to accelerate progress will shape the developing market and value creation.

5.2 A commercial sovereign computing capability

Key Point 50: The UK NQTP is currently a diverse ecosystem of funded R&D, supported technology development in industry and other initiatives including the development of a National Quantum Computing Centre. This has created world class capabilities in QIP and determined efforts are being made to establish a sovereign, full-stack capability

The UK NQTP is currently an ecosystem that includes funded R&D calls, supported industry technology development calls, development of a National Quantum Computing Centre (NQCC), funding for skills and training initiatives, plus numerous academic, industry and international partnerships.

Two QIP platforms, built using trapped ion and superconducting qubits, lead the technology race; the UK is world leading in the former and well positioned in the latter to achieve a globally-leading position by the end of Era 2. Additionally, there is UK capability in other platforms (such as photonics-based QIPs) that could overtake the current leaders or attain significant market share by fulfilling certain requirements (such as a critical need for low size, weight and power).

The UK also has world-leading strengths in software and algorithm development but to maintain the momentum which has been achieved it is critically necessary to ensure future software developments are closely coupled with hardware evolution throughout the software development cycle. Application software is built from algorithms and as

⁸⁰ This is likely to be a significant underestimate, as information on the level of investment made by the IT majors is not publicly available.

⁸¹ While China's overall spend on quantum computing is unknown, the government is investing \$10B in building the world's largest quantum research facility in Hefei. Alongside this, the number of Chinese patents and applications filed in relation to quantum computers has rapidly increased since 2014.

⁸² Technology and system readiness levels (TRLs and SRLs) are used to summarise the maturity of technologies and procurement projects during the development of complex systems. TRLs (which differ slightly between organisations such as NASA, MOD, NATO etc) are more easily quantified and range from 1 (basic scientific principles demonstrated) to 9 (mission-proven technology in routine use). A TRL of about 5 (at which prototype systems have been demonstrated in a relevant environment) indicates the maturity at which development focus typically begins to transfer to industry from research laboratories

⁸³ **Appendices C, D and E** survey the quantum algorithms from which many believe most benefit will be derived from small NISQ machines (Era 1, **Append C**), larger NISQ machines (Era 2, **Appendix D**) and fully scalable, fault tolerant quantum computers (Era 3, **Appendix E**). Applications comprise the use of one or more algorithms to solve specific problems

algorithms mature, hardware re-design may be the key to more efficient – ‘better’ - solutions.

Thus, the UK is well-placed to maintain a leading global position, derive economic and Defence and Security benefit from previous and future investments and ensure future sovereign or assured capabilities, only provided targeted investment continues and keeps pace with global endeavours. Enabling interactions within the QIP community (for instance between research laboratories and industry, software and hardware communities or software developers and end users), will facilitate the development of early machines in the UK will be key to future success.

The UK has significant and engaged user communities, but these users must start building suitably qualified and experienced personnel (SQEP) capabilities now to enable effective exploration of useful applications during Era 1 and build resilience for the future. To meet these needs, users should engage with the NQCC through a programme that is tailored to build SQEP and business readiness as well as links to the QIP community if user needs are particularly niche.

Software and hardware development are closely coupled in quantum computing and necessary in order to explore application areas and speed-ups which could potentially bring forward the date of useful applications. Enabling this co-development will be a key requirement of a future programme, as well as access to simulators and annealers, crucial for the development of compilers and software and skills development.

The UK does not currently have an established integrator that has prototype machines to access. For these reasons it is likely that enabling early access to machines or emulators will mean developing relationships with the IT majors and other major players to complement the UK ecosystem strengths.

The UK has a thriving ecosystem of hardware and software spin-outs and SMEs focussed on developing QCs or exploring applications, built out of the UK’s academic centres of excellence and the NQTP. The UK also has a strong advanced manufacturing base that is already selling components like control systems and lasers to the globally emerging market. The NQCC could provide a focal point to bring these activities together and scale them by seeking to build demonstrator devices based as much as possible on UK-sourced components. This highly interconnected and commercially focussed ecosystem is a significant UK strength and has attracted numerous companies (including IBM, Rigetti, SeeQC, and Google) who have set up UK based activities to benefit from the skills and know-how.

The UK quantum computing ecosystem is still embryonic and fragile and it needs to be stimulated further to reach critical mass given the times to market and the current weak market pull. End users are engaged but not yet willing to invest sufficient resource to develop the full quantum stack⁸⁴ required to ensure QIP reaches the market space and creates a sustainable UK commercial sector. Non-government investment is needed to develop the products and customer demand for systems and services is essential to sustain the new quantum information sector and the British Business Bank is working to de-risk investment opportunities for UK investors and allied states but as the market develops, particularly in the current context, to reconsider whether a quantum-specific fund is required to achieve the scales of funding required at series B and onwards.

⁸⁴ The ‘Full Quantum Stack’ is the full range of quantum computing technologies from the highest product, services, through applications, software systems & assurance, systems integration & scaling, and qubit gates and their control and readout

5.3 Targeted UK translational quantum computing research

Key Point 51: Mid TRL, translational QIP research is supported by the Oxford-led Quantum Computing and Simulation (QCS) Hub

In December 2014, Phase 1 of the UKNQTP established a flagship research entity, the Networked Quantum Information Technology (NQIT) Hub which became the Quantum Computing and Simulation Hub (QCS) in the Phase 2 UKNQTP which began on 1st December 2019.

NQIT encompassed nine universities (Bath, Cambridge, Edinburgh, Leeds, Oxford, Southampton, Strathclyde, Sussex and Warwick) and had connections to five other universities not formally Hub partners (Heriot-Watt, Bristol, Durham, Imperial College London and Sheffield). In addition, NQIT worked with more than 30 commercial companies (including IBM, Lockheed Martin, Raytheon BBN, Google and Toshiba) and government organisations (including the UK's National Physical Laboratory (NPL), Dstl and the US's NIST) plus small and medium-sized enterprises (including Rohde & Schwarz, Covision and Oxford Instruments). The ambitious goal was to understand how to build a universal, scalable quantum computer with error correction. In Phase 1, NQIT focused on ion trap, photonic, solid-state and superconducting platforms as well as quantum algorithm development.

The Phase 2 Hub for Quantum Computing and Simulation is continuing NQIT's work, broadening the consortium to 23 research teams in 16 universities and engaging with 35 commercial and government organisations.⁸⁵ The programme is focussing on:

- Simulation, especially focused on materials discovery;
- NISQ platform development to demonstrate, within the Phase 2 Hub, super-classical performance in areas of relevance to users outside the quantum technology field;
- Universal, scalable, fault-tolerant quantum computer development for general purpose applications.

5.4 Quantum computing sovereign capability development

Key Point 52: IUK is supporting commercialisation of QIP through Wave 3 of the Industrial Strategy Challenge Fund and the Department for Business, Energy and Industrial Strategy is leading a programme to establish a National Quantum Computing Centre which will accelerate translation of QCS Hub R&D into commercialisable technology

In February 2019, IUK announced that the 'Commercialising quantum technologies' challenge had been shortlisted⁸⁶ for funding through Wave 3 of the Industrial Strategy Challenge Fund (ISCF)⁸⁷ and the Autumn Statement of 2019 announced that up to

⁸⁵ <https://gow.epsrc.ukri.org/NGBOVViewGrant.aspx?GrantRef=EP/T001062/1>

⁸⁶ <https://innovateuk.blog.gov.uk/2019/02/05/industrial-strategy-challenge-fund-wave-3-shortlist/>

⁸⁷ Conceived in 2016, the Industrial Strategy Challenge Fund is part of government's Industrial Strategy which aims to raise productivity and earning power in the UK, see <https://www.ukri.org/innovation/industrial-strategy-challenge-fund/>. Funding, currently totalling £4.7B is being released in Waves. The first ISCF funding for quantum technologies comprised 4 'Quantum Pioneer' projects, see <https://www.ukri.org/innovation/industrial-strategy-challenge-fund/quantum-technologies/>. Part of ISCF Wave 2, the 2 year projects, led by industry, who provided funding matching the IUK investment of £20M, were announced in

£153M would be made available by government provided this was matched by at least £205M from industry. The purpose of the ISCF QT Challenge funded projects is to advance readiness levels beyond the Hub demonstrators and de-risk the transfer of technology to industry thereby accelerating the development of pre-production prototypes and commercial products.

The Department for Business, Energy and Industrial Strategy (BEIS) is leading a programme to establish a National Quantum Computing Centre (NQCC) as part of Phase 2 of the NQTP. The NQCC⁸⁸ has a key role to play building the UK's sovereign QIP capability. Announced in the 2018 Autumn Statement and based at Harwell, the NQCC will be a dedicated national centre whose aim is to develop commercially viable, fully scalable, fault tolerant, general purpose quantum computing hardware, software and applications. It is expected to be fully operational by summer 2021 and deliver a NISQ computing capability that, for a range of tasks, outperforms conventional computers by 2025. The initial focus will be developing NISQ machines to demonstrate technologies, give assured and direct access to developers and drive the formation of a sovereign quantum computing supply chain. An onshore, large computer manufacturer which carries out the necessary systems engineering to produce an operating quantum computer is regarded as the ideal model to successfully create a sovereign quantum computing hardware manufacturing sector and establishing such an organisation is a key part of the NQTP strategy.

The UK has a strong record in developing and delivering conventional computer software (London is sometimes called 'Silicon Roundabout' in acknowledgement of this) and has a number of strong research groups developing quantum algorithms. The NQTP Phase 2 Oxford Hub includes more quantum algorithm development work than in Phase 1, but it is essential that industry collaboration is strongly encouraged and thrives. Fortunately, there are signs that this is happening; for instance, the UK has a very strong record in spin-outs and start-ups developing quantum software.

5.4.1 ISCF Wave 3 QIP projects

Key Point 53: In 2020 IUK has made 10 grant awards worth £25.7M for QIP projects under the first Wave 3 ISCF Quantum Technologies Challenge call which address technologies across the full quantum computing stack

The ISCF Wave funding is being released through two calls. The first competition, worth £75M across all quantum technologies (timing, sensing, imaging, communications and computing and simulation), was concluded in the first half of 2020 but at the time of writing (May 2020) IUK has not made public the full competition results.

IUK funded 3 types of project which had to be industry-led:

- Feasibility studies (FSs) providing up to £500K of grant award for projects of duration up to 18 months to carry out proof of concept work to validate novel ideas;

November 2018 and are developing a quantum gravity sensor, a miniature quantum clock and two quantum encryption systems for secure data transmission

⁸⁸ <http://uknqt.epsrc.ac.uk/about/nqcc/>

- Collaborative R&D projects (CRDPs) providing £2 – 10M of grant award over a period of up to 36 months to deliver new products or services, with a focus on end users and spanning the full supply chain;
- Technology projects (TPs) providing £4 – 10M grant award over a period of up to 36 months allowing industry to work together on technology challenges facing commercialisation;
- Investment Accelerators (IAs) providing a mechanism to leverage IUK investment with co-investment by venture capital companies.

The QIP Challenge attracted intense interest and start-ups and spin-outs, of which the UK has many, are expected to be a strongly represented in funded projects developing both software and hardware. £25.7M of grant funding was awarded to 10 QIP projects which, together, span the full quantum computing stack.

Six FSs were funded in total investigating innovative ideas relevant to photonic, silicon and superconducting platforms. Three will be addressing software issues including software systems and assurance, applications and quantum computing services. One will address the development of hardware and two will be focussed on critical underpinning technologies. Taken together, these proof of concept projects could lead to step changes in these three platform technology areas.

One CRDP was funded which addressed the full quantum stack.

Three TPs were funded. One, which will be applicable to all of the principal QIP platforms, will focus on the development of system software and applications and is expected to be closely coupled to hardware implementations of qubit control and input and output of data. Two will address the challenges of engineering scalable hardware for four platform technologies; one project will target superconducting platforms while the other will have applications in trapped ion, neutral atom and photonic platforms which share the need to generate, distribute and detect quantum light to operate their qubits).

5.4.2 The current ISCF Wave 3 quantum computing portfolio

Key Point 54: Hardware projects span the leading platforms and address key challenges including systems engineering and scalability. Software projects address qubit control, operating systems (including for hybrid digital / quantum machines) and application software

Hardware projects span platforms based on ions, neutral atoms, photons, superconductors and silicon but there are none addressing NV centres⁸⁹ or quantum annealing. The practicalities of systems engineering and scalability, of particular concern for superconducting platforms, are being addressed as are a number of essential enabling technologies such as generation of quantum light on demand and its detection.

⁸⁹ NV centres are point defects in diamond; pairs of carbon atoms in the diamond's crystal lattice are replaced by nitrogen atoms (N) plus an adjacent empty lattice site (vacancy V). For technology applications, negatively charged NV centres, formed by applying a voltage to the crystal), are the subject of numerous R&D programmes developing sensitive magnetic and mechanical stress sensors, bioimaging schemes, masers, quantum communications and computing

Software projects span low level control of qubits and quantum processors, foundational work towards a universal compiler and operating system for hybrid quantum/digital computers⁹⁰ and the development of application software for materials design with the ultimate aim of addressing important societal challenges such as climate change.

5.5 A QIP Strategy for Defence and Security

Key Point 55: ISCF Wave 3 projects funded during 2020 span a broad range of QIP technologies not including Quantum Neural Nets (QNNs). These have been intensively studied and could benefit all business enterprises, especially Defence and Security. If action is not taken now, it is possible that the UK may be left behind in this important area. The second tranche of ISCF Wave 3 funding, expected in 2021, could support the development of QNNs for machine learning and managing complex systems

Section 3.3 identified the potential for quantum neural nets (QNNs) to revolutionise machine learning and artificial intelligence. A QNN is a neural net⁹¹ which is executed on a quantum computer. QNNs are tolerant of ‘noise’ and full connectivity of the qubits in the NISQ processor is not essential (although probably desirable).

The development of QNNs for pattern matching applications has attracted significant investment especially for exploitation of the D-Wave quantum annealing machines. Some first reports⁹² described the use of QNNs to classify and search imagery and ultimately identifying features of interest, detecting anomalies and instances of change is expected to be possible almost in real time even with the D-Wave processors available during Era 1.

Other areas where QNNs are expected to be valuable, in addition to pattern matching, include machine learning, artificial intelligence, financial technology, control and validation of complex systems (such as autonomous vehicles) and together these span many critical activities in Defence and Security and early adoption could provide significant, potentially disruptive, increases in capabilities during Era 1 with only limited investment required now. Efficiencies in the deployment of manpower and associated resources could be realised, potentially releasing budgets now used for routine tasks to be reallocated to other priorities. Additional benefits include building a ‘quantum ready’ workforce of suitably qualified and experienced personnel (SQEP) who would allow organisations to benefit from the increasingly powerful QIP platforms, algorithms

⁹⁰ Some tasks such as arithmetic and input/output will not be done well by quantum computers and it is expected that, at least initially, practical systems will comprise hybrids in which the quantum processor acts as a specialised co-processors for specific tasks which they perform much better than digital computers. This is already done with digital machines where coprocessors can perform floating point arithmetic, graphics, signal processing, string processing, cryptography or interface with peripherals

⁹¹ The software is based on an algorithm called a ‘neural net’, so named because it has similarities to the operation of biological neurons. Neural nets are a mature and proven method of pattern-matching, but they impose extremely high loads on a classical digital computer architecture. However, by being intrinsically parallel, quantum annealers can execute a neural net in one machine cycle instead of thousands or millions. Therefore, extra and potentially overwhelming quantum speed-up is clear, and we believe it is this ability to run pattern-matching neural nets that explains significant investment by Google, IBM, Intel and others.

⁹² See eg. Nguyen et al, <https://arxiv.org/pdf/1905.13215>

and applications which are expected to appear in the future, slowly in Era 1 but with increasing speed and diversity through Eras 2 and 3.

However, existing quantum computers are electrically noisy and of modest scale and many believe that scale and noise performance must improve radically before quantum computers can solve valuable problems; this belief encourages potential investors to wait. It is a challenge to show, unambiguously, that this is not true and that existing quantum computers can already solve valuable problems.

Stimulated by the progress in NISQ hardware, QNNs have been the subject of much research during the past two years,⁹³ but the NQTP currently does not address QNNs and unless action is taken now it is possible that this is an emerging important QIP area in which the UK may be left behind. If this happens, it will be unlikely that the National Programme will be able to provide the tools and trained personnel which would be needed by Defence and Security to investigate and exploit QNNs as an early adopter during Era 1. MOD would not be able to migrate selected tasks to quantum protocols and enjoy the benefits which would accrue.

However, £78M of grant remains to be awarded by IUK (attracting further industry investment) and the opportunity exists for Defence and Security together with the NQCC to lead on developing Challenges which could be funded through the next IUK call expected to be in January 2021. The most promising QIP application areas compatible with NISQ platforms available during Era 1 (quantum annealers and the emerging circuit model machines) include:

- Quantum machine learning;
- Software verification and validation;
- Modelling complex systems;
- Optimisation of complex systems performance.

Apart from technological progress, the experience gained as the Challenges progress will include:

- Leadership, vision and governance of QIP systems in Defence and Security;
- An understanding of technology readiness, likely system constraints and future user needs;
- Skills in transitioning science to technology and system engineering;
- Progressive programme and project management;
- Dynamic risk management;
- Management of intellectual property and know how;
- System engineering, interoperability and standardisation.

MOD should begin working with the QIP Challenge process immediately to be certain that it can engage usefully with consortia during the Challenge formulation and bidding and evaluation process and is suitably prepared to embrace the technologies developed by successful consortia.

⁹³ See e.g. <https://www.nature.com/articles/s41467-020-14454-2> and <https://ai.googleblog.com/2018/12/exploring-quantum-neural-networks.html>

5.5.1 An example NISQ-value challenge:

Key Point 56: An exemplar ISCF Wave 3 Challenge in QNNs could be the development of a NISQ algorithm to identify, localise and track arbitrary features in imagery data

Existing quantum annealers, such as the D-Wave machine, and circuit model quantum computers appear able to ‘compute’ Boltzmann networks. This is important because there is a mature signal processing / data analysis algorithm, called the ‘Neural Net’, which can be cast as one version of a Boltzmann machine. As discussed in **Section 3.3.2**, neural nets are a powerful method of performing image analysis (e.g. robot vision) and have been used for many years in automatic financial trading. Currently, they have to be ‘solved’ on digital computers which become slow when the number of neural net nodes becomes large, such as in image processing. This speed problem arises because digital computers are compelled to mimic parallel computation to ‘execute’ neural nets. Quantum computers, however, are truly parallel and compute all the node values simultaneously. This would appear to give quantum computers a huge intrinsic speed advantage, an advantage which potentially scales as the square of the number of nodes in the neural net.

The Challenge is, given an unclassified data set of images, to develop and implement on NISQ platforms, an algorithm that identifies, localises and tracks through a time series of RGB-coloured imagery, features of arbitrary pixel size. The algorithm should be able to identify and track features through the time series. The features may be partially obscured and present at different locations within the field of view, with different sizes and orientations. The algorithm must provide an estimate of the confidence of feature recognition and tracking. Implementation may be on physical hardware or emulators.

6 Summary and conclusions

Key Point 57: NISQ computers are available now and businesses should begin assessing the opportunities and threats expected from large scale machines expected to appear within the decade. Broadly the NQTP spans all QIP technologies but has no work on quantum neural nets (QNNs) which, on existing commercial machines, could create near-term ‘early wins’

Quantum computers are a novel type of analogue computer and have been anticipated for half a century. They will not be the near-magical machines which over-zealous supporters have promised but it is clear from many theoretical studies and computational simulations that there will be a number of tasks at which they will be overwhelmingly superior to digital computers. Commercial enterprises like Google, and other nations like China, have invested heavily in quantum computing anticipating the advent of large scale machines.

However the commonly held view that Quantum Computing will only be important in the distant future, and for niche reasons, appears mistaken. The view of the authors is that MOD needs to take action now as quantum computing has the potential for major impact on key military activities including intelligence analysis, logistics and autonomy, in only 5-10 years from now.

After coordinated and coherent R&D in the UK and elsewhere resulting in rapid progress over the past five years, emerging Noisy Intermediate Scale Quantum machines are now at a maturity level that business enterprises can begin assessing realistically the opportunities and threats presented by Quantum Information Processing and preparing to embrace the change in business models over the next decade which will inevitably follow as this disruptive technology matures by investing in building a suitably qualified and experienced workforce, exploring and developing where necessary quantum application software.

Broadly, the UK National Quantum Technology Programme (NQTP) is addressing the full stack of quantum computing (ie, hardware, software including low level control systems, algorithms and application software) but there are gaps. One such gap is the absence of NQTP supported work on Quantum Neural Nets (QNNs).

The ability of NISQ computers, including D-Wave, to run neural nets, a mature, versatile and extremely powerful Artificial Intelligence (AI) algorithm, possibly at extremely high speeds, is a game changer and could provide ‘quick wins’ before 2025.

Existing commercial D-Wave machines are expected to facilitate better analysis (with pace, precision, pre-emption and predictive power) and could be disruptive by enabling real time autonomous pattern matching tasks. In the longer term the same technology running on improved quantum machines could give strategic advantage by predicting intent at the nation-state scale and usher in a change in ‘business model’ for MOD.

If, as it appears, quantum computers can ‘break’ the problem of neural nets requiring extreme computing power when run on sequential digital computers, then a whole range of possibilities opens up. The combination of quantum computers and neural nets could provide true ‘computer vision’, where a computer can ‘understand’ and break down an image into content. Neural nets can process other classes of data pattern as well, are superb adaptive correlators and are a mainstay of Machine Learning AI research.

7 Recommendations

Key Point 58: MOD should work with NQTP Partners to formulate and propose a QNN Challenge to be supported by the tranche of ISCF Wave 3 funds expected to be released during 2021. MOD should also ensure it has adequate SQEP to derive full early-adopter advantage from the technologies developed through the QNN Challenge

MOD should begin to invest to build SQEP (Suitably Qualified and Experienced Personnel) in QIP (Quantum Information Processing) in preparation for the adoption of QIP machines into its business practices which could begin to happen as soon as 2021.

The first steps should be to work closely with NQTP Partners (National Quantum Technology Programme) to make practical assessments such as benchmarking studies of existing NISQ (Noisy Intermediate Scale Quantum) machines and their digital competitors.

MOD should also work with NQTP Partners, especially IUK, to formulate and propose a Quantum Neural Nets Challenge involving the use of quantum neural nets to solve problems of practical importance for Defence and Security. The same QNNs would have similar game changing value to many other business enterprises.

If adopted by IUK, the QNN Challenge would be supported as part of the second tranche of ISCF Wave 3 funds expected to be made available during 2021, MOD will need to begin urgently to prepare to engage usefully with consortia during the Challenge formulation and bidding and evaluation process and ensure it is suitably prepared to embrace quickly the technologies developed by successful consortia and to derive early adopter advantage.

8 Acknowledgements

This Landscape document has benefited greatly from advice and comments from many key individuals across the UK to whom the authors extend their gratitude. Of particular note are:

Roberto Desimone	BAE Systems (Applied Intelligence)
Viv Kendon	Durham University
Sir Peter Knight	Chair, UKNQTP Strategic Advisory Council
Roger McKinlay	ISCF Challenge Director Quantum Technologies
Prof. D Paul	Glasgow University

9 List of abbreviations

AES	Advanced Encryption Standard
AI	Artificial Intelligence
BBC	British Broadcasting Corporation
C3	Command, Control and Communications
C4ISR	Command, Control, Communications, and Computers, Intelligence, Surveillance, and Reconnaissance
CAD	Computer Aided Design
CAS	Chinese Academy of Sciences
CFI	Canada Foundation for Innovation
CFREF	Canada First Research Excellence Fund
CIFAR	Canadian Institute for Advanced Research
CMD	Cyber Mimic Defence
CMOS	Complementary Metal-Oxide-Semiconductor
CREST	Core Research for Evolutional Science and Technology
DARPA	Defense Advanced Projects Agency
DCDC	Defence, Concepts and Doctrine Centre
DDQCL	Data-Driven Quantum Circuit Learning
DSS	Decision Support System
DES	Data Encryption Standard
DFT	Discrete Fourier Transform
DL	Deep Learning
Dstl	Defence Science and Technology Laboratory
EPSRC	Engineering and Physical Sciences Research Council
FPGA	Field Programmable Logic Array
FRQI	Flexible Representation of Quantum Images
GNSS	Global Navigation Satellite System
GPU	Graphics Processing Unit
IBM	International Business machines
IDEA	International Data Encryption Algorithm
IoP	Institute of Physics
ISCF	Industrial Strategy Challenge Fund
IUK	Innovate UK
LANL	Los Alamos National Laboratory
LNISQ	Low Noise Intermediate Scale Quantum
MIT	Massachusetts Institute of Technology
ML	Machine Learning
MOD	(UK) Ministry of Defence
MOS	Metal-Oxide-Semiconductor
MV	Machine Vision
NATO	North Atlantic Treaty Organisation
NEC	Network Enabled Capability
NEQR	Novel Enhanced Quantum Representation
NIST	National Institute of Standards and Technology
NISQ	Noisy Intermediate Scale Quantum (computer)
NPL	National Physical Laboratory
NQCC	(UK) National Quantum Computing Centre
NQEC	Network Quantum Enabled Capability

UK OFFICIAL

NQI	National Quantum Initiative
NQIT	Networked Quantum Information Technology
NQL	National Quantum Laboratory
NQTP	(UK) National Quantum Technology Programme
NRC	National Research Council
NSA	(US) National Security Agency
NSERC	Natural Sciences and Engineering Research Council of Canada
OCR	Optical Character Recognition
PI	Principal Investigator
PLATO	Programmed Logic for Automated Teaching Operations
PNT	Precision Navigation and Timing
PT	Population Transfer
QA	Quantum Annealing
QAI	Quantum Artificial Intelligence
QAE	Quantum Auto-Encoder
QAOA	Quantum Approximate Optimisation Algorithm
QBIP	Quantum Boolean Image Processing
QCS	Quantum Computing and Simulation (UK NQTP Hub)
QFT	Quantum Fourier Transform
QIC	Quantum Image Compression
QIP	Quantum Information Processing
QKD	Quantum Key Distribution
QML	Quantum Machine Learning
QPU	Quantum Processing Unit
QRMW	Quantum Representation of Multi-Wavelength images
QST	Quantum State Tomography
QuAIL	(NASA) Quantum Artificial Intelligence Laboratory
QUBO	Quantum Unconstrained Binary Optimisation
QulmP	Quantum Image Processing
R&D	Research and Development
RBM	Restricted Boltzmann Machine
RGB	Red-Green-Blue
RISC	Reduced Instruction Set Computer
RQTR	Russian Quantum Technologies Roadmap
RSA	Rivest-Shamir-Adleman
SAPI	Server Application Programming Interface
S&T	Science and Technology
SL	Semantic Learning
SME	Small and Medium-sized Enterprises
SQEP	Suitably Qualified and Experienced Personnel
SQUID	Superconducting QUantum Interference Device
STFC	Science and Technologies Facilities Council
SVM	Support Vector Machine
TCP/IP	Transmission Control Protocol/Internet Protocol
TFQ	Tensor Flow Quantum
TRL	Technology Readiness Level
UK	United Kingdom
UKNQTP	UK National Quantum Technologies Programme
UKRI	UK Research and Innovation

US
VQE

United States of America
Variational Quantum Eigensolver

APPENDIX A An introduction to quantum computers

A.1 Data representation within classical and quantum computers

A classical digital computer has access to data stored in memory which exists only in a finite set of discrete states. The fundamental unit of memory is the bit (binary digit), which can take values 0 or 1, and a state of the memory is represented as a unique string of these bits (physically, in a computer, these bits are represented as easily distinguished voltages corresponding to ‘on’ and ‘off’ states of transistors). Any information – numbers, text, sound or images – can be represented by a sufficiently large collection of bits.

By analogy a quantum computer represents data by a set of quantum states. The simplest states, which can only be properly described by quantum physics, are two level states called qubits (quantum bits). Examples include the spin-up and spin-down states of an electron and the polarization of a single photon, in which the two states can be taken to be the orthogonal polarisations (horizontal- and vertical- or left- and right- polarised). A classical system can exist only as one state or the other but quantum physics allows a qubit to be a superposition of both states. It is this fundamental property which makes quantum computers more powerful than their classical equivalents. Thus information – again, numbers, text, sound or images – can be represented by a sufficiently large collection of qubits.

The *exact* state of a qubit is not directly accessible, although a measurement will return either a ‘0’ (North pole) or a ‘1’ (South pole) with a probability depending on the square of the angle that the original vector makes with the ‘equator’ (see Figure 2). The longitude defines the phase of the wave-function. Qubits hold numerically continuous (not digital) variables. Once a measurement has been made, the quantum nature of the original information cannot be recovered and the readout from a qubit is a binary 0 or 1. When a qubit interacts with the environment the process is, essentially, a measurement which is why the quantum nature of a state is inherently very fragile.




Figure 2: A schematic comparison of bits and qubits. The qubit can be imagined to be a unit vector $|\psi\rangle$ pointing in the direction (θ, ϕ) on a unit sphere (called the Bloch sphere) such that $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. When the qubit is read out, it is found to be in the basis state $|0\rangle$ with probability α^2 and basis state $|1\rangle$ with probability β^2 . Thus, all quantum measurements and calculations are inherently probabilistic in nature.

Bits in classical computers can be corrupted if the transistors storing the bit are subject to extreme heat, ionising radiation, etc. but modern circuitry makes such events uncommon and error rates are very small (ranging from 10^{-10} – 10^{-17} errors/bit/hour). Qubits, in comparison, are very much more fragile (described by the ‘fidelity’ of the qubit); since 2000, focused research to build high-fidelity qubits has seen their error rates fall roughly logarithmically. Trapped ion realisations have historically shown the

best fidelities and currently exceed 99.9%⁹⁴ with error rates $< 5 \times 10^{-2}$. This improvement in fidelity is achieved by ever greater isolation of the qubit from its environment and this poses a fundamental problem: if the qubits are to survive for a sufficiently long time to be usefully manipulated in a computation, how can they be easily initialised at the start, and read at the end, of a computation?

Problems realising qubits which can store quantum information for a sufficiently long time that it can be usefully manipulated using a suitable (quantum) algorithm held back QIP for several decades. However, over the past 5 – 10 years, there has been a great advance in qubit fidelities which has seen devices built with significant numbers of qubits possessing lifetimes which are usefully long before environmental noise degrades stored information. Such devices have recently been called NISQ (Noisy Intermediate Scale Quantum) computers and much research has been directed to devising quantum algorithms which can tolerate the noise. At the time of writing, the largest NISQ machine is Google's Bristlecone device which has 72 qubits comprising superconducting Josephson junctions.

A more insidious problem is that of the internal connectivity. In a classical machine, the information is stored in 'registers' or memory. The bits interact either as part of an algorithm (e.g. via an adder) or they may get shifted around using a 'bus', i.e. a system that is part of a backbone of the machine.

In a quantum computer, the qubits usually need to be in proximity for their working parameter(s) (e.g. spin, charge etc) to become entangled as a necessary step towards computation. This is very challenging on a large scale, since qubits need to be moved around without interacting with the environment (i.e. losing their 'coherence') or, in an adiabatic machine the topology must match characteristics of the problem.

A.2 Quantum computing paradigms

There are many approaches to building a quantum computer although conceptually the easiest, and commonest, paradigm is the circuit model whose principles are derived from classical digital computers. These in turn are derived from the paradigm of Boolean logic, arithmetic, and switches and relays. An alternative paradigm is adiabatic quantum computing. This requires the system is prepared in a state represented by an energy surface; points on this surface are considered to be in a one-to-one correspondence with the system configurations and the surface is then slowly (adiabatically) distorted into the shape that represents the problem; the lowest point on the surface corresponds to the state of the system equivalent to the 'answer'. The best-known examples of this type of quantum computer are the series of computers built by D-Wave Systems.

Given the many different types of quantum processor, it is natural to seek a metric by means of which different machines can be rated. 'Quantum supremacy' (also called 'quantum advantage') was introduced in 2017 by Caltech's John Preskill but only differentiated between those machines able to complete a task not possible using a classical computer. Broadly, the greater the number and fidelity of qubits comprising the quantum processor the 'better' the processor is expected to be.

⁹⁴ Fidelity and error rates are related but defined in different ways. Hence in the example the error rate is NOT 5×10^{-3} but instead 5×10^{-2}

A.2.1 Figures of merit: quantum volume

When comparing QIP hardware, it is useful to have a single figure of merit. In 2018, IBM introduced⁹⁵ the concept of ‘quantum volume’ which is a more sophisticated measure of a quantum computer’s performance. The metric considers

- Number of qubits;
- Connectivity;
- Gate performance;
- Algorithm errors;
- Compilers and software stack performance.

This metric quantifies the performance of a quantum computer by considering how well it can run a complex algorithm. It considers how many operations can be performed before the qubits decohere or a critically fatal number of errors occurs. The number of qubits and the number of operations that can be performed without overwhelming errors accumulating are termed the width and depth of a quantum circuit. The greater the depth, the greater the number of steps that can be carried out by the quantum computer; ‘deeper’ circuits can run more complex algorithms than ‘shallow’ circuits. To fully assess the ‘depth’ of a circuit, the connectivity, gate performance, coherence time, error rates, and the performance of compilers and the software stack must all be assessed and measured. An algorithm produces a single number - the ‘volume’ that can be compared against other quantum processors. IBM validated their approach by testing on their 5-qubit system released in 2017, their 20-qubit system released in 2018 and the Q System 1 it released in 2019. Quantum volumes of 4, 8 and 16 were found. According to IBM, this improvement in performance is similar to the rate predicted by Moore’s Law.

A.2.2 An overview of circuit model quantum computers

The ‘holy grail’ of quantum computing is considered to be the engineering of large scale, fault tolerant, circuit-model quantum computers which could run any quantum algorithm. Ideally, these would comprise architectures of near perfect qubits (fidelities infinitesimally close to unity) operating with vanishingly small gate errors but, in practice, schemes have been devised in which large numbers of physical qubits work together as smaller numbers of logical qubits.⁹⁶ Estimates of the numbers of physical qubits required for these error correction schemes range from 100s to 1000s. In broad terms, academics favour the trapped-ion approach (individual qubits are trapped metal cations, metal atoms from which one or more electrons have been stripped away) while large industry prefer the superconducting approach (individual qubits are superconducting Josephson junctions⁹⁷) Both technologies have their strengths and weaknesses:

- Fidelities: trapped-ion qubits, ~99.9%; superconducting qubits, ~99.4%;

⁹⁵ <https://arxiv.org/abs/1811.12926>

⁹⁶ A physical qubit is one that is realised as part of the structure of the machine, e.g. an ion in a trapped ion machine. As in classical machines, redundancy is introduced as part of a scheme to eliminate errors. Thus, a logical qubit is a qubit that has been ‘cleaned up’ using one or more of these schemes.

⁹⁷ A Josephson junction (see <https://www.scientificamerican.com/article/what-are-josephson-juncti/>) is a tiny conducting loop with a weak section where a current will flow indefinitely if left undisturbed. The current flow may be clockwise or counter-clockwise. One direction represents a ‘1’ and the other a ‘0’. With a very low current to generate a single quantum of magnetic flux, it may behave as a qubit and can exist in a superposition of states.

- Operation times: superconducting qubits, ~10 – 50 nanoseconds; trapped-ion qubits, ~3 – 50 μ seconds;
- Connectivity: trapped-ion qubits, all-to-all; superconducting qubits, nearest neighbours;
- Scalability: CMOS-compatible superconducting qubit devices, with confidence that 1000 qubit devices can be realised by 2025; trapped-ion qubit devices, have significant problems scaling beyond 50 trapped ions;
- Essential enabling technologies: superconducting qubits, cryogenics; trapped-ion qubits, UHV, magnetic shielding, etc.

IBM, Intel, Rigetti, Alibaba and Google are developing superconductor-based quantum computers while Honeywell, Alpine Quantum Technologies and IonQ favour ion-trap technology. The biggest challenge faced by superconductor-based approaches is increasing the qubit connectivity. Ion traps hold up to about 50 ions and the trapped-ion approach faces a major challenge introducing additional traps to increase the qubit number above ~50.

From an application development viewpoint, superconducting quantum computers are attractive because major developers offer cloud access to their machines (especially IBM's Q Experience⁹⁸) allowing development of algorithms concurrently with hardware development.

A.2.3 An overview of adiabatic quantum computers

The Canadian company D-Wave Systems was founded in 1999 and, using ideas from condensed matter physics, created quantum adiabatic processors first demonstrated in 2007. Subsequently, software (quantum algorithms) followed capable of solving a diverse range of practical problems including logistics, artificial intelligence/machine learning, materials science, drug discovery, cyber security, fault detection and financial modelling. The technology underlying the D-Wave quantum processing unit (QPU) comprises superconducting niobium loops through which currents circulate; the (flux) qubit states are implemented as the direction of the current flow (clockwise or counter-clockwise) with corresponding magnetic fields. A multi-qubit processor is built by coupling together individual qubits, using superconducting Josephson junctions (again, using niobium loops), although the connectivity of qubits is less than the ideal in which every qubit is connected to every other qubit.

The first chip, Rainier, in the D-Wave One computer comprised 128 qubits organised into 16 cells of 4 'horizontal' and 4 'vertical' qubits with cells tiled vertically and horizontally, a pattern D-Wave named the 'Chimera'. The maximum connectivity of any qubit is 6 but connectivity is 5 at the edges of the pattern. Subsequent chips (Vesuvius, W1K and W2K in the D-Wave Two, D-Wave 2X and D-Wave 2000Q computers) comprise 512, 1152 and 2048 qubits, respectively, but the Chimera pattern is maintained (figure 2a). D-Wave have been developing their technology at pace with roughly two chip development cycles per year. In October 2018, the D-Wave 2000Q was made available to the public through D-Wave's cloud service⁹⁹ and in February 2019 the next-generation Pegasus chip (figure 2b) was announced claimed to be 'the world's most connected commercial quantum system' with a maximum of 15

⁹⁸ <https://quantum-computing.ibm.com/>

⁹⁹ <https://cloud.dwavesys.com/leap>

connections per qubit, 5640 low-noise qubits and be available in mid-2020. More detail about the structure and operation of D-Wave machines may be found below in **B.7**.

D-Wave provide extensive tutorials and a programming guide¹⁰⁰ and a server application programming interface (SAPI) to those wanting to use the hardware. C/C++, MATLAB and Python libraries are provided to reformulate the user problem in terms of an embedded quantum unconstrained binary optimisation (QUBO) problem which D-wave can solve. Low-level programming, and data input and output, are transparent to the user.




Figure 3: a) the Chimera (left) and b) Pegasus (right) D-wave architectures
(reproduced from www.dwavesys.com with permission)

A.2.4 An overview of measurement based quantum computers (MBQCs)

This genre of quantum computing is also known as the ‘cluster state model’ and has no direct classical paradigm and is difficult to understand. Clusters of qubits in a highly entangled state are generated and used for computation.

The standard circuit model approach assumes the ability to perform any quantum operation from a universal set of gate primitives.¹⁰¹ This is difficult to achieve practically, particularly in multi-qubit configurations. In an extended computation, in realistic conditions, the quantum state being acted upon would most likely be rapidly corrupted and the calculation would fail. Whilst error correction can correct the fault, it introduces a significant overhead in terms of qubits required. The MBQC model avoids these overheads.

¹⁰⁰ https://docs.dwavesys.com/docs/latest/doc_handbook.html

¹⁰¹ There are more of these than in the classical case, where all operations may be constructed using "NAND" gates.

Knill, LaFlamme, & Milburn¹⁰², following Gottesman and Chuang, invented measurement based quantum computing which uses linear optics and requires only single qubit gates plus so-called Bell basis measurements or Bell measurement gates. Technically, it was shown using this approach that universal quantum computation is possible with only linear optical elements and photodetectors. Two Bell measurement gates can move an arbitrary 2-qubit state between two locations using a process known as ‘teleportation’¹⁰³ and, given the assumed ability to store instances of an entangled state, provide the basis for a more complex machine to perform any feasible quantum operation.

Later, an alternative approach was proposed by Raussendorf and Briegel¹⁰⁴, more often known as ‘one way’ quantum computing, and requiring only single qubit measurements. The system is prepared in an initial, highly entangled, state called a cluster state. A set of measurements is made on single qubits and the order and choice of basis for these measurements defines the computation, the path chosen relies on the results of previous measurements. It is a ‘one way’ scheme because, as the computation is performed, time asymmetry is introduced and the computation can only run forwards. The approach is attractive because the technical challenge becomes that of preparing the initial cluster states rather than executing the subsequent single qubit measurements, which are assumed to be straightforward. In reality this may not quite be the case, since the single qubit measurements can affect neighbouring qubits and this limits possible architectures.

A.2.5 An overview of topological¹⁰⁵ quantum computers

Exotic possibilities arise in the quantum physics of particles confined to move in only two dimensions, particularly at very low temperatures and in the presence of very strong magnetic fields. A topological quantum computer is a theoretical system most often employing anyons which are two dimensional quasiparticles (i.e. excitations that usually exist on surfaces) whose world lines (trajectories) form braids in two-dimensional space and time. As time proceeds, the calculation takes place via interactions between these anyons. The system is thought to be comparatively robust due to the stability of these braids (they possess structures similar in concept to knots).¹⁰⁶ Thus, the idea behind topological quantum computing is to encode information into topological degrees of freedom which are intrinsically error free (in a suitable thermal environment) in terms of error avoidance rather than error correction.

Such machines would be similar in computing ability, power and capability to circuit models of computation, although certain problems may map more or less easily on to their structure. None of these machines have yet been demonstrated, although work has commenced on some of their building blocks. Experimental evidence for the

¹⁰² <https://arxiv.org/pdf/quant-ph/0006088v1.pdf>

¹⁰³ Teleportation

¹⁰⁴ See <http://arxiv.org/pdf/quant-ph/0301052v2.pdf>

¹⁰⁵ Topology is a mathematical discipline concerned with geometrical properties which are not affected by continuous deformations including stretching and bending. Hence, a tea cup is considered to be topologically the same as a doughnut because of its handle.

¹⁰⁶ See for example <http://iopscience.iop.org/1367-2630/focus/Focus%20on%20Topological%20Quantum%20Computation>

existence of some types of anyons was observed in 2005. The type of anyons required for topological quantum computing are thought to exist in rotating Bose Einstein condensates, quantum spin systems and superconductors.

Microsoft are investing in this technology, and braiding has been demonstrated by Marcus in Copenhagen. Microsoft say that the qubits are spatially large (they are quasi particles made up of multiple real particles) which means that the errors are much fewer and that qubits are almost naturally fault tolerant.

A.2.6 An overview of emulators of quantum processors

‘Emulator’ and ‘simulator’ are used interchangeably by many. In this document a simulator is taken to mean a physical system, such as an ensemble of atoms held in an optical lattice, which can be used to estimate the properties of a second, more complex quantum system, such as a solid. Emulator is reserved for a piece of software which runs on an HPC and which can computationally describe the quantum states of a many qubit system and how those states can be modified by the action of quantum mechanical operators. Emulators this describe a fully entangled system of quantum objects. If there are N such objects (qubits) the number of possible quantum states is given by 2^N and the computer resources required to completely describe such states (without any approximations) rapidly exhausts even the largest HPCs as N increases and systems comprising about 40 qubits are currently the largest which can be described without any approximations.

The website¹⁰⁷ for Oxford University’s Quantum Exact Simulation Toolkit, Quest, gives useful values for the resources required by Quest to emulate different numbers of qubits: 26 qubits requires a 2 GB machine, 29 qubits require a 16 GB machine and 45 qubits require 1, 048, 576 GB (i.e., 1 petaByte, PB). Not just the storage requirements but also the execution times for the emulation increase exponentially and systems comprising ~45 qubits are likely to remain the upper limit of may be computationally described without approximation.

There are many emulators¹⁰⁸ many of which are open source. In addition to these Microsoft’s QDK (see **Appendix J.2.5**) and ATOS’s (see **Appendix J.3.1**) are available commercially.

¹⁰⁷ <https://quest.qtechtheory.org/about/>

¹⁰⁸ See <https://quantiki.org/wiki/list-qc-simulators> for a list

APPENDIX B Quantum computing hardware

B.1 Introduction

A quantum computer is completely unlike a digital computer to the extent that the term is misleading. It is much more like the form of computer used from the Victorian era onwards, the analogue computer. Analogue computers were usually implemented using mechanical, hydraulic, or pneumatic means and a system was set up that would mimic the real system allowing the calculation of outcomes in advance. The quantum computer is similar in concept but exploits quantum interactions rather than those of cams, wheels, pipes and reservoirs.

Analogue computers were quickly and effectively put to military use. The fire control systems on battleships from about 1905 incorporated successively better analogue computers until fire could be accurately directed over tens of kilometres from moving platforms, taking account of a multitude of factors. Early weapon locating radars (Green Archer also known as Radar, Field Artillery, No. 8) used analogue computers until digital computers allowed cheaper and more compact implementation. Submarines used torpedo data computers to calculate firing angles and early models of the UK economy operated by the UK Treasury were hydraulic analogue computers. Analogue computers died out as they were replaced by more adaptable digital computers once mainframe computers like the IBM 360 became available. Digital computers were not necessarily quicker but they were cheaper, far easier to program/adapt and could be more accurate.

Quantum computers provide a novel implementation of analogue computers using quantum effects such as entanglement. As small NISQ machines appear, they are used together with digital computers in hybrid form since there are some tasks such as simple arithmetic and data input/output which quantum computers will never do well. Quantum computers, at least for the foreseeable future, will remain specialised co-processors used to do specific tasks at which they excel.

Most current quantum computers are NISQ machines and will remain so until it is possible to engineer machines with sufficiently many qubits to implement error correction schemes. Leading circuit-model hardware is being developed by IBM, Intel and Google and the current position is summarised below. D-Wave is prominent developing an adiabatic quantum computer and, given the expectation that D-Wave is a strong candidate for early adoption of QIP by Defence and Security, the hardware is described in some detail below. Xanadu¹⁰⁹, QET Labs¹¹⁰, JQI¹¹¹ and others, are developing an alternative approach using photons described in **Section B.9**.

B.2 IBM

The hardware is based on transmon¹¹² superconducting qubits (invented at Yale in 2007 and engineered from two capacitatively shunted superconductors to have low sensitivity to charge noise). The architecture is scalable and error correction can be

¹⁰⁹ <https://www.prnewswire.com/news-releases/xanadu-receives-4-4m-investment-from-sdtc-to-advance-its-photonic-quantum-computing-technology-300987885.html>

¹¹⁰ <http://www.bristol.ac.uk/physics/research/quantum/>

¹¹¹ <https://jqj.umd.edu/news/semiconductor-quantum-transistor-opens-door-photon-based-computing>

¹¹² transmission line shunted plasma oscillation qubit

incorporated. The hardware is stacked in layers whose temperature decreases from 4 K at the top of the stack to 15 mK at the base.

To-date, over 2.5 million experiments have been on the IBM Q platform¹¹³ and more than 60 research papers published. One landmark publication was a detailed solution of the non-trivial problem of fully entangling 16 qubits.¹¹⁴ Strong user engagement will be important in the future for the efficient, application orientated development of quantum computing. In the UK, Oxford University is currently engaged with IBM (as an IBM Q-Hub regional centre of quantum computing education, research, development, and implementation which provides collaborators online access to IBM Q quantum technology) but there are many more overseas government and industry research organisations engaged as partners.¹¹⁵ Current applications projects include quantum chemistry for drug and materials design and optimization of transportation logistics and finance.

In April 2018, IBM revealed the first start-ups joining the IBM Q Network with cloud-based access to IBM's quantum computers and other resources. These include:

- 1Qbit: (Vancouver, Canada) builds quantum and quantum-inspired solutions for demanding computational challenges. Their hardware-agnostic services allow development of scalable applications. The company is backed by Fujitsu Limited, CME Ventures, Accenture, Allianz and The Royal Bank of Scotland;
- Cambridge Quantum Computing (CQC);
- Zapata Computing: (Cambridge, MA) provides quantum computing, services developing algorithms for chemistry, machine learning and security;
- Strangeworks: (Austin, TX) develops QIP tools for software developers and systems management;
- QxBranch: (Washington, D.C.) provides data analytics for finance, insurance, energy, and security customers. The company is developing quantum tools exploiting machine learning and risk analytics;
- Quantum Benchmark: (Kitchener-Waterloo, Canada) is a venture-capital backed software company seeking to provide solutions which enable error characterization, mitigation and correction as well as performance validation of quantum computing hardware;
- QC Ware: (Palo Alto, CA) develops hardware-agnostic quantum software for Fortune 500 companies including Airbus Ventures, DE Shaw Ventures and Alchemist as well as US government agencies including NASA;
- Q-CTRL: (Sydney) is using its hardware agnostic platform (Black Opal) to improve quantum computer performance and reduce the lead time for QIP tools which can solve real world problems. Q-CTRL is backed by Main Sequence Ventures and Horizons Ventures.

¹¹³ IBM have quantum hardware sites at Tokyo (20 qubits), Melbourne (14 qubits), Tenerife (5 qubits) and Yorktown Heights (5 qubits). Typical clock speeds are ~5 GHz, with T_1 and T_2 times 10 - 70 micro-seconds. Gate and readout errors are $(0.7 - 3.0) \times 10^{-3}$ and $(3.0 - 10.0) \times 10^{-2}$ respectively.

¹¹⁴ Wang et al, <https://www.nature.com/articles/s41534-018-0095-x>

¹¹⁵ <https://www.research.ibm.com/ibm-q/network/members/>. The network includes clients from Fortune 500 companies, academic institutions, and US national research labs, including JPMorgan Chase, Daimler, Samsung, Barclays, Honda, Oak Ridge National Lab, Oxford University and University of Melbourne.

In addition to real quantum hardware, IBM offers high-performance quantum simulation (Qiskit Aer) which can be accessed (through Qiskit or IBM Q Experience, see¹¹⁶). This allows ideal experimental circuits to be tested before running on real hardware, the performance of which can be predicted by adding noise in a controllable way.

B.3 Google

Google's research areas in hardware development are¹¹⁷:

- Superconducting qubit processors with chip-based scalable architectures targeting two-qubit gate errors of <0.5%. Bristlecone, announced in March 2018, is Google's most recent quantum processor with 72 qubits and Google are 'cautiously optimistic' that, with system engineering to achieve optimally low error rates, equal to or better than their previous 9 qubit device¹¹⁸, it will show quantum supremacy;¹¹⁹
- Quantum neural networks research is developing a framework to implement a quantum neural network on NISQ processors available now or in the near future.¹²⁰ The advantages which may be achieved by manipulating superposition of very large numbers of states is a key research objective;
- Quantum-assisted optimisation: Google are developing hybrid quantum-classical machines for optimization problems. These would take advantage of thermal noise to allow tunnelling to globally lowest energy states of the problem Hamiltonian (in much the same way as D-Wave).

B.4 Intel

Intel's declared goal¹²¹ is a complete quantum computing system (hardware, algorithms and software and control electronics) on a chip and has adopted two approaches to quantum computing.

Like many other research groups, they are developing a superconducting qubit approach, exemplified by the Tangle Lake 49-qubit chip announced in January 2018. The launch of the 49-qubit chip happened only a few months after the announcement of the 17-qubit chip developed in conjunction with Intel's Dutch partners, QuTech and Delft University of Technology. The chips, made with a 'flip-chip' processing method, have an architecture allowing improved reliability, good thermal performance and reduced RF interference between qubits while the fabrication process enables smaller features and scalable interconnects (and higher data flow on and off the chip) compared to wire bonded chips.

Intel are also developing a 'spin qubits in silicon' approach which seeks to exploit Intel's many year's-experience in silicon chip technology. Intel liken the technology to existing semiconductor electronics and transistors but differs by exploiting the spins of single electrons, manipulated by low-amplitude microwave pulses. This effort is at a lower TRL than their superconducting technology but may progress more rapidly, perhaps

¹¹⁶ <https://www.research.ibm.com/ibm-q/technology/simulator/>

¹¹⁷ <https://ai.google/research/teams/applied-science/quantum-ai/>

¹¹⁸ Demonstrated readout and single gate errors of 0.1% and 2 qubit gate errors of 0.6%.

¹¹⁹ Google published a Nature article on 23rd October 2019

(<https://www.nature.com/articles/d41586-019-03213-z>) claiming this milestone had been passed using a 53 qubit processor and the consensus view, after intense scrutiny, concurred.

¹²⁰ <https://github.com/quantumlib/cirq>

¹²¹ <https://www.intel.co.uk/content/www/uk/en/research/quantum-computing.html>

even overtaking the superconducting approach. A CMOS-based approach allows a high qubit density, which aids entanglement with neighbouring qubits. In February 2018, QuTech and Intel announced a 2-qubit silicon spin-qubit based quantum device which should be able to operate at ~1 K, less technologically challenging than the ~20 mK necessary for superconducting qubit operation. Progress in other areas includes demonstration of an algorithm, a compiler and control electronics.

B.5 Microsoft

Beginning with the establishment of Station Q in 2006, Microsoft has been working on developing a scalable quantum computer based on topological qubits. These use quasi-particles (called anyons which still have to be demonstrated experimentally). The topological approach is attractive because anyons, like fermions, cannot occupy the same quantum state and so are resistant to errors. Thus, quantum computers based on topological qubits do not need error correction schemes, which for other types of qubit is believed to increase the required numbers of qubits by factors of about 10^3 - 10^4 . Building a topological quantum computer is correspondingly easier than, say, building an ion trap-based device, or will be when anyons can be physically realised.

It was shown in 2002¹²² that topological quantum computers are equivalent to other types of quantum processor and, in particular, are more appropriate for running some types of quantum algorithm (such as those concerning knot theory). Although error resistant, topological quantum computers give a level of accuracy which is directly proportional to the number of anyons comprising the machine. ('Conventional' quantum computers, when error-free, solve problems with absolute accuracy.) Thermal fluctuations in the processor produce random pairs of anyons which interfere with other anyons but this is straightforwardly avoided by ensuring the anyons are physically separated by a distance at which the interaction is effectively zero.

Microsoft envisage a future quantum computing system, from software to hardware, integrated into its Azure cloud service. The Microsoft Quantum Development Kit will be integrated into Visual Studio. The Q# language has been created for quantum code development together with extensive quantum libraries. When the code is complete, the concept is to run a quantum simulation to check for bugs and validate that the solution is ready to be run on a quantum computer.

B.6 Honeywell

In March 2020, Honeywell Quantum Solutions¹²³ announced that it planned to make available over the internet, by Summer 2020, its trapped ion based quantum computer. With just 2 qubits, it claimed a quantum volume of 16 which value it ascribes to the very high fidelities possible with trapped ion qubits. By the time the system is available commercially, the quantum volume is expected to increase to about 64 through further improvements. Honeywell has also pledged to add additional qubits to the system each year until 2025 and believes that it will be able to demonstrate increases in quantum volumes by factors of 10 each time.

Honeywell claim a unique selling point. Because of the long qubit coherence times, the calculation can be halted and a qubit can be interrogated and its value reset depending on the measurement. Honeywell liken the process to the addition of an 'IF' statement to an algorithm.

¹²² <https://link.springer.com/article/10.1007%2Fs002200200645>

¹²³ <https://www.honeywell.com/en-us/company/quantum>

B.7 D-Wave

D-Wave machines use quantum annealing to solve NP-hard, unconstrained binary optimization (QUBO) problems¹²⁴ by mapping the problem onto a graph which represents the topology of the machine's interconnected qubits. Each qubit is a superconducting quantum interference¹²⁵ device (SQUID) fabricated from niobium which becomes superconducting at cryogenic temperatures. Electricity flows without resistance and magnetic fields are set up which can point 'up' or 'down' allowing for the representation of two states.

D-Wave's mathematical treatment of the QUBO problem is to seek values of variables $x = (x_1, \dots, x_N)$ which minimise the quadratic objective function $(x) = \sum_{i=1}^N q_{ii}x_i + \sum_{j=1}^N \sum_{i=i}^{j-1} q_{ij}x_ix_j$, given the set of q_{ij} real coefficients and the constraint that $x_i = 0$ or 1 .¹²⁶ Objective functions can be represented mathematically by graphs which are sets of nodes (qubits representing the variables x_i) connected by edges (the coupling constants q_{ij}).

Complete graphs, K_n , are described in terms of the number of nodes, n , which they have and a complete graph with n -nodes represents the edges of an $(n-1)$ D polytope (a generalisation of a 3D polyhedron). Figure 4 shows the complete graphs K_n together with the $n(n-1)/2$ edges connecting every node to all other nodes. A k -vertex connected graph $K_{n,k}$ is a graph with only $k < n(n-1)/2$ edges connected. The degree of a node of a graph is the number of edges that are incident to the node.

¹²⁴ A well-known QUBO problem is the Graph Colouring problem in which the nodes (or vertices) of a graph are coloured in such a way that no two adjacent nodes have the same colour

¹²⁵ The term 'Interference' refers to the wave behaviour of the electrons flowing in the superconducting niobium which causes quantisation of the magnetic flux created by the electric current

¹²⁶ The QUBO problem is equivalent to the Ising problem in statistical mechanics which seeks the minimum energy, $E(s) = -\mu \sum_i h_i s_i + \sum_{i,j} J_{ij} s_i s_j$, for a lattice of atomic spins μ with $s_i = \pm 1$ subject to external magnetic fields h_i and site-site coupling J_{ij} via the transformation $s_i = 2x_i - 1$.




Figure 4: Complete graphs, K_n (https://en.wikipedia.org/wiki/Complete_graph)

In D-Wave's machines to date, the qubits are not fully connected (which restricts the classes of problem which can be solved) and in the D-Wave 2000Q quantum processing unit, and predecessors, the arrangement of connections is known as a Chimera graph. The next generation processors will have a Pegasus graph architecture and is the first fundamental change in the architecture of D-wave machines since the first appeared (D-Wave One).

The D-Wave Chimera architecture comprises unit cells of $K_{4,4}$ graphs. Unit cells are tiled into $N \times N$ arrays of sparsely connected qubits denoted CN by D-Wave. Each unit cell comprises 2 sets of 4 qubits, one set of 4 'vertical' and one set of 4 'horizontal' qubits making a total of 8 qubits per unit cell. Within the unit cell, each vertical qubit is connected (coupled) to every horizontal qubit and vice versa ('bi-partite' connectivity). Between 'internal' unit cells each vertical/horizontal qubit is connected to the corresponding vertical/horizontal qubit in the previous and next cell giving 6 connections per qubit. Qubits in 'edge' unit cells have only 5 connections to other qubits, see Figure 5. In D-Wave's terminology, Chimera qubits have a length of 4 (internal connections within the unit cell) and a degree of 6 (maximum number of qubit connections internal and external to the unit cell).




Figure 5: The C4 Chimera architecture illustrating qubit connectivity

The D-Wave 2000Q QPU supports a C16 Chimera graph within which 2048 qubits are logically mapped into a 16x16 array of 8-qubit unit cells with a degree of 6. The largest number of qubit couplers achieved in the C16 Chimera is 6000 indicating the substantial lack of connectivity compared to the maximum possible.¹²⁷

During the evolution of the Chimera architecture, D-Wave machines developed solely by tiling more unit cells so that the qubit number, but not the degree, increased. In the Pegasus, as with the Chimera architecture, qubits belong to either vertical or horizontal sets but, while internal couplers connect orthogonal qubits (vertical ↔ horizontal) and external couplers connect colinear qubit pairs (in the same horizontal row or vertical column), in the Pegasus architecture there are three types of coupler: internal, external and odd. Odd couplers connect pairs of vertical or horizontal qubit pairs in (respectively) adjacent columns or rows. Pegasus qubits have a length of 12 and degree of 15 and this greater connectivity allows more complex problems to be solved using the same number of qubits or, conversely, fewer qubits are needed using a Pegasus architecture processor than needed by a Chimera architecture processor to solve the same problem. The time to solution, however, is not reduced as the connectivity increases.

The Pegasus graph, PM, contains $24M(M-1)$ qubits and has a maximum degree of 15¹²⁸. In the first set of Pegasus chips to be announced, Pegasus(0), a PM graph contains $8(3M-1)(M-1)$ non-edge qubits and $8(M-1)$ edge qubits (supporting K_4 and $K_{6,6}$ sub-graphs built from the main processor fabric¹²⁹). The P16 chip announced in February 2019 contains 5760 qubits in total, with 5640 in the main fabric. This is large increase in qubit connectivity compared to the D-Wave 2000Q allows greater qubit

¹²⁷ $\frac{1}{2} \frac{2048!}{2!2046!} = 2096128$

¹²⁸ K Boothby et al, 'Next-Generation Topology of D-Wave Quantum Processors' 14-1026A-C, D-Wave Technical Report Series (2019)

¹²⁹ 'fabric' is a D-Wave term used to denote the network of qubits as nodes of the graph. The fabric of Pegasus graphs, necessarily, includes disconnected nodes and while building their graph representation, D-Wave's graph generators currently take the value of a Boolean variable ('nice_coordinates') to control whether coordinate systems compatible with Chimera addressing should be constructed

entanglement increasing the computational capability of the machine. D-Wave are conducting characterisation studies to quantify the benefit but Boothby et al⁹¹ expect these to include:

- more efficient embeddings¹³⁰ of complete subgraphs and bipartite subgraphs;¹³¹
- improved run times;
- improved error correction schemes.

Through Cloud access, a number of users have used D-Wave's machines for various optimisation tasks. Volkswagen and Japanese car components manufacturer Denso have separately worked with D-Wave to develop traffic flow optimisation models while Tohoku University in Japan has developed tsunami evacuation models following the magnitude 9.0 earthquake off Japan's Pacific coast which devastated parts of eastern Japan in 2011. The P16 chip with 5760 qubits is expected to be brought online for cloud access by mid-2020 and further increase experimentation with D-Wave's evolving hardware. Other modes of access include purchase of a machine for exclusive use but the high cost involved (at least \$15M) have restricted the adoption of this approach to a few organisations including NASA, Google, Lockheed Martin and Los Alamos National Laboratory.

There has been much debate about various aspects of the D-wave series of machines including whether their operation depends in any way on the types of quantum phenomena which underpin the operation of the currently emerging quantum technologies (superposition, entanglement, tunnelling, etc.) and what, if any, speed up these effects confer on D-Wave processors compared to classical computers. One of the most prominent critics is Scott Aaronson who has systematically challenged D-Wave's published assertions that their machine is superior to traditional classical machines at solving certain carefully chosen problems and has yet to demonstrate quantum supremacy of its processor. There now seems to be agreement that D-Wave machines do use quantum tunnelling in the solution of computing problems, but no certainty that they will solve real world problems exponentially faster than classical computers and problems have yet to be found for which D-Wave outperforms all classical competitors.¹³²

The class of problems for which the D-Wave machine is best suited is determined by the architecture of the chip. With currently demonstrated connectivities, D-Wave processors do not allow the machine to act as a universal Turing machine but are well suited to solving quantum unrestrained binary optimisation (QUBO) problems or any NP-complete or NP-hard problem which can be mapped into this form.

¹³⁰ Before solving an optimisation problem using a D-Wave machine, its graphical representation ('problem graph') must be mapped onto some or all of the hardware qubits. This 'minor embedding' is a subgraph of the hardware graph where the qubits represent the nodes (variables) of the graph and the connections between them the edges. It is known to be NP-hard to find an optimal minor embedding of an arbitrary problem graph into an arbitrary hardware graph. There are various algorithms to find minor embeddings; an heuristic algorithm proposed by Cai (see <https://arxiv.org/abs/1406.2741>) is the most versatile so far

¹³¹ The internal connections in the Chimera unit cell, Figure 4, constitute a partite subgraph

¹³² 'Quantum computer gets design upgrade', E Gibney, Nature, **541**, 447 – 448 (2017)

The first step is to recast the problem into an equivalent combinatorial optimization problem so that the solution corresponds to the minimum value of the energy function $E(s)$,

$$E(s) = - \sum_{i=1}^N h_i s_i - \sum_{i<j} J_{ij} s_i s_j$$

where the variables s_i are constrained to the values ± 1 and J_{ij} and h_i are real numbers. N is the number of variables that define the problem. D-Wave computers are designed to find approximations to the solutions of the Ising spin problem described the Hamiltonian (energy function) $H_{problem}$,

$$H_{problem} = - \sum_{i=1}^N h_i \hat{\sigma}_i^z - \sum_{i<j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z$$

To solve a problem, the D-Wave machine is set to an initial state in which all coupling constants J_{ij} are set to zero and all biases, h_i , set to 1 so that $H_{initial} = - \sum_{i=1}^N h_i \hat{\sigma}_i^z$. The coupling constants and biases are then slowly ('adiabatically') changed so that 'initial' configuration is evolved into the 'problem' configuration. Effectively what happens is that the system is slowly perturbed from its ground state over a time period T_{solve} . At any given time, T , the system's Hamiltonian is

$$H(s) = A(s)H_{initial} + B(s)H_{problem}$$

where $s = 1/T$. If the couplings and biases are changed sufficiently slowly, then at any given time T the system is always in its ground state and at the end of the computing run ($T = T_{solve}$) the system Hamiltonian has evolved to that for the problem.

Although the speed at which the D-Wave system evolves is too rapid for it to be regarded as a truly adiabatic process¹³³, it is necessary that the processor is maintained very close to absolute zero to control the amount of thermal noise. There appears to be an optimal processor temperature of the order of 15 – 20 mK because small amounts of thermal noise assist tunnelling and accelerate the system's evolution to its global minimum energy. Refrigeration energy requirements scale almost independently of qubit number N ¹³⁴ (because the principal cooling load is the chip rather than the annealing process) and scaling to architectures with at least 10^4 qubits is practical.

D-Wave processors have been used in many application areas including the financial sector for trading trajectory optimization, in molecular biology to model protein folding in bioscience and for the development of binary classifiers in AI and for computer vision.

Of particular interest in the context of defence and security are autonomous machine tasks including:

- identifying threats in online traffic;
- extraction of information from images.

D-Wave has experimented with machine learning on the chip, setting up a Quantum Boltzmann machine, a type of stochastic recurrent neural network, which D-Wave

¹³³ See for instance 'Molecular Reaction Dynamics' R D Levine, Cambridge University Press (2005)

<https://doi.org/10.1017/CBO9780511614125>

¹³⁴ D-Wave One cooling energy required = $0.00009 N^2 + 0.001 N + 33.217$ kJ

describe as ‘fundamentally different from previous machine learning models and eventually allowing a machine to generate new data that is statistically indistinguishable from the kind of data on which it was trained’. In principle, this could have powerful applications such as creating speech indistinguishable from that of real humans. The company has launched a spin-off called Quadrant to develop this approach, focussing on deep learning using only small amounts of training data.

D-Wave has raised much funding from a wide range of investors, including investment bank Goldman Sachs, In-Q-Tel, Bezos Expeditions (the investment arm of Amazon), BDC Capital, Harris & Harris Group and Draper Fisher Jurvetson, a venture capital firm.

D-Wave has sold only small numbers of machines but continues to attract buyers from government (NASA), industry (Lockheed Martin and VW) and commerce (Google) and attract interest in its technology. Recently, cloud access has begun allowing users such as Oak Ridge National Laboratory to have cloud access to a D-Wave 2000Q system, allowing them to explore hybrid computing as to access the latest generation D-Wave processors.

B.8 Competing digital silicon technologies

While quantum computers are expected to provide the only way of solving certain problems, none of these problems are believed to be critically important to Defence and Security activities. Thus, this Landscape would be incomplete without consideration of the significant progress being made in digital silicon technologies.

Examples where quantum computers already bring a unique capability are in:

- Quantum simulation. Quantum interactions are extremely difficult and computationally costly to model on digital computers, such that only fairly simple cases can be addressed. However, a capability to model large quantum systems would be useful in chemistry, life sciences and materials research;
- Quantum annealing. Quantum computers are said to avoid the issue where a local minimum or maximum is mistaken for the optimum answer. It is unclear how much value this brings in practice, it would only be valuable for extremely complex optimisation problems (such as the protein folding problem).

Quantum computer emulators, running on digital computers, already exist. It could therefore be argued that approaches that ‘only’ a quantum computer can execute can also be performed digitally. However, the huge memory sizes and processing capacity that must be applied to accurately model quite small quantum computers (circa 40 qubits is the current maximum which can be emulated) requires an expensive supercomputer.¹³⁵ At present, however, the vast majority of Intelligence related problems that a quantum computer might tackle can also be tackled by digital computers. However, if the data deluge problem grows as expected, then quantum computers might become the only practical way of managing the information processing load.

B.8.1 Hardware and Software

Digital computers are very ‘general purpose’ and can give solutions for any well understood problem. Diverse examples include a game of chess, control of a robot and

¹³⁵ The commercially available ATOS Quantum Learning Machine runs on Bull HPCs with ~1 million cores running at 25 – 30 GFlops.

the ubiquitous office spreadsheet. But success can be very limited; execution may not be at a useful speed, or be affordable in terms of resources such as memory. Image processing and complex route-finding in 3 dimensions (the generalised 'Travelling Salesman' problem) are examples of everyday but very costly problems.

Computer programming is constrained by the need to align how humans 'think' and how digital computers operate. The programming language must allow humans to express their solution in an understandable way but it also needs to take a form which permits automatic translation into the many sequential steps by which digital computers work. New languages are continually being developed, but evolution is slow. In essence, all seek to improve the alignment between how a human conceptualises a problem and how a computer then executes that solution.

The central issue is that digital computers require that a problem be expressed as a sequence of steps that are executed sequentially. Modern computer languages and systems do allow a degree of parallel activity, whereby parts of a programme are run concurrently rather than sequentially, but in essence all digital computers are sequential.

Quantum computers are truly parallel, and it is very difficult to make them perform a series of steps. For example, there is not yet a viable 'quantum memory' to hold intermediate products of calculation. This means it will not, in general, be possible to take software written for digital computers and re-compile it for a quantum computer.

This means that any way of expressing 'solutions' that can be compiled onto both types of computer becomes especially important. Neural nets and Annealing-type optimisation are good examples, but of the two, neural nets appear by far the most important to Intelligence as they are already being applied to a huge range of 'pattern matching' problems in Defence and Security.

B.8.2 Quantum and Digital Hardware Competition

Digital technology will be 'head to head' with Quantum Information Processing, on a like-for-like basis, when 'running' conceptually parallel software such as Neural Nets or Annealers. Both perform the same task in a conceptually similar way and the key questions become those of relative cost and performance.

There can also be a 'head to head' comparison when digital processors deliver a function and quantum information processors deliver the same function, albeit by a totally different method. This is likely to be a 'chalk and cheese' comparison but in reality, business problems often can be solved in very divergent ways. Examples might be facial recognition, number plate recognition and feature recognition in images. Quantum computers and digital computers might do this differently but could be directly compared based on error rates for feature identification, processing speeds and cost. This quickly converts into a 'benchmarking' exercise, where a problem (either 'real' or a 'test') is solved by the alternative methods and error rates, speeds and costs are compared.

In the proposed follow-on study, the focus is on comparing the performance of digital and quantum computers when running software matched equally to both, such as a neural net. In reality, it is not just quantum computers that are becoming bigger and faster, the same is true of digital computers. It would also appear that mainstream manufacturers such as Intel are developing digital computers optimised to run neural nets. This will compound the rate at which digital competitors improve, and delay the

point at which quantum computers become cost-effective. The anticipated 'digital competitors' to quantum information processing are those that run neural net pattern matching software accurately, quickly and cheaply.

Computer annealers are less of a focus (noting that quantum computers that can anneal, such as D-Wave, can also run neural nets). Digital computers are usually suitable (and fast enough) to solve all but the most complex optimisation problems. Quantum annealers are a little different in that they claim to avoid incorrect local optimisation 'solutions' through quantum tunnelling and would come into their own in cases of extremely complex optimisation surfaces.

B.8.3 Low-cost and Highly Parallel non-Quantum Computers

B.8.3.1 Cloud Computers

Cloud computers comprise very large numbers of interlinked general-purpose computers and so do not offer outstanding processing power per £. However, they are ideal for sharing computer power over many users where processing loads fluctuate significantly with time. In essence each customer can ask for a brief, vast surge in power without facing a huge cost increase. The model is one of 'rent' rather than 'purchase'.

Cloud computers are likely to be well matched to image processing problems provided the images come in batches and not continuously. They will also be well matched to a problem where the expectation is that the algorithms will later be executed by replacement quantum computers.

Neural net solutions are likely to map well onto Cloud computing, which could be used to run either large neural nets or many small ones at far higher speeds than any single computer might achieve. It offers a solution that scales efficiently.

However, the advantage is achieved by harnessing very large numbers of computers which are not individually cheap. Cloud computers are not expected to be able to run neural nets any more cost-effectively than single cores, although expert advice must be sought on this question.

B.8.3.2 GPU (Graphical Processing Unit) based Solutions

'Graphics cards' are purchased in huge numbers for business and consumer purposes, there being roughly one unit per computer purchased. They command a good price as the apparent performance of a gaming or Computer Aided Design (CAD) computer is driven by its graphics card. Performance is growing very rapidly; new generations have been appearing about every 2 years with doubling processing speeds. Moore's law is still operating although the performance gains are not just achieved by improved photolithography.

A modern graphics card consists of a large memory which contains the display image (and usually one or more additional images being processed while the first is displayed) and many special on-chip computers working on the individual pixels. These on-chip computers are RISC (Reduced Instruction Set Computer) devices optimised for high precision multiplication, division and addition. This means that each one occupies a very small area of silicon compared to a general-purpose computer core of comparable power. This kind of RISC processor is highly capable of running both neural net and image processing software where the essential calculation is multiplication to achieve weighting, then cumulative summation.

The RISC processors communicate through shared memory using special architectures that minimise 'clashes' in memory access. It is important when harnessing many RISC processors to moderate memory contention as they interact, which limits the number of RISC cores that can be harnessed effectively. This is analogous to the quantum computer issue of having each qubit 'entangled' with as many others as possible. In essence, GPUs added to a normal computer multiply its processing power many-fold and in principle a single ordinary computer could operate many GPUs.

GPU chips are programmable in C, C+, Python and many other languages. The manufacturing leaders are NVidia and AMD. The problems they can solve efficiently must align to their special architecture and they are widely used to run neural nets and other image processing software; they would be well matched to highly parallel activities like Digital Annealing (see Section **F4.4**).

Based on silicon area as the cost driver, GPUs can be expected to outperform general purpose computers by one or two orders of magnitude (for the same price) if used to 'run' neural net software.

B.8.3.3 Bespoke neural net chips

Intel have been developing special chips optimised to running neural net programmes for several years. These are being made available to Universities and research organisations at low cost. Neural net internal mathematics is very simple so the required computing cores occupy little silicon area. Communication is at relatively low levels of number precision, which creates new design options but, even so, there is a major interconnection problem.

In principle, every node in the hidden layer must communicate with every input node, potentially requiring thousands of crossing 'wires'. Neurons in animals and insects achieve this enormous degree of interconnection by using 3D structures but at present silicon technology is intrinsically 2D; '3D' integrated circuits are realised by vertical stacking of 2D chips connected at a global, intermediate or local level; their compact nature has seen wide adoption for flash memory in mobile devices.

Digital silicon neural net solutions use 'telephone exchange' type switching, moving messages sequentially along shared 'busses' rather than in fully parallel circuitry. This limits execution speeds and so the design trade becomes silicon area for communication vs silicon area for computation. This is very like the trade-off GPUs require but neural nets generally need simpler 'messaging' enabling a reduction of the silicon area assigned to communication.

Based on silicon area as the cost driver, bespoke neural net chips can be expected to out-perform GPUs by perhaps an order of magnitude. GPUs are already much more economical of silicon area than general purpose computer cores, such that special neural net chips could potentially outperform a normal computer by between 2 and 3 orders of magnitude, for the same silicon area. Thus, for Defence and Security applications it is probable that GPU and/or neural net chips will become the primary competitors to quantum information processors adopted at an early stage of QIP development over the next 5 – 10 years.

Intel examples include their neural compute stick¹³⁶ but Intel are also developing quantum options.¹³⁷ AMD are developing a chip based on 'VEGA' graphics architecture¹³⁸ and NVidia is developing GPU based solutions.¹³⁹ Matlab provides a Deep Learning toolbox¹⁴⁰ which can be used to train a convolutional neural network or long short-term memory networks.

B.8.3.4 Niche optimisation solutions

There are very many areas of mathematical and scientific modelling where optimisation is critically important as well as many technology applications including bio-technology, chemical and manufacturing process control, data analysis, engineering design, environmental management, financial planning, logistics, packaging and risk management. The process of optimisation is analogous to finding the overall lowest point on a landscape of hills and valleys. Slightly more rigorously, the function to be optimised is imagined as a high-dimensional rippled sheet or surface; 'hills' and 'valleys' correspond to local extrema and the global optimisation process requires identification of the position of the lowest point on the entire surface. Such problems have been claimed to be solved most rapidly using quantum annealing machines, such as D-Wave (Section 1.2.2), which make use of quantum tunnelling to explore the sheet to find the lowest 'valley'.

Conceptually, the annealing algorithm starts at some point on the multi-dimensional surface and moves the point 'downhill' until a minimum is reached. It is then necessary to determine if this minimum point is the global minimum. If it is, the true solution has been identified. To determine if this the true solution, it is usual to repeat the minimisation (possibly thousands of times for a complex surface) from different starting positions and choose the overall lowest solution found which is hopefully the true optimum solution. Thus, with digital annealing there is a probability, which is difficult to quantify, that a solution is found quickly but that it is not the true solution.

In August 2019, Fujitsu announced it had developed a Digital Annealer.¹⁴¹ Whilst not a quantum machine or algorithm, nonetheless it uses a digital circuit design inspired by quantum phenomena and carries out a similar calculation – that is, performing a global optimisation to identify the global minimum of a multi-variate function. Although not claiming the technology will outperform future quantum computers, Fujitsu emphasise their Digital Annealer is simpler to operate (for instance, cryogenic temperatures are not needed) and programme while still allowing the solution of problems which are difficult for conventional classical computers. It is seen to have wide application to optimisation problems already managed through computational optimisation, including logistics, resource management, materials discovery and network control, as well as new areas including autonomous vehicles and healthcare.

¹³⁶ https://www.intel.com/content/dam/support/us/en/documents/boardsandkits/neural-compute-sticks/NCS2_Product-Brief-English.pdf; <https://software.intel.com/en-us/neural-compute-stick>; <https://www.intel.ai/nervana-nnp/>; <https://newsroom.intel.com/editorials/intel-creates-neuromorphic-research-community/>;

¹³⁷ <https://newsroom.intel.com/news/future-quantum-computing-counted-qubits/>

¹³⁸ <https://www.amd.com/en/products/professional-graphics/instinct-mi25>

¹³⁹ <https://developer.nvidia.com/discover/artificial-neural-network>

¹⁴⁰ <https://www.mathworks.com/help/deeplearning/ug/neural-networks-with-parallel-and-gpu-computing.html>

¹⁴¹ <https://www.fujitsu.com/jp/documents/digitalannealer/services/da-shoukaiEN.pdf>

B.8.3.5 Competition through Software

Currently quantum computers appear able to ‘run’ special purpose and perhaps ‘niche’ algorithms, as well as neural nets. Neural nets are general purpose if the task is pattern matching but they can also provide control loops and are a favourite of the Machine Learning community.

Because neural nets can place huge computational loads on computers, many other ‘ad hoc’ methods have been developed over the years, including hybrid approaches, to recognise features in images. Two mainstream examples are automatic number plate recognition (which now works well without needing a supercomputer) and facial recognition and/or face tracking in crowds. A quantum computer constrained to running neural nets must ‘defeat’ classical computers running this sort of honed, special purpose software. Specially ‘honed’ software may achieve the results a neural net will but with only a small fraction of the computing power.

B.9 Photonic quantum computers

The use of quantum light to represent quantum information is yet another paradigm of quantum computation. It has a number of attractive features including:

- Operation at room temperature without cooling;
- Very low power requirements;
- Chip-scale devices can be mass-manufactured using existing technologies;
- Simple architectures for QIP networks because integrated solid-state technology is common to photonic quantum computers and quantum communications.

Photons can be used and manipulated to represent superpositions of quantum states using linear optical components such as beam splitters, mirrors, waveplates and phase shifters (Mach-Zehnder interferometers¹⁴²). Linear optical elements apply unitary transformations to the qubits (sometimes called ‘flying qubits’) represented by the quantum light and an appropriate circuit of linear optical element can be designed comprising any quantum circuit model. Continuous variable quantum computing is also possible using the linear optics approach.

Although the individual elements in the circuit preserve the statistics of quantum light passing through the circuit, and so photonic quantum computing is not degraded by ‘noise’, there are two significant problems which must be overcome in order to build a practical photonic quantum computer. The first is photon loss due to scatter or absorption as light passes through the circuit. The second arises from the very small interaction of photons with other photons which adds complexity to the practical realisation of quantum gates; one approach to solve this problem is to use non-linear optical elements (such as Kerr cells) in the circuit. It was proved in 2001, however, that universal quantum computers can be built solely using linear optical elements.¹⁴³

¹⁴² A Mach-Zehnder interferometer is commonly used in linear optics to measure or control the phase differences between two beams of light. See

https://en.wikipedia.org/wiki/Mach%E2%80%93Zehnder_interferometer

¹⁴³ The KLM scheme, Knill, Laflamme and Milburn ‘A scheme for efficient quantum computation with linear optics’, Nature. **409** (6816): 46–52 (2001)

APPENDIX C Quantum algorithms of primary importance in Era 1

C.1 Introduction

The strengths and weaknesses of quantum computers are unlike those of digital computers and there is an ongoing search for algorithms that harness the particular strengths of quantum computers with a specific focus currently on noise-tolerant algorithms suitable for NISQ machines and available on a 0 – 5 year timescale.

The following list is not exhaustive but it includes those most often cited as capable of having, potentially, transformational effect and judged to be likely to be relevant to Defence and Security. Note that quantum algorithms can also be executed on emulators¹⁴⁴ running on digital computers, but less quickly.

- Shor's algorithm;
- Grover's algorithm;¹⁴⁵
- Quantum Fourier Transform;
- Quantum annealing;
- Machine Intelligence and Learning.

In the following sections, the nature of each quantum algorithm and its likelihood of impacting on military operations are summarised.

C.2 Shor's Algorithm¹⁴⁶:

Why it matters:

A convenient and robust approach to encoding and decoding digital messages or data is provided by the RSA algorithm¹⁴⁷. In this system, senders and receivers have a private key known only to them and a public key exists which is openly available. RSA is based on the principle that if a large number (the public key) is created as the product of two large prime numbers (the two private keys) then it is computationally impractical to reverse engineer the private keys from the public key.¹⁴⁸

If RSA is 'broken' then old intercepted data becomes de-codable and unsafe. However, it is envisaged that alternatives to Rivest-Shamir-Adleman (RSA) (including hashing, Data Encryption Standard (DES), Advanced Encryption Standard (AES), International Data Encryption Algorithm (IDEA) and Diffie-Hellman algorithms) will come into general use such that communication in the intermediate term will be 'safe' from this particular attack method. For the longer term, various post-quantum cryptographic methods are being developed.

Technical note, how it works:

¹⁴⁴ The ATOS Quantum Learning Machine is an example. It runs on Bull HPCs and is limited to algorithms requiring no more than about 40 (logical) qubits. Emulators can model different amounts of noise and give a good understanding of how NISQ machine performance would differ from that of a 'perfect' machine.

¹⁴⁵ This is a special case of an algorithm which does not need a fault tolerant computer – i.e., it will tolerate some noise.

¹⁴⁶ Biography of Peter Shor at https://wikipedia.org/wiki/Peter_Shor

¹⁴⁷ <https://tools.ietf.org/html/rfc8017>

¹⁴⁸ This is an example of an NP-hard problem. 'NP' denotes a non-deterministic polynomial problem all of which are believed to be difficult to solve unlike P (polynomial) problems which are considered easy to solve. Examples are multiplication (P) and factoring (NP). One of the long-standing problems in computer science is proving or disproving P=NP.

Shor's algorithm is a 'toolkit' of parts that achieve this factorisation by looking for a period in the function $f(x)$ which takes the form $f(x) = a^x \bmod(N)$ described in detail later. It achieves this by:

- Converting the problem to a related problem which, if solved, makes the original problem soluble;
- The first stage uses digital computers;
- The second stage uses a quantum Fourier transform.

Shor's algorithm can be implemented on digital computers, including an equivalent to the quantum Fourier transform.

If the public key is the integer N and its factors (the two private keys) are P_1 and P_2 then

$$N = P_1 \times P_2.$$

While there are many digital computer algorithms than can calculate P_1 and P_2 given N , for very large N they all require extreme levels of computing power and time, which can readily be made more extreme by making N even larger. For every bit added to the number N 's bit length, the required computing power doubles.

Shor's algorithm, instead, factorises $(P_1 - 1) \times (P_2 - 1)$ for which there is a known solution provided by the mathematician Euler. Shor calculates the series of numbers provided by the function $f(x) = a^x \bmod(N)$ for $x = 0, 1, \dots$ until the values of $f(x)$ begin to repeat. a is an arbitrary number (for ease of calculation, chosen initially to be small, e.g. $a = 2$, and then subsequently increased as needed). Not all values of a are acceptable indicated by the absence of periodicity in $f(x)$ over $x = 0 \rightarrow \sqrt{N}$.

The modulo arithmetic ensures $f(x)$ is never larger than the Public Key, allowing each term of the series to be calculated by one multiplication and one division. The value of x when $f(x)$ starts to repeat will be a factor of $(P_1 - 1) \times (P_2 - 1)$.

Shor's algorithm collects these factors (of which there will be many; several of them small, e.g. 2 and 4 will always be factors) and uses the quantum Fourier transform (QFT) to find which combination of factors, C_f , describes either of $(P_1 - 1)$ or $(P_2 - 1)$. Success is apparent because $P_1 = C_f + 1$ and if N/P_1 is an integer, then that integer is P_2 .

This was trialled using a simple BASIC program running on an iPad. The algorithm was capable only of dealing with double precision arithmetic values of N and the programme became slow even with values of P_1, P_2 of a few hundred. Often $f(x)$ directly delivered the value of either of $(P_1 - 1)$ or $(P_2 - 1)$ at an early stage in trials of a .

Running the QFT on a quantum computer gives the possibility of an exponential speed-up important for large N ; however, for large N , practical considerations have led to significant debate about the feasibility of the method. There are many reports of 'stunt' demonstrations; thus Jiang et al¹⁴⁹ claimed to have factorised $376289=571 \times 659$ using the D-Wave 2000Q but the method used is not generally applicable and is unsuitable for breaking RSA (and achievable even with the IBM 16 qubit quantum computer). There have been similar claims, most recently¹⁵⁰ that

¹⁴⁹ <https://arXiv:1804.02733>

¹⁵⁰ F Grier, 2018

$$(383123885216472214589586724601136274484797633168671371) = (618970019642690137449562081) \times (618970019642690137449562091)$$

which can be rewritten as $2^{178} - (13 \times 2^{91}) + 651 = (2^{89} - 21) \times (2^{89} - 31)$. These examples are included to show that, while there are special exceptions, quantum computing is not yet ready to factorise numbers of the sizes used for cryptography.

Defence and Security significance:

Shor's algorithm can break RSA encryption but from our simple experiments this seems unlikely to happen for some decades, if at all. The problem is twofold.

- The $f(x)$ functions are calculated on a digital computer and the computational load becomes very extreme for large N such that it is difficult to reach the point where the QFT can be fed the factors;
- The QFT requires at least \sqrt{N} qubits, one for each value of the $f(x)$ series which may have approaching \sqrt{N} elements and this would appear to demand a quantum computer of minimum (logical) qubit size 2^{512} for a 1024 bit RSA code. In the context that currently the biggest quantum computers still have (physical) qubits measured in tens, this is an extremely large number. The more qubits that are fully entangled the difficult it is to achieve error free operation and qubit number of 2^{512} is unlikely ever to be accomplished.

The issue for Defence and Security is the security of encrypted data and data transfer and advice from GCHQ/NCSC (who continually investigate the crypto-threat from QIP and are the National Authority for cyber security) will be critical to determine future tactics and strategies..

The difficulty in implementing Shor's algorithm using quantum computers of realistic size does not mean that RSA is safe. Grover's algorithm (below) also appears to offer an attack route and needs fewer qubits.

C.3 Grover's algorithm:

Why it matters:

Grover's algorithm searches a quantum database for a value, and returns the location of the data. In this, it is analogous to an internet search engine. It only returns one location; if there are several data entries of the same value then several searches will be needed.

Grover's algorithm is claimed to be very much faster than established search methods for very large databases in that it requires \sqrt{N} steps to search, where N is the number of locations containing data. Practitioners regard this quadratic speed-up claim as an exaggeration; the maximum speed-up is limited by how much serialisation can be tolerated, so a saving significantly greater than 20 or 30 bits of security is difficult to imagine in the near to medium term.

Finding 'needles in data haystacks' is standard problem in data management and so Grover's algorithm addresses an extremely important area but it is probably true that advances in classical computing will prove to be of greater overall significance.

Technical note, how it works:

The simplest version of Grover's algorithm requires that a relationship between the data held at a location and the location index be known¹⁵¹, and it inverts that relationship using a quantum computer to find the index given the data, provided $D_{index} = f_{index}$ is known. (It is not clear if this limitation can be overcome for all real databases but there will be many cases where it can be.)

Defence and Security significance:

The utility of Grover's quantum algorithm is doubtful because good alternative solutions exist based on classical digital computers. The premise that Grover's algorithm is superior is founded on the idea that a classical system will need to search every one of the N locations to find what it needs, which will take N steps rather than \sqrt{N} . In reality this assumption is flawed.

For example, internet searches do not require the search engine to go through every piece of data on the whole web, looking for the search string. Instead an indexing technique is used such that (much smaller) index tables are searched to find the location. Large databases routinely use indexing methods that vastly improve on linear search methods.

In addition, there is an existing electronic solution in mass production, developed precisely for the purpose of finding values in tables. This is the Contents Addressable Memory¹⁵² (CAM) memory, which returns the location given a data value. These are widely used in data routing servers e.g. those of CISCO Systems. Implementation is often by using FPGA (Field Programmable Logic Array) chips and the location is returned in one memory cycle (a few nanoseconds). Current quantum computers operate with a response time of the order of 20 – 100 microseconds (depending on type), so start with 3 - 5 order of magnitude speed disadvantage.

Electronic CAM is much more expensive than other forms of digital memory since it requires a larger area of silicon on the integrated circuit (and hence fewer chips to the wafer). In order for Grover's algorithm to provide this function it will need to be cheaper or smaller than a digital implementation which is already very fast.

Despite the expected future needs to search extremely large data sets, Grover's algorithm is not likely to be pre-eminently important over the coming decade because of these classical methods offering better performance.

C.4 Quantum Fourier transform:

Why it matters:

The quantum Fourier transform is the standard Fourier¹⁵³ data transform but implemented on a general-purpose quantum computer. This requires a number of qubits not less than the number of data samples to be processed.

¹⁵¹ More complex schemes can be used when the relationship isn't known and also schemes which use Shor's algorithm to count solutions when the number of solutions to the problem is unknown.

¹⁵² While there are many different designs of CAM and a number of academic publications detailing particular examples, the generic concept is set out in https://en.wikipedia.org/wiki/Content-addressable_memory

¹⁵³ Jean-Baptiste Joseph Fourier lived from 1768-1830. While electronics and signal processing is often perceived to be a fast moving research area, its present day mathematical mainstays are underpinned by 15th and 16th century scientific study https://en.wikipedia.org/wiki/Fourier_transform

It is hard to overstate the importance of the Fourier transform to signal processing. It is probably the default approach to signal extraction or detection and can be applied to time varying data and/or 2-dimensional data such as images. Modern radars, sonars and communications systems make heavy use of Fourier transforms. They are less used in image processing because of the extremely high computational loads involved.

Short length Fourier transforms (up to 1024 elements) can be executed in hundreds of microseconds using built in hardware and software in single chip classical computers and FPGAs. Special purpose hardware can execute a transform in $\log_2 N$ cycles, where N is the number of data points. For a 1024-point transform that means 10 cycles and a cycle might take 25 nanoseconds. In radars it is common to have dedicated Fourier units processing data in under a microsecond. Usually the system architecture sets the timeframe and the electronics limits the processable number of data points, N . In the authors' opinion, for the QFT to achieve wide adoption, it should be able to process data characterised by large N in one cycle in 20 microseconds or less.

Technical note, how it works:

The Fourier transform re-expresses a data series in terms of a superposition of sine waves of different amplitudes and phases. Historically, electronics has been heavily based on the use of resonant circuits which exhibit sine responses at particular frequencies and electronics engineers tend to think of signals and events in terms of their periodicity. Essentially the Fourier transform converts a time series into a set of frequencies and amplitudes, or the inverse. Usually a special version of the Fourier transform is used, the Discrete Fourier Transform (DFT). This can only deal with data sets whose length is an integer power of 2 e.g. 128, 256, 512, 1024 but is favoured for reasons of computational efficiency.

Defence and Security significance:

To displace existing Fourier analysis, the quantum implementation needs to be either cheaper or faster for a particular size of transform. At least for larger data sets (exceeding 256 elements for a 1-dimensional transform, exceeding 16x16 for 2-dimensional transforms), it is expected that quantum computers will offer faster, large-scale Fourier processing than possible with digital electronics. However, this would require general purpose quantum computers in excess of 256 qubits in size. Progress in the last decade has seen >50 qubit devices demonstrated and although a fully error corrected, 256 logical qubit machine is a significant challenge, construction of a 256 NISQ machine is feasible although it is hard to imagine their SWAP will be in any way similar to existing digital solutions.

In the event of large quantum Fourier transforms becoming feasible, they would be valuable for pattern recognition in large data sets. However, it is difficult to predict future adoption of QFTs because it is critically dependent on the quantum volume of available machines.

C.5 Quantum annealing

Why it matters:


Quantum annealing has become almost synonymous with D-Wave computers; more detail can be found in **Appendix B.7**.

Technical note, how it works:

Quantum annealing algorithms aim to find the global minimum of a function, defined in terms of a given set of states, by a process using quantum fluctuations. Quantum

annealing can be used in optimisation and sampling problems. In the former type of problem, the algorithm searches for the “best” solution, that is the combination of states which define the function which has the lowest cost function. Often one wants to minimise the energy of some arrangement of objects¹⁵⁴ and the “best” solution is that which has the lowest energy and corresponds to the globally lowest point on an energy landscape. The latter type of problem is encountered in machine learning problems where one wishes to build a probabilistic model of the cost function by sampling.

Quantum annealing is similar to simulated annealing (named for the similarity to the metallurgical process in which a metal is heating and cooled to improve the metal's strength) and is a classical probabilistic method of finding a global minimum of a many dimensional function, $E(s)$. The function $E(s)$ to be minimized can be regarded as equivalent to the internal energy of the system in state s and the aim is to identify s corresponding to minimum possible energy. In quantum annealing quantum mechanical tunnelling probability (from one state of the system to another) plays the role of temperature in simulated annealing.



Quantum tunnelling is a well-known phenomenon in which quantum objects have a probability of being observed in areas of space even though their kinetic energy is insufficient for them to be able to reach those areas. A simple classical analogy is the behaviour of a ball rolling across level ground on top of which exists a smooth hillock. If the ball has little kinetic energy (is rolling slowly) it may partly climb the hillock but falls back. Only if the kinetic energy is greater than the

additional potential energy the ball must gain to climb to the top of the hillock does the ball pass over. In contrast, a quantum ball, whatever its speed, will have a non-zero probability of being found on the far side of the hillock. The probability depends on the ball's kinetic energy (e), and the height (h) and width (w) of the hillock. The probability is greater as the ball's kinetic energy increases (constant h and w), the hillock's height reduces (constant e and w) or the hillock's width reduces (constant e and h).

Under circumstances where shallow minima are separated by tall, thin barriers, quantum annealing outperforms simulated annealing on a classical machine. This is because the probability of large thermal fluctuations is small, except at high temperatures, while for quantum tunnelling through a barrier, the probability varies $\sim e^{-w} \cong 1$ if the barrier width w is sufficiently small. It is expected that the computational efficiency of quantum annealing on a quantum computer would be greater than that found running the quantum algorithm on a classical machine.

Defence and Security significance:

Planning and Optimisation is a theme running throughout Defence and Security, not least in logistics. There are never sufficient assets to do all the tasks, or time to ideally assign what you have. Indeed it can be argued that the essence of combat is to concentrate your own assets where your opponent is least able to concentrate against

¹⁵⁴ eg. of spins on a lattice - the Ising model - which was the first formulation of the algorithm in its present form by Kadowaki and Nishimori in 1998 although Finnila and co-workers had described a related algorithm in 1994

you. In that sense Quantum Annealing addresses one of the most pivotal defence and security problems. There are very similar commercial analogies.

However it is no panacea, particularly over the next 5 or 10 years. In order to apply quantum annealing to a problem it must be expressed in the mathematics of complex surfaces. Digital computers have mature software design to translate ‘problems’ into the huge numbers simple operations digital computers can execute at speed. The equivalent software for quantum annealers is, by comparison, both novel and immature; furthermore it needs rare specialist expertise to apply it. It took decades, in which vast improvements were accomplished, to allow digital computers to achieve the broad application range they have. We should expect comparable delays before novel software render quantum annealers widely useful.

The possible exception lies in the ability of at least some quantum annealers to ‘run’ a form of software already well advanced as it can be implemented digitally, the neural net.

C.6 Machine intelligence and learning:

Why it matters:

In machine learning, ‘kernel’ methods are a ubiquitous class of algorithms for pattern analysis and classification, the best known of which are support vector machines (SVMs¹⁵⁵). The general task of pattern analysis is to find relations in datasets (including clusters, rankings, principal components, correlations and classifications). For many algorithms that solve these tasks, the raw data have to be transformed into feature vector representations via a user-specified feature map. Kernel methods, in contrast, require only a user-specified kernel (a similarity function comprising pairs of data points). When the feature space (data sets) becomes large the kernels become computationally expensive to estimate. Quantum computers have exponentially large quantum state spaces (of dimension 2^n for n qubits) and many believe that this will enormously benefit artificial intelligence.

The greater the amount of training data which can be input to the machine learning algorithms used to train AI systems the more ‘intelligent’ is the AI. This implies massive amounts of data must be input to the system, classified and analysed. Additionally, quantum computers are expected to revolutionise machine learning by efficiently sampling, in fine detail, computationally complex feature spaces and making possible the extraction of new insights from the input data. As quantum volumes increase (see Section A.2.1), quantum processors will be able to perform increasingly extensive feature mapping and a key question is at what threshold quantum volume will quantum processors outperform the most powerful classical computers?

How it works:

A recent paper published jointly by the IBM Watson AI Lab and MIT, published in Nature¹⁵⁶ describes the development of a quantum machine learning algorithm suitable for running on quantum computers expected to be available in the next 5 years.

The paper describes two quantum algorithms, suitable for running on NISQ machines, which were implemented on a 5-qubit transmon quantum processor (see Section B.2).

¹⁵⁵ SVMs are non-probabilistic binary linear classifiers. Given a training set of data, each datum marked as belonging to one of two categories, an SVM training algorithm builds a model that assigns new examples to one category or the other.

¹⁵⁶ ‘Supervised learning with quantum-enhanced feature spaces’ V Havlíček, A D. Córcoles, K Temme, A W. Harrow, A Kandala, J M. Chow and J M. Gambetta, Nature, **567**, 209 – 212 (2019)

The algorithms solve a problem of supervised learning: the construction of a classifier. One method, the quantum variational classifier, uses a variational quantum circuit to classify the data in a way similar to the method of conventional SVMs. The other method, a quantum kernel estimator, estimates the kernel function on the quantum computer and optimizes a classical SVM. However, there is one particular algorithm within the Machine Intelligence and Learning collection which the authors believe is of pivotal importance in the context of Quantum Information Processing for Defence and Security. This algorithm is the neural net.

Neural nets are a well-established set of methods hitherto run on digital computers. They identify patterns in data, including images. They are named neural nets because they have similarity to the operation of neurons in the brains of insects, animals and people. Neural nets usually deliver near class leading performance in pattern matching and are 'trained' on real data as opposed to being programmed in the traditional sense. They are a major strand of research in AI/machine learning. Their deficiencies are that they demand very high levels of computing power, are black boxes and require large datasets to avoid the inadvertent introduction of human biases. Once 'trained' neural nets work in a highly complex way and it is hard to determine how they have calculated a finding.

Artificial neural nets, typically, have tens to millions of artificial neurons, arranged in layers, to process and analyse information input to the network. Each layer is connected to those either side of it. The input layer receives information which is passed to subsequent, hidden, layers which transform the input data into meaningful information reported through the output layer. (See Figure 6 below) The 'best' neural networks are fully connected between layers. The connections are weighted and higher weightings correspond to greater influence on connected neurons, similar to the way brain cells trigger each other across synapses. Weights can be positive (the neuron excites those to which it is connected) or negative (the neuron inhibits those to which it is connected).




Figure 6: Schematic diagramme of a simple neural net

When a neural net is being trained or operating after being trained, data is input to the network via the input layer of neurons. This triggers the neurons in successive hidden layers until the signal arrives at the output layer - a feedforward network. Each neuron receives weighted inputs from the neurons to their left. The neurons sum all the inputs and if the sum is greater than a threshold value, the neuron 'fires' and triggers the neurons to which it is connected on its right-hand side.

Learning is accomplished by a feedback process called backpropagation in which the output produced by a network is compared with the output it was intended to produce. Working backwards from the output layer towards the input layer, the differences are used to modify the connection weights of the neuronal

connections. In an iterative process, backpropagation seeks to create sets of weights for which actual and intended output agree; when this is achieved the training is complete

Once trained using sufficiently large and representative data sets, the network will analyse new input data using the learned connection weights. The capability and accuracy of the network is only as good as the data used for learning. If the training data did not span the new input, the analysis will be meaningless.

Although neural nets have complex topologies, the calculations they perform for each node are simple. A node multiplies each of its many inputs by a different factor and sums the results. It then may apply a non-linear function to that summation to create its 'output'. When implemented on digital computers the number of 'weighted sum' calculations varies as the square of the number of nodes and rapidly becomes enormous and often intractable.

Neural nets usually deliver near class leading performance in pattern matching and are 'trained' on very large real data sets as opposed to being programmed in the traditional sense. They are a major strand of research in AI/machine learning. Once 'trained' neural nets work in a highly complex way and it is hard to determine how they have calculated a finding. A 'trained' net may not always be behaving as the trainer assumes.

Neural nets 'run' on quantum computers have the potential to operate at enormously higher speeds than when implemented on a classical digital computer. The size of a neural net can be represented by the number of nodes which make it up, ' N '. If analyzing an image feature, then the size N would be similar to (but larger than) the number of pixels. In a digital computer computational load scales as N^2 operations and leads to very slow execution times. In a quantum computer evaluation requires one operation, potentially complete in nanoseconds.

Defence and Security significance:

Section 3.3.1 describes the use of automated methods to enhance data analysis and identifies ML methods, such as the Automated Statistician (**Section 3.3.1.1**), which can produce simple text- and graphics-based reports autonomously when provided with sufficient data.

The detail and accuracy of these reports depends critically on the volume, and richness, of the data analysed and so quantum processors will be key enablers for accurate, sophisticated autonomous analysis of intelligence data.

The Authors regard the neural net as pivotal in the context of Defence and Security because it is capable of decomposing images into searchable elements, rendering images 'searchable' just as is done currently with text documents. In addition, neural nets allow computers to recognise features within an image e.g. identify the presence of a vehicle or recognise a face. A neural net-based processor could draft a text report of image content, but processing speed would be unacceptable using ordinary computers. Hitherto neural nets have only been applied to relatively simple data such as financial trading patterns and small image features such as letters or symbols e.g. optical character recognition. If it were possible to 'run' neural nets enormously more quickly than currently possible, their impact could be transformational not only to automatic image searching, understanding and filing but also to robotics and autonomy. This is of huge commercial as well as military importance and, in the view of the Authors, eclipses all other algorithms in a Defence and Security context.

APPENDIX D Quantum algorithms of secondary importance in Era 1

D.1 Introduction

In Eras 1 and 2, large scale, fault tolerant quantum computers are not expected to be available and thus practical algorithms must be able to exploit NISQ hardware and require relatively few qubits, few quantum gate operations and little or no error correction. Five quantum algorithms which the authors believe will be of primary importance to Defence and security are described in **Appendix C**. This Appendix describes a further five NISQ algorithms which may have niche value to Defence and Security.

As a group, these algorithms have wide ranging applications of interest to Defence and Security including:

- VQE (**D.1**): general optimization problems, quantum simulations, quantum chemistry;
- QAOA (**D.2**): many applications but of particular importance are machine scheduling, image recognition and the layout of electronic circuits;
- DDQCL (**D.3**): computer vision, speech synthesis, image and text analysis and molecular design for drug discovery;
- QAE (**D.4**): quantum simulation, data distribution across nodes in a quantum network, reducing quantum memory requirements in quantum communication channels and simplifying quantum circuits;
- PT (**D.5**): protein folding.

D.2 Variational quantum eigensolver, VQE

The VQE algorithm can be used to find the eigenvalues of any matrix. In this algorithm a quantum subroutine is run inside a classical optimization loop. The algorithm prepares and measures the expectation value of some operator on the quantum state. A classical non-linear optimizer subsequently minimises the expectation value by varying the parameters describing the quantum state. This is repeated iteratively until convergence is obtained. An implementation for a hybrid classical-quantum photonic computer was developed in O'Brien's group in 2013 and applied to the many (2!) electron molecular ion HeH^+ (published in 2014).¹⁵⁷ Ideally, the algorithm can achieve exponential speed up.¹⁵⁸

Applied to molecules, VQE finds ground state solutions of the Schrödinger equation. In the first step, the molecular Hamiltonian is converted into a qubit Hamiltonian (kinetic plus potential energies) by mimicking electron-electron interactions with entangled photonic (flying) qubits for a specified internuclear separation, R . Next, following the variational approach, a trial wavefunction is adopted and represented by one and two qubit gates. The first step requires more qubits as the size of electronic structure problem increases; the accuracy of the final solution is limited by the number of quantum gates available for the second step. In the third step, the energy of the trial state is evaluated given contributions of the various electron-electron (Coulomb and exchange) interactions defined in step 2 and in a fourth, classical, step these contributions are varied to give the overall lowest energy achievable. Using the wave-

¹⁵⁷ <https://www.nature.com/articles/ncomms5213>

¹⁵⁸ https://slidelegend.com/queue/arxiv171001022v2-quant-ph-9-oct-2017_5a178a491723dd541adfa98a.html

function found, a new molecular potential is calculated and averaged with the last iteration and then steps 2 – 4 are repeated until any change in the solution (energy or wave-function) is smaller than a preset threshold. The Variational Theorem¹⁵⁹ ensures that the solution is the best which can be found for the given flexibility (number of qubits) in the description of the electronic wave-function. In principal, the optimisation (step 4) could be run also on a quantum machine.

A significant limitation of the VQE algorithm for computational chemistry is the need to specify a form of the variational wave-function and there has been work to relax this constraint creating an adaptive algorithm which describes well highly correlated molecules even as they approach the dissociation limit.¹⁶⁰

D.3 Quantum approximate optimisation algorithm, QAOA

Optimization problems are usually formulated as minimization problems where some “cost function” is minimised to find the optimal solution. Optimisation is widely used in physics, mechanics, engineering and economics and as the complexity of the problem increases, ever more efficient solution methods are needed. Of particular importance are machine scheduling, image recognition and the layout of electronic circuits.

The QAOA is a family of hybrid algorithms which give super-polynomial speed up compared with classical algorithms and can be run on noisy quantum machines. In 2014, an algorithm was described which produces approximated solutions for combinatorial optimisation problems.¹⁶¹ Understanding its speed up over classical algorithms is still an active area of research and it is a popular candidate to use on defined problems to demonstrate quantum supremacy.

There is a significant literature devoted to the comparison of QAOA for specified NP-hard problems (such as Max-Cut for graph partitioning or colouring) with state of the art classical algorithms (such as AKMAXSAT). Simply, given a graph of nodes and edges, each node is coloured black or white. Each node next to one of the opposite colour scores a point. The aim is to find the colouring scheme (“cut”) that scores the most points. The classical algorithm requires all partitions of the graph, and their associated node sums, are enumerated.

The quantum algorithm imagines that sets of bit strings exist which correspond to the maximum cut of the graph and are equal to the ground state of a Hamiltonian (cost function) which is determined by first constructing a function, C_{ij} , that returns a 1 if the edge connects white and black nodes or 0 if the nodes are the same colour. Thus $C_{ij} = \frac{1}{2}(1 - z_i z_j)$ with the $\{z_i\}$ taking the values of ± 1 . The total cost is $\sum C_{ij}$. The corresponding Hamiltonian is $\sum_{ij} I - \sigma_i^z \sigma_j^z$. Functions which approximate the many body ground state of this Hamiltonian can be formed and measured. The measurement gives, to a high probability, the bit string corresponding to the maximum cut and the statistics improve as more measurements are made.

D.4 Data-driven quantum circuit learning, DDQCL

Following Turing’s paper “Computing Machinery and Intelligence” in which he posed the question “Can machines think?” research effort evolved towards addressing the

¹⁵⁹ [https://en.wikipedia.org/wiki/Variational_method_\(quantum_mechanics\)](https://en.wikipedia.org/wiki/Variational_method_(quantum_mechanics))

¹⁶⁰ <https://www.nature.com/articles/s41467-019-10988-2>

¹⁶¹ <https://arxiv.org/abs/1411.4028>

question "Can machines do what we (as thinking entities) can do?" and the field was named Machine Learning (ML) in 1959 by Arthur Samuel. Subsequently, ML has attracted increasingly intense research as society has become more data centric. The prospect of large-scale quantum computers has raised the possibility that they may allow more complex analyses to be carried out more quickly. However, big data sets imply the need for quantum machines with large numbers of logical qubits which are not expected to be available for at least a decade.

Accordingly, there has been much effort directed towards the development of quantum algorithms which can run on NISQ-classical hybrid machines and the DDQCL is one such approach which has become a leading research effort. Unlike the VQE and QAOA algorithms, which attempt to minimise a well-defined cost function, the DDQCL algorithm is probabilistic in nature and creates a generative model by sampling input data according to more than one cost function. The algorithm has been implemented on (trapped ion) circuit model and quantum annealing machines. Applications include computer vision, speech synthesis, image and text analysis and molecular design for drug discovery.

When the DDQCL hybrid algorithm is implemented on a circuit model, fully connected, N-qubit NISQ machine, the 2^N amplitudes of the wave-function are used to understand the correlations in the input data.¹⁶² The method assumes a (fixed) set $\mathbf{D} = (\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(D)})$ of D independent, randomly orientated vectors. The input data is mapped to the vectors $\mathbf{x}^{(d)}$ (for a simple black and white pattern the $\mathbf{x}^{(d)}$ map to ± 1) and the algorithm seeks to determine iteratively the statistical distribution, P_D , of the D vectors. The quantum circuit model assumes a wavefunction, $\psi(\vartheta)$, parameterised by the angle ϑ and following the Born interpretation, $|\psi(\vartheta)|^2$ represent probabilities $P_\vartheta(\mathbf{x})$ in terms of which a cost function is written, $C(\vartheta) = -\frac{1}{D} \sum_{d=1}^D \ln(P_\vartheta(\mathbf{x}^d))$. ϑ is then adjusted to minimise the cost using a classical optimiser. At any iteration the cost is approximated using samples from the data and measurements from the quantum circuit, hence the name "data-driven quantum circuit learning". A good approximation to the real distribution is obtained only if the model is sufficiently flexible; this flexibility results from the complexity of the quantum circuit. The drawback is that more flexible models are more challenging to optimize because of their larger number of parameters.

The quantum circuit comprises alternating layers of one- and two-qubit gates parameterised, respectively, by single qubit rotations, $\{\vartheta_i^{(l,k)}\}$, and two qubit entangling rotations, $\{\vartheta_{ij}^{(l)}\}$. The subscript denotes the qubits entangled by the operation and the superscript l denotes layer number of the qubit. The superscript k is needed in the one-qubit gate case as a rotation identifier (the implementation requires that arbitrary single qubit rotations are decomposed into three simpler rotations).

D.5 Quantum Auto-Encoder, QAE

An autoencoder is a type of unsupervised neural network used to learn efficiently representations of data, ie. to learn an encoding for a set of data. This allows the neural network to express the input in a lower dimensional space and ignore "noise". As well as reducing the dimensionality of the data, the ability to reconstruct data is learned as the autoencoder attempts to generate from the reduced encoding a close

¹⁶² <https://www.nature.com/articles/s41534-019-0157-8>

approximation to the original data. The most important application is in information retrieval systems.

The experimental realisation of a hybrid quantum-classical quantum auto-encoder algorithm has been described recently¹⁶³ and outlines a scheme in which photonic qutrits are compressed into qubit states. This permits reversible (lossless) compression when a set of quantum states does not span the full Hilbert space in which they are initially encoded. Potentially, quantum data compression can benefit many applications including quantum simulation, data distribution across nodes in a quantum network, reducing quantum memory requirements in quantum communication channels and simplifying quantum circuits.

In one approach, a 3×3 unitary transformation, U , is characterised by three input modes (the qutrit) and three output modes.

The input qutrits are encoded as superpositions over three single photon modes (one spatial mode supporting two polarization modes plus another spatial mode with fixed polarization). The transformation U is implemented as a sequence of 2×2 unitary transformations (physically realised with half- and quarter-wave plates which provide a simple, stable way of controlling the unitary transformation). Iterative training then seeks to minimize the occupation probability of the third output (“junk”) mode across the (training) set of input states. Lossless compression is achieved when the junk mode is unoccupied; the decoding step is just the inverse of U and the junk mode can be discarded giving a qubit output state. (When the compression is imperfect, there is a nonzero probability of finding photons in the junk output state and this is a measure of the error in the compression.) The cost function can be defined as the average junk mode occupation probability over the different training states.

Thus, the hybrid algorithm comprises:

- Preparing the input qutrit state as a superposition of one photon states;
- Implementing (as an optical circuit) the unitary transformation U from the input qutrit to output qubits, including a “junk” state;
- Calculating the cost function as the average junk mode occupation;
- Iteratively adjusting U to minimise the occupation of the junk state by a classical gradient descent optimisation routine.

D.6 Population transfer, PT

In fundamental science, transferring the state of a quantum system to a given distribution of populations is an important problem with applications in atomic and molecular physics.

In applied science, adiabatic manipulation of quantum states is an essential tool in quantum information processing and requires efficient algorithms. Accordingly, significant effort has been directed towards developing algorithms for population transfer between states, especially for energy minimum searches / annealing optimisations where the states have similar energies.

The adiabatic control of quantum eigenstates involves slowly changing some perturbation parameter (such as magnetic field) which controls the energy of a state adjacent to the rest of the states of the system’s Hamiltonian; if the change is

¹⁶³ <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.122.060501>

sufficiently slow, the Adiabatic Theorem¹⁶⁴ ensures the system remains in the same eigenstate. If the perturbation is applied too rapidly, the eigenstate cannot adapt and so the probability density changes little from initial to final state.¹⁶⁵ These rapid changes, called “shortcuts to adiabaticity” are useful in achieving faster guided evolution of the system toward the desired final state and bypassing the restriction imposed by the adiabatic theorem.

The use of shortcuts to adiabaticity have been proposed in adiabatic quantum computing, quantum annealing and “holonomic” quantum computing as a means of demonstrating quantum advantage. In quantum thermodynamics, the same idea has been used to suggest high-efficiency engines may be possible, by suppressing of quantum transitions during adiabatic cycles and in QIP, the same technique could be used for fast, fault-tolerant universal quantum computers and quantum repeaters in quantum networks.

¹⁶⁴ The Adiabatic, or Born-Fock, Theorem (1928) states “a physical system remains in its instantaneous eigenstate if a given perturbation is acting on it slowly enough and if there is a gap between the eigenvalue and the rest of the Hamiltonian's spectrum”

¹⁶⁵ In the limit of instantaneous changes – the diabatic or sudden approximation - there is no change in the probability density but there are usually no states of the final Hamiltonian with the same functional form of the probability density as the initial state and so the system ends in a linear combination of states of the final Hamiltonian such that the composite probability density equals the initial probability density.

APPENDIX E Quantum algorithms for Era 3

E.1 Introduction

The quantum algorithms described in this Appendix can be executed only on large scale, fault tolerant quantum computers not expected to be available before 2030. They require many qubits with long coherence lifetimes so that quantum information which is encoded among these qubits survives for at least time required to complete the quantum calculation.

The following list includes Deutsch's algorithm for historical interest, since it was the first to be described, but has little practical value plus two other algorithms which may have future value for Defence and security.

- Deutsch's algorithm: historical interest;
- Quantum Simulation: prediction of quantum properties of large numbers of entangled quantum objects;
- Linear equations: determination of simple relationships between an outcome and one or more variables that drive that outcome.

E.2 Deutsch's algorithm

Speed-up: exponential. Following Feynman's 1981 outline of a possible future quantum computer, Deutsch described a simple quantum algorithm in 1985, later improved together with Jozsa. Although of little practical use, it was a milestone in quantum algorithm development because it was the first of such algorithms, because it demonstrated exponential speed-up compared to a classical algorithm (which requires exponentially many calls) and because it is a deterministic algorithm and always produces a correct answer. Deutsch's algorithm was the inspiration for Simon's algorithm and, subsequently, Shor's algorithm. A brief description is included here for historical interest only since it is not expected to have significant value for Defence and Security.

Consider a one-bit function which can take either of two values, $f(x) = 0$ or 1 . Classically, to determine the value of the function requires two function calls but on a quantum machine, a one-qubit measurement (which will give one bit of information) can be chosen so that the one bit is a global property of the function, in this case $f(0) \oplus f(1)$. The algorithm makes use of information "hidden" in the structure of the problem to outperform the classical algorithm.

E.3 Quantum Simulation: Why it matters:

The prediction of the behaviour and properties of a large ensemble of entangled quantum objects is extremely hard to calculate using digital computers. When the objects are atoms in a molecule, the discipline is known as quantum chemistry. The problem grows exponentially with the number of interacting objects (N quantum objects interact in 2^N ways). However, approximate methods which run on classical digital computers have been developed and quantum chemistry¹⁶⁶ has been extremely

¹⁶⁶ The first application of quantum mechanics to a problem of chemical interest was Heitler and London's article in 1927 ('Wechselwirkung neutraler Atome und homöopolare Bindung nach der Quantenmechanik' Zeitschrift für Physik. 1927, **44**, 455–472) which gave the first mathematical description of a chemical bond. Subsequently much of the theory was developed by Born,

successful in predicting the behaviour of even complex molecules, to the extent that new chemicals can be designed to order with a high level of confidence in how they will behave. But this is only true if the quantum interactions are few in number or if approximations are made (for instance in Density Functional Theory, the 2^N non-local interactions between N electrons are replaced by a functional of the local electron density).

Quantum computers sidestep the complexity issue by setting up individual qubits to mimic the behaviour of quantum interactions in the target chemical or system. This allows them to scale up to relatively complex systems with moderate numbers of qubits.

In particular Quantum Simulators have the potential to substantially improve on digital simulations of complex molecules, such as pharmaceuticals, improving property prediction and options for synthesis (to give samples for testing) and they could become a key tool in the armoury of synthetic chemists. However, interviews of synthetic chemists conducted by BEIS's Innovate UK in 2015 as part of their Quantum Roadmapping programme did not expose major discontent with existing digital simulators used for quantum chemistry. In addition, quantum computer-based simulations of molecules are likely to require a much higher programming skill level than required to use digital models therefore the value of Quantum Simulation is uncertain. Understanding the folding and unfolding of proteins, however, is an area of bio- and medicinal chemistry which could benefit from large scale quantum simulators.

How it works:

Qubits are initialised to mimic the quantum interactions of the quantum system of interest such that the behaviour of the Quantum Simulator replicates that of the target system. Rapid 'what if?' experiments can then be conducted, including adjustment of the quantum system configuration to identify an optimal system for the desired use. This is essentially analogue computing, as exploited by the military in battleship fire tables where cams and gears were used to solve the fire control equations prior to digital computers becoming available.

Defence and Security significance:

The effect will be hard to predict as Quantum Simulators would essentially increase quantum chemistry productivity and accuracy. The analysis of pathogens, poisons and explosives and design of chemistries to protect against specified threats, could occur more quickly and accurately. While this is not considered likely to impact critically on Defence and Security activity, it has niche value by allowing better understanding of physical options for conflict. Wider defence could benefit from the discovery and design of strongly correlated materials (heavy fermion materials that include insulators and electronic solids difficult to describe with current quantum chemical methods) that show unusual electronic and magnetic properties.

Oppenheimer, Hartree, Fock, Pauling and Hückel. The development of digital computers in the 1960s enabled its systematic application to materials and processes of chemical interest to begin. Further advances in theory (which saw the development of Density Functional theory by Slater, Pople, Kohn and Sham) and computers now allow application to chemistry and quantum chemistry is widely used by the chemical and pharmaceutical industries.

E.4 Solution of linear equations:**Why it matters:**

Linear equations depict simple relationships between an outcome and one or more variables that drive that outcome. 'Linear' indicates that no quadratic or higher power laws are present in the relationship. Although this might appear absurdly simple for any 'real' problem it is usually possible to depict complex problems as sets of linear equations used within defined bounds. Even very complex problems such as the simulation of electronic circuitry often uses linear models.

Strongly correlated materials are a broad class of heavy fermion compounds that include insulators and electronic materials, and show unusual electronic and magnetic properties. A related topic is linear programming where a set of constraints is expressed in terms of linear equations and a 'zone' is automatically identified which meets all those constraints. Another is linear (or multiple) regression, a statistical technique which seeks to identify driving factors against an outcome from actual event data.

Graphics cards designed to synthesise imagery are based on linear equations. The individual calculations are fairly simple but millions need to be performed to build even a simple image and so computer graphics cards consist of arrays of very fast arithmetic processors, predominantly performing linear operations.

Analysis, modelling and systems design based on linear equations is extremely widespread. In general, linear equations can readily be solved by digital computers, provided the number of variables does not become large.

How it works:

Commonly a linear problem requires that a set of simultaneous (linear) equations be solved. Expressed in matrix mathematics, the core operation required to solve this problem is to invert an N -dimensional matrix, where N is the number of variables.

Subject to a number of constraints on the matrix, the Harrow and Lloyd algorithm¹⁶⁷ provides an efficient method of inverting the matrix using a quantum computer.

Defence and Security significance:

At present quantum computers are not large enough to solve linear equation problems beyond a scale that would be seen as trivial for a digital computer. As quantum computers grow this will become less true but problems involving very large numbers of simultaneous equations are not common.

However, a notable exception is route finding through a complex and dynamic set of obstructions, a problem facing autonomous robots.


It is not envisaged that quantum solution of linear equations will be important for Defence and Security purposes over the next decade. Strongly correlated materials are a wide class of heavy fermion compounds that include insulators and electronic materials, and show unusual electronic and magnetic properties.

¹⁶⁷ <https://arxiv.org/abs/0811.3171>

APPENDIX F Boltzmann, Hopfield, Hamming and Neural Networks

There is a tendency for different academic groups to use alternative terminologies when describing multi-node fully parallel computing networks. For example, neural nets run on D-Wave computers are often described as Restricted Boltzmann Machines (RBMs).

A Boltzmann Machine is a fully connected network, where all nodes connect to all other nodes



In this example the nodes v_1 to v_4 are the input nodes, h_1 to h_3 are the hidden layer. No output node is shown; one of h_1 to h_3 might be used.

In terms of topology the difference between a Boltzmann Machine and a neural net is that a neural net does not have connections between nodes in the same layer. In this example that means a neural net would not have direct links between any of v_1 to v_4 and have no direct links between any of h_1 to h_3 .

A Restricted Boltzmann Machine also has no direct links between nodes in the same layer and is often treated as if equivalent to a neural net.

However, the links depicted in a Boltzmann Machine may have a different meaning to the links depicted in a neural network. In a neural network the link depicts a single numerical value generated by the layer to the left and fed to the layer on the right. It is unidirectional and takes a single numerical value although in a practical implementation it might well provide a noisy signal (a probability distribution) or one made noisy by quantisation (number of bits in the digital signal). In a Boltzmann Machine the 'signal' can be a probability distribution and so is more general.

The question also arises as to whether a link is bi-directional or not, a matter of physical implementation. Some neural nets, such as Hopfield and Hamming nets, feed their outputs back as another input such that one of the nodes v_i is determined by the output state of the network and not externally. However, a Boltzmann topology implies that links might be bi-directional such that the values of the input layer v might have to be externally forced, else the hidden layer h can modify it.

A Boltzmann machine is an obvious way of describing a set of fully entangled qubits and so its adoption by the physics community in the context of quantum computing is

quite natural. However, the physical implementation of quantum computers varies and whether links take values or quantum distributions and whether the influence is bi-directional depends on the computer design.

Neural net design deliberately avoids the ambiguities of Boltzmann Machines (bi-directional links of uncertain meaning). If cross links are permitted in a layer, a neural net becomes much more difficult to implement on a digital (sequential) computer and the resulting problems make slow in operation and hard to train.

An everyday example of this 'cross-link' problem is encountered in Excel spreadsheets. To sum a column of numbers and put the result in a box, the column is selected and the '=Sum()' function used. But if the 'answer' box is accidentally included within the Sum() list, Excel stops and cites a 'circular reference error'. It is refusing to calculate a value not only because there is (probably) a mistake but also because the answer is uncertain when calculated sequentially.

In order to correctly compute a hidden layer in a Boltzmann (as opposed to a Restricted Boltzmann) Machine, the calculation must either be carried out in a truly parallel way or a problem arises. The value of any node cannot be calculated before the ones next to it are known. And so a 'chicken and egg' situation arises; a result will be achieved if the calculation is repeated until a stable answer is derived, but:

- It might not happen. The value could oscillate or become unstable;
- The wrong answer might be obtained, because it converges to a local minimum or maximum;
- It is slow because the calculation has to be repeated many times.

Neural nets avoid this simply by having no crosslinks in any one layer and neural net linkages are uni-directional. Output to input feedback, if present, is made explicit.

Overall a Boltzmann Machine is a more generalised depiction of a single cycle computing framework than a neural net, but the constraints deliberately adopted by neural nets have made them 'computable' on conventional digital computers and also enabled better progress on training methods.

An RBM is commonly used as 'shorthand' to indicate that the Boltzmann Machine has adopted the neural net representation norms i.e. lack of crosslinks in any one layer and (probably) uni-directional links.

APPENDIX G QulmP algorithms for circuit model quantum computers

There are very many quantum image processing algorithms for circuit model quantum computers, the principal ones being:

- Quantum Boolean Image Processing (QBIP);
- Flexible Representation of Quantum Images (FRQI);
- Novel Enhanced Quantum Representation (NEQR);
- Quantum State Tomography (QST).

QBIP algorithms are simple, working in the computational basis states (i.e., $|0\rangle$ and $|1\rangle$), and are particularly useful because of their tolerance to noise and relatively low computational cost.

FRQI algorithms map every pixel into a 4-dimensional basis with pixel position information encoded in a 2-dimensional qubit sequence and greyscale (or colour) information encoded as the probability amplitude of a single qubit. The principal advantage of the algorithm is using the superposition of qubit sequences to store the position information for all pixels allowing operations to be performed on all of them at the same time. However, because of the pixel mapping, FRQI is computationally intensive (scaling quadratically in the image size); other drawbacks with FRQI include the impossibility of accurate image retrieval and very poor image compression.

The NEQR approach is essentially an improved FRQI algorithm where two entangled qubit sequences, rather than a single sequence, are used to store the greyscale (or colour) and position information; the whole image is stored in the superposition of the two qubit sequences.

QST is a very important part of QulmP because it allows estimates to be made from measurements on unknown quantum states, represented in general by a density matrix. For an n -qubit state there are $2^{2n} - 1$ parameters that describe the state which determines the number of measurements required in the state estimation process. As for all quantum mechanical measurements, the observable is described by an Hermitian operator;¹⁶⁸ the realistic representation of the measurement process requires that the operator includes an appropriate term representing inevitable noise present in the input channel.

Although the above algorithms can process tri-colour (Red-Green-Blue, RGB) images, there appears to be only a single quantum algorithm reported to date which can process multi-spectral images. This is the Quantum Representation of Multi-Wavelength (QRMW) images and uses a superposition of basis states to store different wavelength data for each pixel in the image.

Clearly, the greater the amount of information contained in the image the larger is the number of qubits and execution time required by the image processing algorithm. Consequently, attention has been given to the development of quantum image compression (QIC) techniques. Typically, pixels are grouped by colour, position or channel (for multi-spectral images) and processing operations are applied to all pixels

¹⁶⁸ Hermitian, or self-adjoint, operators are represented by (in general, complex) matrices equal to the complex conjugate of their transpose. The eigen-values of an Hermitian operator are all real. Operators representing all physical observables are Hermitian

in a group instead of operating on pixels individually. Compression ratios as large as 50% have been achieved.

APPENDIX H TensorFlow Quantum (TFQ)

TFQ allows the construction of quantum data, quantum models and classical control parameters as tensors in a single computational graph. 'TensorFlow Ops' allows quantum 'measurements' to be made which return as probability distribution functions. Training is performed using the 'Keras' open-source Python neural-network library.

Google give as an example the supervised classification of quantum states using a quantum neural network (QNN). The challenge is to classify 'noisy data'. To build and train a model it is necessary to:

- Prepare the quantum data - quantum data is loaded as tensors specified as quantum circuits (written in Cirq) to create quantum data 'on the fly';
- Evaluate a QNN model – the QNN can be tested using Cirq and later embedded inside a TensorFlow compute graph. Based on the data structure, parameterized quantum models can be selected from drop-down categories. A successful QNN disentangles the input data exposing information hidden in classical correlations, making it available for measurement and processing;
- Sample the data – to extract classical information, TFQ provides ways of averaging over repeated instances of the previous two steps. The probability distributions depend on the quantum states and on the observable required;
- Evaluate a classical NN model – having obtained the classical information it can be further processed with deep NNs to extract information 'hidden' in correlations;
- Evaluate the Cost Function - desired cost functions can be evaluated from the classical processing;
- Evaluate Gradients / Update Parameters - after evaluating the chosen cost function, model parameters can be varied to minimise subsequent costs of processing.

TFQ can simultaneously train and execute multiple quantum calculations by parallelisation across a number of platforms. Google have released an open source emulator 'qsim'¹⁶⁹ (which mimics a 32 qubit quantum circuit with gate depth of 14 in 111 seconds on a single Google Cloud node¹⁷⁰). Combined with TFQ, Google has demonstrated 1 million circuit simulations for 20 qubit quantum circuit at a gate depth of 20 in 60 minutes on a Google Cloud node.¹⁷¹ TFQ will shortly be able to use the Sycamore quantum processor instead of an emulator.

¹⁶⁹ <https://github.com/quantumlib/qsim>

¹⁷⁰ <https://arxiv.org/pdf/1910.11333.pdf>

¹⁷¹ <https://arxiv.org/pdf/2003.02989.pdf>

APPENDIX I UK and other National quantum technology programmes

In February 2019, the Institute of Physics (IoP) journal ‘Quantum Science and Technology’ published an open access edition¹⁷² titled ‘Focus on Quantum Science and Technology Initiatives Around the World’ in which leading quantum physicists from Australia, Canada, China, Europe, Japan, Russia, the UK and the US review their respective national programmes in quantum technology. The comments below are based, in part, on these eight articles.

I.1 **Australia**

More than two decades of funding basic quantum physics and enabling technology development, as elsewhere around the world, laid a sound base for translational development of quantum technologies. Beginning in 2011, the Australian Research Council began supporting two national centres of excellence for quantum technology. That number has now risen to four and comprise the Centres for Engineered Quantum Science (EQUS), Exciton Science, Future Low-Energy Electronics Technologies (FLEET) and Quantum Computation and Communication Technology (CQC²T). Funding runs to 2023 – 2024. The Australian Department of Defence created a Next Generation Technology Fund in 2018 and has quantum technology as one of the priority areas. Total funding from the federal government comprises AUS\$130M with smaller sums provided by some state governments. A National Strategy began to be developed in late 2018.

The CQC²T, director Michelle Simmons, is developing silicon- and photonics-based quantum computers in which the qubits comprise photons of light (‘flying qubits’). The broad R&D programme spans the development of logical qubits from small-scale codes to cluster-states and surface-code¹⁷³ precursors with an overall goal of developing a scalable architecture with error correction (arising from photon loss) and quantum interconnects; integrated optics provides the qubit coupling. Work understanding how the design of quantum algorithms can be exploited in processor architectures is key to scale-up engineering. A particular challenge for photonic computing is the control of entanglement in photon pairs (photons do not interact directly) and two approaches are being developed to address this challenge: implementing a deterministic 2-photon entangling gate using a quantum memory and demonstrating continuous-variable entanglement through squeezing.¹⁷⁴ Ultra-long memories for optical quantum information are vital for many practical applications of QIP. CQC²T has a world leading capability for noiselessly interfacing flying qubits to long-lived atomic states and has demonstrated the highest efficiency quantum memory and the longest measured atomic coherences.

Large scale adoption of a silicon-based quantum processor depends on the successful development of scalable silicon chip engineering including the fabrication of arrays of sources and detectors, controllable electronics allowing interaction between the

¹⁷² https://iopscience.iop.org/journal/2058-9565/page/Focus_on_quantum_science_and_technology_initiatives_around_the_world

¹⁷³ Surface codes are quantum error correcting codes defined on a 2D lattice of qubits

¹⁷⁴ Squeezing a quantum system, defined by two or more non-commuting quantum parameters, moves measurement uncertainty from one parameter to another. In photonics, either photon numbers or phases can be squeezed; the squeezed mode can be measured more precisely but the squeezing strongly affects the other mode which becomes less well defined

quantum light and the classical environment and integrated circuits tailored to different photon entanglement schemes. This work is supported by unique facilities for nano-fabrication, including precise atom-implantation into device structures, ion-beam milling and nano-scale precision cutting and polishing. Similarly, the Centre has world class facilities for materials and device characterisation and assessment.

To commercialise their quantum computing research CQC²T is collaborating with companies launched by government and industry including Silicon Quantum Computing Pty. Ltd. Silicon Quantum Computing is owned by the Federal and State Governments, the Commonwealth Bank of Australia, Telstra (an Australian telecommunications company) and the University of New South Wales. Its goal is to develop a 10-qubit quantum integrated circuit prototype (using CMOS technology with phosphorus donors and silicon quantum dots) by 2022; CQC²T will provide scalable architectures with high frequency multiplexing and error correction.

I.2

Canada

Over the past decade, Canada has invested more than CAN\$1B in quantum science and has strengths in quantum computation and communications (for instance, the BB84 protocol¹⁷⁵) and is well placed to translate this into technology over the next 5 years. The private sector has been a strong driver of progress; Lazaridis and Fregin invested CAN\$150M to support the 2001 creation and operation of the Institute of Quantum Computing in Waterloo, Ontario. A recent investment is the independent, not-for-profit, Quantum Valley Ideas Laboratory for technology development. Canada also has a number of start-ups in quantum technology, most notably D-Wave Systems at Burnaby in British Columbia.

Government funding comes from the Natural Sciences and Engineering Research Council of Canada (NSERC) for university-based discovery research, innovation and training and has awarded CAN\$267M for quantum research between 2006 and 2015. Technology development is funded by the Canada Foundation for Innovation (CFI) and has invested over CAN\$100M, enhancing Canada's ability to attract and retain quantum researchers. The Canadian Institute for Advanced Research (CIFAR) is a private, not-for-profit institution that invests CAN\$25M per year for skills training, part of which is awarded for training in quantum research. The tri-agency Canada First Research Excellence Fund (CFREF) is supported by NSERC which seeks to build world-leading capabilities creating economic advantages for Canada and funds three quantum research programs in Canada at Institut Quantique, Université de Sherbrooke, (CAN\$33.5M in 2015 for quantum information and materials), Stewart Blusson Quantum Matter Institute at the University of British Columbia (CAN\$66.5M in 2015 for quantum materials and future technologies) and Transformative Quantum Technologies at the University of Waterloo (CAN\$76M plus CAN\$68M from partner contributions in 2015 for technologies to advance deployable quantum devices).

Internationally, through CIFAR funded quantum programs, Canada is engaged in many international collaborations whose goals are to advance the frontiers of quantum science and support and maintain Canada's development as a leading quantum nation.

¹⁷⁵ BB84 is the first cryptography scheme to be described and was developed by Bennett and Brassard in 1984 and is provably secure. Information in an intercepted signal can only be obtained by disturbing the signal itself and so BB84 provides a way of securely communicating a private key from one party to another for use in secure communications.

Canada's largest research organisation is the National Research Council (NRC) whose mandate is to support industrial innovation and advance knowledge and technology development to address current and future economic, social and environmental challenges. In 2014, NRC launched a CAN\$50M/5-year Quantum Photonics Sensing and Security R&D program focused on developing quantum technologies relevant to cyber security and sensors for environmental and health monitoring. Other government agencies engaged in quantum technology R&D to promote early adoption include the Communications Security Establishment (CSE) and Defence Research and Development Canada (DRDC). These national labs, respectively, are developing quantum cryptography and quantum sensing technologies, especially for Precision Navigation and Timing (PNT).

Like the UK, Canada has a strong photonics industry which is a key enabler for the development and manufacturing of quantum technology systems. A number of established companies, such as Imperial Oil, are exploring the adoption of quantum sensing systems and D-Wave Systems (quantum annealers) and ISARA (quantum-safe security solutions) are well established OEMs. Recent quantum computing start-ups include 1QBit, Anyon Systems Inc., Xanadu, Quantum Benchmark and RANOVUS.

Canada regards itself as the initiator of the race to develop a large-scale quantum computer. In the words of Raymond LaFlamme, sometime director of the Institute for Quantum Computing at the University of Waterloo, Canada 'fired the starting gun' with the launch of the Institute for Quantum Computing in 2002, years before large scale initiatives in China, the UK and the US began, and was the first country with a quantum computer manufacturer (D-Wave). The Institute for Quantum Computing is developing photonic- and nano-electronics-based QIP and is undertaking quantum algorithm research. D-Wave is described briefly in Section 1.2.2 and more detail in B.7 above. Xanadu is developing photonic QIP. University of Waterloo based Quantum Benchmark provide solutions for optimizing hardware design and quantum computing performance. 1Qbit, in Vancouver, is developing QUBO algorithms which run on quantum annealers as well as algorithms for quantum chemistry which run on circuit-based quantum processors. Anyon Systems is working with Google and its principal focus is to develop design tools for quantum electronics. RANOVUS is an OEM of Quantum Dot multi-wavelength lasers and digital and photonics integrated circuit technologies.

I.3 China

The IoP article,⁵⁶ co-authored by Jian Wei-Pan, estimates that through successive five-year plans, beginning in 2006, central and local government funding of Chinese R&D in quantum information science over the past decade totals about \$987M, although given the pace of Chinese progress many in the West believe the true figure is much larger. (Although titled 'Quantum information research in China' the article describes quantum metrology as well as quantum computing and communications.)

Since 2017, China has significantly increased the pace of its quantum R&D. In 2016, President Xi Jinping established a national quantum strategy whose aim is for China to become technologically self-reliant, surpass the United States and become the global high-tech leader. Following this, the creation of an \$11B, 37 hectares, National Quantum Laboratory (NQL) in Hefei was announced. China had already led the way with quantum communications in 2017 with a 2,000-kilometer long quantum network

linking Beijing, Shanghai, Jinan and Hefei, initially for banking, but there are plans for a global quantum network to which China is expected to transition its military communications. The Micius quantum satellite was launched in 2016 as part of a secure ground-space-ground quantum communications link and in 2018 was used for a secure 75-minute videoconference between the Chinese Academy of Sciences (CAS) in Beijing and the Austrian Academy of Sciences in Vienna during which 2 GB of data was exchanged.

The NQL has a 30-month, \$11.4B budget which is intended for both scientific research and technology development 'of immediate use to the Chinese armed forces' according to Jian-Wei Pan. In addition to targeting stealthy submarines, a large-scale quantum computer is planned, potentially targeting Western encryption systems. China is collaborating widely with the west, including with the United States, and, for instance, the 2013 development of quantum algorithms to solve linear equations was done collaboratively with Canada and Singapore.

The ultimate goal for quantum computation is to realize a programmable universal quantum computer and several technologies are being investigated with the objective of high-precision fault tolerant quantum logic gates. The developments might take a relatively long time, but China has a good track record of setting targets over long periods of time and adhering to plans. Quantum simulators (using unspecified technology) are planned for some special but important applications including the design of artificial photosynthesis systems and materials showing superconductivity at high temperatures. Over the next 5 years, China expects to demonstrate a quantum simulator which solves problems faster than can any classical computer.

Compared to their quantum communications efforts, China's quantum computing R&D has much greater private sector investment, similar to the US. The Alibaba Quantum Computing Lab, a collaboration between Alibaba's cloud computing arm, Aliyun, and the Chinese Academy of Sciences (CAS) established in 2015, is perhaps the best-known Chinese quantum computing programme and is carrying out research on systems believed to have the greatest potential for practical applications. Aliyun has experience in classical algorithms, architectures and cloud computing while CAS has experience in quantum computing.

Significant progress has been made in understanding the fundamental physics of cold atom and cold molecule systems, including the creation of topological quantum states in ultra-cold bosonic and fermionic quantum gases, and the work is an enabler for quantum simulators. For quantum photonic computing, InAs/GaAs quantum dots were developed which give single photons efficiently on demand and Wang *et al* demonstrated a 5-photon boson sampling machine which achieved an efficiency > 24 000 times greater than any previously reported machine. In basic science, Jian-Wei Pan's group has demonstrated quantum light with 5-photon entanglement, 6-photon CAT states¹⁷⁶ and up to 18-photon hyper-entanglement in spatial, polarization, and orbital angular momentum degrees of freedom. Optical quantum computing algorithms (Shor's, linear equation solvers, machine learning, teleportation and quantum cloud computing) have been developed for applications in code-breaking, big data and quantum simulation.

¹⁷⁶ A cat state, named after Schrödinger's cat, is a quantum state which simultaneously satisfies two diametrically opposite conditions

In the field of superconducting quantum computing, groups from the University of Science and Technology of China, Zhejiang University and the CAS are developing superconducting qubits. 10-qubit entanglement was achieved in 2017 and 18-qubit in 2019. Recently, quantum walks of strongly correlated microwave photons were demonstrated in a 1D array of 12 superconducting qubits. Objectives include: by 2020, to achieve the coherent manipulation of 30 qubits; by 2025, to develop quantum simulation with calculation speeds matching today's fastest supercomputers and by 2030, 'comprehensive realization of common-use quantum computing functions' through a quantum computer prototype with 50 to 100 qubits. If the work with topological quantum states can be applied to fabricate topological qubits, then a quantum computer with 100 topological qubits might be the equivalent of a superconducting qubit machine which has $10^5 - 10^6$ qubits.

In artificial intelligence, Google is working closely with China and the insight from this prompted Eric Schmidt, former chairman of Alphabet, to warn that China will overtake the United States in ML by 2025. China is investing in, and deploying, AI on a scale no other country is doing. It has announced \$Bs in funding for start-ups, launched programmes to attract researchers from overseas and streamlined its data policies so that researchers have access to large data sets (for instance, health data). News-reading and other robots are becoming ubiquitous and AI-powers foreign relations strategy but perhaps of most concern is targeted R&D to incorporate it into its military. Many of the ML algorithms which power these AI applications will use neural nets and, as noted above (Section 2.2), neural net software runs with little modification on quantum processors. Thus, success in China's quantum computer programme will almost certainly further promote Chinese AI dominance.

I.4 European Union

The European commission has invested more than €500M over many years in quantum physics research through the Future and Emerging Technologies programme (collaborative projects), the European Research Council (individual researchers) and the Marie Skłodowska-Curie programme (early career training). This activity has informed a European quantum technology Roadmap¹⁷⁷ and, together with the Quantum Manifesto,¹⁷⁸ led directly to setting up the Quantum Flagship, a 10 year, €1B programme to encourage Europe's strengths and overcome its weaknesses. The programme began in 2017 and its goal is to create a federated effort from the EU member states with three key objectives:

- Expand European leadership in quantum research, including training and skills;
- Ensure a competitive industry to position the EU as a future global leader in quantum technologies;
- Attract and support innovative research, businesses and investments in quantum technology.

The objectives of the Flagship for quantum computing are:

- In 3 years – demonstrate technologies for fault tolerant quantum processors;
- In 6 years – demonstrate an error-corrected or fault tolerant quantum processor;

¹⁷⁷ <https://iopscience.iop.org/article/10.1088/1367-2630/aad1ea>

¹⁷⁸ https://qt.eu/app/uploads/2018/04/93056_Quantum-Manifesto_WEB.pdf

- In 10 years – demonstrate quantum algorithms which achieve quantum supremacy.

The first call resulted in 20 funded projects (from 141 proposals) of which about a third address basic science or underpinning enabling technologies.

Two projects (using ion-trap and superconducting technologies aiming at up to 100 qubits) have as their objective the development of quantum computers that are competitive with state-of-the art conventional machines. Two quantum simulation projects will develop an atom / ion based programmable simulator and an ultracold atom device for quantum cascade frequency combs whose entangled modes will be used as qubits.

In addition to the Flagship, a 'quantum fleet' of other funding instruments are planned. The next framework programme, 'Horizon Europe', will support the Flagship through a new concept called 'Missions' awarded through a new organisation called the European Innovation Council. QuantERA, comprising 32 funding organisations from the 26 EU countries, will continue to support quantum technology (a second call worth €20M ended in February 2019; note this is coordinated investment from existing S&T allocations in EU countries, not new money).

Other important funding will come from:

- the European Space Agency which is already flying quantum technology demonstrators (examples being secure quantum communications and next generation time and frequency transfer);
- a follow on to the European Metrology Programme for Innovation and Research is planned by the European Association of National Metrology Institutes.

I.5

Japan

Japan has traditionally invested heavily in quantum technologies but mainly in the civil research sector. About ¥27B has been invested by the Japanese government in quantum research projects over the past 15 years across basic quantum science and quantum technology development, initially with principal foci being QIP and quantum communications. From 2016 onwards technologies such as quantum gyros and optical lattice clocks have been added and have received significant funding as have important enabling technologies including quantum cybernetics, control theory and quantum-classical hybrid systems. Space based technologies, such as constellations of satellites for communications, are also being developed.

In 2003, the Japanese Science and Technology Agency began the Core Research for Evolutional Science and Technology (CREST) project for QIP, metrology and sensing. Photonic, superconducting, neutral atom, atomic ensemble, and continuous variable approaches were explored. Subsequently (up to 2010), the use of spins on dopants in semiconductors, trapped ions and molecular vibrations and rotations were also explored. CREST was restarted in 2016 aiming for the 'Creation of an innovative quantum technology platform based on the advanced control of quantum states' and will run up to 2022 focusing on quantum simulators, sensing, imaging and quantum repeaters and state control.

Work during 2009 – 2013 claimed to demonstrate the extreme inefficiency of universal quantum computers (ignoring overheads due to quantum error correction).

Reproduced below is Table 1 from the paper by Yamamoto¹⁷⁹ in the IoP special issue referenced in footnote 61.

Table 1. Time-to-Solution for NP-hard MAX-CUT Problems.

Problem size	Universal quantum computation ^a	Heuristic		
		QAOA machine ^b	Quantum annealer ^c	Quantum neural network ^d
$N = 20$	4×10^{-3} (s)	600 (s)	1.1×10^{-5} (s)	1.0×10^{-4} (s)
$N = 50$	6×10^2 (s)	—	5.0×10 (s)	3.7×10^{-4} (s)
$N = 100$	2×10^{10} (s) (~700 years)	—	($\sim 10^{17}$ (s))	2.5×10^{-3} (s)
$N = 150$	6×10^{17} (s) (~20B years)	—	($\sim 10^{32}$ (s))	5.4×10^{-2} (s)

^a Theoretical limit (no decoherence, no gate error, all-to-all connections, 1 ns gate time).

^b Rigetti Computing 19 bit machine (Quantum Approximate Optimization Algorithm, Dec. 2017).

^c D-WAVE 2000Q with sparse connection (Experimental: $N = 20, 50$, Extrapolated: $N = 100, 150$).

^d NTT 2000 CIM with all-to-all connections (Experimental: $N = 20, 50, 100, 150$).

The results in this table appear to be not yet confirmed by other groups but, *if correct*, one implication is that quantum neural nets are the only practical method for large, computationally hard, problems.¹⁸⁰ A second implication is that a real-world circuit-model quantum computer requires large resources to correct errors; it was shown that a quantum computer with one and two qubit gate errors of less 10^{-3} requires at least 10^8 – 10^9 physical qubits and computational time of a few days to factor a 1024-bit integer number using Shor's algorithm.

As a consequence of these findings, quantum neural nets have been developed for special purpose optical-quantum computers available globally through a cloud service since 2017. Four systems are available: (i) a Coherent Ising Machine for solving NP-hard Ising problems; (ii) a coherent Boolean satisfiability problem (SAT) solver for solving NP-complete k -SAT problems (SAT problems in k variables); (iii) a coherent XY machine for solving continuous optimization problems; (iv) a coherent cryptomachine for solving problems on encrypted codes. Future work aims to understand the limitations of the computational power of quantum neural nets.

A new 10 year, ¥22B quantum information science and technology programme (Q-LEAP) managed by the Japanese Science and Technology Agency began in 2018 and covers quantum simulation and computation, quantum sensing, and ultrashort-pulse lasers. The first thrust of the programme consists of a flagship project for a superconducting quantum computer supported by 6 basic research projects developing hardware and software for quantum simulation and computation. The second thrust consists of a flagship project on solid state quantum sensors and 7 basic

¹⁷⁹ <https://iopscience.iop.org/article/10.1088/2058-9565/ab0077/pdf>

¹⁸⁰ An NP problem of size n is one for which the number of steps needed to check the answer is smaller than the value of some polynomial function of n . An NP-complete problem is an NP problem such that if a solution can be found requiring a number of steps which varies as a polynomial function of problem size, then solutions to all NP problems can be found similarly in polynomial times. NP-complete decision problems are the hardest problems in NP ('NP-hard') and no solutions in a polynomial number of steps have been found yet. The unsolved problem $P = NP$ asks if polynomial time algorithms exist for NP-complete and, therefore, all NP problems. It is widely believed that $P \neq NP$. The Travelling Salesman problem is an NP-complete problem

research projects developing quantum sensing technologies for different applications. The third thrust consists of a flagship project on advanced laser innovation supported by 4 basic research projects on attosecond¹⁸¹ pulse laser technologies and applications.

I.6

Russia

Quantum technologies in Russia are on the country's list of strategically important cross-cutting technologies in its National Technology Initiative and Digital Economy National Programme. Like the UK, the focus includes quantum computing and simulation, quantum communications, quantum metrology and sensing.

The Russian Quantum Center (RQC) was founded in December 2010 by Belousov (Parallels Inc), Lukin (Harvard) and Demmler (Harvard) funded through grants and privately (totalling about 2B rubles, ~£24M at 2019 exchange rates) and physically located within the Skolkovo Innovation Center in Moscow. Additional quantum technology research centres are the Kazan National Research Technical University in Tatarstan, the M.V. Lomonosov Moscow State University and the NTI Center for Quantum Communications at the National University of Science and Technology which together are supported with a further ~2B rubles from government and industry. A 5-year Russian Quantum Technologies Roadmap (RQTR) has been prepared and a budget of ~EUR1B requested.

The RQTR's objective is a cloud-accessible NISQ computer available within the next few years. Technologies being studied include superconducting qubits, neutral atoms, trapped ions, photons and polaritons; all of the experimental work is supported by theoretical studies. Software/quantum (and quantum inspired) algorithms under development include quantum error correction codes, quantum error suppression methods and large-scale emulators.

In 2017, Lukin announced the RQC had developed the then most powerful functioning quantum computer. The device comprised 51 cold atom qubits which was observed to show many-body dynamics corresponding to persistent oscillations of crystalline order and was used to realise novel quantum algorithms. The quantum processor was tested at Harvard and was claimed to solve problems which were difficult for conventional supercomputers solve.

One of the leading academics developing the mathematical foundations of quantum information theory is Professor A S Holevo at the Steklov Mathematical Institute of the Russian Academy of Sciences (RAS). Other work includes Dukhov's group who are developing methods to check quantum resource availability in NISQ devices and researchers at the Valiev Institute of Physics and Technology RAS are developing methods to control quantum systems, especially using quantum machine learning (QML). Also addressing QML, as well as quantum enhanced optimization and quantum enhanced simulation of electronic structure, is the Deep Quantum Laboratory at Skoltech which, in addition to research, has developed education programmes in quantum technology and quantum information processing.

¹⁸¹ 1 attosecond is 10^{-18} seconds.

I.7 United Kingdom

Quantum Technology is one of the 'Eight Great Technologies'¹⁸² plus 2 (Quantum Technology and Internet of Things) identified by the UK government to 'propel the UK to future growth.'

Starting in 2014, the government created a National Quantum Technology Programme (NQTP) and committed £270M to exploit decades of investment by EPSRC in fundamental quantum physics. By translating this basic science into technology, the government wishes to create a new quantum industry in the UK. Following a Blackett Review¹⁸³ and an House of Commons S&T Select Committee report,¹⁸⁴ a follow-on funding announcement in November 2018 brought the total committed funding from government and industry to about £1B. Part of this £1B, was a 2014, investment of £36M over 5 years by UK MOD to develop quantum sensors for PNT. The progress of the UK quantum technology programme in 2016 is reviewed [here](#). The NQTP comprises low TRL work (\leq MOD TRL 4¹⁸⁵) in universities clustered into 4 Quantum Hubs (Sensors and Timing led by Birmingham University, Imaging led by Glasgow University, Communications led by York University and QIP led by Oxford University) plus industry led, higher TRL projects funded by Innovate UK (IUK) through UK Research and Innovation's (UKRI's) Industrial Challenge Strategy Fund (ISCF) aiming to develop technologies towards commercialisation.

The UK regards its quantum computing research as world leading and has a rich mix of academic research groups pursuing almost all types of quantum processors plus significant effort in quantum algorithm development. The flagship research entity was the Networked Quantum Information Technology (NQIT) Hub in Phase 1 of the UK NQTP which became the Quantum Computing and Simulation Hub (QCS) in Phase 2 NQTP which began on 1st December 2019.

NQIT encompassed nine universities (Bath, Cambridge, Edinburgh, Leeds, Southampton, Strathclyde, Sussex and Warwick) and had connections to five other universities not formally Hub partners (Heriot-Watt, Bristol, Durham, Imperial College London and Sheffield). In addition, NQIT worked with more than 30 commercial companies (including IBM, Lockheed Martin, Raytheon BBN, Google and Toshiba) and government organisations (including the UK's National Physical Laboratory (NPL), Dstl and the US's NIST) plus small and medium-sized enterprises (including Rohde & Schwarz, Covision and Oxford Instruments). The ambitious goal was to understand how to build a universal, scalable quantum computer with error correction. In Phase 1, NQIT focused on ion trap, photonic, solid-state and superconducting platforms as well as quantum algorithm development.

Internationally, IBM has selected Oxford University as a partner in its Q-Hub (see Section B.2). In its first major collaboration in the UK, with a £5.5 million prosperity

¹⁸² <https://www.gov.uk/government/publications/new-eight-great-technologies-quantum-technologies>

¹⁸³ <https://gov.uk/government/publications/quantum-technologies-blackett-review>

¹⁸⁴ <https://www.parliament.uk/business/committees/committees-a-z/commons-select/science-and-technology-committee/inquiries/parliament-2017/quantum-technologies-17-19/>

¹⁸⁵ <https://data.gov.uk/data/contracts-finder-archive/download/713667/5f5887bd-8b04-4ab0-9388-c47f37514c1d>

partnership over five years with Bristol University and University College London, Google is developing quantum software for modelling and simulation.

Phase 2 will continue NQIT's work, broadening the consortium to 23 research teams in 16 universities and engaging with 35 commercial and government organisations. The programme will focus on:

- Simulation, especially focused on materials discovery;
- NISQ platform development to demonstrate, within the Phase 2 Hub, super-classical performance in areas of relevance to users outside the quantum technology field;
- Universal, scalable, fault-tolerant quantum computer development for general purpose applications.

In the UK, advancing readiness levels further beyond the Hub demonstrators usually involves technology transition to industry who, with varying degrees of financial support from IUK, develop technologies into pre-production prototypes and towards commercial products. The UK lacks large computer manufacturing companies to stimulate the commercialisation of UK R&D in quantum computing but UKRI is leading a programme to establish a National Quantum Computing Centre (NQCC) as part of phase 2 of the NQTP.

The NQCC¹⁸⁶ aims to build the UK's capability to be at the forefront of quantum computing, delivering greater prosperity and security advantages for the UK, as announced in the Budget in November 2018. Based at Harwell, the NQCC will be a dedicated national centre with the aim of working towards fully scalable, fault tolerant, general purpose quantum computing. The initial focus will be developing NISQ machines to demonstrate technologies, give assured and direct access to developers and drive the formation of a sovereign quantum computing supply chain and a large computer manufacturer which carries out the necessary systems engineering to produce an operating quantum computer. The UK already has a flourishing photonics sector, a key enabler for quantum technology systems including quantum computers, but there is a need for more a more diverse supply chain. The recent establishment of the Quantum Technology Leadership Group¹⁸⁷ is expected to facilitate development this. Encouraging a multi-national computer manufacturer to establish a UK manufacturing capability must also be achieved since the UK currently lacks this. The NQCC is expected to be fully operational by summer 2021 and deliver a NISQ computing capability that, for a range of tasks, outperforms conventional computers by 2025.

The UK has a strong record in developing and delivering conventional computer software (London is sometimes called 'Silicon Roundabout' in acknowledgement of this) and has a number of strong research groups developing quantum algorithms. The NQTP Phase 2 Oxford Hub includes more quantum algorithm development work than in Phase 1, but it is essential that industry collaboration is strongly encouraged and thrives. Fortunately, there are signs that this is happening.

¹⁸⁶ <http://uknqt.epsrc.ac.uk/about/nqcc/>

¹⁸⁷ <https://www.teledyne-e2v.com/news/teledyne-e2v-hosts-uks-quantum-technology-leadership-group/>

I.8 United States

The National Quantum Initiative (NQI) Act (H.R. 6227) was passed nearly unanimously by both houses of Congress and signed into law by President Trump on 21st December 2018. It authorised \$1.2B to be invested in quantum information science over five years; this funding will go to NIST, the National Science Foundation (NSF) Multidisciplinary Centers for Quantum Research and Education, the Department of Energy Research and the National Quantum Information Science Research Centers. An executive order established an NQI Advisory Committee comprising experts from industry, research and federal agencies. Almost immediately, the Department of Energy increased its quantum research funding by \$80M.

Many across the US government are beginning to consider advances in quantum science to be as important as previous national priorities, such as the arms and space races, and the council on foreign relations has labelled quantum science as ‘a race the United States can’t [afford to] lose. However, many view US government efforts to develop quantum technologies, especially quantum computing, as inadequate¹⁸⁸ and are looking to the private sector to make good federal deficiencies.

These concerns continue; in October 2019,¹⁸⁹ Scott Aaronson (David J. Bruton Centennial Professor of Computer Science at The University of Texas at Austin, and director of its Quantum Information Center) commented ‘China is ahead of the US right now in quantum communications, simply because they decided to invest a lot in that area while the US decided not to. I think that the US retains a strong lead in quantum computation with other important centers being Canada, the UK, the EU, Australia, Singapore, Israel.’ Similarly, the U.S. House Committees on Science, Space, and Technology, and its subcommittees on Research and Technology and Energy, all voiced concerns that the quantum sector in the United States was falling behind international competition and China, for example, is believed to be investing thirty times more than the US government in quantum technology.

In QIP, US private sector R&D companies (Intel, Google, IBM, etc.) are leading efforts to develop a scalable, fault-tolerant quantum computer based on superconducting qubit technology. Microsoft’s approach uses intrinsically error-free topological qubits that allows naturally scalable systems. Also making significant progress are start-ups Rigetti Computing, IonQ and Quantum Circuits. (See Sections **B.2** – **B.5** respectively for an overview of IBM, Google, Intel and Microsoft’s research.) Google claimed it had demonstrated ‘Quantum Supremacy’ in October 2019 although this was later contested by IBM.

I.9 Summary: UK prospects

Three factors that are necessary to bring quantum technologies out of the laboratory are:

- relevant use-cases with significant market potential;
- professional engineering on a large scale;
- significant research to overcome current scientific and technological limitations.

¹⁸⁸ <https://thehill.com/opinion/technology/419810-the-united-states-needs-better-quantum-science-as-a-national-policy>

¹⁸⁹ <https://www.forbes.com/sites/moorinsights/2019/10/10/quantum-usa-vs-quantum-china-the-worlds-most-important-technology-race/>

The UK is strongly positioned compared to the rest of the world, especially in lower TRL work although its ability to engage in professional engineering on a large scale is a matter for concern. The establishment of the NQCC will help develop demonstrators to higher TRLs but to achieve its stated aim 'to be at the forefront of quantum computing, delivering greater prosperity and security advantages for the UK' it is essential that large computer companies engage with the NQTP. Although the UK has a healthy supply chain for some components of QIP systems, it lacks sovereign computer manufacturers. It is encouraging, therefore, that IBM and Google have both entered into collaboration agreements during the past year (with Oxford and Bristol Universities and with University College London respectively) to pursue quantum computer development. To achieve the government's aim, however, it will be necessary to create and retain sovereign capabilities for manufacturing and software development. The former will be harder to achieve than the latter.

APPENDIX J Industry, Supply Chains, UK Capability

J.1 Introduction

Commercial providers of QIP may be classified according to their roles across the four layers of the full quantum computing stack: hardware, systems, system software and applications layers. Some organisations provide end-to-end capability (such as IBM, Google, Microsoft, Rigetti, D-Wave Systems and others) while some adopt niche positions, say in hardware or software.

At the current, early, stage of software and hardware development, the software development remains closely coupled to the hardware on which it is executed and the co-development of these technologies is essential for any nation which wishes to establish a leading position.

The UK does not have an established integrator but it does have flourishing spin-outs and start-ups developing software and hardware technologies and so the emerging relationships with IT major players (such as that between IBM and the Oxford Hub), together with the strong advanced manufacturing base already selling components such as vacuum systems, lasers and control systems, potentially positions the UK to be a future key player in quantum computing. The UKNQCC could be critically important in building a new, global-scale, quantum computing industry from these components.

J.2 End-to-end capability

The UK has no companies which operate across the full quantum stack

J.2.1 D-Wave

D-Wave and its activities have been discussed elsewhere in this document. See for instance **Section 2.2.2** and **Appendix B.7**.

J.2.2 Google

Google is developing quantum simulation and quantum machine learning (QML) algorithms to support academia, industry and national labs as part of the Google AI business. In terms of their own business, published information (unsurprisingly) implies Google plans to target markets open to data rich digital information business challenges such as financial services, portfolio optimisation, risk analysis, health care, logistics and data analytics and does not see a strong business driver for encryption breaking, partly because of the introduction of post-quantum cryptographic methods. The company sees quantum-assisted optimization and inference techniques as critical for machine-learning and artificial-intelligence systems which, they believe, could improve the management of renewable power generation, remote-sensing and early-warning systems as well as having value in managing on-line services and goods and warehousing. In line with Google's high-profile activities developing driverless cars, self-driving vehicles are seen as an important business area for QML.

Research areas are¹⁹⁰:

1. Superconducting qubit processors;
 - a. with chip-based scalable architectures targeting two-qubit gate errors of < 0.5%. Bristlecone, announced in March 2018, is Google's most recent

¹⁹⁰ <https://ai.google/research/teams/applied-science/quantum-ai/>

quantum processor with 72 qubits and Google are “cautiously optimistic” that, with system engineering to achieve optimally low error rates, equal to or better than their previous 9 qubit device¹⁹¹, it will allow demonstration of quantum supremacy;

2. Quantum simulation;
 - a. the focus is on quantum algorithms for modelling systems of interacting electrons with applications in chemistry and materials science¹⁹²;
3. Quantum neural networks;
 - a. developing a framework to implement a quantum neural network on noisy intermediate-scale quantum processors available now or in the near future.¹⁹³ The advantages which may be achieved by manipulating superposition of very large numbers of states is a key research objective;
4. Qubit metrology;
 - a. Google believe a two-qubit loss less than 0.2% is critical for effective error correction and are working to demonstrate quantum supremacy experiment;¹⁹⁴
5. Quantum-assisted optimisation;
 - a. Google are developing hybrid quantum-classical machines for optimization problems. These would take advantage of thermal noise to allow tunnelling to globally lowest energy states of the problem Hamiltonian (in much the same way as D-Wave).

J.2.3 IBM

IBM was among the first major companies to make available early QIP service for business, engineering and science, IBM Q, which comprises the entire quantum computing technology stack plus hardware accessible over the cloud through Qiskit from May 2016.

The hardware is based on transmon (a portmanteau from transmission line shunted plasma oscillation qubit) superconducting qubits (invented at Yale in 2007 and engineered from two capacitatively shunted superconductors to have low sensitivity to charge noise). The architecture is scalable and error correction can be incorporated. The hardware is stacked in layers whose temperature decreases from 4 K at the top of the stack to 15 mK at the base.

To-date, over 2.5 million experiments have been on the IBM Q platform¹⁹⁵ and more than 60 research papers published. One landmark publication was a detailed solution of the non-trivial problem of fully entangling 16 qubits.¹⁹⁶ This strong user engagement

¹⁹¹ Demonstrated readout and single gate errors of 0.1% and 2 qubit gate errors of 0.6%.

¹⁹² <https://github.com/quantumlib/OpenFermion>

¹⁹³ <https://github.com/quantumlib/cirq>

¹⁹⁴ On 20th September 2019, an article in the Financial Times reported a Google research paper temporarily posted on line which claimed "To our knowledge, this experiment marks the first computation that can only be performed on a quantum processor" - a calculation requiring 10,000 years on IBM's Summit HPC could be run in 3 minutes on Bristlecone. Google later declined to comment on the newspaper article.

¹⁹⁵ IBM have quantum hardware sites at Tokyo (20 qubits), Melbourne (14 qubits), Tenerife (5 qubits) and Yorktown Heights (5 qubits). Typical “clock speeds” are ~5 GHz, with T_1 and T_2 times 10 - 70 micro-seconds. Gate and readout errors are $(0.7 - 3.0) \times 10^{-3}$ and $(3.0 - 10.0) \times 10^{-2}$ respectively.

¹⁹⁶ Wang et al, <https://www.nature.com/articles/s41534-018-0095-x>

will be important in the future efficient, application orientated development of quantum computing. In the UK, only Oxford University is currently engaged (as an IBM Q-Hub regional centre of quantum computing education, research, development, and implementation which provides collaborators online access to IBM Q quantum technology) but there are many more overseas government and industry research organisations engaged as partners.¹⁹⁷ Current applications projects include quantum chemistry for drug and materials design and optimization for transportation logistics and finance.

In April 2018, IBM revealed the first start-ups joining the IBM Q Network with cloud-based access to IBM's quantum computers and other resources. These include:

1. 1Qbit: (Vancouver, Canada) builds quantum and quantum-inspired solutions for demanding computational challenges. Their hardware-agnostic services allow development of scalable applications. The company is backed by Fujitsu Limited, CME Ventures, Accenture, Allianz and The Royal Bank of Scotland;
2. Cambridge Quantum Computing (CQC): see 6.1.2.1;
3. Zapata Computing: (Cambridge, MA) provides quantum computing, services developing algorithms for chemistry, machine learning and security;
4. Strangeworks: (Austin, TX) develops QIP tools for software developers and systems management;
5. QxBranch: (Washington, D.C.) provides data analytics for finance, insurance, energy, and security customers. The company is developing quantum tools exploiting machine learning and risk analytics;
6. Quantum Benchmark: (Kitchener-Waterloo, Canada) is a venture-capital backed software company seeking to provide solutions which enable error characterization, mitigation and correction as well as performance validation of quantum computing hardware;
7. QC Ware: (Palo Alto, CA) develops hardware-agnostic quantum software for Fortune 500 companies including Airbus Ventures, DE Shaw Ventures and Alchemist as well as US government agencies including NASA;
8. Q-CTRL: (Sydney) is using its hardware agnostic platform (Black Opal) to improve quantum computer performance and reduce the lead time to QIP tools which can solve real world problems. Q-CTRL is backed by Main Sequence Ventures and Horizons Ventures.

In addition to real quantum hardware, IBM offers high-performance quantum simulation (Qiskit Aer) which can be accessed (through Qiskit or IBM Q Experience, see¹⁹⁸). This allows ideal experimental circuits to be tested before running on real hardware, the performance of which can be predicted by adding noise in a controllable way.

J.2.4 Intel¹⁹⁹

Intel's declared goal is a complete quantum computing system (hardware, algorithms and software and control electronics) has adopted two approaches to quantum

¹⁹⁷ <https://www.research.ibm.com/ibm-q/network/members/>. The network includes clients from Fortune 500 companies, academic institutions, and US national research labs, including JPMorgan Chase, Daimler, Samsung, Barclays, Honda, Oak Ridge National Lab, Oxford University and University of Melbourne.

¹⁹⁸ <https://www.research.ibm.com/ibm-q/technology/simulator/>

¹⁹⁹ <https://www.intel.com/content/www/us/en/research/quantum-computing.html>

computing: like many other research groups they are developing a superconducting qubit approach, exemplified by the Tangle Lake 49-qubit chip announced in January 2018. The launch of the 49-qubit chip happened only a few months after the announcement of the 17-qubit chip which was developed in conjunction with Intel's Dutch partners, QuTech and Delft University of Technology. The chips, made with a 'flip-chip' processing method, have an architecture allowing improved reliability, thermal performance and reduced RF interference between qubits while the fabrication process enables smaller features and scalable interconnects (and higher data flow on and off the chip) compared to wire bonded chips.

Intel are also developing a "spin qubits in silicon" approach which seeks to exploit Intel's many year's experience in silicon chip technology as a route to chip-scale quantum computers with many (millions) of qubits. Intel liken the technology to existing semiconductor electronics and transistors but differs by exploiting the spins of single electrons, manipulated by low-amplitude microwave pulses. Investment levels are modest (~\$50M) and this effort is at a lower TRL than their superconducting technology but may progress more rapidly, perhaps even overtaking the superconducting approach. A CMOS-based approach allows a high qubit density, which aids entanglement with neighbouring qubits. In February 2018, QuTech and Intel announced a 2-qubit silicon spin-qubit based quantum device which should be able to operate at ~1 K, compare to the ~20 mK necessary for superconducting qubit operation. Progress in other areas includes demonstration of an algorithm, a compiler and control electronics. And they foresee 'commercial' quantum computing as soon as 2025, a decade earlier than my estimates.

Intel do not link the two technology areas but have a strong interest in neuromorphic computing²⁰⁰ for AI. Strong AI might be possible in the future if Intel is successful in developing quantum processors with millions of qubits

J.2.5 Microsoft²⁰¹

In 1997 Kitaev²⁰² introduced the idea of topological qubits which can be used, conceptually, to build hardware which is immune to decoherence and thus does not require large resources devoted to error correction. Particular types of quasi-particle²⁰³ called 'anyons' cannot interact and so quantum states comprising anyons do not decohere. In 2005, experiments with gallium arsenide (GaAs) semiconductor devices in high magnetic fields and at temperatures close to absolute zero topological qubits have been claimed to be seen.²⁰⁴

²⁰⁰ <https://www.intel.co.uk/content/www/uk/en/research/neuromorphic-computing.html>

²⁰¹ <https://www.microsoft.com/en-us/quantum>

²⁰² <https://www.sciencedirect.com/science/article/abs/pii/S0003491602000180?via%3Dihub>

²⁰³ Elementary excitations from the ground state of a many particle system are usually called 'quasiparticles' if they are fermions or 'collective excitations' if they are bosons. They are a convenient way of simplifying the description of large systems. The charge carriers implied by the effective mass theory of semiconductors are well-known quasi-particles. Electrons travelling through the semiconductor are perturbed in a complex way by the surrounding atomic nuclei and electrons but the charge carriers of effective mass theory, treated simply as electrons with renormalized masses moving in free space, give a good description of many semiconductor properties

²⁰⁴ <https://arxiv.org/pdf/cond-mat/0502406v1>

Microsoft have been working for about 15 years towards the demonstration of a full quantum stack machine based on topological qubits working in a hybrid system as part of the Microsoft Azure environment. Their concept is to:

- Work in Visual Studio, using Microsoft Quantum Development Kit tools;
- Using Q#, write the solution code using Microsoft quantum libraries;
- Run a quantum simulation to debug and validate the solution;
- Once validated, run the solution on Microsoft's classical/quantum hybrid computer within the Azure environment.

Until a large scale topological quantum processor is built, Microsoft (in partnership with 1QBit) offer the open source Quantum Developer Kit (QDK) allowing Q# code to be written and run on an emulator (see **Appendix A.2.6**). QDK contains standard libraries (providing the basics of a computer programme such as addition, multiplication, etc. as well as the .NET and Python libraries), libraries for computational chemistry (such as coupled cluster methods for electronic structure calculations), machine learning and numerics. The numerics library contains the operations and functions that relate complicated arithmetic functions to the native machine operations; this simplifies the implementation of Oracular functions, such as Schor's algorithm. Utilities are provided including a resource estimator, which evaluates the number of qubits likely to be needed during a calculation, and a probability outcome estimator, which evaluates the expected probability of a measurement.

The full state simulator can rapidly exhaust available HPC resources if more than about 30 qubits are required and so the less capable Toffoli simulator is available; this can only simulate quantum codes requiring no more than X, CNOT, and multi-controlled X quantum gates but can be used with millions of qubits.

Microsoft's key application areas include optimisation, machine learning, simulation of quantum systems (including simulations of chemicals and chemical reactions) and cryptography.

J.2.6 Oxford Quantum Circuits²⁰⁵

The company, OQC, was set up in 2017 and claims to have the most advanced quantum computer in the UK and offers the full quantum stack. The hardware is based on the Coaxmon which has a 3D architecture, rather than the 2D structures used in other machines. This innovative approach makes possible much simpler systems engineering and allows scalable structures to be designed. Together with Royal Holloway, University of London, (RHUL) commercial qubits are fabricated using RHUL's Superfab and measured and operated in OQC's laboratory, currently hosted by Oxford University but soon to move into in-house facilities.

The superconducting Coaxmon²⁰⁶ structure is fabricated in tiled arrays on sapphire chips; each component comprises a coaxial circuit quantum electro-dynamics (QED) structure with the qubit and aluminium microwave resonator opposite each other on the chip. Copper-beryllium wiring for control and readout circuitry are co-axial and perpendicular to the chip plane.

²⁰⁵ <https://oxfordquantumcircuits.com/>

²⁰⁶ <https://arxiv.org/pdf/1703.05828>

Recently,²⁰⁷ a consortium of led by OQC (and including SeeQC UK, Oxford Instruments, Kelvin Nanotechnology, University of Glasgow and RHUL) was awarded £7M by IUK to industrialise the design, manufacture and test of Coaxmons as part of the concerted ISCF-backed effort to establish a UK manufacturing capability in quantum computers. The project will exploit know how and IP developed under Phase 1 of the UKNQTP and launches in August 202 with first deliverables expected in 2021.

J.3 **Quantum software for the systems and application layers**

J.3.1 **ATOS²⁰⁸**

Atos is a French multinational information technology service and consulting company headquartered in Bezons, France. It specialises in hi-tech transactional services, end-to-end communications and cloud, big data and cybersecurity services. ATOS launched their Quantum Learning Machine programme in 2016. It has four priority areas:

- Quantum algorithms for machine learning;
- Next generation hybrid architectures incorporating quantum co-processors;
- Quantum safe cryptography for the post-quantum age;
- The Quantum Learning Machine (QLM) which provides a complete programming and simulation environment for quantum software development, education and training running on Bull high performance computers (HPCs).

The Atos QLM emulates a physical quantum circuit model computer and can simulate up to about 40 qubits depending on the HPC memory available. The interface is written in Python and the Atos Quantum Assembler (AQASM) allows quantum gates to be defined or mixed with predefined gates/quantum programs imported from other developers.

The ATOS website lists six major national laboratories who have purchased, and operate, the QLM:

- The Hartree Centre (UK STFC);
- Argonne National Laboratory (USA);
- Technical University (Denmark);
- Oak Ridge National Laboratory (USA);
- CEA – Atomic Energy Commission (France);
- Campus Hagenberg (Austria).

J.3.2 **Cambridge Quantum Computing²⁰⁹**

Cambridge Quantum Computing Limited (CQC) is based in the UK but describe itself as a global quantum computing company. It was founded in 2014 and has a focus on quantum software for quantum-age cybersecurity.

CQC has developed a platform-independent quantum compiler (t|ket>), with a Python interface allowing access to quantum hardware from Google, IBM, ProjectQ, PyZX and Rigetti. This is being used to create application software for quantum chemistry

²⁰⁷ <https://thequantumdaily.com/2020/04/23/oxford-quantum-circuits-led-consortium-wins-grant-to-boost-quantum-technologies-in-the-uk/>

²⁰⁸ <https://atos.net/en/products/quantum-learning-machine>

²⁰⁹ <https://cambridgequantum.com/>

(targeting the design of pharmaceuticals and other chemicals, materials and agrochemicals) and quantum machine learning (targeting deep learning for time-series modelling, decision-making and optimization).

In 2019, CQC announced their photonics-based Ironbridge quantum machine. The website claims Ironbridge provides device-independent communications security guaranteed by the laws of quantum mechanics which can be used for post-quantum encryption algorithms, cached entropy generation for IoT devices, key generation for certificates and quantum watermarking.

J.3.3 Riverlane²¹⁰

Riverlane was founded in 2016 and is developing software which can exploit the capabilities of quantum co-processors in hybrid computers to accelerate the simulation of quantum systems. Applications include the development of new battery materials and pharmaceuticals.

In June 2017, the start-up raised €3.7 million in seed funding, principally from venture capital investors Cambridge Innovation Capital, Amadeus Capital Partners and Cambridge Enterprise. It will use this funding to demonstrate its technology across a range of quantum computing hardware platforms, targeting early adopters in materials and drug discovery. It will also plans to expand its team of quantum software researchers and computational physicists

J.4 Commercial providers of quantum hardware

²¹⁰ <https://www.riverlane.com/>

Initial distribution

	KIS	Dstl	electronic
--	-----	------	------------

Report documentation page

v5.0

* Denotes a mandatory field

1a. Report number: *	DSTL/TR121783	1b. Version	FINAL
2 Date: *	June 2020	3. Number of	vii+123
4a. Report UK protective	UK OFFICIAL		
4b. Report national caveats: *	NONE		
4c. Report descriptor: *	NONE		
5a. Title: *	Quantum Information Processing Landscape 2020: Prospects for UK Defence and Security		
5b. Title UK protective marking: *	UK OFFICIAL		
5c. Title national caveats: *	NONE		
5d. Title descriptor: *	NONE		
6. Authors: *	Andrew Middleton, Stephen Till and Lt Cdr Matt Steele RN		
7a. Abstract: *	<p>Technologies belonging to the 'First Quantum Revolution' came from understanding quantum mechanics; they are ubiquitous and critical to daily life. The 'Second Quantum Revolution' will see the introduction of technologies that exploit subtler quantum effects. Dstl has reviewed the quantum computing landscape, to identify potential opportunities for UK Defence and Security. This landscape report is one of the products produced.</p> <p>As was done with the UK Quantum Technology Landscape document prepared by Dstl in 2013 (DSTL/PUB75620), this version is being released to UK colleagues and Stakeholders for comment and additional input. The document currently considers primarily the academic QIP landscape with little about either the enabling technologies or the QIP UK industrial sectors. The authors would particularly welcome input and advice from colleagues in industry concerning these areas, especially given their dynamic nature and the need to protect UK intellectual property and know how.</p> <p>After review, key findings and recommendations will be added and the document will be finalised and issued.</p>		
7b. Abstract UK protective	UK OFFICIAL		
7c. Abstract national caveats: *	NONE		
7d. Abstract descriptor: *	NONE		
8. Keywords:	Quantum information processing		

* Denotes a mandatory field

UK OFFICIAL

9. Name and address of publisher: *	10. Name and address of funding UK Strategic Command HQ Northwood Middlesex HA6 3HP
11. Funding source	
12. Dstl project number: 710251	
13. Programme:	
14. Other report numbers:	
15a. Contract start date:	15b. Contract end 31 st March
16. IP conditions for report:	
17a. Patents: NO	
17b. Application number:	
18. Release authority role:	

Guidance on completing the report documentation page can be found on the [Gov.UK website](https://www.gov.uk).

UK OFFICIAL

UK OFFICIAL

THIS PAGE INTENTIONALLY LEFT BLANK

UK OFFICIAL

UK OFFICIAL



UK OFFICIAL