H2020-FETHPC-2014

Coordination of the HPC strategy

EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D4.5

Second report on crosscutting issues and disruptive innovation strategies

Final

Version: 1.1
Author(s): Carlo Cavazzoni, Cineca
Date: 24.08.2017

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. №: FETHPC-671558**	
	Project Title: European eXtreme Data and Computing Initiative	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: D4.5	
	Deliverable Nature: Report	
	Dissemination Level: PU	**Contractual Date of Delivery:** 31 / 08 / 2017
		Actual Date of Delivery: 24/ 08 / 2017
	EC Project Officer: Evangelia Markidou	

* - The dissemination level are indicated as follows: **PU** – Public, **CO** – Confidential, only for members of the consortium (including the Commission Services) **CL** – Classified, as referred to in Commission Decision 2991/844/EC.

Document Control Sheet

Document	**Title: Second report on crosscutting issues and disruptive innovation strategies**	
	ID: D4.5	
	Version: 1.1	**Status: Final**
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): D.4.5-final.doc	
Authorship	**Written by:**	Carlo Cavazzoni, Cineca
	Contributors:	Rosa Di Felice, University of South California and CNR-NANO Institute Matteo Paris, Università di Milano
	Reviewed by:	Guillaume Colin de Verdière, CEA Florian Berberich, JUELICH
	Approved by:	MB/TB

Document Status Sheet

Version	**Date**	**Status**	**Comments**
alpha	31/07/2017	Draft	
1.0	10/08/2017	Final version	
1.1	24/08/2017	Revision of final version	Fix according to review

Document Keywords

Keywords:	High Performance Computing (HPC), Extreme Computing, Exascale, Disruptive Technologies, Advanced Technologies, Quantum Computing.

Table of Contents

References and Applicable Documents

1 Barends, R. *et al.* Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits. *Phys Rev Lett* **111**, doi:ARTN 080502 10.1103/PhysRevLett.111.080502 (2013).

2 Grover, L. K. A fast quantum mechanical algorithm for database search. *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing. New York, New York, USA: ACM Press*, 212-219 (1996).

3 Mohseni, M. *et al.* Commercialize early quantum technologies. *Nature* **543**, 171-174 (2017).

4 Bonesteel, N. E., Hormozi, L., Zikos, G. & Simon, S. H. Braid topologies for quantum computation. *Phys Rev Lett* **95**, doi:ARTN 140503 10.1103/PhysRevLett.95.140503 (2005).

5 Gibney, E. Billion-euro quantum project takes shape. *Nature* **545** (2017).

6 Shor, P. W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *Siam J Comput* **5** (1997).

7 Feynman, R. P. Simulating Physics with Computers. *Int J Theor Phys* **21**, 467-488 (1982).

8 Lloyd, S. Universal Quantum Simulators. *Science* **273**, 1073-1078, doi:10.1126/science.273.5278.1073 (1996).

9 Olson, J. *et al.* Quantum Information and Computation for Chemistry. *arXiv:1706.05413* (2017).

10 Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-Gordon, M. Simulated Quantum Computation of Molecular Energies. *Science* **309**, 1704-1707 (2005).

11 Babbush, R. *et al.* Exponentially more precise quantum simulation of fermions II: Quantum chemistry in the CI matrix representation. *arXiv:1506.01029* (2015).

12 Babbush, R. *et al.* Exponentially more precise quantum simulation of fermions I: Quantum chemistry in second quantization. *New J Phys* **18** (2016).

13 Romero, J. *et al.* Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *arXiv 1701.02691* (2017).

14 Shen, Y. C. *et al.* Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure. *Phys Rev A* **95**, doi:ARTN 020501 10.1103/PhysRevA.95.020501 (2017).

15 Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602-606, doi:10.1126/science.aag2302 (2017).

16 Reiher, M., Wiebe, N., Svore, K. M., Wecker, D. & Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proc Natl Acad Sci U S A*, doi:10.1073/pnas.1619152114 (2017).

17 Haner, T. & Steiger, D. 0.5 Petabyte Simulation of a 45-Qubit Quantum Circuit. *arXiv:1704.01127* (2017).

18 Wecker, D., Hastings, M. B. & Troyer, M. Progress towards practical quantum variational algorithms. *Phys Rev A* **92** (2015).

19 Shen, Y. *et al.* Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure. *Phys Rev A* **95** (2017).

20 McClean, J. R., Kimchi-Schwartz, M. E., Carter, J. & de Jong, W. A. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys Rev A* **95** (2017).

21 Farhi, E. & Harrow, A. W. Quantum Supremacy through the Quantum Approximate Optimization Algorithm. *arXiv:1602.07674* (2016).

22 Lloyd, S., Mohseni, M. & Rebentrost, P. Quantum algorithms for supervised and unsupervised machine learning. *arXiv:1307.0411* (2013).

23 Rebentrost, P., Mohseni, M. & Lloyd, S. Quantum Support Vector Machine for Big Data Classification. *Phys Rev Lett* **113**, doi:ARTN 130503 10.1103/PhysRevLett.113.130503 (2014).

24 Brereton, R. G. & Lloyd, G. R. Support Vector Machines for classification and regression. *Analyst* **135**, 230-267, doi:10.1039/b918972f (2010).

25 Lau, H. K., Pooser, R., Siopsis, G. & Weedbrook, C. Quantum Machine Learning over Infinite Dimensions. *Phys Rev Lett* **118**, 080501, doi:10.1103/PhysRevLett.118.080501 (2017).

26 Lloyd, S., Garnerone, S. & Zanardi, P. Quantum algorithms for topological and geometric analysis of data. *Nature communications* **7**, 10138, doi:10.1038/ncomms10138 (2016).

27 Hernandez, M. & Aramon, M. Enhancing Quantum Annealing Performance for the Molecular Similarity Problem. *Quantum Inf. Process.* **16**, 133, doi:10.1007/s11128-017-1586-y (2017).

28 Choi, V. Minor-embedding in adiabatic quantum computation: I. The parameter setting problem. . *Quantum Inf. Process.* **7**, 193-209, doi:10.1007/s11128-008-0082-9 (2008).

29 Cai, J., Macready, B. & Roy, A. A practical heuristic for finding graph minors. *arXiv:1406.2741* (2014).

30 Perdomo-Ortiz, A., Dickson, N., Drew-Brook, M., Rose, G. & Aspuru-Guzik, A. Finding low-energy conformations of lattice protein models by quantum annealing. *Sci Rep* **2**, 571, doi:10.1038/srep00571 (2012).

31 Babbush, R., Love, P. J. & Aspuru-Guzik, A. Adiabatic Quantum Simulation of Quantum Chemistry. *Sci Rep* **4**, doi:10.1038/srep06603 (2014).

32 Curtarolo, S. *et al.* The high-throughput highway to computational materials design. *Nature Mat.* **12**, 191-201, doi:10.1038/nmat3568 (2013).

33 Sumpter, B. G., Getino, C. & Noid, D. W. Theory and Applications of Neural Computing in Chemical Sciences. *Annu Rev Phys Chem* **45** (1994).

34 Lorenz, S., Gross, A. & Scheffler, M. Representing high-dimensional potentialenergy surfaces for reactions at surfaces by neural networks. *Chem Phys Lett* **395**, 210-215 (2004).

35 Lucas, A. Ising formulations of many NP problems. *Frontiers in Physics* **2**, doi:10.3389/fphy.2014.00005 (2014).

36 Garnerone, S., Zanardi, P. & Lidar, D. A. Adiabatic Quantum Algorithm for Search Engine Ranking. *Phys Rev Lett* **108**, doi:Artn 230506 Doi 10.1103/Physrevlett.108.230506 (2012).

List of Acronyms and Abbreviations

AISBL	Association Internationale Sans But Lucratif (International Non-for-Profit Association)
AQC	Adiabatic Quantum Computation
ARC	Australian Research Council

ASCR	Advanced Scientific Computing Research Program
BDEC	Big Data and Extreme-scale Computing
BDV	Big Data Value
CBQC	Circuit-Based Quantum Computing aka gate-model QC or GMQC
CECAM	Centre Européen de Calcul Atomique et Moléculaire
CoE	Centres of Excellence for Computing Applications
cPPP	contractual Public-Private Partnership
CSA	Coordination and Support Action
CWI	Centrum Wiskunde & Informatica
D	Deliverable
DG	Directorate General
DoW	Description of Work
DW	D-Wave
EC	European Commission
ECMWF	European Centre for Medium-range Weather Forecasts
EESI	European Exascale Software Initiative
ENES	European Network for Earth System modelling
EPOS	European Plate Observing System
EQuS	ARC Centre of Excellence for Engineered Quantum Systems
EsD	Extreme scale Demonstrators
EU	European Union
FET	Future and Emerging Technologies
FP7	Framework Programme 7
GDP	Growth Domestic Product
GMQC	Gate-Model Quantum Computation
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HPC	High Performance Computing
IARPA	Intelligence Advanced Research Projects Activity
IDC	International Data Corporation
IESP	International Exascale Software Project
INVG	Istituto Nazionale di Geofisica e Vulcanologia (National Institute of Geophysics and Volcanology)
ISV	Independent Software Vendor
IT	Information Technology
KPI	Key-Performance Indicator
LANL	Los Alamos National Laboratory
LHC	Large Hadron Collider
M	Month
NQIT	Networked Quantum Information Technologies
NSA	National Security Agency
NSF	National Science Foundation

OS	Operating System
PM	Person Month
Q	Quarter
QAOA	Quantum Approximate Optimization Algorithm
QC	Quantum Computing
QEO	Quantum Enhanced Optimization
QPU	Quantum Processing Unit
QuAIL	NASA Ames Quantum Artificial Intelligence Laboratory
QuArC	Microsoft’s Quantum Architectures and Computation
QUBO	Quadratic unconstrained binary optimization
R&D	Research and Development
R&I	Research and Innovation
RFP	Request for Proposal
ROI	Return On Investment
SHAPE	SME HPC Adoption Programme in Europe
SHS	Social and Historical Sciences
SME	Small and Medium Enterprise
SRA	Strategic Research Agenda
SWOT	Strengths, Weaknesses, Opportunities and Trends
TRL	Technology Readiness Level
UCSB	UC Santa Barbara
US	United States
USC	University of Southern California
USRA	University Space Research Association
UvA	University Of Amsterdam
VQE	Variational Quantum Eigensolver
VU	Vu University Amsterdam
WG	Working Group
WP	Work Package

1 Introduction

Following EXDCI deliverable D4.1 “First Report on cross-cutting issues and disruptive innovation strategies” and taking into account recent scientific advances on quantum computing, we have chosen to put the focus on this topic.

The reminder of this report gives an analysis of the current state of the art of quantum computing. We conclude this report by a set of potential actions for the HPC community.

2 Quantum Computing

Quantum computing (QC) is essentially harnessing and exploiting the amazing laws of quantum mechanics to process information. A traditional computer performs operations on long strings of “bits”, which encode either a zero or a one. A quantum computer, on the other hand, uses quantum bits, or qubits. A qubit is a quantum system that encodes the zero and the one into two indistinguishable quantum states. The qubit takes its final value (0 ***or*** 1) when read. But, because qubits behave quantumly, we can capitalize on the phenomena of "superposition" and "entanglement". Superposition is essentially the ability of a quantum system to be in multiple states at the same time — that is, something can be “here” and “there,” or “up” and “down” at the same time. Entanglement is an extremely strong correlation that exists between quantum particles — so strong, in fact, that two or more quantum particles can be inextricably linked in perfect unison, even if separated by great distances. Thanks to superposition and entanglement, a quantum computer can process a vast number of calculations simultaneously. Think of it this way: whereas a classical computer works with ones and zeros, a quantum computer will have the advantage of using ones, zeros and “superpositions” of ones and zeros (see Fig. 1 Upper panel: digital computer gate can act on data one bity at a time. Lower panel: quantum computer gate can act on all possible states of the 8 Qbits at the same time.).

The intrinsic property of the superposition of states of a quantum system is the key feature that allows a quantum computer to compute solutions of complex problems at a speed which is beyond comparison with a typical digital computer, e.g. a quantum computer can in principle solve a problem with factorial complexity with a single instruction. Given that, a full featured quantum computer with a complete instruction set can in theory outperform any digital computer, but the reality is quite far from that.

These properties make QC very attractive for HPC, where quantum computing could be used in synergy with digital computers to speed-up complex applications otherwise unaffordable on digital supercomputers (of today and in the foreseeable future). Traditional HPC systems could take care of all the instructions which cannot be executed on the quantum computer (typically I/O, and control code) offloading to the quantum engine by mean of a dedicated library of computational kernels that can be reformulated using the quantum logic.

On the other hand, small quantum computers can be simulated today on large supercomputers, allowing for the development of the methods and software in a user friendly environment, well ahead of the availability of a real machine.

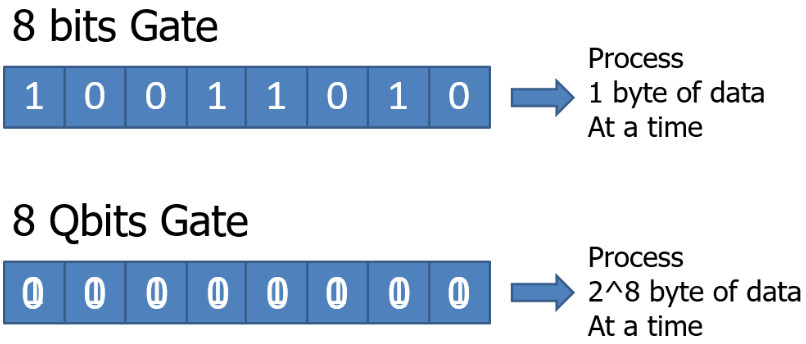

Fig. 1 Upper panel: digital computer gate can act on data one bity at a time. Lower panel: quantum computer gate can act on all possible states of the 8 Qbits at the same time.

2.1 Paradigms for quantum computation

As introduced above, quantum computers make use of the peculiar characteristics of quantum systems, such as superposition and entanglement, to improve performances. To make the point clearer: QCs **directly** exploit quantum effects and the word directly is important here, since ordinary semiconductor-based computers also use quantum effects, such as tunneling, whereas quantum computers must have accurate control over quantum states and operations, as well as an architecture such to prevent decoherence within the timescale of computation. **Controlling or removing quantum decoherence is one of the greatest challenges to fabricate quantum hardware**. This usually means isolating the system from its environment as interactions with the external world cause the system to decohere. However, other sources of decoherence also exist. Decoherence is irreversible, as it is effectively non-unitary, and is usually something that should be highly controlled, if not avoided. Decoherence times for candidate systems typically range between nanoseconds and seconds at low temperature. Currently, some quantum computers require their qubits to be cooled to 20 millikelvins in order to prevent significant decoherence.

Quantum Computing Architectures. As a matter of fact, different architectures have been suggested for realistic, programmable, arbitrary-scale quantum computers, with different approaches to the main issue: providing low-overhead fault-tolerant quantum computation for application of large-scale quantum computation to real-world problems. Three main architectures have been shown to be formally equivalent. On the other hand, their underlying elementary concepts and the requirements for their practical realization can differ significantly. In the following we will briefly summarize the three main paradigms of quantum computing architectures also providing some elements to assess their advantages and limitations: circuit-based QC (CBQC – aka gate-model QC or GMQC), measurement-based QC (MBQC), adiabatic QC (AQC).

Notice that in order to assess architectures and implementations one should take into account that practical implementation of quantum algorithms may require a large number of qubits and many operations on them. Besides, since errors arise during communication, qubit operations, data storage, and measurement, any assessment must properly account for these sources of errors.

CBQC has been considered **the** only possible architecture of quantum computing for a long time. Indeed, in order to solve a particular problem, computers, be it classical or quantum, follow a precise set of instructions that can be mechanically (or quantum-mechanically) applied to yield the solution to any given instance of the problem.

CBQC uses qubits, i.e. the intrinsic spin-1/2-like degree of freedom of any bistable quantum system, to encode information and unitary operations to process them. The interest raised by

CBQC and its development had been fostered by the precise identification of criteria that should be fulfilled by any architecture suitable for a scalable quantum computer, i.e.: (i) It should be possible to initialize an arbitrary N-qubit quantum system to a known state; (ii) A universal set of quantum operations must be available to manipulate the initialized system and bring it to a desired correlated state; (iii) The technology must have the ability to reliably measure the quantum system; (iv) It must allow much longer qubit lifetimes than the time of a quantum logic gate. The second requirement encompasses multi-qubit operations; thus, it implies that a quantum architecture must also allow for sufficient and reliable communication between physical qubits.

Ordinarily, in a classical computer, the logic gates other than the NOT gate are not reversible. In CBQC quantum logic gates are reversible. However, classical computing can be performed using only reversible gates. For example, the reversible Toffoli gate can implement all Boolean functions. This gate has a direct quantum equivalent, showing that **quantum circuits can perform all operations performed by classical circuits.**

Available physical platforms for CBQC. In principle, we know how to build a quantum computer; we start with simple quantum logic gates and connect them up into quantum networks. Trapped atomic ions and superconducting circuits have been suggested as promising candidates for implementations. Whereas the superconducting system offers faster gate clock speeds and a solid-state platform, the ion-trap scheme is characterized by better qubits' quality (coherence times) and reconfigurable connections.

Advantages, limitations, perspectives of CBQC. The performances of CBQC implementations actually reflect the topology of connections in the base hardware. As we will see, this is true also for devices implementing adiabatic QC, thus supporting the idea that quantum computer applications and hardware should be co-designed.

Selected references on CBQC. D. P. DiVincenzo, *The Physical Implementation of Quantum Computation*, Fort. Phys. 48, 771 (2000); A. Steane, *Quantum Computing*, Rep. Prog. Phys. 61 117 (1998).

MBQC is a relatively recent proposal compared to gate-based QC. In MBQC the processing of quantum information takes place by rounds of simple measurements on qubits or qudits[1], which are prepared in advance in a highly entangled resource state usually a cluster state or graph state. It is an instance of one-way computation because the resource state is destroyed by the measurements. In MBQC the outcome of each individual measurement is random, but they are related (and correlated) in a way that the computation always succeeds. In general the choices of basis for later measurements need to depend on the results of earlier measurements, and hence the measurements cannot all be performed at the same time.

Any Measurement Based computation can be made equivalent to a quantum circuit by using quantum gates to prepare the entangled resource state. For cluster and graph resource states, this requires only one two-qubit gate for each link, and thus the scheme is efficient. The converse is also true: Any quantum circuit can be made equivalent to a MB scheme.

Available physical platforms for MBQC. MB quantum computation has been demonstrated by running the 2 qubit Grover's algorithm [2] on a 2x2 cluster state of photons. A linear optics quantum computer based on one-way computation has been proposed and several sources of photonic cluster states have been demonstrated as well. Cluster states have also been created in

---

[1] A qudit is a quantum unit of information that may take any of *d* states, where *d* is a variable. Namely, it ist he generalization of the qubit, which corresponds to *d=2*.

optical lattices but were not used for computation, mostly because the atomic qubits were too close together to be addressed and thus measured individually.

Advantages, limitations, perspectives of MBQC. MBQC is in principle scalable since sources of highly entangled states are available and may be used offline. Usually those sources are not deterministic and the rate of production scales badly with the number of qubits, but these is not usually seen as a limitation since the time needed to produce resource states is not part of the computation. A more serious issue concerns the life-time of highly entangled states, which should be compared with the time needed to perform measurements.

Selected reference on MBQC. H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, M. Van den Nest, *Measurement-based quantum computation* Nat. Physics **5**, 19-26 (2009); Raussendorf; D. E. Browne & H. J. Briegel *Measurement based Quantum Computation on Cluster States* Phys. Rev. A **68** 022312 (2003).

The principles of Adiabatic QC are rooted in the so-called quantum annealing protocol, suggested for finding the global minimum of a given objective function over a given set of candidate solutions by exploiting quantum fluctuations. In quantum annealing, the system is initialized in an equal-weight superposition of all possible states and then left free to evolve according to its, usually time-dependent, Hamiltonian. Annealing is obtained introducing a slow transverse-field, slow enough for the system to stay close to the ground state of the instantaneous Hamiltonian, i.e. to evolve adiabatically. If the rate of change of the transverse-field is then accelerated, the system may leave the ground state temporarily but is likely to arrive in the ground state of the final problem Hamiltonian, i.e., adiabatic evolution. The transverse field is finally switched off, and the system is expected to finally lands in the ground state of the classical model corresponding to the solution to the original optimization problem.

In AQC the Hamiltonian of interest is that with a ground state describing the solution to the problem of interest. This Hamiltonian may be involved but another system with a simple Hamiltonian is prepared and initialized to the ground state. Then, the simple Hamiltonian is adiabatically evolved to the desired Hamiltonian. Since the system remains in the ground state (so the term adiabatic), at the end the state of the system describes the solution to the problem.

AQC has been shown to be polynomially equivalent to conventional quantum computing in the circuit model and it is robust against dissipation since the system is always in its ground state.

Available physical platforms for AQC. An experimental demonstration of the success of quantum annealing for random magnets was reported immediately after the initial theoretical proposal. Current implementations of AQC are the only commercial devices available (D-wave). They are based on Josephson junctions qubits and contain CPUs made of approximately 512 qubits in the first generation (now 2000), the number of functional qubits varying significantly from chip to chip, due to flaws in manufacturing.

Advantages, limitations and perspectives of AQC. Whether or not D-Wave computers offer a concrete enhancement of performances is yet to be tested conclusively. Every problem that have tested so far can still be solved faster on classical computers. Currently, AQC computers are not general purpose, but rather are designed for quantum annealing. Specifically, the computers are designed to use quantum annealing to solve the so-called quadratic unconstrained binary optimisation problem.

Selected references on AQC. J. Brooke, D. Bitko, T. F. Rosenbaum and G. Aeppli, *Quantum annealing of a disordered magnet*, Science **284** 779 (1999); P Ray, BK Chakrabarti, A Chakrabarti *Sherrington-Kirkpatrick model in a transverse field: Absence of replica symmetry breaking due to quantum fluctuation*s Phys. Rev. B **39**, 11828 (1989).

2.2 Funding and effort: academy versus industry

Since its inception in the early nineties, research in theoretical and experimental quantum computation was essentially funded by public bodies. In Europe, the European Commission funded several projects and coordination actions in the sector information and communication technologies. In the USA, funding mostly came from the National Science Foundation and the Army Research Office. While these efforts allowed a terrific progress in both quantum information theory and in searching for qubit technologies (see 3.1.1), we are not yet at the stage of having reliable and powerful quantum hardware for testing all the theoretical predictions.

Few years ago, we witnessed a revolution that promises to lead to commercial quantum hardware of interesting size for real-world application: this revolution consists of a steering in funding sources, with industry playing now a major role.

From fintech, to big data, to hardware design, cybersecurity, general analysis services, information and systems modelling, biotechnology, and a host of other sectors, once quantum computing gains sufficient traction, we seem to be looking at, alongside artificial intelligence, the face of the next wave of disruption. Thus, industrial funding and venture capitals play a major role.

In investment terms, there are host of quantum computing companies with significant venture capital support, but few are backed by public companies. Three buck this trend: D-Wave, backed by, amongst others, Harris & Harris, Canada-based Pender Growth Fund, and Goldman Sachs (NYSE:GS), Post-Quantum, supported in part by London's Barclays (NYSE:BCS), and 1QBit, which have received support from CME Group (NASDAQ:CME) and the Royal Bank of Scotland (NYSE:RBS). In terms of already listed IT giants making moves in the quantum world, Intel, IBM (NYSE:IBM), Microsoft, and Google (NASDAQ:GOOG) (NASDAQ:GOOGL), for instance, are also increasing their quantum investments.

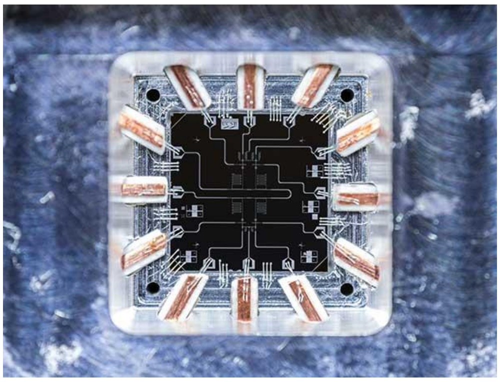

Fig. 2. A 2014 prototype of a Google qubit (0.6 cm by 0.6 cm) known as a transmon[1], based on superconducting circuits. From aps.org. Google's quantum computing test will use 49 updated versions of these qubits.

Google is venturing on the superconducting qubit technology [1] (Fig. 2), pursuing both digital and analog quantum computation, doing both theoretical and experimental research. They are pushing the field to the commercialisation of new devices and services for users. Their “manifesto” for quantum technologies was recently published in Nature [3]. In this comment, Google’s Quantum Artificial Intelligence team [https://research.google.com/pubs/QuantumAI.html], directed by Hartmut Neven, identifies three applicative priorities (quantum simulation, quantum-assisted optimization, quantum sampling, on which we expand in the next subsection), technical hurdles and business opportunities. The main hurdle is scaling the number of qubits while maintaining coherence; error correction is a big issue, because it needs many additional qubits (one possible way out is the design of error-prone algorithms, an example of which is the variational quantum eigensolver – VQE – approach, see subsection 2.3). Concerning business opportunities, they start from the principle that in the digital era even a modest increase in product power/quality can boost user numbers and revenue; they note that the markets that are most open to such disruptions are information-rich and digital, and involve business challenges that rely on many variables: for instance, financial services, health care, logistics and data analytics. Specifically, quantum simulation applications can aim at stronger polymers for aircrafts, more effective catalytic converters for cars, more efficient materials for solar cells, better pharmaceuticals, and more breathable fabrics. Related business models could supply quantum simulators for users: laboratories might pay a subscription for access, computing companies and HPC centers could act as consultants. In 2014, Google hired John Martinis (University of California Santa Barbara) and his entire group, to intensify efforts towards fabricating quantum hardware of useful size in less than a decade. In recent years, Google has also hired a number of scientists coming out of academic institutions highly involved in quantum information research, such as University of Southern California Information Sciences Institute, Massachusetts Institute of Technology and Harvard University.

Microsoft's Quantum Architectures and Computation (QuArC) Group was founded at the end of 2011 [https://www.microsoft.com/en-us/research/group/quantum-architectures-and-computation-group-quarc/] and is managed by Krysta Svore. The group is a team of leading quantum computer scientists and engineers dedicated to developing real-world quantum algorithms, understanding their implications, and designing a comprehensive software architecture for programming such algorithms on a scalable, fault-tolerant, quantum computer. Microsoft, however, is hoping to encode its qubits in a kind of quasiparticle: a particle-like object that emerges from the interactions inside matter. Such quasiparticles are called non-abelian anions and the firm aims to exploit their topological properties to build topological quantum computers [4]. Detailed plans were revealed by Alex Bocharov in an interview published in Nature News&Comments on 21 October 2016 [http://www.nature.com/news/inside-microsoft-s-quest-for-a-topological-quantum-computer-1.20774]. Bocharov identifies topological quantum computation as a strategy to bypass extensive and expensive error correction. On the Microsoft Blog on 20 November 2016, there was the announcement that Microsoft is doubling down on its commitment to the tantalizing field of quantum computing, making a strong bet that it is possible to create a scalable quantum computer using what is called a topological qubit [https://blogs.microsoft.com/next/2016/11/20/microsoft-doubles-quantum-computing-bet/]. To boost the production of topological qubits of usable size, also Microsoft hired faculty members from top academic institutions, already long active in quantum technologies and fundamental properties: Leo Kouwenhoven from Delft University, Charles Marcus from the Niels Bohr Institute at the University of Copenhagen, and Matthias Troyer from ETH Zürich.

IBM Q [https://www.research.ibm.com/ibm-q/] is an industry-first initiative to build commercially available universal quantum computers for business and science. Similarly to Google, IBM is also targeting a superconducting qubit technology. As of 17 May 2017, IBM Q has successfully built and tested two of the most powerful universal quantum computing processors. A 16-qubit device is for public use by developers, researchers and programmers via the IBM Cloud at no cost (more than 300,000 quantum experiments have been run by users on the IBM Cloud). The first prototype commercial processor, with 17 qubits and architecture innovations, is the most powerful built by IBM to date, roughly twice as powerful as the free version in the cloud. IBM was one of the first private companies that invested in the development of quantum hardware.

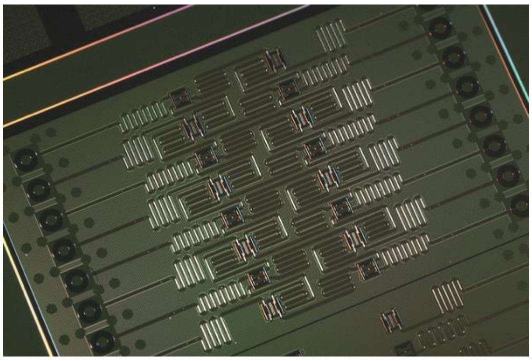

Fig. 3. The IBM 16-qubit processor that is available to users via the IBM Cloud. Credit: IBM Research.

Intel is the other big computer company that is not lagging behind the aforementioned giants. While IBM, Microsoft and Google are all trying to develop quantum components that are totally different from those present in present computers, Intel is trying to adapt the workhorse of the existing computer technology, namely the silicon transistor, for the task. Intel has a team of quantum hardware engineers in Portland, Oregon, who collaborate with researchers in the Netherlands, at TU Delft's QuTech quantum research institute, under a $50 million grant established in 2015. In December 2016, Intel's quantum group reported that they can layer the ultra-pure silicon needed for a quantum computer onto the standard wafers used in chip factories. This strategy makes Intel an outlier among industry and academic groups working on qubits. Jim Clarke, who leads Intel's project as director of quantum hardware, argues that silicon qubits are more likely to get to a satisfactory scaling trend. The expertise and equipment used to make conventional chips with billions of identical transistors should allow work on perfecting and scaling up silicon qubits to progress quickly. Intel's silicon qubits represent data in a quantum property called the "spin" of a single electron trapped inside a modified version of the transistors in its existing commercial chips. Another reason to work on silicon qubits is that they should be more reliable than the superconducting equivalents, which are error-prone.

The race between Google and IBM to sell true quantum computers before they really exist was covered by the journal Wired on 6 March 2017 [https://www.wired.com/2017/03/race-sell-true-quantum-computers-begins-really-exist/]. Industrial investments in quantum technologies have large coverage in public media; we point out in particular The New York Times, with articles entitled "IBM Wants Everyone To Try A Quantum Computer" (4 May 2016), "Microsoft Spends Big To Build A Computer Out Of Science Fiction" (20 November 2016), "Researchers Report Milestone In Developing Quantum Computer (4 May 2015 – on Google+UCSB), "Microsoft Makes Bet Quantum Computing Is Next Breakthrough" (23 June 23 2014), a video featuring Justin Trudeau speaking on Quantum Computing (25 April 2016).

On the other hand, the following is a non-exhaustive list of government and non-profit agencies that perform research or provide funding for Quantum Computing:

- The ARC Centre of Excellence for Engineered Quantum Systems (EQuS) started in 2010 under the Australian Research Council (ARC) of the Australian Government;
- The Centre for Quantum Computation & Communication Technology has 17 coordinated programs under the ARC, involving facilities at six different universities in Australia; the team is mostly active in Silicon-based qubits;
- The DNA-SEQ Alliance is an inter-disciplinary, cross-organizational effort that is designed to revolutionise cancer treatment and drug discovery, focusing on protein kinase inhibition; they have a partnership with D-Wave and 1QBit to leverage quantum computation to perform rapid and inexpensive genomic sequencing, creating bioinformatics profiles that enable detailed crystallographic analysis;
- The U.S. government agency IARPA (Intelligence Advanced Research Projects Activity) has a program called Quantum Enhanced Optimization (QEO) that funds research into enhanced methods for quantum annealing to solve hard combinatorial optimization problems;
- Los Alamos National Laboratory (LANL) formally organized a Quantum Institute in 2002 to perform research into Quantum Computing and Quantum Cryptography;
- NASA Ames Quantum Artificial Intelligence Laboratory (QuAIL) explores applications of quantum information processing to NASA missions, particularly quantum algorithms for hard computational problems that arise in aeronautics, Earth and space sciences, and space operations and explorations; in collaboration with Google and USRA (University Space Research Association) they operate the D-Wave 2X computer installed at Nasa Ames Research Center in Mountain View, California;
- NSF (National Science Foundation) has a Quantum Information Science (QIS) group that supports theoretical and experimental proposals that explore quantum applications to new computing paradigms or that foster interactions between physicists,

mathematicians, and computer scientists that push the frontiers of quantum-based information, transmission and manipulation;

- NSA (National Security Agency) has an obvious interest in utilizing quantum computation to break codes; most of the NSA-funded work on quantum computer development takes place at University of Maryland's College Park Campus;
- Networked Quantum Information Technologies (NQIT) is the largest of four hubs in the UK National Technology Program;
- The Perimeter Institute for Theoretical Physics is a non-profit center for scientific research, training and educational outreach in foundational theoretical physics based in Waterloo, Ontario, Canada;
- QuSoft, located in Amsterdam, the Netherlands, is a new research center formed in December 2015 dedicated to quantum software; it is a join initiative between CWI, UvA, and VU and will be focusing its research in the areas of few-qubit applications, quantum testing and debugging, quantum architectures, and quantum cryptography;
- QuTech, located in Delft, the Netherlands, is a joint public-private partnership organization with partners TU Delft and TNO and others that develops quantum technologies based on superposition and entanglement aimed at scalable quantum networks and quantum computers; QuTech's research is guided by three roadmap paths including Fault Tolerant Quantum Computing, Quantum Internet, and Topological Quantum Computing;
- USRA (Universities Space Research Association) is an independent, nonprofit research corporation where the combined efforts of in-house talent and university-based expertise merge to advance space science and technology;
- The European Commission (EC) has funded several collaborative projects during the past two decades, mostly under the future and emerging technologies (FET) initiative of the Information and Communication Technology (ICT) sector, with a cumulative investment of around €550M; the EC is now running coordination actions to shape the €1 billion "Quantum Technologies" Flagship [5];
- European countries are also investing at the national level and high-quality research is being pursued in Austria, Germany (QUTEGA, http://www.laserfocusworld.com/articles/2017/06/german-quantum-initiative-qutega-starts-with-optical-single-ion-clock.html), Italy, Spain, France (https://scholar.google.fr/scholar?q=esteve+quantum+computing+programming&hl=fr&as_sdt=0&as_vis=1&oi=scholart&sa=X&ved=0ahUKEwiz68DRvJXVAhWKL8AKHQw9AhwQgQMIKDAA) and practically all European countries.

2.3 Applications

To compile a list of possible applications of quantum computation, we refer to the distinction between the two main paradigms, gate-model quantum computation (GMQC) and adiabatic quantum computation (AQC) – see section 3.1.1. We remark that AQC is the only paradigm for which exists a commercial exploitable device; therefore AQC applications are more mature than GMQC applications.

We point out two relevant documents that were compiled as the outcome of two workshops aimed at exploring funding opportunities and real-world applications of quantum computation:

- ASCR Report on Quantum Computing for Science, workshop sponsored by the USA DoE, held in Bethesda (MD), February 15-16, 2015 [https://science.energy.gov/~/media/ascr/pdf/programdocuments/docs/ASCRQuantumReport-final.pdf];

- NSF workshop report on quantum information and computation for chemistry, workshop sponsored by the USA NSF, held in Arlington (VA), November 14-15, 2016 [http://aspuru.chem.harvard.edu/blog/wp-content/uploads/2017/06/NSF-report.pdf]

Another event closely related to this deliverable is the CECAM workshop “Synergy between quantum computing and high-performance computing”, to be held in Zürich on August 22-24, 2017 [Link to the workshop https://www.cecam.org/workshop-0-1464.html].

GMQC Applications. GMQC should be suitable to any algorithms and should manifest its advantages over classical computation especially on problems that have exponential classical complexity. Among the most famous problems that should scale polynomially on a quantum computer are Grover’s [2] and Shor’s [6] algorithms: the former searches for a specified entry in a non-ordered database, the latter should perform prime factorization of integers in essentially polynomial time. Note that a number of algorithms for public-key cryptography depend on the fact that there is no known efficient classical algorithm to factor integers into prime numbers; thus, there are implications of quantum computing power for the field of cybersecurity.

The idea of a quantum computer was first proposed in 1981 by Nobel laureate Richard Feynman, who pointed out that accurately and efficiently simulating quantum mechanical systems would be impossible on a classical computer, but that a new kind of machine, a computer itself “built of quantum mechanical elements which obey quantum mechanical laws” [7], might one day perform efficient simulations of quantum systems. Classical computers are inherently unable to simulate such a system using sub-exponential time and space complexity due to the exponential growth of the amount of data required to completely represent a quantum system. Quantum computers, on the other hand, exploit the unique, non-classical properties of the quantum systems from which they are built, allowing them to process exponentially large quantities of information in only polynomial time. Thus, a natural field of application of quantum computation is quantum chemistry and many-body systems. Alan Aspuru-Guzik (Harvard University) is the most active group leader in this respect, though Seth Lloyd (MIT) opened this avenue in 1996 [8], demonstrating the feasibility of Feynman’s proposal. The scientific community is interested in both what quantum computing can do for chemistry and what chemistry can do for quantum computing[9].

Building from tools developed by Seth Lloyd and co-workers, Aspuru-Guzik *et al* developed a quantum algorithm for computing molecular electronic ground state energies in polynomial time, paving the way for a variety of subsequent developments in quantum computing for chemistry [10]; relevant issues to be solved are quantum state preparation and how to handle multi-reference states. A crucial aspect of developing quantum algorithms for quantum chemistry is the estimate of the resources required to perform relevant simulation tasks, namely understanding the scaling in terms of number of electrons and spin-orbitals, as a function of the number of qubits. Aspuru-Guzik’s group members developed algorithms for the simulation of sparse Hamiltonians, which have a favorable scaling as a function of molecular size; they were used for quantum chemistry formulation in first [11] and second quantization[12]. For testing, quantum algorithms for molecular energies can be executed on classical computers [13].

Besides intense efforts at Harvard University in collaboration with Google, other groups are looking at the molecular electronic structure with quantum algorithms. The quantum implementation of the unitary coupled cluster for simulating molecular electronic structure was published in February 2017 by a Chinese partnership [14]. In Europe, Matthias Troyer (ETH Zürich, recently moved to Microsoft Research USA) has also been a major player to devise quantum algorithms for quantum chemistry and many-body systems and investigate the scaling [15,16], encompassing GMQC and AQC and bridging the HPC community.

We remark that a related field of investigation is the simulation of quantum circuits on existing classical supercomputers, which is an important communication bridge with the HPC community. The record in this respect is the "0.5 petabyte simulation of a 45-qubit quantum circuit" by Damian Steiger and Thomas Häner (ETH Zurich), announced in April 2017 and realized on the NERSC (Berkeley) supercomputer Cori, utilizing 8192 nodes and 0.5 petabytes of memory. The impressive achievement was allowed by highly-tuned kernels combined with reduced communication requirements [17].

Within the context of quantum chemistry, but possibly extendable beyond, hybrid quantum-classical algorithms are attracting attention. Hybrid strategies could use classical CPUs to select classically hard part of an algorithm and feed these instances into quantum processing units (QPUs). This field is particularly suitable to match existing HPC platforms and competence. Hybrid algorithms have already been proposed. VQE [13,18-20], for instance, parametrises a quantum state through some polynomial number of classical parameters, uses a quantum device to evaluate the expectation value of an objective function which depends on them and updates the parametrisation through a classical non-linear minimization procedure. VQE has been applied to the simulation of molecules using quantum optics and superconductors. Another representative hybrid algorithm is the quantum approximate optimisation algorithm (QAOA) [21], developed by Edward Farhi and collaborators, which is promising to approximately solve combinatorial optimization problems.

The other foremost field of application for GMQC is supervised and unsupervised machine learning. Seth Lloyd proposed quantum versions of machine learning approaches [22], including support vector machine [23,24]. Traditional machine learning and deep learning (neural networks) are a promising channel towards artificial intelligence as it is nowadays conceived: note that the Google group that is devoted to quantum computation, directed by Hartmut Neven, is labelled "Quantum Artificial Intelligence" [https://plus.google.com/+QuantumAILab]: in fact, many learning tasks rely on ASCR solving hard optimisation problems or performing efficient sampling.

Most proposals for the quantum version of machine learning utilise the finite-dimensional substrate of discrete variables. Within a Europe-USA-Canada collaboration, the framework was lately generalized to the more complex, but still remarkably practical, infinite-dimensional systems [25].

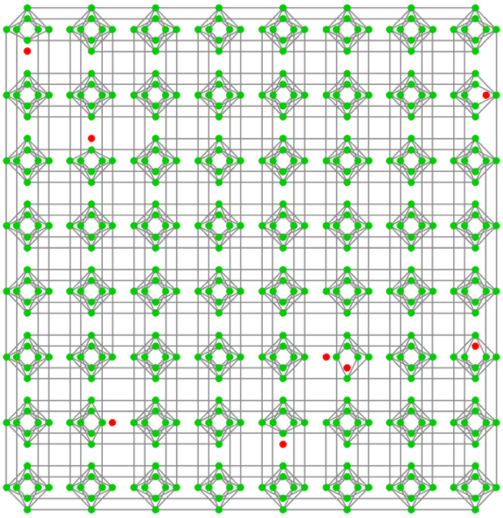

Fig. 4. Schematic representation of the "Chimera" hardware graph of the D-Wave Two X (DW2X) housed at the Information Sciences Institute (ISI) at the University of Southern California (USC). Green circles represent active qubits, red circles represent inactive qubits and lines represent couplings between qubits. Each qubit can be coupled to a maximum of six other qubits.

Topological and geometric analysis of data, to identify patterns in big data, is another important topic for applications of quantum computing [26].

AQC Applications. First we focus on ongoing applications of the commercial quantum annealer D-Wave (DW) and then we generalize. Scientific work on DW devices is currently carried out on two different processors, DW2X (1098 qubits) and DW200Q (2000 qubits); early attempts conducted on DW1 (128 qubits) and DW2 (512 qubits) have been published. The Chimera connectivity graph of a functional DW2X device is shown in Fig. 4.

The most promising applications of DW devices are as optimisers and simulators. As simulators, DW computers can tackle materials science and statistical mechanics. Optimisation problems based on machine learning can efficiently be mapped on DW devices, but also other optimisation schemes are feasible. The conference AQC 2017, held in Tokyo on 26-29 June, gave an exhaustive overview of the state-of-the-art.

An impressive achievement was presented by Richard Harris (D-Wave Systems Inc.), who reported on the simulation of three-dimensional spin glasses on DW2000Q. It has been demonstrated that the quantum simulator can efficiently describe the phase transitions, e.g., from antiferromagnetic to spin glass. This result was achieved by investigating the susceptibility and the order parameter as a function of doping. No optimisation of the annealing time was attempted. This is a milestone in the demonstration that DW can be exploited to do computations in materials science and statistical mechanics.

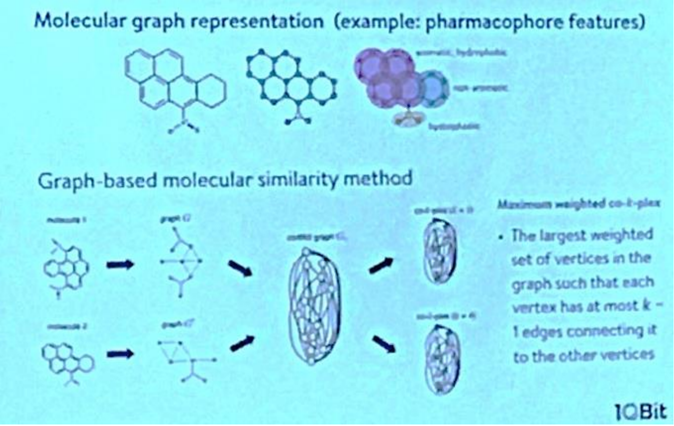

Molecular graph representation (example: pharmacophore features)

Graph-based molecular similarity method

Maximum weighted co-k-plex

- The largest weighted set of vertices in the graph such that each vertex has at most k − 1 edges connecting it to the other vertices

1QBit

Fig. 5. Graph-based molecular similarity: a molecule can be mapped to a graph with different levels of structural complexity; the graph can be ported onto DW. Photo from presentation at AQC 2017, taken by Rosa Di Felice.

Companies that do services for customers are starting to use DW. The company 1QBit, based in Vancouver (Canada), solves industry most demanding computational challenges by recasting problems to harness the power of quantum computing. It was heavily represented at AQC 2017 with oral presentations and posters. Maritza Hernandez-Gaete reported on the use of DW2X for the molecular similarity problem, which has huge relevance in drug design [27]. The strategy is to map the molecular structure onto a graph and then porting the graph onto the Chimera graph that is typical of DW processors. In the graph-based molecular similarity method, different levels of complexity of the molecular structure can be retained: e.g., atoms and bonds can be mapped to vertices and connections, respectively; or cycles and their interactions can be mapped to vertices and bonds (Fig. 5). In porting the graph to the Chimera graph, minor embedding [28,29] is necessary, because of the limited DW connectivity between superconducting devices [2]. Actually, the limited connectivity is one of the major bottlenecks for real-world applications. Different research groups world-wide are working to design quantum devices with more complete connectivity. For example, Andrea Rocchetto presented efforts ongoing at the University of Oxford to use stabilisers as design tools for new forms of the Lechner-Hanke-Zoller annealer; this strategy, though, is criticised for not being quantum. Sruti Puri (University of Sherbrooke) presented a quantum Ising machine based on oscillators, also putatively not quantum.

---

[2] A minor embedding maps a logical problem qubit to a set of physical qubits such that for every coupling between pairs of logical qubits in the logical problem there exists at least one physical coupling between the corresponding sets of physical qubits. A minor embedding is found by performing a series of edge contractions, which effectively joins vertices together, thereby allowing for a graph with fewer vertices but a higher degree of connectivity to be obtained.

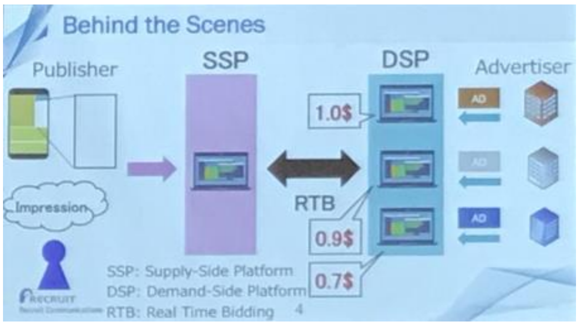

Fig. 6. Display advertising at Recruit (Japan). Photo from presentation at AQC 2017, taken by Rosa Di Felice.

In addition to the exploitation of DW as an optimizer, 1QBit is also exploring the use of DW to study phase transitions, e.g., Shunji Matsuura. Another company that is making progress in using DW for real-world applications is Recruit, a communication company based in Japan. Shinichi Takayanagi (Recruit) showed that DW can treat the problem of optimising display advertising, namely the advertising messages that appear on websites (the advertiser pay a publisher when the ad is clicked – Fig. 6). The problem can be cast into a QUBO formulation, which is suitable for DW chips; an example on a campaign with 14 advertisements and 24 users was shown, but the real-world data had to be pruned to fit the size and connectivity constraints of present quantum annealers. Kotaro Tanahashi (Recruit) presented feature selection by quantum annealing, important in problems tackled by machine learning.

Academic examples of real-world problems span the fields of physics and biology and rely on porting a machine learning formulation onto a QUBO problem. Such efforts are carried out by Daniel Lidar at the University of Southern California (USC). In collaboration with Maria Spiropulu (Caltech), he uses DW2 to process data from the Large Hadron Collider (LHC) and to predict plasma instabilities in future fusion experiments; a specific application concerns the classification of data related to the Higgs boson discovery, which is affected by the large Standard Model background and small Higgs signal (this work is currently under review for publication). In collaboration with Remo Rohs at USC, Lidar proves that DW2X can be exploited to tackle the problem of transcription factor-DNA binding specificity, both to classify and rank genetic sequences to match experimental data; furthermore, advantages to use a quantum annealer rather than classical approaches, emerge when the experimental data set for training is fairly small (this work is under review for publication). In these applications, as highlighted above for the molecular similarity problem, minor embedding[28,29] is necessary and the embedding strategy should be chosen properly.

Early work on DW included a crude model of protein folding [30]. Although most approaches to the quantum simulation of quantum chemistry were described above in the subsection on GMQC, adiabatic formulations exist [31].

Work that is still in the planning phase addresses the use of DW in materials science, to predict an optimized material for a chosen performance property or to characterize catalysis processes. For instance, one may want to select a binary compound with high thermoconductivity, or with high resistance to mechanical stress. Based on exhaustive materials genome databases constructed with results from density functional theory calculations[32], in principle it is possible to write the search for the optimal material as a QUBO problem portable on DW chips. The materials genome initiative has been hugely funded by the DoE in the USA (http://materials.duke.edu; http://ceder.berkeley.edu) and by the EC in Europe (http://aiida.net, http://www.max-centre.eu, https://www.nomad-coe.eu) and it would be very valuable to have tools to interrogate the produced databases. In addition to materials optimization, quantum annealing could also be exploited to model potential energy surfaces of reactants on surfaces [33,34].

AQC is also explored for software validation and verification, a hot topic for aerospace corporation, such as Lockheed Martin.

Andrew Lucas published the Ising formulation of NP-hard problems[35], opening the way to AQC application to many general situations. NP is the set of decision problems *solvable* in polynomial time by a theoretical non-deterministic Turing machine; NP-hard means "at least as hard as any NP-problem," although it might, in fact, be harder. Classes of such Ising-portable instances are: (i) partitioning problems; (ii) binary integer linear programming; (iii) covering and packing problems; (iv) problems with inequalities; (v) coloring problems; (vi) Hamiltonian cycles; (vii) tree problems; (viii) graph isomorphisms. The applications mentioned above represent a subset of these classes.

Without proposing an application on real data, Garnerone, Zanardi and Lidar presented in 2012 an adiabatic quantum algorithm which prepares a state containing the same ranking information as the PageRank vector [36]. The latter is a central tool in data mining and information retrieval, at the heart of the success of the Google search engine.

In summary, adiabatic quantum annealers implement a physical Hamiltonian and are therefore suitable to simulate physical systems represented by that Hamiltonian (e.g., spin glasses for DW); optimisation problems are the other big target of AQC; other possible, yet unexplored applications to NP problems are suggested in a published article [35].

3 Conclusion

We conclude this document by proposing a set of actions for the HPC community:

- Participate in conferences/schools/workshops to get in touch and create a common scientific language with the QC community;
- Promote the use of high-end HPC systems of the PRACE infrastructure to emulate QC and provide tool and environment to allow developer to run and develop applications on the QC emulator;
- Hire expert people on QC from top institutions and train HPC people by visiting exchanges and other instruments;
- Propose coordination actions at the European level;
- Run experiments on existing D-Wave devices;
- Run experiments on the IBM cloud;
- Experiment with European solutions such as AQLM (see https://atos.net/en/2017/press-release/general-press-releases_2017_07_04/atos-launches-highest-performing-quantum-simulator-world) to prepare next generation of software using quantum computers.