

Impact case study (REF3b)

Institution: PHYESTA

Unit of Assessment: UoA 9 – Physics

Title of case study: A New Generation of Supercomputers results from the Co-Design of a Computer Chip for Lattice QCD Calculations

1. Summary of the impact

Impact: Economic gains

PHYESTA designed 8% of the area of the computer chip for IBM's most recent BlueGene/Q supercomputer product. Global install base of design exceeds \$500M.

Significance:

Unique experiment in co-design at the cutting edge of technology. Adopted by both IBM and Fujitsu, who have led in Green500 energy efficiency and top500 supercomputer rankings.

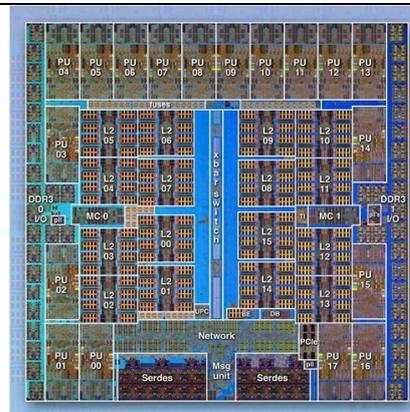
Reach:

This supercomputer architecture has been installed in labs in the UK, the US, the EU, and Japan and is accelerating computational science and advanced manufacturing around the globe. In the UK the BlueJoule system installed in the Hartree center at Daresbury is driving HPC uptake in the advanced manufacturing sector.

Beneficiaries:

IBM, Fujitsu, computational science and the HPC community worldwide.

Attribution: This work was led by Dr Peter Boyle (School of Physics & Astronomy, University of Edinburgh) in collaboration with Columbia University and IBM.



2. Underpinning research

In 2000, the Lattice QCD research group in the Institute for Particle and Nuclear Physics entered into a collaboration with Columbia University and the IBM T J Watson Research Center to jointly develop QCDOC, a supercomputing architecture customised for Quantum Chromodynamics simulations [R1, R2]. The QCDOC project combined the use of then nascent system-on-a-chip technology with a relatively slow but very low power processor, and a large amount (4MB) of on-chip embedded DRAM. This enabled us to integrate a six dimensional torus communications network, to accelerate application performance with very high bandwidth memory, and to optimise the overall system price, performance and power efficiency. We therefore developed a machine with very good computational efficiency in an early application of hardware-software co-design [R3, R4].

In December 2007 Boyle was invited to lead an international team designing the memory prefetch engine for IBM's next generation BlueGene/Q architecture. This was a unique academic-industrial collaboration on core IBM technology. The component designed by Boyle comprises 8% of the die area and the project was the subject of a legal Collaboration Agreement between University of Edinburgh, Columbia University and IBM, whereby Boyle worked during the research and development phase as an external contractor to IBM.

This has been a unique experiment in co-design at the cutting edge of technology, using advanced QCD software and silicon design skills to feed back and ensure best performance. Many architectural decisions have been influenced by our QCD codes and by the academic design team members [R5, R6]. The prefetch engine is a key performance differentiator, which was under our design control, particularly during the bringing-up of the prototype chip and the debugging of the system. The chip design is included in BlueGene/Q, and has led to four USPTO patents. It has

Impact case study (REF3b)

also been included in the Fujitsu K supercomputers. Boyle received the Ken Wilson award (Lattice 2012), and has also received a Gauss award for contributions to supercomputing.

Personnel:

The key PHYESTA researcher involved was Dr Peter Boyle (Academic staff, 2000-present)

3. References to the research

| | |
|------|--|
| [R1] | C. Allton et al., “2+1 flavor domain wall QCD on a $(2\text{ fm})^3$ lattice: light meson spectroscopy with $L_s = 16$ ”, Phys.Rev.D 76, 014504 (2007), DOI: 10.1103/PhysRevD.76.014504, [58] |
| [R2] | D. Antonio et al., “Neutral Kaon Mixing from (2+1)-Flavor Domain-Wall QCD”, Phys.Rev.Lett. 100 , 032001 (2008), DOI: 10.1103/PhysRevLett.100.032001, [45] |
| [R3] | P. Boyle et al., ‘The QCDOC project’, Nuclear Physics B - Proceedings Supplements, 140 , p. 169, (2005), DOI: 10.1016/j.nuclphysbps.2004.11.179, URL: tinyurl.com/lmgq34s , [19] |
| [R4] | P.A. Boyle et al., ‘Overview of the QC DSP and QCDOC Computers’, IBM Journal of Research and Development, 49 , p. 351, (2005), DOI: 10.1147/rd.492.0351, URL: tinyurl.com/le6kz9x , [35] |
| [R5] | P.A. Boyle, ‘The BAGEL assembler generation library’, Computer Physics Communications, 180 , p. 2739, (2009), DOI: 10.1016/j.cpc.2009.08.010, URL: tinyurl.com/om8y2vw , [16] |
| [R6] | R.A. Haring et al., ‘The IBM Blue Gene/Q Compute Chip’, IEEE Micro, 32 , p. 48, (2012), DOI: 10.1109/MM.2011.108, URL: tinyurl.com/nypgv3l , [32] |

References R1, R3 and R6 best illustrate the underpinning research quality. [Number of citations]

4. Details of the impact

One of our collaborators on QCDOC moved from Columbia to become the Chief Architect of the BlueGene/L project, which IBM developed in collaboration with Lawrence Livermore National Laboratory. This built on many of the design concepts initiated in QCDOC, and the intellectual connection was recognised by IBM when they included a QCDOC paper in the BlueGene/L edition of their journal for research and development. The BlueGene computer designs remained the fastest in the world for a record duration, from September 2004 until June 2008, and have been a workhorse for accelerating computational science around the world. Our role in developing these designs is acknowledged by IBM: *“My interaction and familiarity with the efforts of the department of physics at the University of Edinburgh cover many years. Physicists at the University of Edinburgh, in their pursuit of fundamental physics through QCD have developed hardware and software resulting in pioneering architectures and methods in supercomputing. These were so successful that they are now part of the mainstream supercomputer offerings of leading vendors such as IBM. A perfect example of this is the BlueGene generation of machines that are now sold worldwide. My current position is the chief architect of all generations BlueGene. In addition I have participated in multiple generations of QCD computers in collaboration with physicists at the University of Edinburgh.”* [F1]

In 2009 Fujitsu adopted the partitioning approach developed by us for their K-series supercomputers, with citation [S1]. Their Tofu network architecture for the 1 billion dollar K-computer design forms the basis of Japan’s national computational science strategy. The continuation of our collaboration with IBM included direct design contributions to the latest BlueGene/Q design. This intellectual contribution is supported by four joint patent applications for novel jointly created intellectual property [S2], and by several press releases [S3, S4]. The BlueGene/Q design is internationally recognised as having taken top place in the Green500 supercomputing list since November 2010, demonstrating a step change in energy efficiency. The design led the top500 in June 2012, becoming the fastest computer in the world, and overtaking the K-computer which also contained some of our original ideas in its network design.

Impact case study (REF3b)

BlueGene/Q is now IBM's premier HPC product, with large multi-Petaflop/s installations in LLNL (USA), Argonne (USA), Cineca (Italy), Edinburgh (UK), Daresbury (UK) and Juelich (Germany). A 1.2 Pflop/s BlueGene/Q prototype has been installed in Edinburgh in November 2011 as part of STFC's DiRAC facility (presently the 23rd fastest computer in the world and the joint fastest single science domain system). A 35M, 1.4 Pflop/s system is installed in Daresbury as part of the Hartree centre which aims to enhance the use of HPC in advanced manufacturing and thus directly impact UK industry. An aggregate of 40Pflop/s was installed in 2012 across the world and is now contributing to the advance of all computational science disciplines. The total worldwide sales of the new generation of supercomputers exceed \$500M. *[text removed for publication]*

5. Sources to corroborate the impact

| | |
|------|---|
| [F1] | Factual statement by the IBM BlueGene Chief Architect |
| [S1] | Web description of Fujitsu K-series architecture I.Ajima, S.Sumimoto and T.Shimuzu http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5331902 <i>Corroborates Edinburgh role in original design and cites [R3] as first reference</i> |
| [S2] | Joint US Patent Applications: US patent application: 20110219208, "MULTI-PETASCALE HIGHLY EFFICIENT PARALLEL SUPERCOMPUTER", URL: tinyurl.com/o7zjt5 US patent award 8,327,077 "Method and apparatus of parallel computing with simultaneously operating stream prefetching and list prefetching engines", URL: tinyurl.com/q7hu569 US patent award 8,347,039 "Programmable stream prefetch with resource optimization", URL: tinyurl.com/ot6qzwt US patent award 8,255,633 "List based prefetch", URL: tinyurl.com/p6bfgvs <i>Corroborate the IP of the prefetch engine design</i> |
| [S3] | STFC Press Release, 26 November 2010 "STFC supported work leads to IBM's 'world's greenest supercomputer'" www.stfc.ac.uk/2107.aspx <i>Corroborates Edinburgh role in supercomputer design</i> |
| [S4] | IBM Press Release 19 November 2010 "Report: IBM Supercomputers Are Most Energy Efficient in the World" www-03.ibm.com/press/us/en/pressrelease/33046.wss <i>Corroborates Edinburgh contribution to BlueGene chip design</i> |
| [S5] | <i>[text removed for publication]</i> |