Institution: University of Bristol		
Unit of Assessment: 11) Computer Science and Informatics		
Title of case study: Isambard – revolutionising supercomputers through mobile technology		
Period when the underpinning research was undertaken: 2013-present		
Details of staff conducting the underpinning research from the submitting unit:		
Name(s): Simon McIntosh-Smith	**Role(s) (e.g. job title):** Professor of High-Performance Computing	**Period(s) employed by submitting HEI:** 04/2009-present
Period when the claimed impact occurred: 2018-July 2020		
Is this case study continued from a case study submitted in 2014? N		

1. Summary of the impact

Bristol's Isambard project has revolutionised the high-performance computing space, taking technology originally designed for mobile devices and using this to build faster, cheaper supercomputers. Isambard's impact has led to a new breed of Arm-based supercomputers, creating a new market worth hundreds of millions of dollars today. Isambard is being used by the Met Office for weather and climate research and has given birth to a whole generation of Arm-based supercomputers, including the fastest supercomputer in the world, Fugaku, at RIKEN in Japan.

2. Underpinning research

The High Performance Computing (HPC) processor market has been dominated by one main CPU vendor, Intel, for the last twenty years. During this period of domination, as often happens during a monopoly, the rates of improvement in processor performance and in performance per dollar have been slowing precipitously. This has resulted in recent supercomputers struggling to deliver the steps forward in performance that were required to tackle important scientific challenges, such as modelling climate change, or simulating how viruses work at the molecular level. To tackle this problem, researchers in the HPC space were considering a shift to Arm-based supercomputers, an architecture previously designed primarily for mobile devices, but progress was hampered by three major roadblocks: 1) a lack of real, production Arm-based machines; 2) scientifically rigorous performance comparisons providing evidence that Arm-based systems would be competitive with the current state of the art; and 3) enough of the software ecosystem (libraries, applications etc) ported to Arm to make the shift possible. In 2016 Professor McIntosh-Smith formed the Isambard research programme to solve these three problems by developing the world's first production supercomputer based on the Arm architecture. A key insight behind adopting a licensable architecture such as Arm is that it enables multiple chip companies to license the core technology before customising it with their own 'secret sauce', resulting in a vibrant market where multiple processor vendors could compete over performance, price and other features, while all providing solutions that are compatible with one another, crucial when targeting a software-driven market such as HPC. Multiple licensable architectures were evaluated, and the Arm architecture was chosen because of the mature software stacks which were available, and its existing broad adoption among silicon vendors, ranging from Apple and Google to Fujitsu, Marvell and NVIDIA. The Bristol team's background in high performance embedded computing, along with its deep experience in porting and optimising complex scientific codes to new architectures, meant it was ideally placed to recognise and exploit this opportunity to leverage Arm's architecture into production HPC environments for the first time.

Prototype development

During 2013-2016, McIntosh-Smith collaborated with 11 international project partners in the European Commission-funded Mont-Blanc 2 project **[i]** to develop prototype Arm-based

supercomputers. These early prototypes were the first to show that Arm-based computers were able to run scientific applications, but they were still much slower than the Intel systems available at the time. During the same period, McIntosh-Smith explored related techniques within two EPSRC projects: the ASEarch Collaborative Computational Project (2013-2014), and the responsive mode project, ENPOWER (2013-2016).

Isambard: breaking through the roadblocks to adopting Arm in HPC

Once the prototypes and basic principles had been demonstrated, McIntosh-Smith formed the Isambard research programme. Backed initially by GBP3 million in EPSRC funding **[ii]**, before a GBP4.1M upgrade **[iii]**, he designed Isambard to enable scientifically rigorous and reproducible results by integrating new, state-of-the-art Arm-based processors from Marvell (a US-based semiconductor company) and Fujitsu, uniquely alongside a heterogeneous collection of best-in-class processors from Intel, AMD, IBM and NVIDIA. Isambard also included the best software suites and established code optimisation best practice, enabling the first ever "best to best" comparisons to be made between all the main competing computer architectures, including Arm. Isambard's results, published between Nov 2017 and May 2019, were the first ever to prove that Arm-based supercomputers could be performance competitive with the state of the art (Intel-based supercomputers), both at the single server level **[3.1, 3.2]** and at the system level **[3.3, 3.4]**, while potentially delivering up to three times the performance per dollar at the processor level (Isambard's ThunderX2 CPUs cost USD1,000 each, but delivered the same performance as Intel CPUs which cost USD3,000 each at that time). The team's system-level papers **[3.3, 3.4]** were also the first to show scaling of scientific software across multiple different Arm-based supercomputers – the Isambard Cray system and the HPE Catalyst system, the latter of which was a GBP1M donation to the University of Bristol on the strength of the Isambard work. The research undertook rigorous benchmarking, comparing a range of scientific software applications across computational chemistry, biochemistry, materials modelling, engineering, and climate modelling.

Figure 1 – Isambard, the world's first production Arm-based supercomputer

The success of the Isambard project led to McIntosh-Smith being awarded the prestigious HPCwire Readers' Choice Award for "Outstanding Leadership in HPC" in November 2018 and a best paper award at CUG 2019 **[3.3]**.

3. References to the research

3.1 **McIntosh-Smith S**, Price J, Deakin T and Poenaru A (2018), *Comparative Benchmarking of the First Generation of HPC-Optimised Arm Processors on Isambard*, CUG 2018, Stockholm https://uob-hpc.github.io/assets/cug-2018.pdf

3.2 **McIntosh-Smith S**, Price J, Deakin T and Poenaru A (2019). A Performance Analysis of the First Generation of HPC-Optimised Arm Processors, *Concurrency and Computation: Practise and Experience (CC-PE),* **31.16**, p.e5110*,* https://doi.org/10.1002/cpe.5110

3.3 **McIntosh-Smith S**, Price J, Poenaru A and Deakin T (2019), *Scaling Results From the First Generation of Arm-based Supercomputers*, CUG 2019, Montreal, Canada (Best Paper Award) http://uob-hpc.github.io/assets/cug-2019.pdf

3.4 **McIntosh-Smith S**, Price J, Poenaru A and Deakin T (2019). Benchmarking the first generation of production quality Arm-based supercomputers, *Concurrency and Computation: Practice and Experience (CC-PE)*, **32.20**, p.e5569, https://doi.org/10.1002/cpe.5569

Funding information

[i] Barcelona Supercomputer Centre (coordinator), **McIntosh-Smith S** (CI), *Mont-Blanc 2 Project*, European Commission FP7, grant agreement no. 610402, 13 partners, 2012-2015, GBP15 million, (Bristol share GBP387,000)

[ii] **McIntosh-Smith S** (PI), *GW4 Tier 2 HPC Centre for Advanced Architectures (Isambard 1)*, EPSRC EP/P020224/1, 2016-2021, GBP3M
Augmented by GBP1.7M from the GW4 Alliance and Isambard project partners: Met Office, Arm, Cray, CFMS and Zenotech (GBP4.7M total)

[iii] **McIntosh-Smith S** (PI), *GW4 Tier-2 HPC Centre for Advanced Architectures (Isambard 2)*, EPSRC EP/T022078/1, 2020-2023, GBP4.1M
Augmented by GBP2.5M from the GW4 Alliance and Isambard project partners: Met Office, Arm, Cray/HPE and UKAEA/CCFE (GBP6.6M total)

4. Details of the impact

In October 2018, the Isambard project delivered the world's first Arm-based supercomputing service, tackling the three major roadblocks to exploiting Arm in HPC detailed in Section 2, and thus opening up a new multi-million pound market for UK-based Arm Ltd, directly leading to hundreds of millions of dollars worth of Arm-based systems being deployed across the US, Europe, Japan and China. Isambard's results were highlighted in the press as being highly disruptive to the HPC market **[5.1]**.

Creation and growth of the HPC market for Arm Ltd (UK) and Marvell (US)

Isambard was instrumental in opening up the high-performance computing market for Arm, as their Senior Director for HPC states: '*Isambard's early results **[3.1-3.4]** were central in helping to kick-start this whole market for Arm, given they were the first independent, scientifically rigorous results proving Arm's competitiveness in this space*.' **[5.2]** Top industry analysts, Hyperion, have quantified the size of this new market **[5.3]**, publishing forecasts that more than 50,000 Arm-based processors would be sold in HPC in 2019 for a total systems revenue of USD145 million, growing to 233,000 processors sold in 2020, for USD675M in revenue. The market for Arm-based CPUs in HPC is predicted to grow to 610,000 processors in 2024 for a revenue of USD1.6 billion, both metrics growing at over 61% CAGR.

Isambard has also crucially enabled the scientific software ecosystem to be rapidly ported to the Arm architecture, a vital step to bringing the cost/performance benefits of the Arm ecosystem to the HPC market. Arm's Senior Director for HPC states that this key impact has: '*resulted in much more highly optimising compilers, better math libraries, and much more mature scientific applications running on Arm; Isambard has often been the first Arm-based supercomputer to which a scientific application is ported. Important examples include: the Unified Model atmospheric code from the Met Office, the application behind much of the weather prediction and climate modelling in the UK and around the world; VASP, the electronic structure chemistry code from the University of Vienna, one of the most widely used codes on supercomputers*

around the world, and the most heavily used code on Archer, the UK's national system; and OpenFOAM, a widely used open-source computational fluid dynamics code.' **[5.2]**
Isambard has resulted in significant benefit to Marvell, the US-based semiconductor company which developed the Arm-based ThunderX2 CPUs used in the project. Marvell's VP/GM of their Server Processor Business Unit states: '*Isambard was a breakthrough project for Marvell, being the world's first production Arm-based supercomputer service. … Following on from Isambard, Marvell has managed to win many more Arm-based supercomputers, including the first Arm-based system to appear in the Top500 league table of the fastest supercomputers in the world (Astra at Sandia National Laboratory in New Mexico) **[5.4]**, the three Catalyst systems at the universities of Edinburgh, Leicester and Bristol in the UK, two systems at Los Alamos National laboratory, also in New Mexico, and other systems at CEA in France and the Barcelona Supercomputing Center in Spain. All of these followed a similar approach to Isambard, using Marvell ThunderX2 CPUs, and together they represent tens of millions of dollars of business*.' **[5.5]**
Isambard has had such significant impact for both Arm and Marvell that in 2019 they contributed substantial monetary investments into the follow-on project, Isambard 2 **[iii]**, totalling GBP750,000 between them **[5.2, 5.5]**.

Isambard's impact in enabling the next generation of supercomputing systems
Isambard continues to have impact with the next generation of Arm-based systems being installed in 2020. It has enabled the Arm-based Fugaku system at RIKEN in Japan, which in June 2020 was confirmed as the new fastest supercomputer in the world. This system achieved a High Performance Linpack result of 415.5 petaflops, beating the second-place system by a factor of 2.8. It is powered by Fujitsu's 48-core A64FX CPU and is '*the first number one system on the list to be powered by Arm processors*.' **[5.4]** According to Arm, '*Tens of partners in our ecosystem are currently designing Arm-based server chips with more being added each quarter, and several designs have already come to market*.' **[5.2]**
Isambard has also been crucial to the Met Office's (UK) evaluation of advanced computer architectures, enabling them to port widely used weather and climate codes to Arm-based CPUs for the first time **[5.6]**. According to their Principal Fellow in Supercomputing, '*The Met Office has recently had outline funding approval for up to GBP1.2B **[5.7]** to develop and deliver weather prediction and climate modelling systems over the next ten years, and what we have learned from Isambard will be a vital input into this process.'* **[5.6]** Isambard has also been an important training platform for Met Office staff via hackathons and courses. He continues: '*These events have brought essential understanding and skills in utilising Arm CPUs and Nvidia GPUs to Met Office staff*.'
Similarly to Arm and Marvell, in 2019 the Met Office invested a further GBP50,000 into Isambard 2 **[5.6]**.

5. Sources to corroborate the impact

5.1 Next Platform – Article (November 2017), *Arm benchmarks show HPC ripe for processor shakeup* [Accessed 2 November 2020]
5.2 Arm Ltd – Corroborating Statement (June 2020), Senior Director High-Performance Computing and Director Research Ecosystem, UK and Europe
5.3 HPC market size and Arm's addressable fraction of this market: https://www.slideshare.net/insideHPC/hyperion-research-hpc-market-briefing-breakfast-at-sc19, slide 11
5.4 The Top500 list of the fastest supercomputers in the world: https://www.top500.org/lists/top500/2020/06/ [Accessed 2 November 2020]
5.5 Marvell – Corroborating Statement (May 2020), VP/GM Server Processor Business Unit
5.6 Met Office – Corroborating Letter (August 2020), Principal Fellow in Supercomputing
5.7 Met Office – Press Release (February 2020), *Up to £1.2billion for weather and climate supercomputer* [Accessed 2 November 2020]