

Impact case study (REF3b)

Institution: Swansea University
Unit of Assessment: 9 - Physics
Title of case study: A benchmark tool for high performance computing
1. Summary of the impact

This case study describes the development, application and commercialisation of an open source tool, **BSMBench** that enables **supercomputer vendors and computing centres** to benchmark their system's performance. It comprehensively informs the design and testing of new computing architectures **well beyond** other benchmarking tools on the market, such as Linpack.

The significance of our code is that, **unlike other benchmarking tools**, it **interpolates** from a communication- to a computation-dominated regime simply by varying the (physics) parameters in the code, thus providing a perfect benchmark suite to test the response of modern multi-CPU systems along this axis. The impact of this work has great reach: a **start-up company**, **BSMbench Ltd**, has been founded to develop and commercialise the software; adopters have included **IBM** – one of the giants of the supercomputer world (where it uncovered errors in their compilers); it has been deployed by **Fujitsu** to validate its systems, by HPC Wales, a multi-site, commercially focussed **national computer centre** and by **Transtec**, an HPC company employing over 150 staff; and **tutorial articles** about BSMBench have appeared in magazines such as Linux Format.

This software tool spawned from our research into “**Beyond the Standard Model**” (**BSM**) physics which aims to understand the Higgs mechanism in particle physics at a fundamental level. This involved simulating quantum field theories using bespoke code on some of the fastest supercomputers on the planet.

2. Underpinning research

The Standard Model of Particle Physics has at its centre the paradigm of electroweak symmetry breaking which explains the masses of fundamental particles. However it is incomplete because, e.g. it does not include gravity and there are issues surrounding the stability of the Higgs mass. Obtaining a more complete theory which answers these questions is a **Grand Challenge** problem. The tool described in this case study was born from Prof **Lucini's** research in ‘walking technicolor’ BSM models which seek to explain electroweak symmetry breaking via a composite Higgs. Due to these theories’ properties, progress can only be made using **numerical simulations**. This follows the same approach used to study the interaction of quarks via the strong force in which physicists have 30 years of computing experience.

Lucini initiated this research soon after his appointment to Swansea in October 2005. He and his group are still actively pursuing this programme and they are recognised as **world-leaders** in this field. The impact in this case study is based on research from **March 2008** [R1-R5]. These papers identified the first numerical signatures validating technicolor as a potential mechanism of electroweak symmetry breaking.

In order to simulate the wide range of parameters of the putative technicolor theories, a parallel, scalable and efficient simulation code is needed, but such a code was not available in 2008. Lucini therefore developed a suitable code, written so that the physical parameters representing different realisations of technicolor (essentially the number and representations of the “quarks” and “gluons” in each theory) could be **freely adjusted** as input parameters. This was very challenging conceptually and technically since varying these parameters leads to very different algebraic structures and hence different coding frameworks.

The performance of modern supercomputers depends on both the raw **computational power** of each processing unit and the **communication speed** between them. The communication component is crucially important since all supercomputers rely on linking many (typically hundreds

Impact case study (REF3b)

of thousands of) CPUs to obtain the maximum possible overall speed. The balance between these two parameters is a crucial factor in supercomputer design and testing.

The key ingredient of Lucini's BSM code which led to this research's impact is that the **user can easily adjust the physics parameters** corresponding to the choice of the technicolor theory being simulated. However, due to the ingenious way the software is written, this causes fundamental changes deep in the code which **independently stress the communication and computational** aspects of the computer system.

Lucini's leadership in this research programme was recognised by the award of a **Royal Society University Research Fellowship** [G1,G2], an **EPSRC software grant** [G3] and an **STFC Special Programme Grant** [G4] which funded a PDRA. Swansea University, recognising the potential impact of this work, granted Lucini sabbatical leave to develop this research programme. Other contributors to this project included the Swansea PDRAs Drs. Patella (at Swansea 2007-2009) and Rago (Swansea 2008-2010).

3. References to the research

Publications:

Swansea authors in bold. R1 and R2 best reflect the quality of the research. Reference [R1] explains the physics background of the research. The research code was the crucial tool used in [R2].

- [R1] **B. Lucini**, "Strongly Interacting Dynamics Beyond the Standard Model on a Spacetime Lattice", *Phil. Trans. Roy. Soc. Lond A* **368**, 2010, 3657, DOI:10.1098/rsta.2010.0030 [Impact Factor 3.1] Cited by 8 Google Scholar
- [R2] L. Del Debbio, **B. Lucini**, **A. Patella**, C. Pica and **A. Rago**, "Conformal Versus Confining Scenario in SU(2) with Adjoint Fermions", *Phys. Rev. D* **80** (2009) 074507, DOI: 10.1103/PhysRevD.80.074507 [Impact Factor 4.2] Cited by 95 Google Scholar
- [R3] Luigi Del Debbio, **Biagio Lucini**, **Agostino Patella**, Claudio Pica, Antonio Rago, "Mesonic Spectroscopy of Minimal Walking Technicolor", *Phys. Rev. D* **82** (2010) 014509 DOI:10.1103/PhysRevD.82.014509 [Impact Factor 4.2] Cited by 52 Google Scholar
- [R4] Luigi Del Debbio, **Biagio Lucini**, **Agostino Patella**, Claudio Pica, Antonio Rago, "The Infrared Dynamics of Minimal Walking Technicolor", *Phys. Rev. D* **82** (2010) 014510, DOI: 10.1103/PhysRevD.82.014510 [Impact Factor 4.2] Cited by 74 Google scholar
- [R5] Francis Bursa, Luigi Del Debbio, David Henty, Eoin Kerrane, **Biagio Lucini**, Agostino Patella, Claudio Pica, Thomas Pickup, Antonio Rago, "Improved Lattice Spectroscopy of Minimal Walking Technicolor", *Phys. Rev. D* **84** (2011) 034506, DOI: 10.1103/PhysRevD.84.034506 [Impact Factor 4.2] Cited by 23 Google scholar

Grant Income:

[G1] **Lucini (PI)** "Symmetry Aspects of Colour Confinement in QCD and Related Models", Royal Society University Research Fellowship, October 2005 to September 2010, £301k.

[G2] **Lucini (PI)** "Non-Perturbative Dynamics Beyond the Standard Model", Royal Society University Research Fellowship (Renewal), October 2010 to September 2013, £325k.

[G3] **Lucini (PI) Aarts, Allton, Hands** "Software Support for UKQCD", EPSRC software programmer support grant, October 2007 to March 2008, £15.6k. Grant awarded for supporting the coding aspects of the research project.

[G4] **Lucini (PI) Armoni** "Lattice Gauge Theories Beyond QCD", STFC Special Program Grant, April 2008 to March 2010, £228k. Grant supported PDRA (Dr. Patella).

Other grants are the STFC rolling and consolidated grants to the theoretical particle physics group of £3.06m (2008) and £1.3m (2011) and a £6m HPC grant supporting the research programme of the UKQCD collaboration (2009), of which this research is an important part.

4. Details of the impact

Simulations using **High Performance Computers (HPC)** are established as the third mode of scientific investigation after experiment and theory and it is inevitable that their impact will grow. HPC systems are best viewed as tightly interconnected collections of individual computers. Therefore, two main factors affect their overall performance: the **individual CPU power** and the **communication throughput** between these CPUs. Understanding both aspects is crucial for designing, developing and testing HPC platforms.

In order to rank the performance of HPC machines, benchmarking software is used. By far the most popular is **Linpack** which is the only performance tool used in the universally accepted list of the world's fastest computers, www.top500.org. However, Linpack measures only computational speed (via dense matrix multiplications) and it *“does not stress local bandwidth”* and *“does not stress the network”* according to Jack Dongarra, who first introduced this benchmark in the 1970's. Thom Dunning, director of the National Center for Supercomputing Applications (University of Illinois at Urbana-Champaign), agrees with this sentiment: *“almost anyone who knows about it will deride its utility”*.

The unique property of our BSMBench code – the **capacity to vary computational and communication** aspects simply by adjusting the physical parameters of the theories it simulates – makes it an exemplary alternative benchmark tool. In 2011, Lucini realised the important impact that his software would have in the design and appraisal of HPC and therefore began packaging the code into a benchmark tool, **BSMBench**. He utilised his IBM contacts to initiate a formal **collaboration with IBM** to implement this work. This project was finalised by the **internship of Swansea PhD student**, Ed Bennett, in the **IBM Research Center** in Cambridge (MA) (from 25/6/11 to 17/8/11) formalised under a Joint Study Agreement between Swansea University and IBM [C1]. Lucini had overall design lead of the benchmark tool and remains the maintainer of the code. This project resulted in a conference presentation, co-authored by Lucini and **Dr Kirk Jordan of IBM Research** at the industrial-facing International Supercomputing Conference, ISC12 [C2]. As a result of this **commercial collaboration**, BSMBench was released as an **open source project** in June 2012, and can be downloaded free of charge from www.bsmbench.org. The code has been designed to run easily on a variety of systems, from small symmetric multi-processor machines to the most powerful supercomputers and it has been tested on systems from 64 to 8192 cores with no degradation in performance.

IBM has been using BSMBench since August 2011 to inform the design of its supercomputers and to promote their cutting-edge **BlueGene/Q** systems. These are currently some of the planet's most powerful HPC architectures: they occupy four of the top ten places in the June 2013 world ranking (www.top500.org). The BSMBench product has had **significant impact on IBM** and, importantly, **helped track some errors in their multi-million dollar BlueGene/Q systems**. From the IBM HPC Associate Program Director [C3]:

*I write to acknowledge the impact that BSMBench, the benchmarking tool derived from your research on Strongly Interacting Dynamics beyond the Standard Model, has had on the evaluation of our BlueGene/Q architecture. While the www.top500.org ranking of these systems is based on the Linpack benchmarking tool, the **capabilities of our architecture were confirmed by the more stringent BSMBench tests**. BSMBench establishes a new concept in benchmarking High Performance Computing systems. **We have used it to test our recent IBM BlueGene/Q systems prior to their commercialization** through the work Ed Bennett did, and we have found the results very informative helping to identify some issues with the compiler and support software.*

BSMBench is being developed and commercialised by **BSMbench Ltd**, a **start-up company founded in 2012 specifically created** to promote, market, and utilise the BSMBench code suite. Lucini was a founder and is Chief Technical Officer and a shareholder. This company has won **significant EU convergence funding worth £180k** to advance and refine the software [C4]. Currently, it is developing general-purpose parallel software based on BSMBench that can be used in HPC applications in sectors such as finance, aerospace engineering, weather forecasting and oil

Impact case study (REF3b)

extraction. BSMbench has recorded downloads of its software from companies such as **Fujitsu, Research in Motion (Blackberry), Jaguar Land Rover, KLM and Microsoft** [C5]. From the BSMbench Company Director [C5]:

“this tool has been a crucial enabler for us and it is thanks to it and to its open source nature that we were able to set up our business, BSMbench Ltd. Basing its activity on BSMBench and on further developments of it, our company is constantly growing its business and attracting a larger base of customers.”

Since the summer of 2012, **HPC Wales** – a £40m national computing centre – has used BSMBench to benchmark and monitor the performance of its systems [C6]. In April 2013, BSMBench was used by **Fujitsu** (a major IT systems company with annual turnover of £1.8bn) to validate and benchmark their computers at HPC Wales’s latest Hub prior to its launch [C7]. BSMBench has also been used by **Transtec**, a European HPC company with annual revenues of €45m, to benchmark their latest commercial products [C8]. Furthermore, Lucini was asked to perform monitoring of **Fujitsu’s Sandybridge clusters** using BSMBench [C9]. The Chief Executive of HPC Wales said [C6]:

“BSMBench has been used extensively on our systems, to monitor their performance and detect potential problems at an early stage. In all our applications, BSMBench has proved to be robust and reliable, and has been an invaluable asset to enable us to provide a high standard of service to our customers.”

The software has attracted the interest of the **Open Source community**. An article reviewing its features and describing its use in an industrial environment was **published in Linux Format**, the UK’s best selling Linux magazine with a monthly circulation of 24,000 [C10], and in issue 124 of the Italian magazine *Linux Pro* [C11]. From the editor of *Linux Format* [C10]:

“I write to acknowledge the impact that BSMBench has had for our company. We carefully select the contributions we include in our publications, for which we consider the packages that are of high interest for our readers. Following this guideline, a 4-page tutorial on using BSMBench was included in Linux Format issue 163 and the software package itself was part of the bonus DVD in issue 164. Moreover, we have been using BSMBench for reviews of hardware systems published in subsequent issues of our magazine.”

5. Sources to corroborate the impact

[C1] Joint Study Agreement (No. W1157505) between Swansea University and IBM

[C2] E. Bennett, B. Lucini, K. Jordan et al., “BSMBench: A HPC Benchmark for Beyond the Standard Model Lattice Physics”, presentation at the *International Supercomputing Conference, ISC12*, Hamburg, June 17-21 2012, tinyurl.com/nydae34

[C3] Emerging Solutions Executive, Computational Science Center, IBM T.J. Watson Research, Cambridge Research Center, Cambridge, MA USA

[C4] EU Convergence Funding to BSMBench, European Regional Development Fund Project Funding HPCW-0033, 24th October 2012. Total value £180k

[C5] Company Director, BSMbench Ltd., Dylan Thomas Centre, 1 Somerset Place, Swansea, SA1 1RR, www.bsmbench.com

[C6] Chief Executive, HPC Wales, Ty Menai, Ffordd Penlan, Parc Menai Business Park, Bangor, Gwynedd LL57 4HJ

[C7] Benchmarking report from Principal Researcher, Fujitsu Laboratories of Europe

[C8] HPC Solution Architect, Transtec AG, Tübingen www.transtec.co.uk

[C9] Project Manager – Hosting, Network and Security, Fujitsu www.fujitsu.com/uk

[C10] Linux Format, 30 Monmouth Street, Bath BA1 2BW www.linuxformat.com

[C11] Linux Pro, Via Torino 51, I-20063 Cernusco sul Naviglio (MI), Italy facebook.com/LinuxPro.it