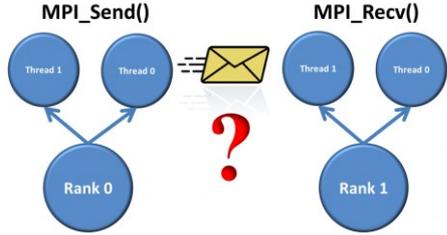


Institution: PHYESTA	
Unit of Assessment: UoA 9 – Physics	
Title of case study: The Message Passing Interface (MPI): An International Standard for Programming Parallel Computers	
<p>1. Summary of the impact</p> <p>Impact: Economic gains The Edinburgh Parallel Computing Centre (EPCC) made substantial contributions to the development of MPI and produced some of its first implementations; CHIMP/MPI, and CRI/EPCC MPI, for Cray T3D and Cray T3E supercomputers.</p> <p>Significance: MPI is the ubiquitous de-facto standard for programming parallel computers. Software written to use MPI can be transparently run on any parallel system, from a multi-core desktop computer to a high-performance supercomputer.</p> <p>Reach: Hardware vendors including Cray, IBM, Intel, and Microsoft all support MPI. The world's 500 most powerful supercomputers all run MPI. Hundreds of companies use MPI-based codes.</p> <p>Beneficiaries: Hardware vendors, software vendors, scientific, industrial and commercial ventures. Specific examples include Cray Inc., <i>[text removed for publication]</i> and Integrated Environmental Systems.</p> <p>Attribution: This work was led by Professor Arthur Trew and Dr Lyndon Clarke.</p>	
<p>2. Underpinning research</p> <p>Most parallel computing systems utilise message passing to communicate information between different processes running on multiple processors. Standard interfaces rather than vendor-specific communication mechanisms are essential for the development of portable libraries, toolkits and applications, thereby giving users the flexibility and security to choose the hardware that best satisfies their operational and computational needs and financial constraints.</p> <p>Inspired by portability issues arising in the computational physics work of UKQCD (led by Kenway [R1,R2], and that of Clarke within the UCKP consortium [R3,R4], Trew led research whose output was the Common High-level Interface to Message Passing (CHIMP) [R5]. CHIMP was a portable parallel software library which provided a set of standard interfaces that ran on parallel hardware from multiple vendors. ([R5] was among Trew's outputs in Edinburgh's physics RAE1996.)</p> <p>In 1994, the Message Passing Interface (MPI) 1.0 standard [R6], was released by the Message Passing Interface Forum (MPIF), a collaboration body of over 40 organisations from academia and industry including leading hardware vendors Cray, IBM, Intel, Meiko, NEC and Thinking Machines. CHIMP's influences on this standard were substantial, and facilitated by Clarke's engagement with the MPIF, where he served as a coordinator of the "Groups, Contexts and Communicators" sub-group. Prior to CHIMP, most message passing interfaces used a single global communications context and collective communications across all processes in a system. CHIMP introduced message contexts to isolate communications between independently communicating processes, and collective communications within specific process groups. This exploits the fact that not all computational processes need to receive all communications within the system. These ideas are directly encapsulated within MPI's defining concepts of groups, contexts and communicators [R6].</p> <p>In 1994 PHYESTA researchers produced one of the first implementations of the MPI standard, CHIMP/MPI, followed in 1995 by CRI/EPCC MPI, a native implementation developed by EPCC and Cray Inc. for the Cray T3D and Cray T3E supercomputers. CRI/EPCC MPI included an implementation of single-sided communications for remote memory access, which had been proposed for MPI and was subsequently adopted as part of the MPI-2 specification, released in 1997. Under Clarke's influence, MPIF had identified these as important upgrade areas. MPI-3 was</p>	

Impact case study (REF3b)

published in September 2012 and continues to acknowledge the important contributions made by EPCC, who in turn continue to contribute to the work of the MPI Forum. Scopus, ScienceDirect, CiteSeerX and Google Scholar yield 7200, 2900, 3900 and 22100 articles respectively, featuring the terms “MPI” and “message passing interface”, published between 1993 and the present.

Personnel:

Key PHYESTA researchers involved were Professor Richard Kenway (1993-present), Professor Arthur Trew (1993-present) and Dr. Lyndon Clarke (EPCC research staff, 1993-1999). Kenway led EPCC as its Director until 1997 and remains its Chairman. Trew was its Consultancy Manager while directly involved in this research [R5], then EPCC Director from 1997 to 2008. Trew was returned as Category A in Physics in RAE1996 and RAE2008, since when he has been Head of the School of Physics and Astronomy in Edinburgh.

3. References to the research

The quality of the work is best illustrated by [R2], [R4], [R5] and [R6]. *[Number of citations]*

[R1]	N. Stanford (for UKQCD collaboration), Portable QCD codes for Massively Parallel Processors, Nucl. Phys. Proc. Suppl. 34, 817-819 (1994); http://www.arXiv.org/abs/hep-lat/9312010 [0]
[R2]	C. R. Alton et al (UKQCD collaboration), Light hadron spectrum and decay constants in quenched lattice QCD, Phys. Rev. D 49, 474-485 (1994); http://prd.aps.org/abstract/PRD/v49/i1/p474_1 [41]
[R3]	L. J. Clarke, Parallel Processing for Ab Initio Total Energy Pseudopotentials, Theoretica Chimica Acta 84, 325-334 (1993) http://dx.doi.org/10.1007/BF01113271 [3]
[R4]	A. de Vita, I. Stich, M. J. Gillan, M. C. Payne and L. J. Clarke, Dynamics of dissociative chemisorption: Cl ₂ /Si(111)-(2x1), Phys. Rev. Lett. 71, 1276 (1993) http://link.aps.org/doi/10.1103/PhysRevLett.71.1276 [49]
[R5]	R. A. A. Bruce, S. Chapple, N.B. MacDonald, A.S. Trew, S. Trewin. CHIMP and PUL: Support for portable parallel computing. Future Generation Computer Systems, 11, 211-219 (1995) http://dx.doi.org/10.1016/0167-739X(94)00063-K [2 (Google Scholar)]
[R6]	J. Dongarra, et al., “MPI - A Message-Passing Interface Standard”. International Journal Of Supercomputer Applications And High Performance Computing, 8(3-4) P. 165-416. September 1994. http://hpc.sagepub.com/content/8/3-4.toc [290 (Google Scholar)] Bibliographical Note /Role statement for Clarke: This special issue represents the initial full publication of the MPI Standard, v1.0 as a research output. This is listed on WoS as a single article with 64 authors, and normally cited as such; WoS lists 35 citations to it. (Scopus lists authorship as 'Anon' and captures only 3 citations despite there being 290 on Google Scholar.) Clarke is among 21 authors who held 'positions of responsibility' for the creation of MPI. (The remaining 43 authors were 'active participants' in the development of MPI.) Specifically he is the second among 5 authors responsible for the "Groups Contexts and Communicators" aspects of MPI (pp 311-356). These roles are explained in the open-access edition of v1.0 (which is otherwise the same as the Journal publication) at http://www.mpi-forum.org/docs/mpi-1.0/mpi-10.ps and also in the 1995 html version (v1.1, with 'minor' changes from v1.0) at http://www.mpi-forum.org/docs/mpi-1.1/mpi-11-html/mpi-report.html

4. Details of the impact

MPI is the ubiquitous de-facto standard for programming parallel systems. The role of EPCC in its development and implementation is acknowledged by independent authorities [S1, S2]. MPI can be used to run parallel applications on large-scale high performance computing (HPC)

Impact case study (REF3b)

infrastructures, local clusters or multi-core desktop computers. To be compliant with the MPI standard, implementations *must* support the operations for point-to-point and collective communications, groups and communicators, which are the parts of MPI to which Clarke directly contributed. There are now myriad MPI implementations, both open source and commercial, many of which were produced (and many continually used) within the REF impact window. For example, MVAPICH, an open source derivative of MPICH, for use with high-performance networks, has recorded over 182,000 downloads and over 2,070 users in 70 countries, including 765 companies [S3]. Among these we cite IBM as an example of a leading HPC supplier using MPI [S4].

The 'top500' is a list of the most powerful super-computers published twice yearly [S5]. The benchmark used to rank these systems is Linpack, which is parallelised using MPI. Consequently, all top500 ranked supercomputers run MPI by definition. More significantly, the Linpack benchmark is itself chosen in large part because so many of the world's most powerful machines run MPI-based software applications on a frequent or continuous basis.

EPCC serves on the HPC Advisory Council [S6], which includes over 300 hardware and software vendors, HPC centres and selected end-users, including 3M, Bull, CDC, Dell, EPCC, Hitachi, HP, IBM, Intel, Microsoft, NEC Corporation of America, NVIDIA, Schlumberger, SGI, STFC, and Viglen. In the Council's online survey of best practices, covering 45 commercial packages, all but one uses MPI. These commercial packages span a range of applications including finite element analysis (ABAQUS-Dassault Systemes, MSC NASTRAN for MEA-MS Software), material modelling (ABAQUS-Dassault Systemes), computational fluid dynamics (AcuSolve-Acusim, CFX and FLUENT-ANSYS, FLOW-3D-FLOW Science, OpenFOAM-OpenCFD Ltd, STAR-CCM+-CD-adapco), molecular dynamics (D.E. Shaw Research), oil and gas reservoir simulation (ECLIPSE-Schlumberger), and crash simulation (LS-DYNA-Livermore Software Technology Corporation).

Without MPI, hardware vendors would face an inefficient marketplace for their hardware, as users would need to develop their software to use vendor-specific infrastructure which would limit the portability of their software, and so limit the ability of users to migrate to other hardware platforms. For software vendors, MPI opens up possibilities for users to seamlessly exploit multi-core, multi-processor, cloud or supercomputer architectures, and so more readily access increased computing power. Since MPI's inception in 1994, EPCC/PHYESTA has collaborated with hardware and software vendors to exploit MPI to deliver these benefits within a commercial context. We next present some recent examples of this, which show how MPI has delivered economic impact. These are a small sample of the ways in which MPI is used daily in commerce and industry around the world (mostly with no PHYESTA involvement beyond its role in the development of MPI itself).

Cray Inc. have been a leading hardware vendor for over 40 years. In the early 90s Cray contributed to the development of the MPI specification. EPCC and Cray together developed one of the first MPI implementations, CRI/EPCC for the Cray T3D machine [R6]. Cray's Manager for Exoscale Research Europe states: *"Today, a Cray MPI library is distributed with every Cray system and is the dominant method of parallelisation used on Cray systems. Cray reported total revenue of over \$420 million in 2012, so this is a large industry which is heavily reliant on MPI and the work that EPCC contributed in this regard"* [F1]. Cray and EPCC continue to collaborate on programming model development and research, on MPI and other programming models, for example, enhancing MPI for exascale as part of the FP7 funded EPIGRAM project. Facts in this paragraph are confirmed in [F1].

[text removed for publication]

In 2013, EPCC worked with Integrated Environmental Systems (IES). (This was done as part of SuperComputing Scotland, a joint EPCC-Scottish Enterprise programme [S7].) IES is the world's leading provider of software and consultancy services on energy efficiency within the built environment. IES's SunCast software allows architects and designers to analyse the sun's shadows and the effects of solar gains on the thermal performance and comfort of buildings. SunCast calculates the effect of the sun's rays on every surface at every hour of a design day, a total of 448 separate calculations. Porting SunCast to run over Microsoft MPI allows for the parallel processing of surfaces, one per processor. When a processor has completed a surface, it notifies the controlling processor of its results and that it is ready to be assigned another surface. With MPI, SunCast can now run on a supercomputer with order-of-magnitude time savings for analyses

Impact case study (REF3b)

– from 30 days to 24 hours in one example. These facts are confirmed in [F3] where IES Director, Craig Wheatley, also comments: "Additionally, using the MPI in IES Consultancy has increased the efficiency and therefore profitability of our own consultancy offering and to date we have used it with 4 live projects with an average analysis time of under 12 hours. These particular projects were very large and complex and would otherwise have taken several weeks."

5. Sources to corroborate the impact

[S1]	M. Snir, S. Otto, S. Huss-Lederman, D. Walker, J. Dongarra. "MPI – The Complete Reference". Scientific and Engineering Computation Series, MIT Press, 1996. ISBN 0262691841. [1144 cites] <i>Corroborates EPCC role in development and implementation of MPI</i>
[S2]	W. Gropp, E. Lusk, N. Doss, A Skjellum. A high-performance, portable implementation of the MPI message passing interface standard. Parallel Computing 22(6) September 1996, pp789-828. http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.29.4569 [748 cites] <i>Corroborates EPCC role in development and implementation of MPI</i>
[S3]	MPICH. Collaborators, http://www.mpich.org/about/collaborators/ OpenMPI. The Open MPI Development Team, http://www.open-mpi.org/about/members/ . MVAPICH2. Current users http://mvapich.cse.ohio-state.edu/current_users/ <i>These listings corroborate ubiquity of MPI usage in both academia and commerce</i>
[S4]	IBM Platform MPI V8.3 delivers high performance application parallelization, IBM United States Software Announcement 212-203, 4 June 2012, http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=897/ENUS212-203&infotype=AN&subtype=CA . <i>Corroborates IBM implementation of MPI</i>
[S5]	Top 500 supercomputer sites, http://www.top500.org/ <i>Corroborates that supercomputers are benchmarked by MPI-Linpack</i>
[S6]	HPC Advisory Council, http://www.hpcadvisorycouncil.com/council_members.php . <i>Corroborates EPCC role in HPC advisory council</i>
[F1]	Factual statement from Manager Exascale Europe at Cray about MPI
[F2]	[text removed for publication]
[S7]	Scottish Enterprise services, SuperComputing Scotland. http://www.scottish-enterprise.com/services/develop-new-products-and-services/supercomputing-scotland/overview msn news. Supercomputers 'good for business'. 22 November 2011. http://news.uk.msn.com/science/supercomputers-good-for-business-35 SuperComputing Scotland. News. Using Supercomputing to Design Energy Efficient Buildings, 18 June 2013. http://www.supercomputingscotland.org/news/using-supercomputing-to-design-energy-efficient-buildings.html <i>Corroborate EPCC partnership with Integrated Environment Systems</i>
[F3]	Factual statement from Director of IES about use of MPI for SunCast