Dr David Ham

Qualifications

PhD Applied mathematics "On techniques for modelling coastal and ocean flow with unstructured meshes" (2006), TU Delft.

BSc. Bachelor of Science (first class honours) mathematics (2001), The Australian National University.

LLB Bachelor of Laws (2003), The Australian National University.

Career

Current

Senior Lecturer (equivalent to Associate Professor) Department of Mathematics, Imperial College London.

Natural Environment Research Council (NERC) Independent Research Fellow Jointly appointed in the Departments of Mathematics and Computing, Imperial College London (2013-2016) Department of Mathematics (2016-)

- Research into code generation based on the composition of **high level mathematical abstractions** for the creation of complex parallel finite element software, particularly for ocean and atmosphere simulation.
- Development of symbolic computer mathematical techniques to automate the generation of **inverse (adjoint) simulation software**.
- Development of **new numerical algorithms** and **parallel finite element software**.
- Creation of **mathematical compiler technology** for the generation of finite element code optimised for emerging vectorised hardware.
- Building strategic research collaborations, and winning competitive research funding.
- Lecturing in computational science, mathematics and simulation software development.
- Supervision of doctoral and masters students.

Previous

Grantham Research Fellow Department of Computing (2012-2013), Department of Earth Science and Engineering (2008-2012), Imperial College London.

Research associate Department of Earth Science and Engineering, Imperial College London (2005-2008).

Research assistant Fluid Mechanics Section, Faculty of Civil Engineering and Geosciences, Delft University of Technology (2001-2005).

Research funding

Total funding £4.8M, since 2010.

- Natural Environment Research Council (NERC) Independent Research Fellowship NE/K008951/1 "Abstracting the environment: automating geoscientific simulation" £485 994, 2013-2018
- PI, NERC grant NE/K006789/1 "Gung Ho Phase 2" - developing a new UK weather and climate model. £294 332, Feb 2013-Jan 2016.
- PI, NERC grant NE/I001360/1 "Automated adjoints: how much do we really know about the source of the Indian Ocean Tsunami?" £59 278, Oct 2010-Sep 2012
- PI, NERC grant NE/I021098/1 "Abstracting the hardware: Assembly algorithms for numerical weather prediction on emerging massively parallel architectures" £101 387, May 2011-May 2013
- Co-I Intel Parallel Computing Centre: Intel Parallel Computing Centre: Performance portablE SeismiC Imaging (PESCI): £145 000, Jan 2015-May 2016
- Co-I Engineering and Physical Sciences Research Council (EPSRC) grant EP/M011054/1 "A new simulation and optimisation platform for marine technology" £434 711, 2013-2018
- Co-I EPSRC grant EP/L000407/1 "Platform: underpinning technologies for finite element stimulation" £1 271 599, 2013-2018
- Co-I, EPSRC grant EP/K008730/1 "PAMELA: a Panoramic Approach to the Many-CorE LAndsape - from end-user to end-device: a holistic game-changing approach." £1.4M, Mar 2013-Feb 2018
- Co-I, EPSRC grant EP/I00677X/1 "Multi-layered abstractions for PDEs" £450 000, Oct 2010-Mar 2014
- Co-I, NERC grant NE/I000747/1"Unstructured mesh dynamical core for atmospheric modelling using geophysically-optimal finite elements" £111 519, Oct 2010-Sep 2011
- Co-I, Imperial/EPSRC grant "Bridging the gaps in computational science and engineering" £26 880 Oct 2010-June 2011
- Winning team, Imperial College Research Excellence Award 2010 £100 000

Invited presentations in last 5 years.

- **Plenary talk** *Higher order DF methods and finite element software for modern architectures* Bath (2016)
- **Keynote lecture** *14th International workshop on Multi-scale (Un)-structured mesh numerical Modeling for coastal, shelf, and global ocean dynamics* Portland, Oregon (2015)
- **Plenary talk**, CliMathNet conference, Bath (2015)
- ICMS workshop *Galerkin methods with applications in weather and climate forecasting* Edinburgh (March 2015)
- **Keynote lecture** PDESoft, Heidelberg (2014)
- University of Edinburgh (2014)
- Americal Geosciences Union Fall Meeting, San Francisco (2013)
- Department of Computer Science, TU Darmstadt (2013)
- NOAA Geophysical Fluid Dynamics Laboratory, Princeton (2013)
- *High-order numerics and high performance computing* Department of Aeronautics, Imperial (2012)
- Simula Research Laboratory, Lysaker, Oslo (2012)
- Partnership for Advanced Computing in Europe (PRACE) workshop on community climate codes, Barcelona, Spain (2012)
- *Dynamical cores for climate models*, hosted by is-enes (Infrastructure for the European Network for Earth System Modelling), Lecce, Italy (2011)
- *Simulating the Spatial-Temporal Patterns of Anthropogenic Climate Change*, run by Los Alamos National Laboratory (2011)
- ENUMATH minisymposium on Domain-Specific Languages and Code Generation for PDE Problems (2011)
- US National Conference on Computational Mechanics (2011)

Esteem and international visibility

- Member of the technical programme committee, Supercomputing (2016)
- Member of the Research Paper Committee, International Supercomputing Conference (ISC 2013 - 2016).
- Member of the scientific committee, PDESoft 2016.
- **The Wilkinson Prize for Numerical Software 2015**. The world's leading numerical software prize, awarded only every 4 years.
- Best poster award, Siam Conference on Computational Science and Engineering (2015).
- Chair, scientific committee and local organising committee. FEniCS '15 workshop.
- External assessor, Swiss Platform for Advanced Scientific Computing (PASC) Grid-tools project (2014-2015).
- External PhD examiner, Division of Geology and Geophysics, University of Sydney (2014 & 2015).
- Member of the Programme Committee, International Conference on Parallel Computing (ParCo 2013 & 2015).
- Member of the editorial board of Geoscientific Model Development (2010-present).
- Organiser and visiting fellow, Isaac Newton Institute for Mathematical Sciences programme *Multiscale numerics for the atmosphere and ocean* (2012).
- Convener, Ocean Modelling Session, European Geosciences Union (2010-2012).
- Member of the scientific advisory committee for the 9th International workshop on unstructured mesh numerical modelling of coastal, shelf and ocean flows, MIT (2010).
- Briefed Thailand's minister for science and technology on simulating the Indian Ocean Tsunami and aspects of tsunami risk (2009).
- Guest editor, special issue of Ocean Modelling (2008-2009).

Staff supervised

Dr Lawrence Mitchell, research associate (June 2013-)

Doctoral supervision

Multilayer abstractions for PDEs
Graham Markall, Department of Computing, Imperial, completed 2013

Productive and efficient computational science through domain-specific abstractions
Florian Rathgeber, Department of Computing, Imperial, completed 2014
Awarded 2nd prize, best thesis in Data Science 2015, Imperial DSI.

Compatible finite element methods for atmospheric dynamical cores
Andrew McRae, Department of Mathematics, Imperial completed 2015
Awarded Department of Mathematics Dorris Chen merit award 2014

Automated optimal vectorisation of unstructured mesh simulations
Fabio Luporini, Department of Computing, Imperial (2012-)

Topology-aware abstractions for PDE solvers
Gheorghe-teodor Bercea, Department of Computing, Imperial (2012-)

Solving the big data problem in climate simulation
Miklòs Homolya, Department of Computing, Imperial (2014-)

Lecturing

Undergraduate courses

Introduction to law for computer scientists introductory lectures on copyright, contract and data protection law for second year computing students (Summer 2013-2016).

Numerical methods I Introductory scientific computing for second year science students (Spring 2011 & 2012)

Maths methods II ordinary differential equations, complex arithmetic and infinite series for first year geophysics students (Spring 2010).

Student online evaluation comments include:

- "Exceptionally impressed - David's explanation of even challenging topics was consistently excellent, and his sense of humour made his lectures a pleasure to attend."
- "*The best lecturer I have had so far at Imperial College.* He explains concepts very well, is very approachable and willing to help."
- "Very good lecturer. Very well structured and brought things together well. Included some history behind the maths, which was interesting."
- "If only all modules had handouts this good...!"

Postgraduate courses

The finite element method: analysis and implementation applied mathematics masters level module. I created and taught the implementation part of the module, combining implementing advanced mathematical algorithms with learning to apply professional software engineering tools and techniques to a scientific computing problem (Spring 2015 & 2016).

The finite element method lecture course in the theory and implementation of the finite element method given to PhD students and RAs in the departments of Computing, Mathematics and Earth Science and Engineering. Lectures are recorded and **available online**. The course was followed by external viewers in Oxford, the UK Met Office, Louisiana State University, and anonymous viewers from the UK, US, Germany, Mexico, India and Saudi Arabia among others. Lectures from the course currently receive about **2000 views per month**. (Autumn 2013–Spring 2014).

Implementing the finite element method day course in scientific software development for PhD students, academics, industry and government scientists. Given at Imperial (annually 2009–2012) and, by invitation, at the University of Oxford (2008) and the UK Met Office (2011).

CSCS-USI Summer School on Computer Simulations in Science and Engineering software engineering crash course for geoscientific modellers. Swiss national supercomputing centre. (2013)

NCAS Climate Modelling Summer School guest lecture on next generation climate models. (2013)

Management and service

- Co-Director of Doctoral Studies, EPSRC Centre for Doctoral Training in Mathematics of Planet Earth (2016-present)
- Associate Director of the **Centre for Computational Methods in Science and Engineering**, Imperial College London (2012-present)
- Responsible for coordinating software licensing between various institutions and individuals developing finite element packages. Liaising with university commercialisation departments and providing science and business cases for the open source licensing of simulation software (2006-present).
- Registrar of the World Universities Debating Council (2004-2007).

Scientific software developed

Scientific software to which I am or have been a core contributor, with the dates of involvement.

Firedrake system for automatically generating finite element simulations from high level mathematical specifications. (2012-) `http://firedrakeproject.org`

Dolfin-adjoint system for automatically generating and executing adjoint simulations to forward finite element models written in FEniCS and Firedrake. (2011-) `http://dolfin-adjoint.org`

Fluidity adaptive mesh finite element fluid simulation package. Used by and for companies and government organisations in the UK and overseas. (2005-2012)

Spud high level system for generating input files and graphical interfaces for simulation software. (2007-2010)

Journal publications

Ford, R., M. J. Glover, D. A. Ham, C. M. Maynard, S. M. Pickles, G. D. Riley, and N. Wood (2015). "Gung Ho: A code design for weather and climate prediction on exascale machines". *Accepted for publication Advances in Engineering Software*.

Heinis, T. and D. A. Ham (2015). "On-the-Fly Data Synopses: Efficient Data Exploration in the Simulation Sciences". *SIGMOD Record* 44.2, pp. 23–28. ISSN: 0163-5808. DOI: 10.1145/2814710.2814715. URL: http://doi.acm.org/10.1145/2814710.2814715.

Luporini, F., A. L. Varbanescu, F. Rathgeber, G.-T. Bercea, J. Ramanujam, D. A. Ham, and P. H. J. Kelly (2015). "Cross-Loop Optimization of Arithmetic Intensity for Finite Element Local Assembly". *ACM Trans. Archit. Code Optim.* 11.4, 57:1–57:25. ISSN: 1544-3566. DOI: 10.1145/2687415. URL: http://doi.acm.org/10.1145/2687415.

Hill, J., E. E. Popova, D. A. Ham, M. D. Piggott, and M. Srokosz (2014). "Adapting to life: ocean biogeochemical modelling and adaptive remeshing". *Ocean Science* 10.3, pp. 323–343. DOI: 10.5194/os-10-323-2014. URL: http://www.ocean-sci.net/10/323/2014/.

Farrell, P. E., D. A. Ham, S. W. Funke, and M. E. Rognes (2013). "Automated derivation of the adjoint of high-level transient finite element programs". *SIAM Journal on Scientific Computing* 35, pp. C359–393. DOI: 10.1137/120873558.

Markall, G. R., A. Slemmer, D. A. Ham, P. H. J. Kelly, C. D. Cantwell, and S. J. Sherwin (2013). "Finite element assembly strategies on multi- and many-core architectures". *International Journal for Numerical Methods in Fluids* 71, pp. 80–97. DOI: 10.1002/fld.3648.

Rognes, M. E., D. A. Ham, C. J. Cotter, and A. T. T. McRae (2013). "Automating the solution of PDEs on the sphere and other manifolds in FEniCS 1.2". *Geoscientific Model Development* 6.6, pp. 2099–2119. DOI: 10.5194/gmd-6-2099-2013.

Du, J., F. Fang, C. Pain, I. Navon, J. Zhu, and D. Ham (2012). "POD reduced-order unstructured mesh modeling applied to 2D and 3D fluid flow". *Computers and Mathematics with Applications* 65, pp. 362–379. DOI: 10.1016/j.camwa.2012.06.009.

Hill, J., M. D. Piggott, D. A. Ham, E. E. Popova, and M. A. Srokosz (2012). "On the performance of a generic length scale turbulence model within an adaptive finite element ocean model". *Ocean Modelling* 56, pp. 1–15. DOI: 10.1016/j.ocemod.2012.07.003.

Cotter, C. J. and D. A. Ham (2011). "Numerical wave propagation for the triangular P1DG-P2 finite element pair". *Journal of Computational Physics* 230.8, pp. 2806–2820. DOI: 10.1016/j.jcp.2010.12.024.

Farrell, P. E., M. D. Piggott, G. J. Gorman, D. A. Ham, C. R. Wilson, and T. M. Bond (2011). "Automated continuous verification for numerical simulation". *Geoscientific Model Development* 4.2, pp. 435–449. DOI: 10.5194/gmd-4-435-2011.

Cotter, C. J., D. A. Ham, and C. C. Pain (2009). "A mixed discontinuous/continuous finite element pair for shallow-water ocean modelling". *Ocean Modelling* 26.1-2, pp. 86–90. DOI: 10.1016/j.ocemod.2008.09.002.

Cotter, C. J., D. A. Ham, C. C. Pain, and S. Reich (2009). "LBB stability of a mixed Galerkin finite element pair for fluid flow simulations". *Journal of Computational Physics* 228.2, pp. 336–348. DOI: 10.1016/j.jcp.2008.09.014.

Ham, D. A., P. E. Farrell, G. J. Gorman, J. R. Maddison, C. R. Wilson, S. C. Kramer, J. Shipton, G. S. Collins, C. J. Cotter, and M. D. Piggott (2009). "Spud 1.0: generalising and automating the user interfaces of scientific computer models". *Geoscientific Model Development* 2, pp. 33–42. DOI: 10.5194/gmd-2-33-2009.

Ham, D. A., C. C. Pain, E. Hanert, J. Pietrzak, and J. Schröter (2009). "Special Issue: The sixth international workshop on unstructured mesh numerical modelling of coastal, shelf and ocean flows. Imperial College London, September 19-21, 2007". *Ocean Modelling* 28.1-3, pp. 1–1. DOI: 10.1016/j.ocemod.2009.02.005.

Ham, D. A., S. C. Kramer, G. S. Stelling, and J. Pietrzak (2007). "The symmetry and stability of unstructured mesh C-grid shallow water models under the influence of Coriolis". *Ocean Modelling* 16 (1-2), pp. 47–60. DOI: 10.1016/j.ocemod.2006.05.008.

Pietrzak, J., A. Socquet, D. Ham, W. Simons, C. Vigny, R. J. Labeur, E. Schrama, G. Stelling, and D. Vatvani (2007). "Defining the source region of the Indian Ocean Tsunami from GPS, altimeters, tide gauges and tsunami models". *Earth and Planetary Science Letters* 261 (1-2), pp. 49–64. DOI: 10.1016/j.epsl.2007.06.002.

Ham, D. A., J. Pietrzak, and G. S. Stelling (2005a). "A scalable unstructured grid 3-dimensional finite volume model for the shallow water equations". *Ocean Modelling* 10, pp. 153–169. DOI: 10.1016/j.ocemod.2004.08.004.

Ham, D. A., J. Pietrzak, and G. S. Stelling (2005b). "A streamline tracking algorithm for semi-Lagrangian advection schemes based on the analytic integration of the velocity field". *Journal of Computational and Applied Mathematics* 192, pp. 168–174. DOI: 10.1016/j.cam.2005.04.055.

Journal publications in review

Homolya, M. and D. A. Ham (2015). "A parallel edge orientation algorithm for quadrilateral meshes". *Accepted subject to revisions, SIAM Journal on Scientific Computing*.

McRae, A. T., G.-T. Bercea, L. Mitchell, D. A. Ham, and C. J. Cotter (2015). "Automated generation and symbolic manipulation of tensor product finite elements". *Submitted*. arXiv: 1411.2940.

Rathgeber, F., D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. McRae, G.-T. Bercea, G. R. Markall, and P. H. Kelly (2015). "Firedrake: automating the finite element method by composing abstractions". *Accepted subject to revisions, ACM TOMS*. arXiv: 1501.01809.

Refereed conference and workshop proceedings

Computer science has a tradition of publishing in refereed conference and workshop proceedings. The following are the papers I have authored since beginning to work closely with more traditional computer scientists.

Bertolli, C., A. Betts, G. Mudalige, N. Loriant, D. Ham, M. Giles, and P. Kelly (2013). "Compiler optimizations for industrial unstructured mesh CFD applications on GPUs". *Languages and Compilers for Parallel Computing, 25th International Workshop (LCPC 2012)*. Vol. 7760. Lecture Notes in Computer Science, pp. 112–126. DOI: 10.1007/978-3-642-37658-0_8.

Markall, G. R., F. Rathgeber, L. Mitchell, N. Loriant, C. Bertolli, D. A. Ham, and P. H. Kelly (2013). "Performance-Portable Finite Element Assembly Using PyOP2 and FEniCS". *28th International Supercomputing Conference, ISC, Proceedings*. Ed. by J. M. Kunkel, T. Ludwig, and H. W. Meuer. Vol. 7905. Lecture Notes in Computer Science. Springer, pp. 279–289. DOI: 10.1007/978-3-642-38750-0_21.

Rathgeber, F., G. R. Markall, L. Mitchell, N. Loriant, D. A. Ham, C. Bertolli, and P. H. Kelly (2012). "PyOP2: A High-Level Framework for Performance-Portable Simulations on Unstructured Meshes". *High Performance Computing, Networking Storage and Analysis, SC Companion:* Los Alamitos, CA, USA: IEEE Computer Society, pp. 1116–1123. ISBN: 978-1-4673-3049-7. DOI: 10.1109/SC.Companion.2012.134.

Markall, G. R., D. A. Ham, and P. H. J. Kelly (2010). "Towards generating optimised finite element solvers for GPUs from high-level specifications". *Procedia Computer Science* 1.1. ICCS 2010, pp. 1815–1823. DOI: 10.1016/j.procs.2010.04.203.

Technical report

Ford, R., M. J. Glover, D. A. Ham, C. M. Maynard, S. M. Pickles, and G. D. Riley (2013). *GungHo Phase 1 Computational Science Recommendations*. Tech. rep. Forecasting Research No: 587. UK Met Office. URL: http://www.metoffice.gov.uk/media/pdf/8/o/FRTR587Tagged.pdf.

Details of selected papers

Automated derivation of the adjoint of high-level transient finite element programs

P. E. Farrell, D. A. Ham, et al.
SIAM Journal on Scientific Computing (2013)
DOI: 10.1137/120873558

The ability to numerically solve the adjoint to a given PDE is of great benefit in optimisation, error estimation and data assimilation problems across a wide range of science and engineering fields. However, the existing approaches to developing adjoint models are labour intensive, error-prone, require very advanced mathematical and programming skills, and produce very inefficient adjoint model software. In this work, we employ a high-level mathematical description of a finite element problem to automatically generate a highly efficient adjoint simulation. This is a very high impact development: using this result, adjoint simulations for finite element models move from infeasible in many circumstances to routinely available. This publication was instrumental in Patrick Farrell and me winning five year EPSRC and NERC research fellowships respectively, and led to the award of the Wilkinson prize.

Cross-Loop Optimization of Arithmetic Intensity for Finite Element Local Assembly

F. Luporini et al.
ACM Trans. Archit. Code Optim. (2015)
DOI: 10.1145/2687415
This paper demonstrates the power of employing domain-specific information in a finite element compiler to automatically achieve very high performance on complex modern parallel hardware. Using the characteristics of the loop nests which occur in finite element assembly, we employ aggressive loop refactorisation, data alignment and, where necessary, explicit low level vector operations. The whole process is automated and hidden from the user, facilitating the composition of this complex performance technology with sophisticated numerical techniques.

LBB stability of a mixed Galerkin finite element pair for fluid flow simulations

C. J. Cotter, D. A. Ham, C. C. Pain, and S. Reich
Journal of Computational Physics (2009)
DOI: 10.1016/j.jcp.2008.09.014

This publication demonstrates the importance of designing numerical methods informed by the physics of the system to be simulated. In this case the critical role of the Earth's rotation in the physics of the ocean (and, indeed, the atmosphere) is correctly represented in the model numerics. This paper led on to two others (Cotter, Ham, and Pain 2009; Cotter and Ham 2011) and ultimately to our grants to work on rewriting the Met Office Unified Model.

Research agenda: automated composable abstractions for sophisticated and high performance finite element simulation

Mathematicians have traditionally used computers in two distinct ways. Symbolic computation packages (Mathematica, Maple, and their more specialised cognates) act as aids to the core mathematical task of reasoning about abstract mathematical objects. Conversely, the conventional approach to numerical computation is to write computer code embodying a (possibly complex, parallel) algorithm which was derived with pen and paper, or possibly using an unrelated symbolic computation package.

My research focuses on the creation of systems of mathematical abstractions, embodied in production-quality computer code, which unify symbolic reasoning and numerical computation in a single seamless process. At each level of abstraction, from discretised mathematics down to parallel execution, the problem is expressed symbolically, and the representation of the problem at the next level down the hierarchy is achieved by automatic code generation. Rather than manually deriving and implementing algorithms, the mathematicians and computer scientists working on the system create processes which reason about the higher level mathematical objects at each stage. This approach has a number of benefits enabling the creation of more sophisticated, higher performance simulation systems for ever-more complex problems and exploiting the fine-grained parallelism and hierarchical memory models of current and emerging hardware. Critically, code generation creates a separation of concerns in which experts in each layer of the process are able to apply their skills efficiently and accurately, while benefiting directly from the expertise and efforts of those working at different levels of the abstraction.

The class of problems on which my research is concentrated is the simulation of physical systems by numerically solving partial differential equations (PDEs) using the finite element method. My focus to date has been the simulation of flows in atmospheric and ocean, although the mathematics and software technology I have developed is applicable much more broadly, and I am always interested in exploring new numerical techniques and new application areas in collaboration with other mathematicians or with domain experts.

Automatic code generation for finite element simulations

The current revolution in hardware towards very fine-grained parallelism and deep memory hierarchies presents a huge challenge to the scientific computing community. To make effective use of emerging architectures, the low level implementation of scientific code must change. However regularly rewriting sophisticated numerical software for new platforms using **conventional software development is infeasible**: many scientific codes take millions of pounds and many years to perfect.

By generating rather than hand-writing the finite element code, the application of finite element as an abstraction breaking the link between discretisation and implementation becomes much more complete. By **generating code at run time**, the system not only becomes flexible and easy to use, but the mathematical form of the original problem is able to be exploited during execution to enable advanced reasoning about caching, linearity of operators and so on.

The project builds on the Unified Form Language (UFL), a domain-specific language for the finite element method developed by the Simula Research Laboratory, with whom I have a strong collaboration. A finite element-specific programming language provides an ideal abstraction above which **numerics for computational science and can be developed essentially independent of their implementation** and below which compiler technology, scheduling algorithms and different approaches to concurrency can be applied. The resulting model generation system, **Firedrake**, is now available online (`http://firedrakeproject.org`). Firedrake is now in use as a platform for numerical research and simulation science in three departments at Imperial, and a growing list of other instituttions including Oxford, Bath Leeds, Houston, Portland Oregon, and Aachen.

Communicating mathematics with software as well as papers

Developing mathematical abstractions for simulation software creates a mechanism for bridging the gulf which exists in computational science between advances in numerical mathematics and computer science on the one hand, and simulation science practice on the other. The traditional position has been that research in numerical methods and in parallelisation and optimisation of algorithms is conducted on idealised simplifications of real science problems.

Computer scientists and mathematicians often feel that their job is done when they have developed a new algorithm and published a paper demonstrating proof-of-concept for the idealised problem. The assumption is that scientists will implement these advances for their own complex simulation requirements. In reality, these scientists are usually experts in their own disciplines and must devote their energies to applying models, not developing them. The result is that **the numerical and computational techniques employed in practical simulation science frequently lag decades behind the state of the art**.

One computational field in which this is far less true is that of matrix solvers. Since the mathematical abstractions of sparse matrices and vectors are readily expressed as encapsulated objects in programming languages, libraries such as PETSc and Trilinos offer the latest algorithms in a readily composable form. **My research employs mathematical abstractions of discretisations to bring the same composability to advanced numerical techniques for scientific simulation, and their parallel implementations**. This closes the circle between developers of techniques and computational science practice and enables investment in advanced numerical and computational techniques to be realised in practice.

Reversing causality: a new development method for inverse problems and mathematical optimisation

Inverse problems are pervasive across scientific computing, from data assimilation in weather forecasting to medical tomography, non-invasive biomedical diagonstics, and the design of tidal turbine arrays. However, **simulation software capable of solving the inverse problem is very rare**. This is the result of two mathematical software challenges: the need to differentiate the operators of the primary model and the need to record the actions and data of the primary model so that these can be unwound in the adjoint model, which runs backwards in time.

Under a code generation approach, instead of using conventional automatic differentiation which is complex and fragile, the high level input language is differentiated. This is a **much more simple and robust process**. I secured NERC funding (grant NE/I001360/1) to develop automated adjoint code. With colleagues at Imperial, I established a collaboration with Dr Marie Rognes of the Simula Research Laboratory which has resulted in the dolfin-adjoint (`http://dolfin-adjoint.org`) system. This package executes the tangent linear or adjoint to a finite element simulation **almost automatically and with near optimal efficiency** (Farrell, Ham, et al. 2013).

This is a seminal example of the use of the **right mathematical abstraction** to design software which makes a previously intractable problem tractable. For example, researchers at the Simula Research Laboratory in Oslo are using Dolfin-adjoint to attempt to invert heart attack observations to find the actual heart defect. Dolfin-adjoint has been employed to develop Open Tidal Farm (`http://opentidalfarm.org`). This software package **combines a dolfin-adjoint shallow water model with optimisation routines** to produce a tidal turbine arrangement which maximises power yield. **My co-authors and I were awarded the 2015 Wilkinson Prize for Numerical Software for developing Dolfin-adjoint**. We have recently been awarded an EPSRC grant (grant EP/M011054/1) to expand the capabilities of the tidal turbine system to full three-dimensional simulations.

Gung Ho: Developing the next UK weather and climate model.

Funded by NERC (grants NE/I021098/1and NE/K006789/1) and in collaboration with the UK Met Office, I am developing **automatic code generation mechanisms, new data models, and automated parallelisation** for use in the next generation of UK weather forecasting and climate simulation software.

It is essential that software, which is due to become operational in the 2020s, can make effective use of whatever supercomputers the Met Office has access to then: even though they will be of a hardware class yet to be designed. Breaking the traditional link between the numerical algorithms and their implementation for a particular platform is key achieving this goal.

The success of **this project is critical** to the ability of the Met Office to meet the requirement from governments and other policy-makers for **regional-scale climate predictions**.

The core of current computational science approach adopted by the Gung Ho project, and in particular the data model and parallelisation model, was designed by me and is documented in Ford, Glover, Ham, Maynard, Pickles, and Riley (2013) and Ford, Glover, Ham, Maynard, Pickles, Riley, and Wood (2015).

FInAT - extending the abstraction to high order finite elements

A limitation of the current Firedrake implementation is that the internal mathematical structure of finite element basis functions is not captured symbolically. This prevents the exploitation of the most efficient algorithms for finite element assembly for higher order function spaces. By supplementing the existing basis function tabulator, FIAT, with a new symbolic abstraction for basis functions, it will be possible to employ compiler techniques such as loop invariant code motion in order to generate algorithmically optimal assembly code. This improvement is essential to generating competitive code for the high order spectral element approaches which are well suited to fields such as seismic inversion, and this has resulted in the award of an Intel Parallel Computing Centre grant to target that particular application. Capturing the mathematical structure of the basis functions in this way also opens the possibility of generating algorithms for more exotic basis functions for example with increased continuity between elements, and exploring algorithmic advances such as the use of Bernstein polynomials.

A mathematical query language for massive climate simulation data

In addition to constructing and executing simulations, an increasingly important scientific methodology is to conduct investigations by calculating diagnostic quantities from the stored results of very large simulations. Possibly the best example of this approach is climate science conducted by working with the IPCC Coupled Model Intercomparison Project (CMIP) data. The current approach to this task is for a researcher to download the archived data for many different models, and write a bespoke query program to calculate the particular quantity of interest. This approach is **labour-intensive, highly error-prone and destroys the data provenance audit trail**.

A far better approach is for data scientist to write a diagnostic query in a **high-level mathematical language**. Code would then be generated automatically to apply the diagnostic to the output of a given model. As well as being far easier than current approaches, the query would be sufficiently concise and expressive that it could be published alongside the results. This is a particularly good example of what can be achieved when **numerical mathematics informs software design**. The future potential of this approach is to allow mathematically sophisticated **diagnostics to be calculated** in a rigorous and accurate manner over multiple diverse and **massive data sets remotely**. The effective utilisation of the huge data sets which are being produced at vast public expense depends on the emergence of such technology. I have been working on this project with a PhD student for the last year, extending our existing Firedrake framework to support the required operations and data sources.

Research funding plans

The core mechanism through which I have acquired research funding is by partnering with application scientists who have a simulation requirement which is difficult or impossible using existing tools due to the complexity of the problem or methods involved. My involvement in the tidal turbine optimisation project, and in the Gung Ho project are key, funded, examples of this form of engagement. This form of partnership has proven attractive to national funding bodies since it combines advancing scienctific computing with concrete impact in a cross-disciplinary partnership. I anticipate continuing to develop and build on this form of partnership both within CWI and with new partners in the Netherlands, as well as to exploit my international network to bid for funding at the European level.

Specifically, I am currently involved in writing a Marie Curie Innovative Training Network proposal with partners from Norway, Belgium, Sweden, France and Germany. The core of that proposal is providing excellent PhD-level training for creators of future generations of scientific computing tools. Looking forward, the Horizon 2020 FETHPC-01-2016 Co-design of HPC systems and applications and FETHPC-02-2017 Exascale HPC ecosystem development calls present particularly strong opportunities to secure funding for code generation research, since performance portability and high productivity high performance programming environments are critical to the effective exploitation of exascale systems. The FETHPC-02-2017 call places emphasis on high productivity programming models for extreme scale computing and is therefore a key opportunity.

Teaching experience

I have experience of teaching mathematics, especially numerical mathematics and scientific computing, to a wide variety of student cohorts, from first year geoscience students through to postgraduate mathematicians. For a full list of courses taught, see page 4. I would be able to teach much of the general applied maths programme, although clearly I would see my strengths in teaching numerical analysis and scientific computing modules.

Curriculum development

The rapid changes which have been underway in the teaching of computational mathematics have presented me with opportunities to develop and deliver new modules from scratch.

I developed *Numerical Methods 1*, a second year scientific computing for geophysicists module from scratch. In the absence of a suitable textbook employing the department's chosen language of Python, I developed the entire curriculum, wrote comprehensive lecture notes and set all assessment. I worked closely with the other mathematics and computing lecturers in the department to ensure the module fits seamlessly into the mathematics component of the degree.

Most recently, Dr Colin Cotter and I collaborated to create a masters level module in finite elements. The module combines strong computational and theoretical components, including advanced topics such as understanding and implementing the generation of finite element basis functions from Ciarlet triples.

Innovation in teaching

Teaching the "lab technique" of scientific computing

Whether in further study or in industry, the ability to translate mathematical theory efficiently into effective and correct computer software is a key skill for many maths graduates. Producing reliable results effectively demands more than just strong mathematical reasoning and the ability to program. Effective software creation demands that students understand and deploy software engineering techniques such as revision control, code review and systematic testing. These practices play an essential role in actually achieving results in computational mathematics, in a manner analogous to the critical importance of good lab technique in the physical sciences.

I integrate these practices into my teaching so that the students have the opportunity to practice and internalise them alongside the mathematical content of the modules. Coursework is accessed and submitted via professional online revision control systems, feedback is provided via code review and the students are taught to understand the mistakes in their work by exploiting the comprehensive test suites which I provide.

YouTube lectures

In 2013 it became increasingly clear to me that a number of my PhD students, as well as collaborators at other universities and the Met Office did not have the background in finite element theory which was needed for us to work most effectively. I therefore taught a series of lectures on the finite element method which I recorded and published on YouTube. The resulting videos continue to receive around 2000 views per month from around the world. I would present this as an example of creatively using teaching to both solve advance research and as an outreach opportunity.

Teaching out of the classroom

Enabling students to understand the history and social context of their subject can increase both understanding and motivation.

I have used the flexibility of tutorial teaching to engage student interest by, for example, conducting tutorials on the streets of London and in the Science Museum: the latter on the subject of mechanical scientific computing devices.