

The Hartree Centre: Experience in Addressing Industrial, Societal and Scientific Challenges

Prof Vassil Alexandrov Chief Science Officer



Hartree Centre

Mission: To transform UK industry by accelerating the adoption of high performance computing, data-centric computing and AI.



What we do

- Collaborative R&D
 Address industrial, societal and scientific challenges.
- Platform as a service
 Give your own experts pay-as-you-go access to our compute power
- Creating digital assets

License the new industry-led software applications we create with IBM Research

Training and skills

Drop in on our comprehensive programme of specialist training courses and events or design a bespoke course for your team







Hartree Research Themes



Our track record

Martiree Centr

HARTRE

Case study | Code optimisation for aero engines

Collaborative R&D

Hartree Centre experts optimised codes for modelling component design, speeding up run times by approx 20%

"Working with the Hartree Centre, we have quickly made significant improvements to our code, delivering faster turnaround and more capability to our engineers."

- Matthew Street, Rolls-Royce







Case study | Computer aided formulation

Collaborative R&D



Faster development process for products like shampoo, reducing physical testing "The Hartree Centre's high performance computing capabilities help us achieve better design solutions for our consumers, delivered by more efficient, cost-effective and sustainable processes."

- Paul Howells, Unilever







Case study | Building the cognitive hospital

Transforming the patient experience using cognitive technology and data analytics

"Helping our patients and their families prepare properly for coming into hospital will really reduce their anxiety and could mean they spend more meaningful time with doctors so we are able to get them better faster."

- Iain Hennessey, Alder Hey Children's Hospital

Collaborative R&D











Our Platforms

ABBAN





Intel platforms

Bull Sequana X1000 (840 Skylake + 840 KNL processors) - One of the largest

supercomputers in Europe focusing primarily on industrial-led challenges.

IBM big data analytics cluster | 288TB

IBM data centric platforms

IBM Power8 + NVLink + Tesla P100

IBM Power8 + Nvidia K80

Accelerated & emerging tech

Maxeler FPGA system

ARM 64-bit platform

Clustervision novel cooling demonstrator



Detailed Case Studies

Met Office - PSyclone

<u>Rupert Ford</u>, Andrew Porter, Sergi Siso, Iva Kavcic, Joerg Henrichs



Case Study: *Solution* Met Office

- **Performance:** Needs to get the results in time for forecast, ever-increasing accuracy goals for climate simulations.
- Productivity: hundreds of people contributing with different areas of expertise, 2 million lines of code (UM)
- Portability: Very risky to chose just one platform: may not be future-proofed, hardware changes more often than software, procurement negotiation disadvantage if you can only rur Difficult to compromise on Hartree Cen

one

Science & Technology Facilities Council

Domain Specific Languages:

Embedded Fortran-to-Fortran code generation system used by the UK MetOffice next-generation weather and climate simulation model



Given domain-specific knowledge and information about the Algorithm and Kernels, PSyclone can generate the Parallel System layer.



PSyclone S

PSyclone Motivation

- Maintainable high performance software
 - Productivity: Single source science code
 - Performance and Performance portability
- Complex parallel code + Complex parallel architectures + Complex compilers = Complex optimisation space => no single solution ... single source optimised code is not attainable
- 'Maintenance issue with directives' Thomas Schulthess @ ENES HPC Workshop 2018
- So ... separate science specification/code from code optimisation



Separation of Concerns

Separate the Natural Science from the Computational Science (performance):



PSyclone: Generate the PSy Layer (and modify existing code)

- A domain-specific compiler for embedded DSL(s)
 - Finite Difference/Volume, Finite Element (configurable)
 - Currently Fortran -> Fortran or OpenCL
 - Supports distributed- and shared-memory parallelism
- A tool for use by HPC experts
 - Hard to beat a human (debatable)
 - Work round limitations/bugs
 - Optimisations encoded as a 'recipe' rather than baked into the scientific source code
 - Different recipes for different architectures





Using/evaluating with ...

• Met Office LFRic model

–Uses PSyclone. Successor to Unified Model. Model is in development. Currently performing aquaplanet tests. Plans to go operational ~2025.

NEMO Ocean model

–NEMO is a community ocean model used by the major European Weather and Climate centres. Working with Met Office and NVIDIA to demonstrate OpenACC performance in NEMO with no change to the original code.

• ROMS Ocean model and MOST Tsunami model

–Australian Bureaux of Meteorology evaluating PSyclone for ROMS and MOST. Good OpenMP performance demonstrated for both codes.



PSyclone structure



Possible ECP ST integration

- Share experience of developing and using DSLs for code performance, portability and productivity (PARSEC, Autotuning Compiler Technology etc).
- Use PSyclone to target upcoming Exascale hardware and associated software.
- Compare source code translation and code generation approaches (low level DSL to high level DSL).
- Compare our approach with related performance-portable ECP software (Kokkos, Legion, RAJA, Autotuning, PROTEAS etc.)
- Use and evaluate ECP software to develop and improve PSyclone backends (Kokkos, Legion, RAJA, OMPI-X, SOLLVE etc.)
- Extend PSyclone functionality by integrating support for ECP profiling (HPCTOOLKIT, EXA-PAPI etc) and library (Trilinos, PETSc etc) software



More Details

- •PSyclone 1.7.0
- •BSD 3-clause
- https://github.com/stfc/PSyclone
- https://psyclone.readthedocs.io
- sudo pip install psyclone









Intel[®] Parallel Computing Center



Chemistry and Materials

Rick Anderson and team



Chemistry and Materials Team @ Hartree

- Skills covering application of simulation techniques, software development and machine learning
- Remit is helping to solve industry challenges by the development and application of software tools
- Deployment of highly optimized and automated methods for industry uptake



Scientific application areas:

- Phase behaviour of complex mixtures: colloids, polymers, and surfactants;
- Hydrodynamics, diffusiophoresis, and other microscale transport processes;
- Noise and stochasticity in biological systems: subcellular processes; microbial ecology (neutral models); epidermal homeostasis; etc;
- Polymer simulation
- Exploration of behavior of granular media
- Parameterization of models for the study of alloys (embedded atom methods)



Dr Rick Anderson – STFC Hartree Centre

Challenge: What if I could try a new formulation on my phone in the same time it takes to get a coffee?



The Formulation Industry

- Formulation industry is worth £1,000 billion per annum and growing
- Significant value-adding step, with value multipliers of up to 100 times
- Formulated products are everywhere; laundry liquids, toothpaste, shampoo, vehicle oils, foods
- Most contain 20+ individual chemicals species
- Drive for sustainable products
- Digitization to reach market earlier





Scientific validation: Development of new techniques



Science & Technology Facilities Council

Insight at Scale: Virtual Experiments





Image courtesy of University of Sofia

Developing virtual analogues to those experiments performed in the laboratory

In the images we see worm like micellar structures from Cryo-TEM (Right) and are able to simulate these (left)

These micelles are responsible for building viscosity in a number of formulated products.



Science & Technology Facilities Council



Automation and scale: Fast and efficient



Developed 100+ digital assets to enable automation of chemistry and material simulation and data analytics

Intimately coupled to the Hartree Software Stack



Optimized existing codes to extract exceptional performance on latest generation of hardware



Deployment: Challenge Met



Mapping the behavior of a formulation takes 3 days in the laboratory (Source: FMCG Company) 85-90% reduction in effort to prepare a formulation and see results FORMERIC®







Porting DL_MESO (DPD) on Nvidia GPUs

Jony Castagna



What is DL_MESO (DPD)

 DL_MESO is a general purpose mesoscale simulation package developed by Michael Seaton for CCP5 and UKCOMES under a grant provided by EPSRC.

It is written in Fortran90 and C++ and supports both Lattice Boltzmann
 Equation (LBE) and Dissipative Particle Dynamics (DPD) methods.

https://www.scd.stfc.ac.uk/Pages/DL_MESO.aspx



...similar to MD

- Free spherical particles which interact over a range that is of the same order as their diameters.

- The particles can be thought of as assemblies or aggregates of molecules, such as solvent molecules or polymers, or more simply as carriers of momentum.



Fi is the sum of conservative, drag and random (or stochastic) pair forces:

$$\vec{F}_i = \sum_{j \neq i}^N \left(\vec{F}_{ij}^C + \vec{F}_{ij}^D + \vec{F}_{ij}^R \right)$$



Examples of DL_MESO_DPD applications

Vesicle Formation



Phase separation







Lipid Bilayer

Polyelectrolyte

DL_MESO: highly scalable mesoscale simulations Molecular Simulation **39** (10) pp. 796-821, 2013









	0 ⁰ 0	0 0	0	0 0	0	
0	00	e 9 3	(B) (B) (B) (B) (B) (B) (B) (B) (B) (B)	(B)	0	0
0	00	a 9 8	0 0 0	() () () ()	0 0	0
	0	ر چ څ	(3) (3) (3) (3) (3) (3) (3) (3) (3) (3)	0 0	0	0 0
0		0	0	0	0	
	(ΓPI	12			•

		GP	U 1	_		
	0 ⁰ 0	0 0	0 0	0	0	
0	00	() () () ()	(B) (B) (B) (B) (B) (B) (B) (B) (B) (B)	ର ଓ ଓ	0	0
0	00	a a	0 0 0	© ()	0	0
	0	ی چ ق	6) (1) (1)	0 0	0	000
0		0	0	0	0	







Use of communication between GPUs:

1) to exchange particles positions using ghost cells (ideally while computing the internal cells!)



3) gathering statistics

4) transfer info for FFTW (not implemented yet!)



exchange particles positions using ghost cells step 1: exchange x-y planes

(2 communications)



Х



exchange particles positions using ghost cells



exchange particles positions using ghost cells



step 3:

exchange y-z planes (with halo data!)

need only 6 communications instead



Weak scaling

1.2 billion particles for a mixture phase separation



Strong scaling

1.8 billion particles for a mixture phase separation



clustering!

Energy and Environment

Xiaohu Guo



Software Development for Energy and Environment

IMPORTANCE: Hartree Centre key technologies, align with STFC global challenge schemes.



Scalable Algorithms Development

Large scale application software development, advanced computational methods development.



Advanced Monte Carlo Methods for Linear Algebra

Vassil Alexandrov



Advanced Monte Carlo Methods for Linear Algebra on Advanced Accelerator Architectures

Anton Lebedev (Institute for Theoretical Physics, University of Tuebingen, Germany), anton.lebedev@uni-tuebingen.de Vassil Alexandrov (Hartree Centre STFC, UK and ICREA, Spain) vassil.alexandrov@stfc.ac.uk

Hybrid Monte Carlo vs Hybrid MSPAI

Speed-ups V100 vs Intel Xeon 8160

Speed-ups Intel Xeon vs V100 and K80















EBERHARD KARLS

ICREA





OpenMP vs GPU implementation





OpenMP vs GPU implementation



OpenMP vs GPU implementation



Possible avenues of collaboration

Overall possible contributions to the collaboration:

- particle methods, hybrid stochastic/deterministic methods for Linear Algebra and multilevel multiscale methods, runtime systems for functional parallelism, domain specific compilers with several other themes in support.
- Apply these methods and algorithms and software technologies into the fields of weather and climate modelling, chemistry and materials, Industry 4.0, cosmology and astrophysics, visual computing.
- Demonstrate clear vertical integration and co-design approach.

Leading to:

- 1. Novel and advanced particle simulations at scale.
- 2. Novel and Advanced scalable algorithms tackling large scale problems, problems with number of degrees of freedom of the order $O(10^{10}) O(10^{14})$ and beyond.
- 3. Contributions to the ECP software stack developments including libraries developments.



Thank you

Vassil Alexandrov vassil.alexandrov@stfc.ac.uk



Find out more:

- @ hartree@stfc.ac.uk
- hartree.stfc.ac.uk
- Im /company/stfc-hartree-centre
- 🥑 @hartreecentre