

Driving discovery and innovation at Supercomputing Wales

ABSTRACT

Researchers across Wales capitalize on the high performance computing resources of Supercomputing Wales to drive discovery and innovation. Some of this research is highlighted in this booklet, which presents just a small sample of the scientific investigations facilitated by Supercomputing Wales.

November 2021

TABLE OF CONTENTS

RESEARCH AT SUPERCOMPUTING WALES	1
RESEARCH BRIEFS	2
A next-generation conjugate gradient benchmark from computational particle physics	2
Biophysical modelling for sustainable expansion of the shellfish industry in Wales, UK	3
Computational design of photocatalyst for wastewater purification	4
Computational Models of Group Behaviour	5
Covate: Forecasting the number of cases for lineages of SARS-CoV-2	6
High Performance Solution of the Boltzmann Equation.	7
Impact of Sea Level Rise on UK Coasts	8
Living Wales: Near realtime monitoring of land use in Wales	9
Modelling the spread of COVID-19 in Wales	10
Optimisation of realistic simulation models in synchrotron microtomography.	11
Performance of infinitely long tidal stream turbine arrays.	12
Quantum Criticality in the 2+1d Thirring Model.	13
Recognising specific foods in MRI scans using deep learning	14
Vectorising OpenQCD	15
ABOUT SUPERCOMPUTING WALES	16
A center of excellence	16
TO LEARN MORE	16

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying and distribution of any software described in this publication require an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, Dell Technologies, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.

Dell Technologies believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Published in the USA 11/21.

RESEARCH AT SUPERCOMPUTING WALES

Supercomputers make it possible for researchers to address questions that couldn't be easily answered by other means, including investigations conducted in laboratories and on desktop computing systems.

This is the case at Supercomputing Wales, where researchers use high performance computing clusters to investigate the types of problems that require massive amounts of compute power. From modelling the power of the sea as a low-carbon energy source to improving the speed and accuracy of tuberculosis testing, researchers draw on the resources of Supercomputing Wales to drive discovery, innovation and scientific breakthroughs.

Some of this research is highlighted in this booklet, which presents just a small sample of the wide-ranging scientific investigations facilitated by Supercomputing Wales.

A next-generation conjugate gradient benchmark from computational particle physics

SOMBRERO

A next-generation conjugate gradient benchmark from computational particle physics



Ed Bennett, Biagio Lucini, Michele Mesiti (Swansea University), Jarno Rantaharju (Helsinki University)

Summary of the problem

- Lattice Field Theory calculations consume vast amounts of HPC resource
- How can we know that the machines being procured give best value for money?
- Develop a benchmark showing timings for the performance-critical operations for representative problems of interest

Beyond $N=3$




Figure 1: Lattice field theories are most frequently constructed on a 4-dimensional grid (or "lattice").

- Lattice QCD studies the quarks and gluons that make up atomic nuclei
- Lattice QCD is a special case of lattice field theory, with a specific choice of the number of quark colors
- As the number of colors changes, so do the computational demands, so benchmarking QCD isn't enough

Introducing SOMBRERO [1]

- Developed from HiRep [2] research code
- Tests CPU performance for a variety of beyond the standard model theories
- Times CG inversion of a large $((10^7)^2)$ sparse matrix

Sample Results

- SOMBRERO was used for technical commissioning of the new DiRAC 3 facilities for the UK physics community




Figure 2: SOMBRERO strong scaling results for the $Sp(4)$ gauge theory with fundamental matter.




Figure 3: Timings for various compiler choices on DiRAC facilities, on 64, 1, and 64 nodes respectively.

- Identified optimal compiler choices for each machine.
- Identified problems with compiler installation on COSMA, allowing rectification before entering service
- Identified on CSD3 (with Ice Lake CPUs) that using 72 CPU cores per node gave relatively small speedup over 48

Current and Future Work

- Investigating porting to GPUs and hybrid multithread-MPI parallelism using SYCL

[1] <https://github.com/sa2c/sombrero>
 [2] <https://github.com/clauidopica/HiRep>

Supercomputing Wales
November 2021



Biophysical modelling for sustainable expansion of the shellfish industry in Wales, UK



Sophie Ward, Peter Robins, Lewis Le Vay (Bangor University)

The Shellfish Centre

The Shellfish Centre is a research & innovation initiative supporting sustainable development of the shellfish sector in Wales.

Using biophysical models, it is possible to simulate larval transport and dispersal of both native and non-native shellfish species around the Welsh coast.




Figure 1: Hydrodynamic model domain for the Irish Sea (high-resolution grid in South Wales).

Biophysical Modelling




Figure 2: Example trajectory (orange line) of a particle (larvae), advected for five days and released from the North Wales coast (blue dot = release location, green = land).

- High resolution ocean models are used to simulate the hydrodynamics of the Irish Sea (tides, winds, waves).
- Outputs from the hydrodynamic models are used to run an “offline” particle tracking (biophysical) models which simulate larval transport and dispersal (for example, from known shellfish beds).

Biophysical modelling projects include:

- Hydrodynamic and biophysical modelling to inform native oyster restoration.
- Validating particle tracking models using bivalve larvae distribution data.
- Optimisation of location of native oyster broodstock and settlement ground for a source-sink enhanced fishery.
- Predictive modelling of possible Pacific Oyster habituation.




Figure 3: Unstructured hydrodynamic model (grid cells illustrated by grey triangles) and residual current flows around Holyhead, North Wales (coloured arrows).

Current and Future Work

These biophysical modelling tools are being applied to a number of other research projects, including:

- Angel Shark Project Wales, in collaboration with ZSL Marine and Natural Resources Wales.
- Understanding dispersal of material (e.g. larvae) in the nearshore region (Ecostructure project).
- Man-made structures as marine stepping stones of non-native species (Ecostructure project).

Supercomputing Wales



Computational design of photocatalyst for wastewater purification

Sachin Nanavati, David Willock (Cardiff University), S. Mandal and A. Rajakumar (Indian Institute of Technology, Kharagpur)

Summary of the problem

- Aim: To gain insights in the design and functioning of a Silver based Coordination Polymer (CP) photocatalyst which will harness sunlight for water splitting.
- Methodology: Large scale simulations on HPC systems to accurately calculate the electronic structure of CP based on density functional theory (DFT).
- Green Chemistry application: usage of CP photocatalyst for degradation of organic pollutants from industrial wastewater.
- Project in tune with UN Sustainable Development Goal (SDG 6.3) *better water quality*




Figure 1: Unit cell of Ag_2bdc CP [1]. Functionalising with an amine group ($-NH_2$) reduces the bandgap of CP from 3.8 to 2.9 eV, thus making it a visible light driven photocatalyst [2].

Experimental partner:

Indian Institute of Technology, Kharagpur

Results

- Reduction in bandgap (E_g) of CP from 3.8 eV (UV range) to 2.9 eV (visible range).




Figure 2: Density of States (DOS) of $Ag(I)NH_2$ bdc CP (top) and pristine $Ag(I)$ bdc CP (down) [2]. Reduction in bandgap is due to amine ($-NH_2$) functionalisation.




Figure 3: Energy band structure and charge density plot of (left) pristine and (right) amine functionalized CP. Valence Band Maxima (VBM) and Conduction Band Minima (CBM) are shown. $Ag(I)NH_2$ CP can absorb visible light.

- UV light constitutes only 4% of solar spectrum while, visible light make 44% of solar spectrum. Increase in CP light absorption efficiency in visible range.

Reference : [1] S. Mandal *et al.*, J. Phys. Chem. C **123**, 23940 (2019), [2] S. Mandal *et al.*, Applied Catalysis B: Environmental (under revision).



Computational Models of Group Behaviour



R.M Whitaker, W. Colombo, (Cardiff University), Y. Dunham, D.G. Rand (Yale University).

- Individual human behaviour can be understood studying the behaviour of groups.
- The *social brain hypothesis* explains human evolution as a consequence if the human ability to organize and interact within groups.
- A simple *reputation based agent based model* where agents are asked to commit to simple *donation acts* is effective to psychological concepts associated to intra- and inter-group behaviour.
- *Social comparison* plays a crucial role in explaining the behaviour of individuals between each other and the group they represent.

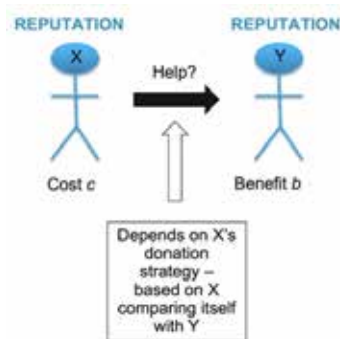


Figure 1: The donation game

- Collaboration with social anthropologists and psychologist from Oxford and Yale universities.
- Role of supercomputing resources is crucial to allow a large exploration of the parameter space to conduct 1000s of diverse simulations.

Identity fusion theory

- The *identity fusion theory* [1] states that group identity does not replace that of individuals. Instead the two are merged into a fused identity the extend of which is represented of the degree of identification with the group.

- Psychological concepts as *prejudice, stereotyping and extremism* can be all explained within the concept of group identity.




Figure 2: Identities of individuals cannot be considered as fully independent from the groups they belong to.

Evolution of prejudicial groups

- *Out-group reputation* is reduced by a prejudice variable that can evolve independently from cooperation.
- Agents can copy the cooperation strategies and prejudice levels of others.
- *Learning (P)* and/or *interacting* with a wider population (*S*) helps to diminish prejudice.
- Natural phenomenon easily triggered but hard to reverse.




Figure 3: Correlation between cooperation and prejudice under different rates of mixing (*S*) and learning (*P*)



SUPERCOMPUTING WALES
LWCHGYFRIFIADURA CYMRU

Supercomputing Wales



Covate: Forecasting the number of cases for lineages of SARS-CoV-2



Anna Price, Alex Southgate, Thomas Connor (Cardiff University), COG-UK Consortium.

Introduction

- As of the end of September 2021, COG-UK [1] has sequenced more than 900,000 SARS-CoV-2 genomes, providing valuable information on the lineages of SARS-CoV-2 circulating in the UK.
- Covate [2] uses the COG-UK metadata to analyse the relationship between lineages observed in both Wales and England, forecasting the time series of sequenced cases for each lineage.

Workflow

Figure 1 shows the workflow. Covate automates the selection of either a Vector AutoRegression (VAR) or Vector Error Correction Model (VECM) for each lineage.

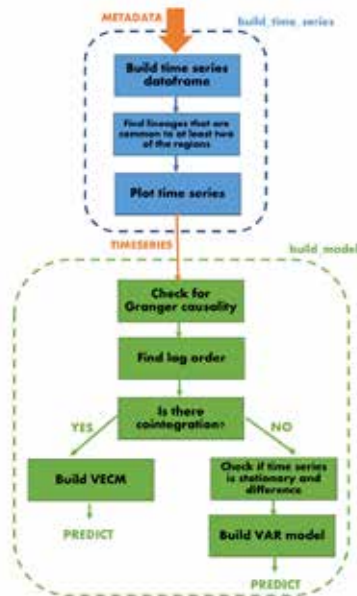


Figure 1: Workflow for Covate

Sample Results

Figure 2 shows a forecast for the Delta variant against the actual number of sequenced cases for Wales and England.

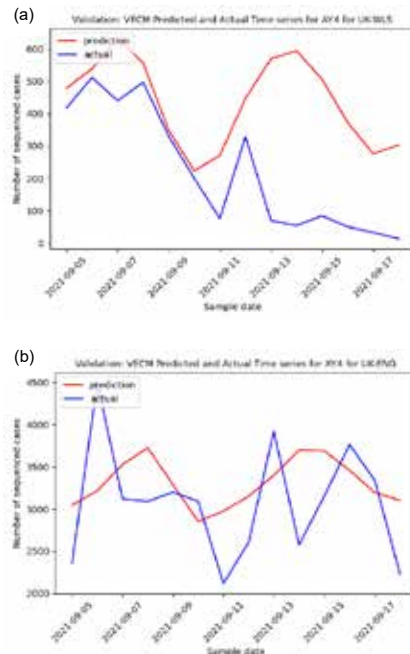


Figure 2: Predicted and actual number of sequenced cases for the Delta sub-lineage AY.4 for (a) Wales and (b) England for the period 5th Sept 21 – 18th Sep 21

Discussion

- There is a strong causal relationship between cases of SARS-CoV-2 in England and Wales.
- Covate can be used to actively monitor the circulation of lineages in Wales.

[1] COVID-19 Genomics UK (COG-UK) (2020) 'An integrated national scale SARS-CoV-2 genomic surveillance network', *The Lancet. Microbe*, 1(3), pp. e99–e100. doi: 10.1016/S2666-5247(20)30054-9

[2] <https://github.com/Pathogen-Genomics-Cymru/covate>

Supercomputing Wales
November 2021



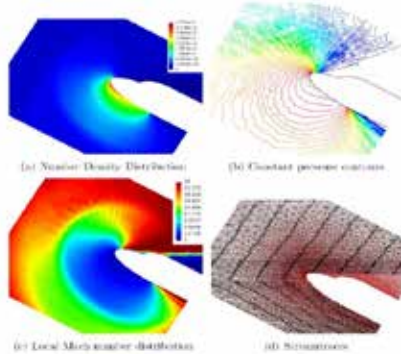
High Performance Solution of the Boltzmann Equation



B. Evans, M. Hanna, M. Dawson, M. Mesiti (Swansea University)

Simulating Hypersonic fluid dynamics

- Numerical simulations of **fluid dynamics at hypersonic speeds** are an extremely hard problem, where the fluid is **not in thermodynamic equilibrium**. This requires working with the **Boltzmann Equation** directly instead of the **Navier-Stokes** equation.




- The Boltzmann Equation describes the transport of the molecular velocity distribution function, f , that describes the **probability** of finding a molecule for a given **point in physical space** for a given **velocity state** at a given time. From f , **pressure** and **temperature** can be determined.

The **BGK approximation** to the Boltzmann equation reads

$$\frac{\partial(nf)}{\partial t} + \mathbf{c} \cdot \frac{\partial(nf)}{\partial \mathbf{r}} = \nu(\mathbf{r}, t)((nf_0) - (nf))$$

where f_0 is the *equilibrium* distribution function.


Numerical Method




The physical space is discretized with a Discontinuous Galerkin (Finite Element) method, and the velocity space with a Discrete Velocity Model

approach.

The whole domain (the *phase space*) is simulated on HPC using **Domain Decomposition** both in the physical space and in the velocity space.



The phase space parallelisation implemented is **significantly more effective than previous implementations**, allowing the efficient use of a larger number of processors [1].



[1] B. Evans, M. Hanna, M. Dawson & M. Mesiti - "High order parallelisation of an unstructured grid, discontinuous-Galerkin finite element solver for the Boltzmann-BGK equation" International Journal of Computational Fluid Dynamics, Volume 33, 2019 - Issue 8

Supercomputing Wales
November 2021



Living Wales: Near realtime monitoring of land use in Wales



Colin Sauze, Richard Lucas, Carole Planque, Sebastien Chognard, Suvarna Punalekar
(Aberystwyth University)

Summary of the problem

- Producing Land Cover and change of use maps from satellite images to:
 - Estimate biomass quantities and measure biodiversity
 - Assess wildfire and flooding damage and recovery
 - Monitor the spread of plant diseases such as ash dieback.
- Using images from the European Space Agency Sentinel and NASA Landsat Satellites going back to the 1980s
- These are very large datasets (75TB), which need a high performance computing system to process them.

How Supercomputing Wales Helps

- RSE Expertise
 - Producing better software, training, technical advice.
- High Performance Computing facilities for data processing
- Hosting additional computing facilities for running web services and data storage.

Applications

- Ongoing project with Welsh Government to provide land use change maps.
- Maps access for all via the web portal
- Useful for other researchers, farmers and local government

Results

- Web portal with data from 2017 to 2020 available online.



Fig 1: Land cover classification map for Wales

Current and Future Work

- Processing 1985 to 2017 data
- Data Cube for Wales
- Continuous near real-time updates to maps from the latest satellite images
- Integrating ground based observations

References

[1] Living Earth: Implementing national standardised land cover classification systems for Earth Observation in support of sustainable development. Owers, C.J., Lucas, R.M., Clewley, D., Planque, C., Punalekar, S., Tissot, B., Chua, S.M.T., Bunting, P., Mueller, N., Metternicht, G., 2021. Big Earth Data 5, 368–390. <https://doi.org/10.1080/20964471.2021.1948179>

[2] National scale mapping of larch plantations for Wales using the Sentinel-2 data archive. Punalekar, S.M., Planque, C., Lucas, R.M., Evans, D., Correia, V., Owers, C.J., Postajko, P., Bunting, P., Chognard, S., 2021. Forest Ecology and Management 501, 119679. <https://doi.org/10.1016/j.foreco.2021.119679>

[3] Living Wales: automatic and routine environmental monitoring using multi-source Earth observation data. Planque, C., Punalekar, S., Lucas, R., Chognard, S., Owers, C.J., Clewley, D., Bunting, P., Sykes, H., Horton, C., 2020. Earth Resources and Environmental Remote Sensing/GIS Applications XI. Presented at the Earth Resources and Environmental Remote Sensing/GIS Applications XI, SPIE, pp. 41–49. <https://doi.org/10.1117/12.2573763>



Supercomputing Wales
November 2021



Modelling the spread of COVID-19 in Wales



Dr Ed Bennett, Dr Mark Dawson, Prof. Mike Gravenor, Prof. Biagio Lucini, Dr Benjamin Thorpe (Swansea University)

Context

- COVID-19 began spreading in the Wales in March 2020
- Welsh Government approached Swansea for modelling of proposed policy interventions in June 2020
- Rapidly assembled team of domain experts and Supercomputing Wales RSEs

Model

- Adapt the model developed by the London School of Hygiene and Tropical Medicine [1]
- Compartmentalise Wales by local authority and by age group
- Model flow from susceptible to infected and onwards based on local prevalence and contact rates




Figure 1: The two possible flows from susceptible to removed.

Feature additions


Implemented additional functionality to respond to Welsh Government requests for specific scenarios

- Vaccination
- Economic indicators
- Hospital and intensive care occupancy
- Differing vaccination effect on transmission vs. on severity of disease

Outcomes

- Fortnightly discussions with Welsh Government
- Regular requests for updated models for changing situation and new policies

Figure 2: An example output of the model, showing projections for hospital admissions



based on a number of assumptions.

- Modelling enabled:
 - Timely implementation of Welsh firebreak
 - Preparation of correct numbers of ICU beds, ambulance service staff, contact tracers during second wave
 - Planning of school reopening in February–March 2021

Current and Future Work

- Better modelling of travel within and outside Wales
- Better accounting for network effects and clustering
- Allow for multiple vaccine brands with different efficacy
- Track multiple viruses simultaneously

[1] <https://github.com/cmimid/covid-uk>

Supercomputing Wales
November 2021



Optimisation of realistic simulation models in synchrotron microtomography



Franck P. Vidal, Iwan T. Mitchell (Bangor University), Jean Michel Létang (INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220).

Summary of the problem

- ISO 10360-11 standard: defines characteristics and methods for measuring object surfaces in X-ray CT.
- Strong imaging artefacts in microtomographic X-ray data: makes the CAD modelling process difficult (Figure 1).
- Solution: Register the CAD models by deploying a realistic X-ray simulation on GPU in an optimisation framework (Figure 2) [1].




Figure 1: CT image containing strong artefacts.




Figure 2: Image registration as an optimisation algorithm.



[1] F.P. Vidal, I.T. Mitchell and J.M. Létang (2021) Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts. *Precision Engineering*. Submitted.

Results

- A registration corresponds to optimising 23 unknowns values.
- On average each registration takes 1 hour and 35 minutes.
- Correlation between real CT image and simulated CT images is >94% (Figure 3).
- Figure 4 shows an example of CAD model.




Figure 3: Checkerboard comparison between the real CT slice and simulated ones. A checkerboard image alternates tiles from the real and the simulated CT slices. Red squares are two examples of tiles from the real CT slices; blue squares from the simulated slices.




Figure 4: Corresponding CAD models.

Future Work

- Model based CT reconstruction to correct imaging artefacts and improve the value of linear attenuation coefficients.

Supercomputing Wales



Performance of infinitely long tidal stream turbine arrays



Pablo Ouro, Takafumi Nishino (University of Oxford),

Tim Stallard (University of Manchester), Peter Stansby (University of Manchester).

- In 2050 about 80% of our energy demand needs to be covered from renewable resources. However, wind and solar energy generation relies on the daily weather conditions.
- Tides are predictable and so is the energy to be harnessed from fast-flowing currents.
- Tidal turbines operate in harsh environments with high turbulence levels challenging their survivability.
- High-resolution simulations allow to study how these turbines operate under highly turbulent conditions allowing to best design tidal turbine arrays and maximise power extraction.

DOFAS: state-of-the-art solver

The in-house code DOFAS (Digital Offshore Farms Simulator [2]) is fully parallelised with MPI and OpenMP that permits to run high-fidelity simulations in hundreds to thousands of CPUs capturing the complex phenomena present in turbulent flows [3]. Tidal turbines are represented with an actuator line method and irregular bathymetry can be represented with an immersed boundary method

Infinitely long tidal array

Our study focused on studying the impact of having an aligned or staggered layouts with different streamwise and spanwise spacing between turbines [1]. Simulations involved a single turbine up to 48 devices and required 864 CPUs to run over a period of 4 days. A total of simulate 1,000 turbine revolutions were computed to obtained converged flow statistics and power performance estimates. The main findings are:

- Aligned layouts increased the power performance for larger streamwise spacing.
- Turbines in staggered layouts achieved larger power coefficients when decreasing the streamwise spacing as a result of the bypass flow impinging the turbines.
- Theoretical results from a linear momentum actuator disc theory agree qualitatively well with DOFAS'.
- LES is becoming a reliable tool for micro-sitting capable of capturing governing phenomena such as wake meandering.
- Power spectral density analysis revealed that the energy decay varies in staggered layouts depending on the spanwise spacing, likely correlated to the constrain of the wake meandering motion.

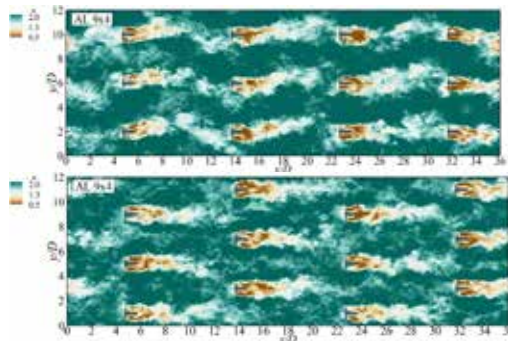


Figure 1: Flow field generated in arrays of turbines with aligned (top) and staggered (bottom) layouts.

[1] Ouro, P. and Nishino, T. (2019) Performance and wake characteristics of tidal turbines in an infinitely large array. *Journal of Fluid Mechanics*. 925: A30.

[2] Ouro, P., Ramirez, L. and Harrold, M. (2019) Analysis of array spacing on tidal stream turbine farm performance using Large-Eddy Simulation. *Journal of Fluids and Structures*. 91: 102732.

[3] Ouro, P., Lopez-Novoa, U. and Guest, M. (2021) On the performance of a highly-scalable Computational Fluid Dynamics code on AMD, ARM and Intel processor-based HPC systems. *Computer Physics Communications*. 269: 108105.

Supercomputing Wales
2021


Quantum Criticality in the 2+1d Thirring Model



Simon Hands (Liverpool University), Michele Mesiti, Jude Worthy (Swansea University)

Phase transitions in 2D relativistic fermions

- Two-dimensional relativistic fermions occur frequently in effective models of layered condensed matter systems (e.g., electrons in **graphene** on a honeycomb lattice are described at low energy by a relativistic Hamiltonian – the **Dirac** Hamiltonian)



- The **Thirring Model** is a Quantum Field Theory which captures the essential features of the Coulomb interaction between electrons. The index l runs over N fermion species.

$$\mathcal{L} = \bar{\psi}_l(\not{\partial} + m)\psi_l + \frac{g^2}{2N}(\bar{\psi}_l\gamma_\mu\psi_l)^2$$

For $N < N_c$ and sufficiently large g^2 , the model has a **phase transition** between gapless **semimetal** and gapless **insulator**. This phase transition correspond to the following symmetry breaking:

$$U(2N) \rightarrow U(N) \otimes U(N)$$

Challenges in the numerical simulations

Numerical studies are made using **Markov Chain Monte Carlo** techniques and a **Lattice Field Theory** formulation.

To improve over earlier attempts [3], we need a numerical formulation of the theory that breaks the **U(2N)** symmetry as little as possible (compatibly with computational costs and implementation challenges).


For this, we use the **Domain Wall Fermion** formulation, which recovers $U(2N)$ as the distance L_s separating “domain walls” in a fictitious third direction x_3 is made large (and for vanishing m). Other approaches are possible [1]. The **main computational challenge** is due to the need to **reduce m** and **drastically increase L_s** .

Results

The value of the condensate $\langle \bar{\psi}\psi \rangle$ is measured, and the equation of state

$$m = A(g^{-2} - g_c^{-2})(\bar{\psi}\psi)^{\delta-1/\beta} + B(\bar{\psi}\psi)^\delta$$

is fit to the data. The success of the fit for $N=1$ data is a strong evidence for spontaneous **U(2)** symmetry breaking [2].



Equation of state fit for $N = 1$ on 16^3 in the $L_s \rightarrow \infty$ limit.

This, together with evidence that there is no such breaking for $N=2$ [3], suggests that $1 < N_c < 2$.

This is in tension with results from [1], suggesting $N_c \sim 0.8$.

Current and Future Work

- Conclude numerical study for $L_s=80$
- Study the 2-point correlation functions

[1] J. Lenz, B.H. Wellegehausen and A. Wipf, Phys. Rev. D 100 (2019) no.5, 054501.

[2] S. Hands, M. Mesiti and J. Worthy, Phys. Rev. D 102 (2020) no.9, 094502.

[3] S. Hands, Phys. Rev. D 99 (2019) no.3, 034504.

Supercomputing Wales
November 2021



Recognising specific foods in MRI scans using deep learning



J. Gardner (Bangor University), S. Al-Maliki (Bangor University, Basrah University, Iraq), E. Lutton (INRA-AgroParisTech, France), F. Boué and F.P. Vidal (Bangor University).

Summary of the problem

- Part of an experimental project aiming at understanding the kinetics of human gastric emptying.
- MRI images of the stomach of healthy volunteers have been acquired using a state-of-art scanner with an adapted protocol (Figure 1).
- The challenge is to follow the stomach content (food) in the data.
- Frozen garden peas and petits pois have been chosen as experimental proof-of-concept as their shapes are well defined and are not altered in the early stages of digestion (Figure 2).
- The food recognition is performed as a binary classification implemented using a deep convolutional neural network (CNN) (Figure 3).
- Exhaustive search of optimal combination of hyperparameters using Supercomputing Wales.




Figure 1: Typical example of MRI slice of the stomach.

[1] Gardner, J. Al-Maliki, S. Lutton, E. Boue, F. & Vidal, E. 2020, *Recognising specific foods in MRI scans using CNN and visualisation*, in *Proceedings of Computer Graphics & Visual Computing 2020 (CGVC 2020)*, The Eurographics Association. <https://doi.org/10.2312/cgvc.20201145>




Figure 2: Data preparation.




Figure 3: Supervised Learning.

Results

- Automatically label the content of the stomach even when the amount of training data is low and the data imbalanced (Figure 4).
- Food recognition in MRI scans produced an accuracy of 97%, precision of 91%, both close to 1 (Table 1)




Figure 4: Examples of labeled stomachs. Circles in magenta depict True Positives (TPs), crosses in magenta True negatives (TNs), blue circles depict false positives (FPs), and blue crosses false negatives (FNs).

Table 1: Classification results over 10 MRI datasets.

TP	TN	FP	FN	Accuracy	Precision	Recall	F1-score
346	2217	33	55	0.97	0.91	0.86	0.89

Supercomputing Wales



Vectorising OpenQCD



Ed Bennett, Mark Dawson, Michele Mesiti, Jarno Rantaharju (Swansea University).

Introduction


The **OpenQCD** code is used to simulate the strong nuclear force, which binds quarks into nuclei and other particles. It is used to make accurate predictions for particle accelerators and to calculate properties of neutron stars. The computationally intensive core of the program is the inversion of a large sparse matrix, the Dirac operator.

We present an extension of the **OpenQCD** to the current generation of Intel HPC processors, Skylake and Knights Landing. The processors support a set of 512-bit vector instructions (**AVX 512**). To take full advantage of these capabilities we implement the most computationally intensive part of the code using Intel's intrinsic instructions.

The algorithm is **memory bandwidth bound** and we therefore optimise for memory management at the expense of computation.

What is SIMD:

Modern CPUs operate in vector units executing a Single Instruction for Multiple Data (SIMD). This allows more operations per cycle (or transistor) with fewer memory operations.



But:

- Requires data in continuous blocks
- Somewhat limited instruction sets

AVX2:

- 256-bit wide vectors, 4 doubles or 8 floats
- 16 floating point registers

AVX 512:

- 512-bit wide vectors, 8 doubles or 16 floats
- 32 floating point registers

Intrinsics

Intrinsic instructions offer a straight-forward interface to control vectorisation and memory management.

Intrinsics are

- Native (C-)functions
- Usually replaced by a single instruction
- Generalisable:
 - Can be compiled for different architectures
 - Compiler handles register allocation

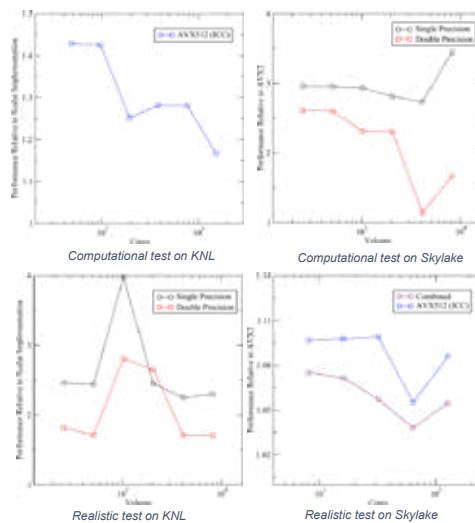
Test Cases

1. Computational load test:

Run the Dirac operator with random data on a single core, removing data dependency and MPI communication.

2. Realistic / Memory Bandwidth test:

Run the full conjugate gradient algorithm on multiple cores. Each iteration depends on the previous result. In addition to the Dirac operator each iteration includes vector operations and reductions.



Supercomputing Wales
October 2018

ABOUT SUPERCOMPUTING WALES

Supercomputing Wales is a strategic program of investment in the university sector in Wales led by Cardiff University, in a consortium with Aberystwyth, Bangor and Swansea universities. Part-funded by the European Regional Development Fund through Welsh Government, this program provides researchers across Wales with access to powerful computing facilities for science and innovation projects.

Academic areas supported by the Supercomputing Wales facilities include science, technology, engineering and mathematics, including computer science, physics, chemistry and biology, as well as environmental, e-science, health, medical and the social sciences.

The Supercomputing Wales Research Software Engineers (RSEs) work with established research teams across a range of disciplines within the consortium universities to help deliver software engineering solutions and algorithms that harness multi-core computing for data-driven simulation, translating specific research needs into high performance computing applications, either by porting existing code or designing and developing new software from scratch, as appropriate for each specific problem.

A center of excellence

At the heart of the Supercomputing Wales initiative is a center of excellence (CoE) spearheaded by two global leaders in digital transformation, Atos and Dell Technologies. This CoE provides Welsh researchers with leading-edge HPC equipment, software and services, including a supercomputer, two Atos BullSequana S Datalake appliances, and on-demand cloud-based supercomputing simulation.

With the HPC systems and expertise of Supercomputing Wales and the Atos Dell Technologies Supercomputing Centre of Excellence, researchers have access to the resources they need to simulate and solve complex scientific problems — and keep the nation of Wales at the forefront of theoretical science and technical innovation.

TO LEARN MORE

To learn more about the exciting research under way at Supercomputing Wales, visit supercomputingwales.com or send an email message to enquiries@supercomputing.wales.

For a closer look at the high performance computing systems in the Supercomputing Wales environment, see the Dell Technologies case study “[World-Class Research](#).”