

The Use of Supercomputers for High Fidelity, Actionable Predictions

Peter V. Coveney

Centre for Computational Science,

Department of Chemistry

University College London

Computational Science Laboratory

Institute for Informatics

Faculty of Science

University of Amsterdam

SURFsara Super Day, 17 December 2019

Acknowledgements

People (1)

Derek Groen (*Brunel University London*)

Robert Sinclair (*University College London*)

Maxime Vassaux (*University College London*)

Krishnakumar Gopalakrishnan (*University College London*)

Robin Richardson (*University College London*) → *Netherlands e-Science Centre*

Dave Wright (*GTN, London*)

Olivier Hoenen (*Max Planck Institute for Plasma Physics*)

David Coster (*Max Planck Institute for Plasma Physics*)

Bruce Boghosian (*Tufts University, Boston*)

Hongyan Wang (*Facebook, Menlo Park*)

...and many others!

Acknowledgements

People (2)

Prof Alfons Hoekstra (*University of Amsterdam, The Netherlands*)

Prof Peter Sloot (*Director of UvA Institute for Advanced Study*)

Prof Daan Crommelin (*Centrum Wiskunde & Informatica, Amsterdam*)

Dr Marco Verdicchio (*SURFsara*)

Dr Walter Lioen (*SURFsara*)

Dr Peter Michielse (*SURFsara*)

Acknowledgements



EuroHPC
Joint Undertaking



COMPAT
Computing Patterns
for High Performance
Multiscale Computing

الصندوق الوطني للبحوث العلمي
Qatar National Research Fund

Member of Qatar Foundation



Distributed
European
Infrastructure for
Supercomputing
Applications



Funding & Projects

Contents

- What can supercomputers do?
- Predicting the properties of advanced materials
- Computational biomedicine
 - Molecular medicine, drug discovery and personalised medicine
 - The virtual human
- Verification, validation and uncertainty quantification (VVUQ)
- The limits of computers
 - a new pathology of the floating point numbers

What can supercomputers do?

- Bigger, better, faster, new and different science
- They consume and dissipate huge amounts of energy. We need solutions to reduce this

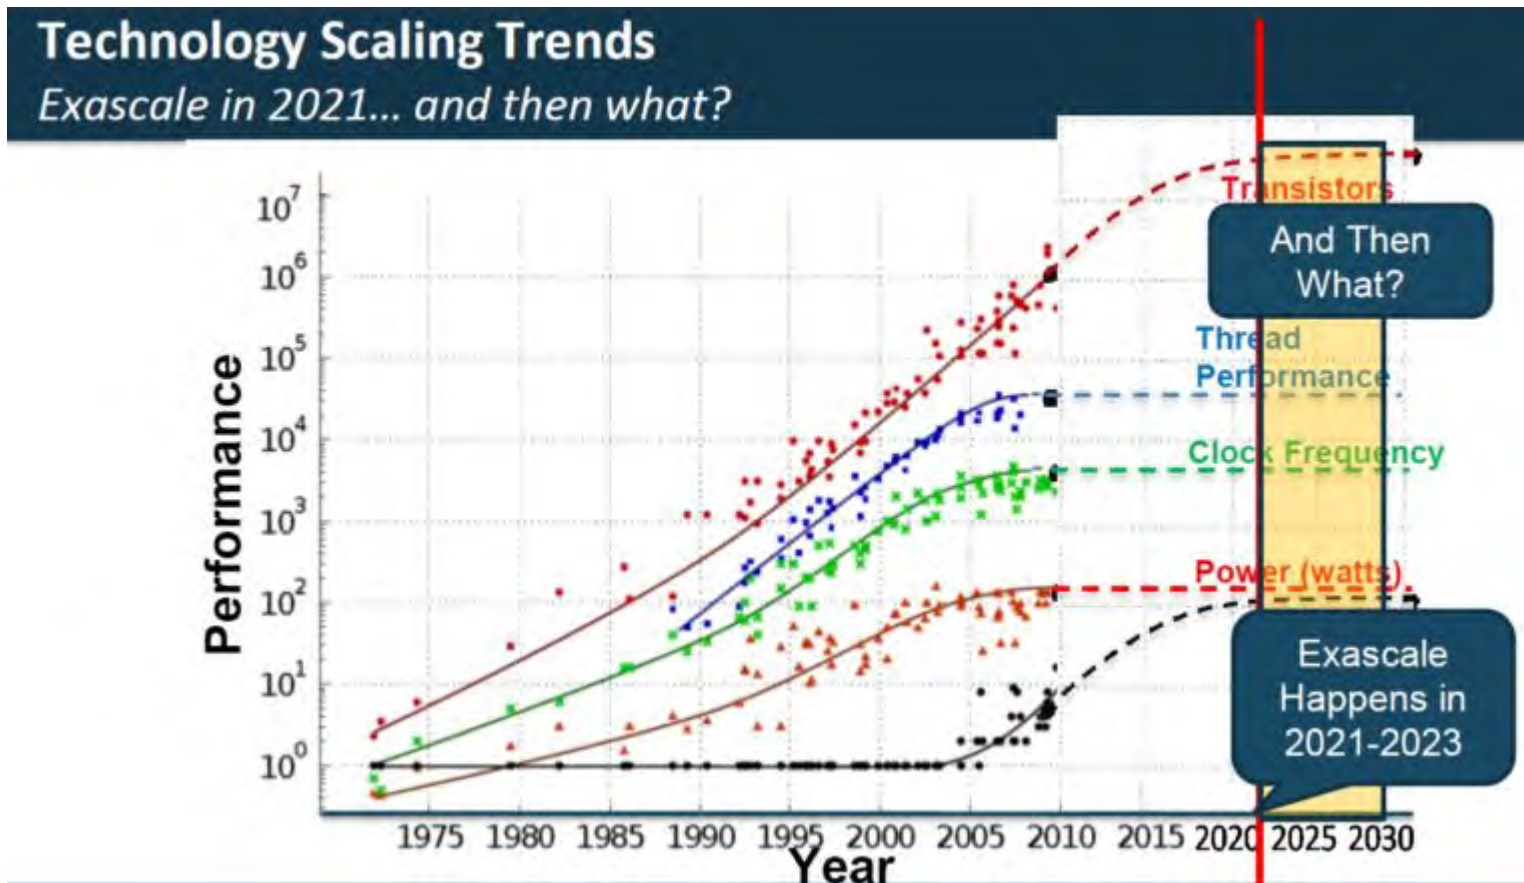


Image source: 3Dincites

<https://www.3dincites.com/2019/09/iftle-426-exascale-computing-is-near-incandescent-lightbulbs-get-a-reprieve/>

HPC resources we use (1)




Summit – 200 PetaFlops (peak)



Frontera – 38.75 PetaFlops (peak)



SuperMUC-NG – 26.8 PetaFlops (peak)




Mare Nostrum 4 – 11.1 PetaFlops (peak)



Piz Daint – 27.1 PetaFlops (peak)

HPC resources we use (2)

- SURFsara (Netherlands)
 - Cartesius
- EPCC (University of Edinburgh)
 - Archer
 - Cirrus
- Barcelona Supercomputing Centre (Spain)
 - MareNostrum (also through PRACE)
- Leibniz Supercomputing Centre (Germany)
 - SuperMUC/SuperMUC-NG
- Cyfronet (Poland)
 - Prometheus
- PRACE
 - Piz Daint (CSCS)
- National Center for Supercomputing Applications (USA)
 - Blue Waters
 - Titan/Summit
- Posnan Supercomputing and Networking Centre (Poland)
 - Eagle



Advanced Functional Materials

- Light, strong, tough, sustainable
- Aerospace and automotive
- Takes *ca.* 20 years for acceptance into commercial operations
- Virtual certification to accelerate rate of commercial adoption

Clay and Graphene Nanocomposites


- To enhance the properties of polymers, nanoscale particles are added to create new type of hybrid material, with improved mechanical performance even at low clay volume fractions.
- Original patent held by Toyota (late 1980s – with clay nanoparticles).
- Used today in several industries: aviation, automotive, energy, medical




Picture from www.nanowerk.com

Simulating Nanocomposites

- **Aim:** Develop accurate multiscale models of 2D-nanocomposites (graphene, graphene oxide, clay and polymer)
- Uses:
 - Predict microstructure
 - Predict materials properties
- We require:
 - Accurate potentials
 - Realistic structures
 - Orchestrating many diverse simulations
 - Couple simulations with different length and time scales



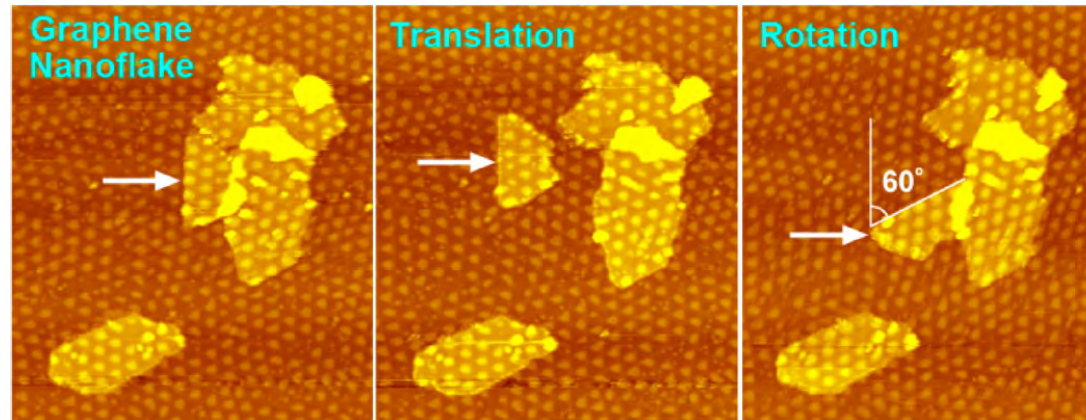
Our Multiscale Nanocomposites Models



DFT: Density Functional Theory
 MD: Molecular Dynamics
 FEM: Finite Element Method

Graphene: nanoscale properties



Graphene flakes slide in a superlubric state over graphite. We have shown through multiscale simulation that friction on the nano-scale manifests through alignment events which dissipate energy to the substrate.



Sinclair, Suter and Coveney. *Adv Mat*, 30 (13), 2018



What is Graphene Oxide's Nano-Structure?



Reactivates of 52
graphene-oxide
structures from
DFT calculations

Sinclair & Coveney, *J. Chem. Inf. Model.*, **59** (6), 2741-2745 (2019)



Random Forest
model predicts
reactivity of all
possible structures



github.com/velocirobbie/make-graphitics

github.com/velocirobbie/graphene-oxide-percolation

Graphene Oxide: Nano-Structure


Graphene Oxide

Using accurate graphene oxide structures we can predict its bulk behavior with coarse grained simulation.

Graphene aggregates in most solvents.

Graphene oxide disperses in polar solvents.


Depending on extent of oxidation, flakes can aggregate into extended 3D structures.





Concurrent Multiscale Modelling (1)

- We have coupled a finite element method with molecular dynamics simulation.
- Instead of constitutive equations, the FEA calculates the material's stress-strain response by calling molecular dynamics simulations.
- We will be able to predict history/time dependent dynamics, like hysteresis and crack-propagation.
- This is based on Heterogenous Multiscale Modelling (HMM)



Scales within the heterogeneous multiscale – A finite element based solution of continuum mechanics at the macroscale coupled to nanoscale simulations through molecular dynamics.

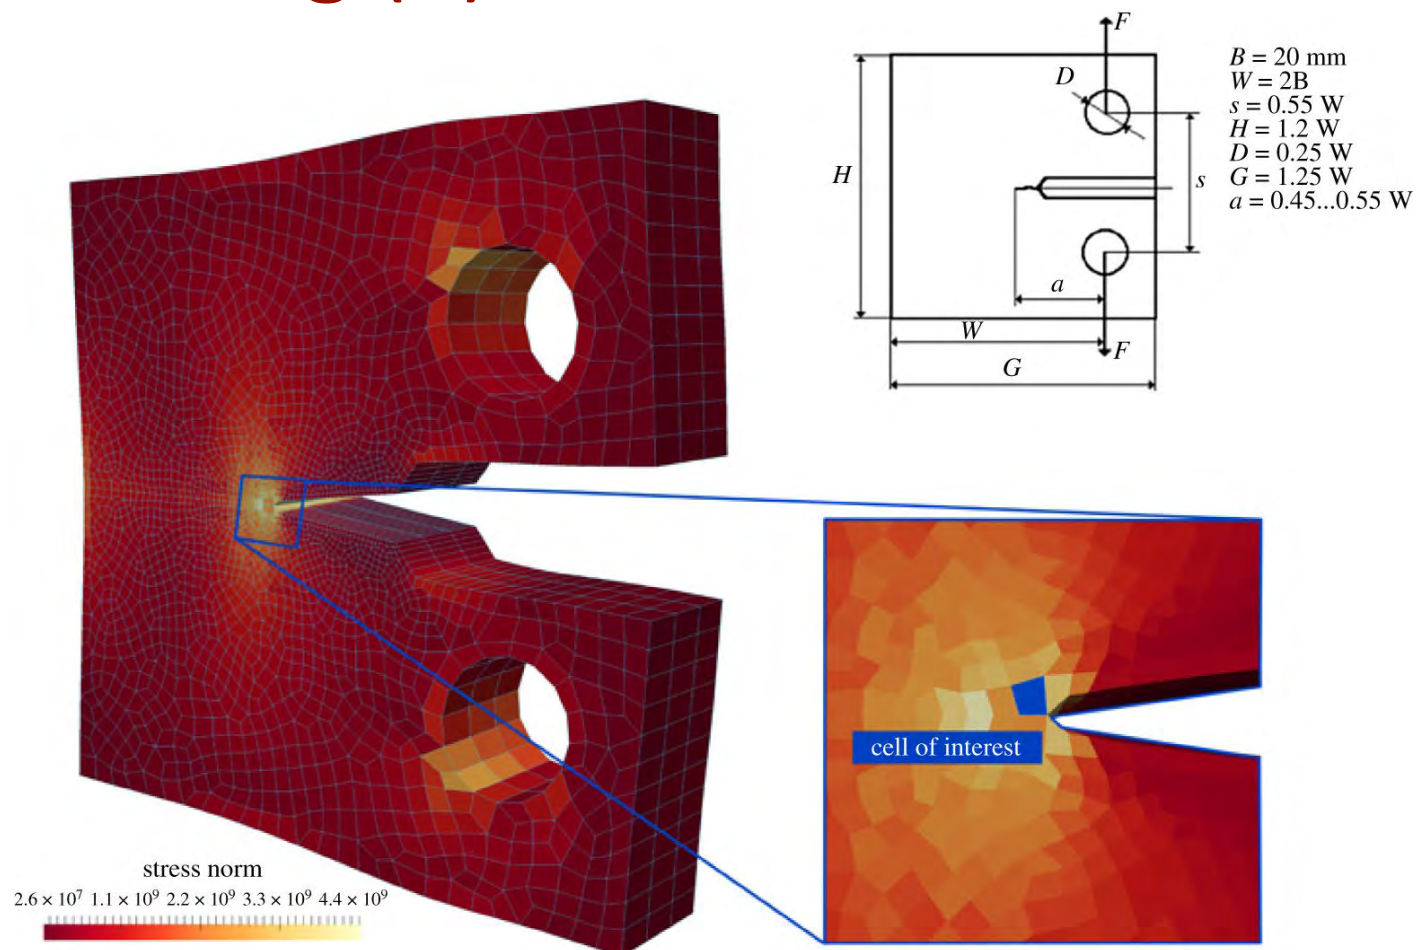
Vassaux, Sinclair, Richardson, Suter and Coveney, *Advanced Theory and Simulations*, **1900122** (2019)

Concurrent Multiscale Modelling (2)

Compact tension tests allow engineers to test the fracture properties of a material

It induces a different mode of failure compared to a dogbone test.

This cannot be simulated using molecular dynamics alone.



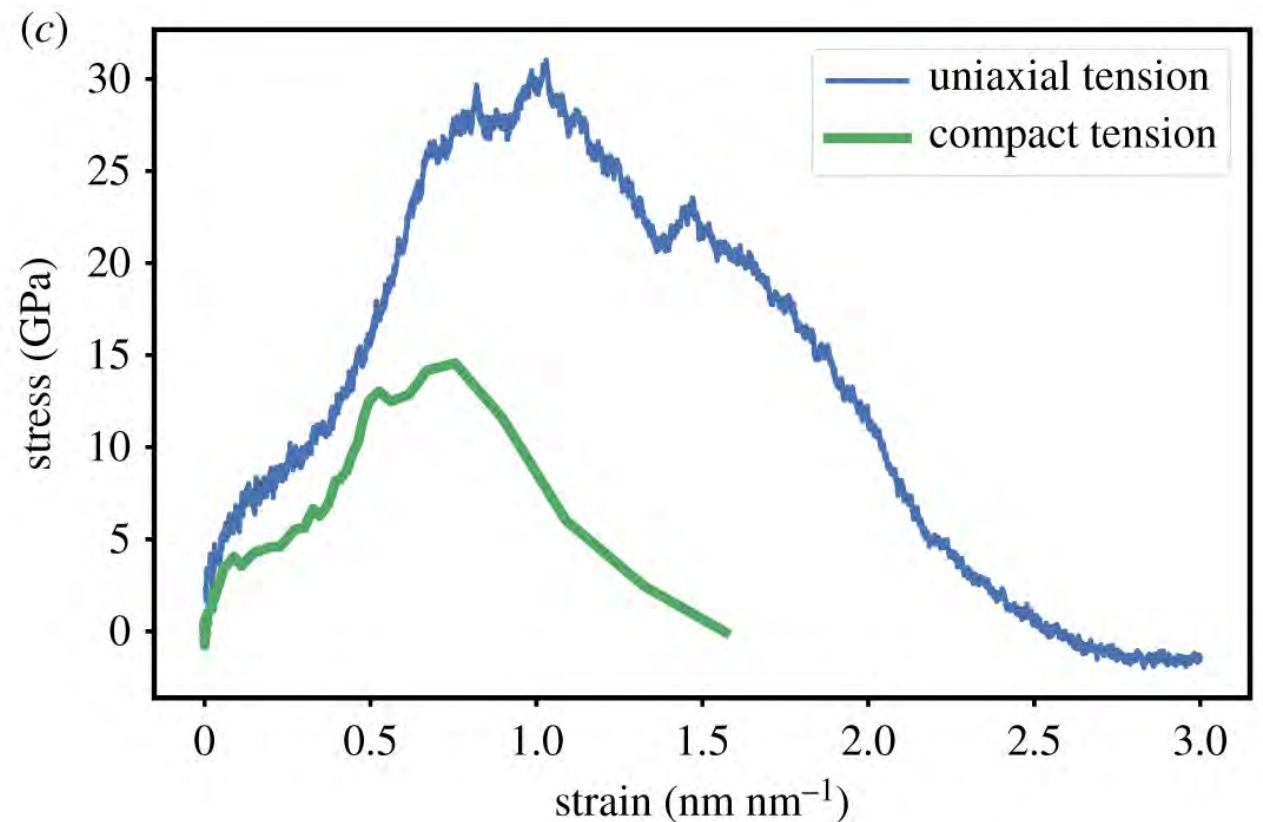
Vassaux, Richardson, Coveney, Phil Trans R Soc A, **377** (2019)

Concurrent Multiscale Modelling (3)

Depending on the continuum level boundary conditions we get significantly different results in the calculated material properties.

Here, a compact tension test induces a different mode of failure compared to a dogbone test.

This allows us to compare more closely with experiment



Vassaux, Richardson, Coveney, Phil Trans R Soc A, **377** (2019)

IMAX Virtual Humans film

Computational Biomedicine

EU Centre of Excellence in computational Biomedicine




- User-driven biomedical modelling and simulation

- Biomedical science across all scales



Molecular Systems




Organisms (The virtual human)




Organs


- Targeting new and emerging biomedical research areas



Cardiovascular medicine



Molecularly-based modelling



Neuro-muscoskeletal medicine



This project has received funding from the European Union's Horizon 2020 research and innovation programme.



Patient Specific Modelling

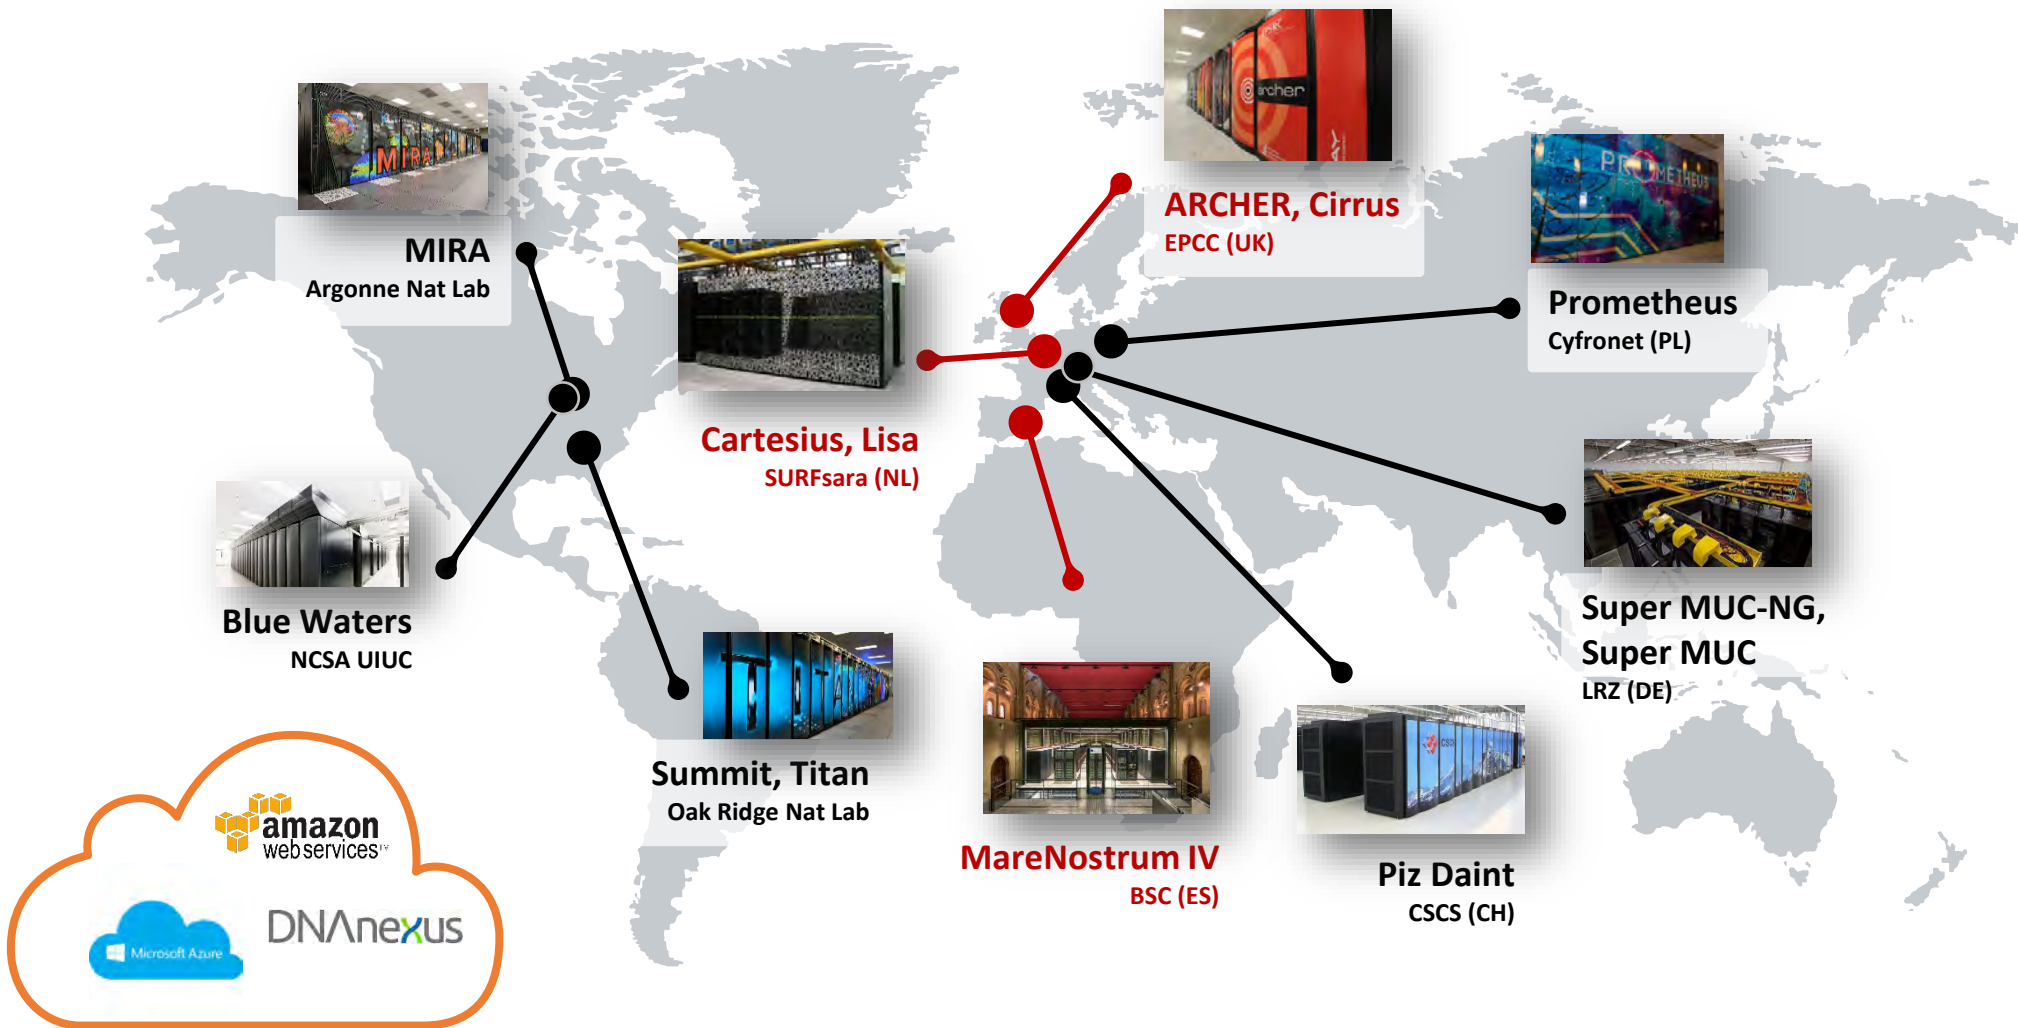
Answer: Assist clinicians' ability to understand how a course of treatment will impact a given individual

How: Simulating the patient using a personalised virtual human

Why:

- ▶ Multiple options for treatment can be investigated and the optimal course chosen for an individual
- ▶ Clinicians can treat a patient with greater confidence
- ▶ Virtual human models can provide non-traditional information about the patient
- ▶ Patients can have a clearer understanding of how a treatment will impact them

HPC in CompBioMed






Cloud Computing in CompBioMed



Software as a Service (SaaS) offers the customer the capability to use services pre-deployed by the provider running on cloud infrastructure without the need to control directly any components of the stack


Containers for biomedical applications



Personalised Medicine


“One Size Does Not Fit All”

PERCENTAGE OF THE PATIENT POPULATION FOR WHICH A PARTICULAR DRUG IS INEFFECTIVE, ON AVERAGE



Sequencing

```
cctcaaatcactctttggcaacgaccatcgtcacaataaagat
aggagggcaactaagggaagctctattagatacaggagcag
atgatacagattagaagacataaattaccaggaagatggaa
accaaaaatgataggggaattggaggtttgccaaagtaaga
cagtatgatcagataccogtagaaatctgtggacataaagttat
aggtagcagtagtagtaggaccacacctttcaacataaattggaa
gaaatctgttgactcagctggctgcaacttaaatttccattagtc
ctattgaaactgtaccagtaaaatgaagccaggaatggatgg
ccc aaaagttaacaatggccattgacagaagaaaaataa
aagcaytagtagaaattgtacagaactgaaaaggaggaa
aaattcaaaaatggcctgaaaatccatacaatactccagtat
ttgccataaagaaaaargacggtagtaaatggagaaaatagt
agatttcagagaacttaataagagaactcaagactctgggaa
gtcaattaggaataccacatcctgcagggttaaaaagaaaa
aatcagtaacagtagctg
```



NATURE | OUTLOOK



Therapy: This time it's personal


Lauren Gravitz

Nature 509, S52–S54 (29 May 2014) | doi:10.1038/509S52a
Published online 28 May 2014

[PDF](#)
[Citation](#)
[Reprints](#)
[Rights & permissions](#)
[Article metrics](#)

Tailoring cancer treatment to individual and evolving tumours is the way of the future, but scientists are still hashing out the details.

Subject terms: [Cancer therapy](#) · [Drug development](#) · [Targeted therapies](#) · [Chronic myeloid leukaemia](#) · [Leukaemia](#) · [Cancer genetics](#) · [Personalized medicine](#)




Brendan Monroe

Drug Discovery


Searching for a needle in a haystack

Virtual screening:

Systematic computer-based prediction of binding affinity of compounds to proteins



You were right: There's a needle in this haystack...




High-Throughput Screening (HTS)

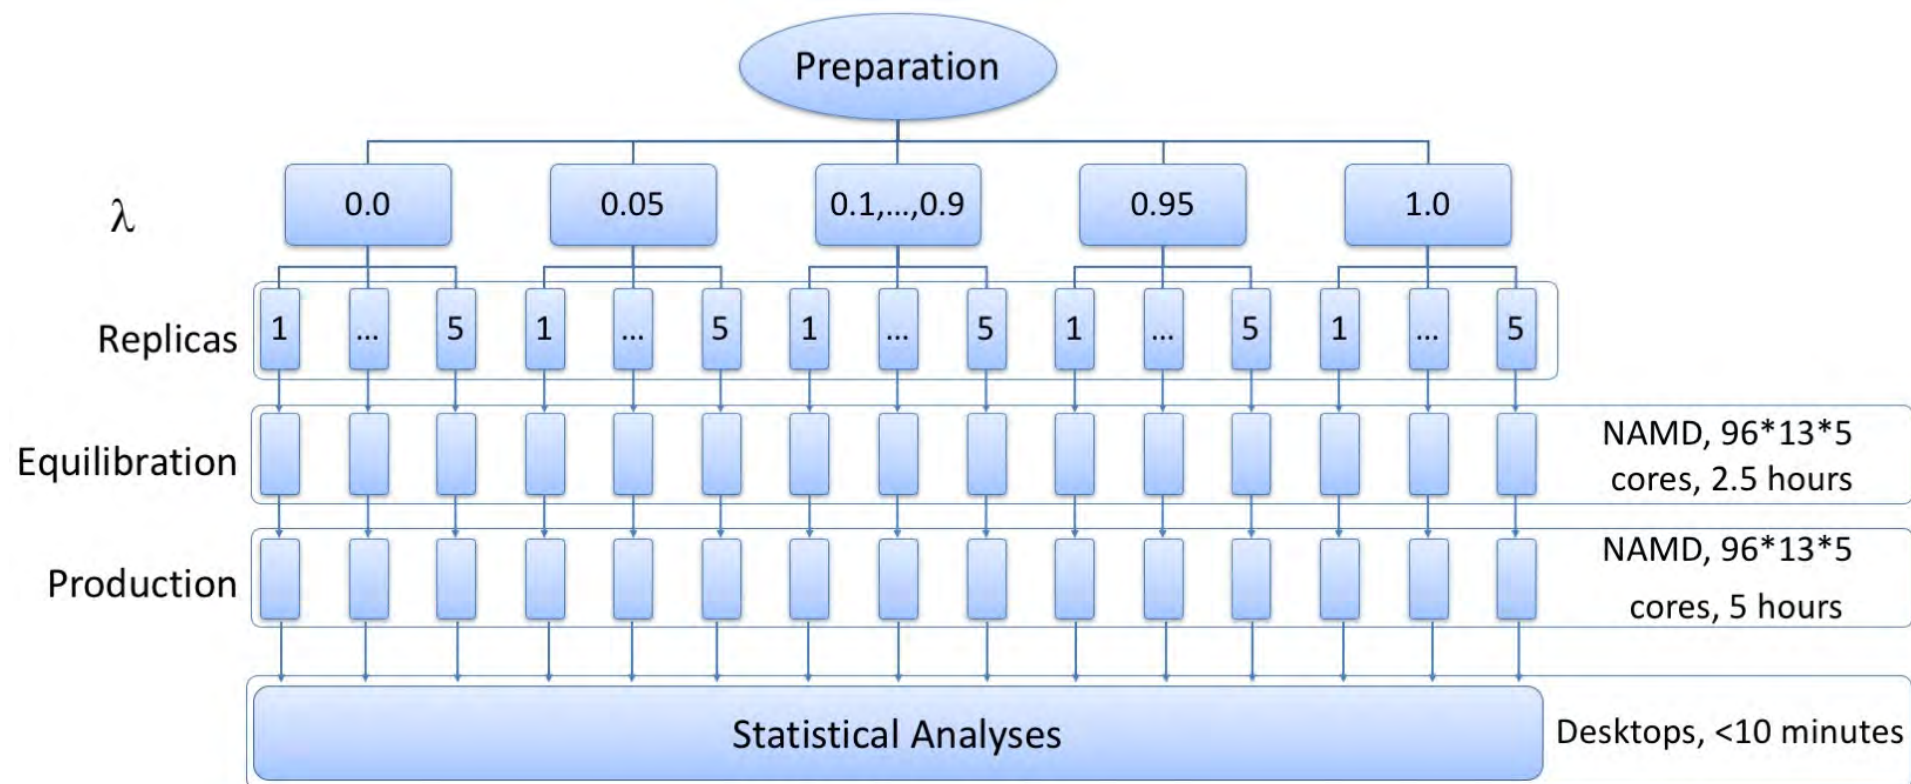
- millions of samples of chemical compounds
- HTS can test 100,000 compounds a day
- Cost of HTS is substantial: 1-10\$/compound

Binding Affinity Calculator (BAC)

BAC can reliably predict binding affinities of compounds with their target proteins, and be used potentially as a drug ranking tool in clinical application or a virtual screening tool in pharmaceutical lead discovery.



Binding Affinity Calculator (BAC) is a software toolkit which automates the implementation of TIES (and ESMACS) methods for binding affinity calculations







A.P. Bhati, S. Wan, D.W. Wright & P.V. Coveney, *J. Chem. Theory Comput.*, **2017**, 13, 210–222.

Pharmaceutical Drug Discovery

Experiment



Computation
blind study

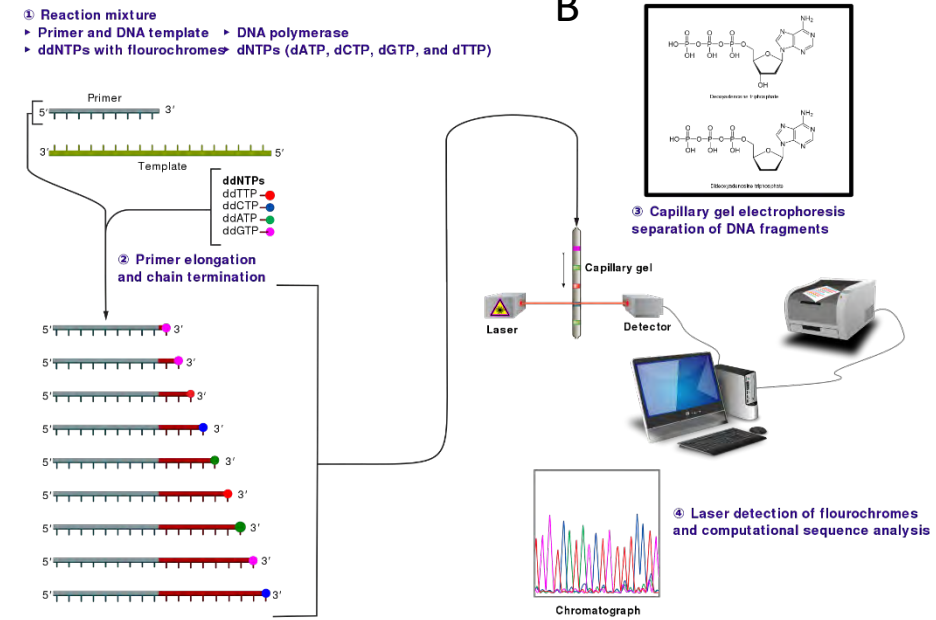
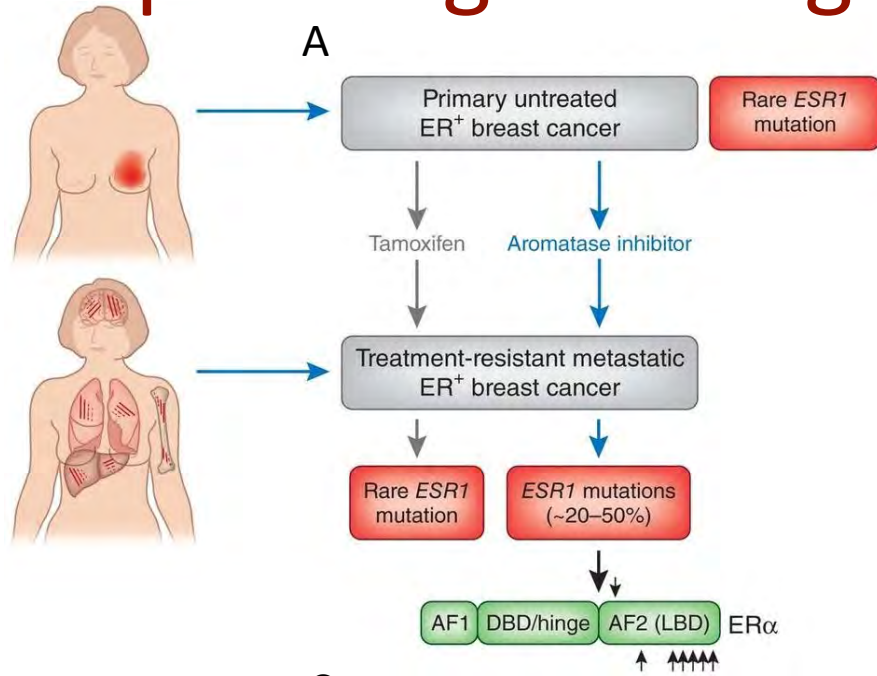


S. Wan, A. Bhati, S. Skerratt, K. Omoto, V. Shanmugasundaram, S. Bagal, P. V. Coveney, *J. Chem. Inf. Model.*, **2017**, *57*, 897–909.


S Wan, AP Bhati, SJ Zasada, I Wall, D Green, P Bamborough, PV Coveney, *J. Chem. Theory Comput.*, **2017**, *13*, 784–795.

DW Wright, S Wan, C Meyer, H van Vlijmen, G Tresadern, PV Coveney, *Sci. Rep.* **2019**, *9*, 6017.


Sequencing – Estrogen Receptor



C Structural Organization of Nuclear Receptors



- ESR1 gene plays a vital role in breast cancer development (A).
- Sanger sequencing and computational analyses to identify mutations in targeted gene (B).
- Identified mutations used for receptor-drug interaction study (C).



S. Wan, D. Kumar, V. Ilyin, U. Al Homsy, G. Sher, K.R.A. Knuth and P.V. Coveney, "From genome to personalised medicine: cancer treatment and discovery of novel variants in Qatar", submitted 2019.


Modelling – Mutations Induce Resistance

- Both L384V and L387R mutations induce resistance
- L387R mutation is likely to block the bindings for all of these drugs and the native estradiol

ESMACS binding affinities for six ligands with three ER receptors: WT, L384V and L387R.


	Afi	Edo	Est	Ral	Tam	Tor
WT	-38.4	-36.2	-29.1	-45.1	-36.9	-37.0
L384V	-37.1	-36.0	-27.5	-41.2	-36.3	-36.8
L387R	-37.4	-35.3	-27.8	-41.5	-33.4	-32.6

Relative TIES-PM binding free energies $\Delta\Delta G$ for six ligands with two ER mutations: L384V and L387R. The energy difference is calculated as $\Delta\Delta G = \Delta G_{binding}^{mut} - \Delta G_{binding}^{WT}$.



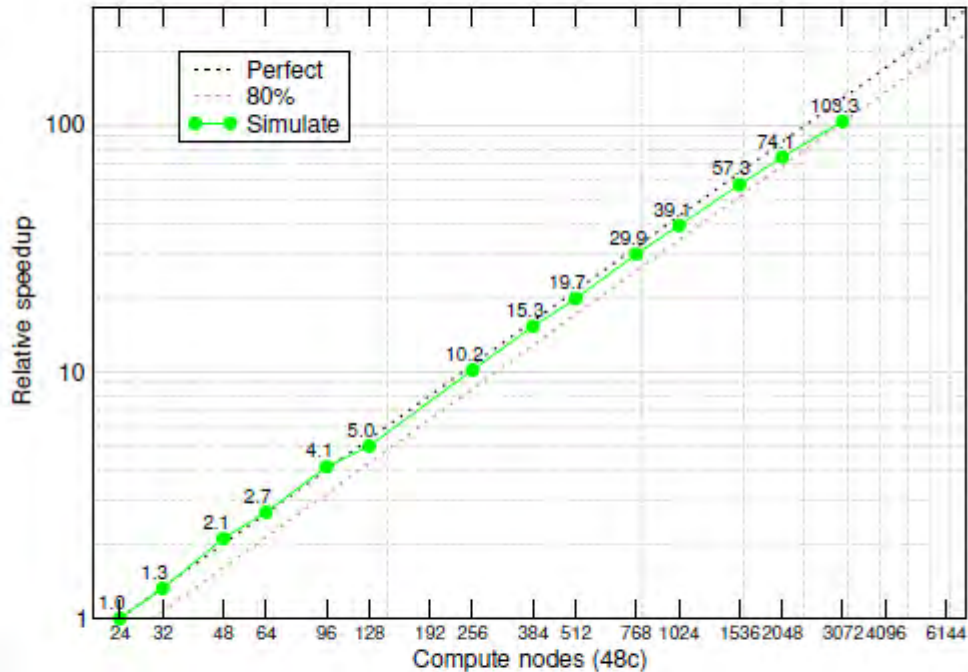
Ligand	L384V		L387R	
	$\Delta\Delta G$	error	$\Delta\Delta G$	error
Afi	2.18	0.42	5.11	0.45
Edo	2.00	0.40	4.78	0.52
Ral	2.17	0.42	6.07	0.92
Tam	2.19	0.41	4.66	0.55
Tor	2.20	0.53	4.58	0.60
Est	2.24	0.31	5.22	1.84

The Virtual Human




- ▶ Develop numerical models to predict behaviour of bodily components - organs, blood, nerves, ...
- ▶ Behaviour needs to be captured at multiple length scales
- ▶ Codes need to communicate at physiological interfaces
- ▶ Resolution needs to be high enough for clinical accuracy
- ▶ Ultimately, calculation time needs to be as short as possible

HemeLB Large-Scale Performance




Scaling performance on SuperMUC-NG
(1 node = 48 cores)



- ▶ Circle of Willis geometry - 10,154,448,502 lattice sites for 5000 steps
- ▶ Strong scaling observed up to 147,456 cores. Full machine scaling underway

Current progress



- ▶ Arteries - 1 inlet, 38 outlets; Veins - 494 inlets, 1 outlet
- ▶ Resolution - $50\mu m \approx 1.46$ billion lattice sites

Observation - geometry generation is difficult, execution is (relatively) straightforward


Challenge - Visualisation of exascale data to be efficient and meaningful

Verification, Validation and Uncertainty Quantification (VVUQ)

- Uncertainty quantification (UQ) common in engineering and applied mathematics
- One typical example is in weather and climate forecasting, which involve coupled models
- Historically, low activity in domains at shorter length and time-scales
- UQ for multiscale problems largely open
- Ensembles play a central role



What is VVUQ?



B.H.Thacker, et al., "Concepts of Model Verification and Validation." 2004. DOI: 10.2172/835920.

- **Verification**
 - Does the computational model fit the mathematical description?

- **Validation**
 - Is the model an accurate representation of the real world?

- **Uncertainty Quantification**
 - How do variations in the numerical and physical parameters affect simulation outcomes?

VECMA Toolkit

- The current release is made up of the following tools:
- **FabSim3** - a python-based automation toolkit for scientific simulation and data processing workflows.
- **EasyVVUQ** – a python library to facilitate verification, validation and uncertainty quantification (VVUQ) for a wide variety of simulations.
- **QCGPilotJob** – a pilot job manager
- **EasyVVUQ-QCGPilotJob** - enables efficient execution of critical parts of EasyVVUQ workflows on HPC machines.
- **MUSCLE3** – a Multiscale Coupling Library and Environment.



Includes



Open source development: <http://www.vecma-toolkit.eu>

EasyVVUQ

- Examples:
 - **Advanced Materials** - **Climate prediction**
 - **Fusion energy** - **Human migration**

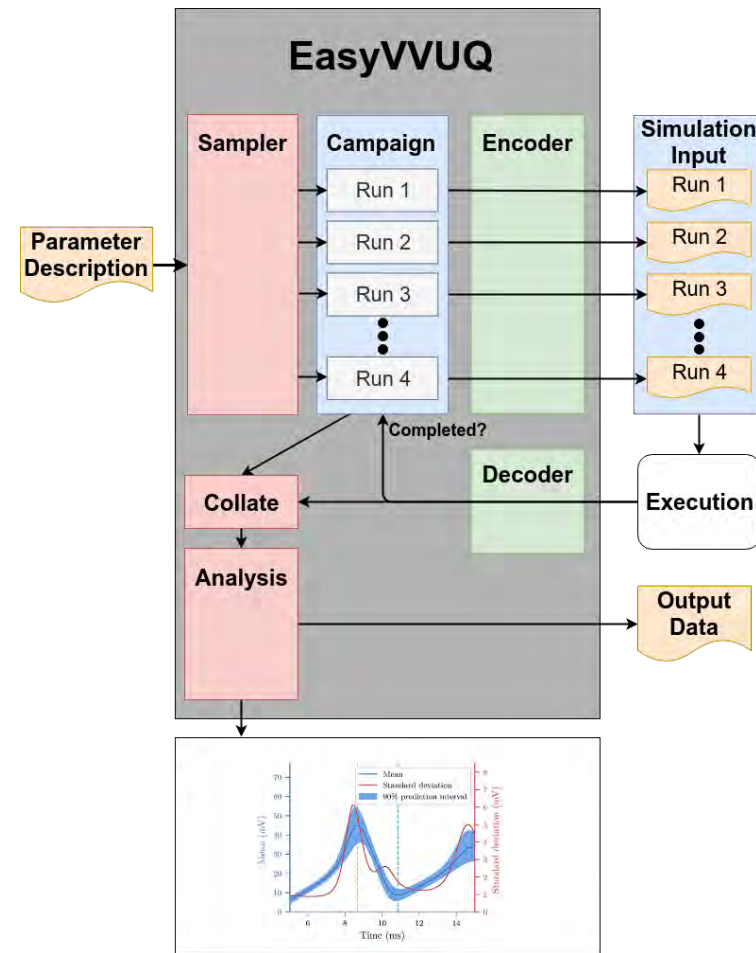


Robin R Richardson



Dave W Wright


- A framework for VVUQ aimed at complex multiscale workflows, with very large and wildly varying resource requirements
- Breaks algorithms down into most basic elements
- Allows rapid prototyping of such multiscale VVUQ workflows
- Python library. Makes use of existing libraries and codes wherever possible.




“Building confidence in simulation: Applications of Easy VVUQ”; authors: D W Wright, R R Richardson, P V Coveney et al (preprint 2019)

EasyVVUQ: A library for verification, validation and uncertainty quantification in high performance computing, Robin A. Richardson David W. Wright, Wouter Edeling, Vytautas Jancauskas, Jalal Lakhilili, Peter V. Coveney. (software release, 2019)

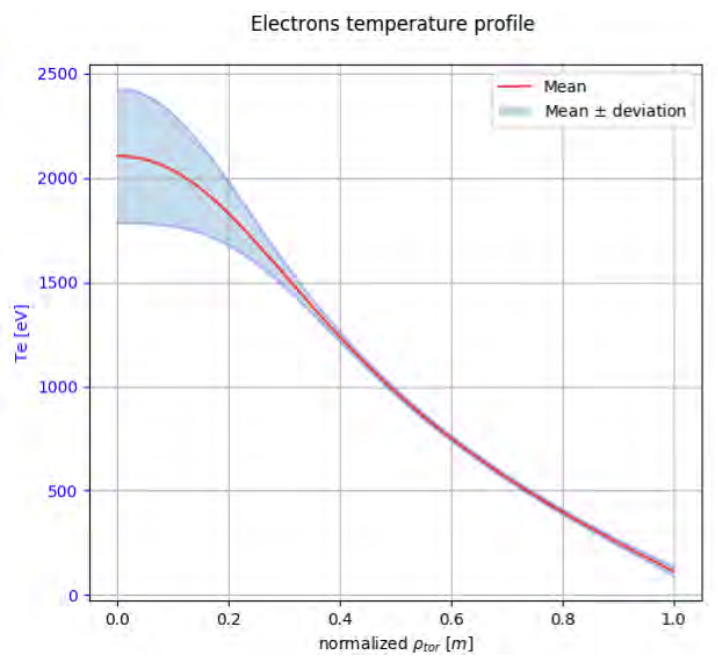
Applications of EasyVVUQ: II





Materials properties



Fusion energy



Climate change



The Limits of Computers

The statistical properties of chaotic systems

A new pathology in the simulation of chaotic dynamical systems

Bruce M. Boghosian (Tufts University)
Peter V. Coveney (UCL)
Hongyang Wang (Tufts University)



Boghosian, Bruce M., Peter V. Coveney, and Hongyan Wang. "A New Pathology in the Simulation of Chaotic Dynamical Systems on Digital Computers" *Advanced Theory and Simulations* (2019): 1900125.

The generalised Bernoulli map

Bernoulli map: a simple dynamical system which exhibits chaotic behavior

$$x_{t+1} = 2 x_t \bmod 1 \quad x \in [0,1)$$

The generalised Bernoulli map also known as the β shift:

$$x_{t+1} = \beta x_t \bmod 1 \quad x \in [0,1),$$

- a one-parameter map where β is either an integer or a rational non-integer (> 1)



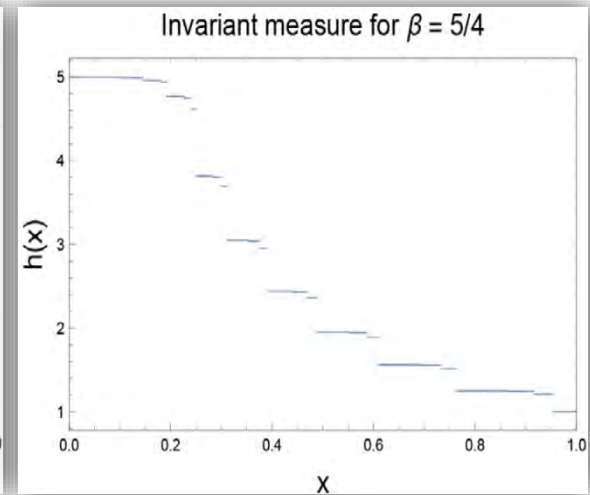

Many things are known about the behaviour of this map using continuum mathematics

Expectation values can be obtained from its unstable periodic orbit (UPO) spectrum

Ruelle, David. *Thermodynamic formalism: the mathematical structure of equilibrium statistical mechanics*. Cambridge University Press, 2004.

Properties of the map

- The invariant measure has discontinuities at a dense set of points in $[0,1)$
- Examples for four non-integer values shown
- Caution: Graphs are less smooth than they appear

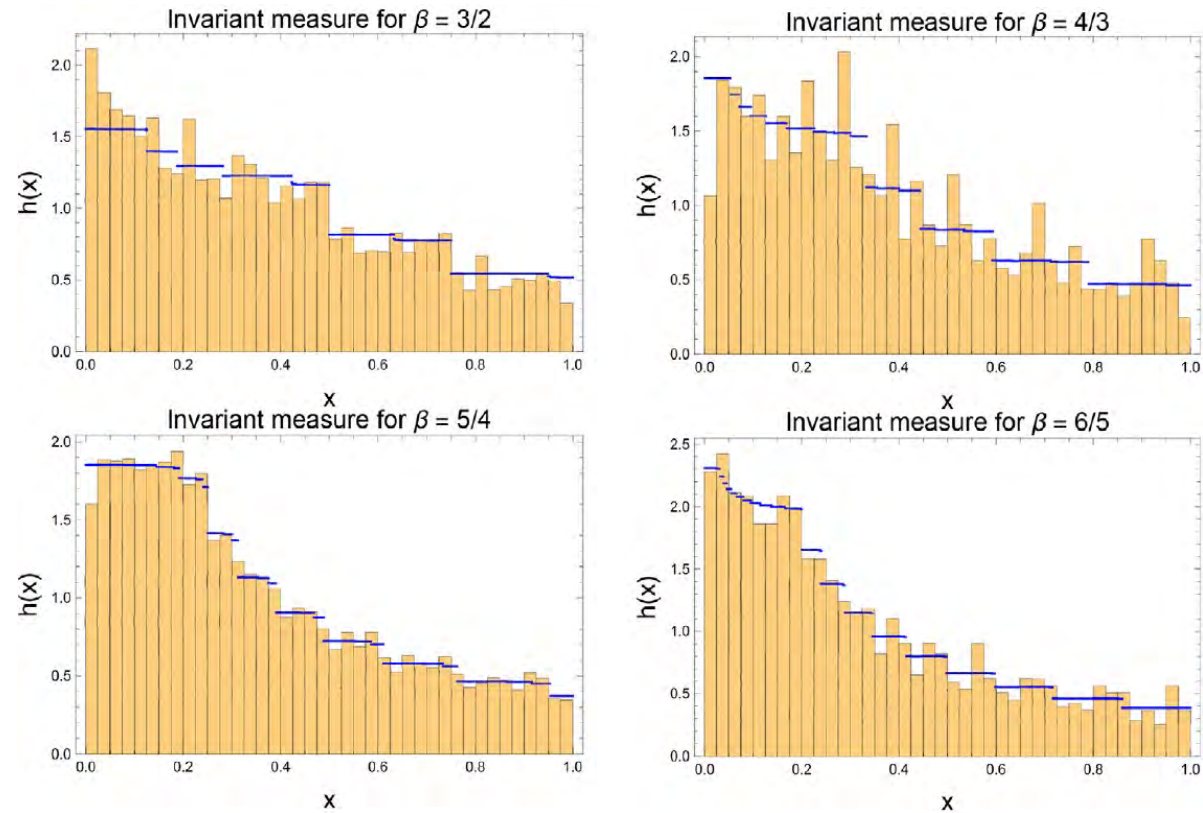


Invariant measures of the generalized Bernoulli map f_β for $\beta = \frac{6}{5}, \frac{5}{3}, \frac{4}{3}, \frac{3}{2}$.
These are normalized so that $h_\beta(1) = 1$, which corresponds to $C = 1$ in the Hofbauer series

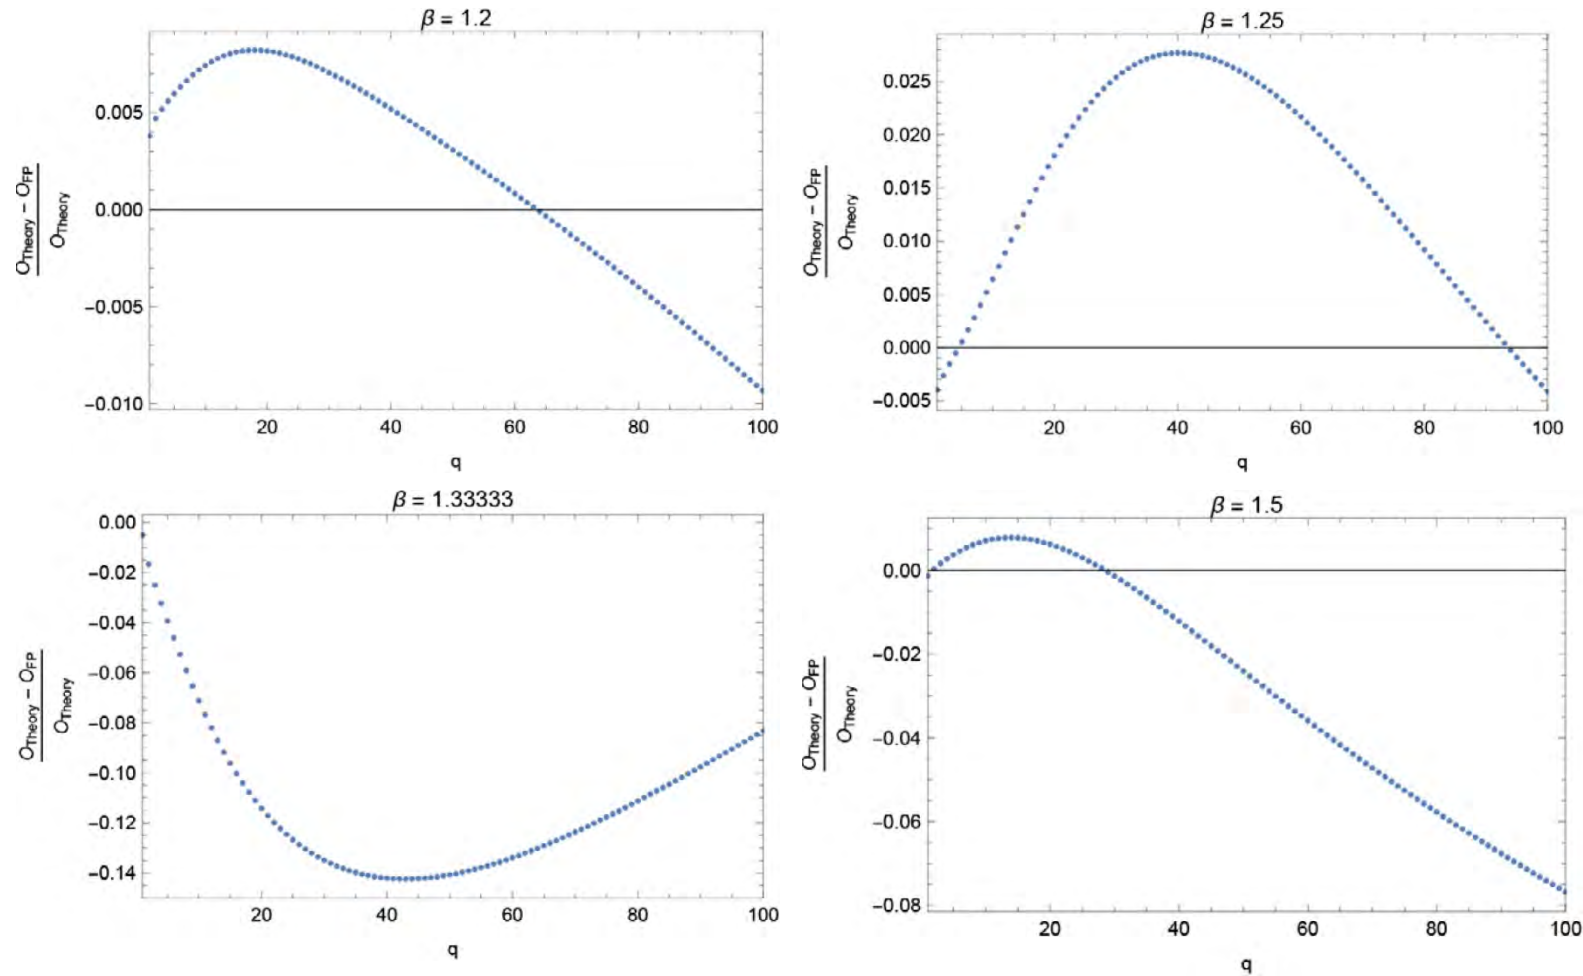
The β shift – floating point representation

- The map can be represented & simulated on digital computers using standard IEEE floating-point numbers
- Single-precision IEEE floating-point numbers consist of 32 bits, of the form $\sigma, e_1, e_2, \dots, e_8, m_1, m_2, \dots, m_{23}$ where σ is the sign bit, e_j are the exponent bits & m_j the mantissa bits.
- Similar construction for double-precision numbers, but using 52 mantissa bits and 11 exponent bits.
- All floating point numbers are dyadic (numbers whose denominators are powers of two)
- However, this is a very poor representation of the rational numbers

- Floating-point arithmetic causes highest damage to the dynamics for even values of β
- Consider $\beta = 2$
- The binary digits shift one place to the left with each iteration
- 1 iteration \rightarrow left shift bits by 1 place \rightarrow loss of 1 bit of precision with each application of the map
- Result will be zero:
 - after 23 iterations for single-precision arithmetic
 - after 52 iterations for double-precision arithmetic
- The invariant measure will be a Kronecker delta at $x = 0$
- In the hypothetical limiting case of number of mantissa bits approaching ∞ , the Kronecker delta would effectively approach a **delta distribution** at $x = 0$
- f.p. arithmetic's exact time-asymptotic result will never be a uniform measure, the correct answer for the real-valued dynamics



- Discrepancy between the exact (blue) and numerical (histogram) invariant measures for the generalized Bernoulli map f_β for $\beta = 3, 5, 7, 9$ and for $\beta = \frac{3}{2}, \frac{4}{3}, \frac{5}{4}, \frac{6}{5}$
- This simulates the average we would obtain if we could run over both an infinite length of time and an infinite ensemble size.
- While the agreement is good for odd integer β (though still greater than round off), it is seen to be very poor for non-integer β



Relative error of the floating-point calculation of the expectation value of x^q for the generalized Bernoulli map f_β for $\beta = \frac{6}{5}, \frac{5}{4}, \frac{4}{3}, \frac{3}{2}$ simulating the average we would obtain if we could run over both an infinite length of time and an infinite ensemble size

Summary of floating point analysis

- Floating point numbers have a strongly detrimental influence on the map due to
 - their discrete and finite nature, and
 - the delicate structure of the attracting set of chaotic dynamical systems
- For even integer values of the parameter ($\beta = 2, 4, 6, \dots$), the long time behaviour is completely wrong
- For non-integer β , relative errors in observables can reach *ca* 15%
- For odd integer β values, relative errors two orders of magnitude larger than those attributable to round-off.

Conclusions

- We have discovered a new pathology in floating point numbers
- The errors persist for arbitrarily high finite precision floating point numbers.
- It is distinct from round off, loss of precision, loss of significance etc., errors.
- The errors arise from the discreteness of floating point numbers.
- There is no reason to think that substantial errors of this sort are not present in chaotic systems of interest in science and engineering.