



HPC at the University of Oxford

Steven Young
Acting Head of Advanced Research Computing

email: steven.young@it.ox.ac.uk

HPC User Forum
Thursday 29 September 2016

As the oldest university in the English-speaking world, Oxford is a unique and historic institution. There is no clear date of foundation, but teaching existed at Oxford in some form in 1096 and developed rapidly from 1167, when Henry II banned English students from attending the University of Paris.



Oxford ranked No 1



World University Rankings 2016-2017

The *Times Higher Education* World University Rankings 2016-2017 list the 980 top universities in the world, making it our biggest international league table to date. It is the only global university performance table to judge world class universities across all of their core missions – teaching, research, knowledge transfer and international outlook.

Read more...

IN PARTNERSHIP WITH ELSEVIER

2017

Show me universities in any country that offer any subject

Or, find a specific university by name

RANKINGS ONLY		PERFORMANCE BREAKDOWN					KEY STATISTICS *
Rank	Title	Teaching	International Outlook	Research	Citations	Industry Income	Overall
1	University of Oxford United Kingdom	89.6	94.5	99.1	99.2	62.5	95.0
2	California Institute of Technology United States	95.5	63.4	95.7	99.8	90.8	94.3
3	Stanford University United States	92.6	76.5	95.9	99.9	60.9	93.8

Read more about the World University Rankings 2016-2017

ACADEMIC INSIGHTS

- World University Rankings 2016-2017: results announced
- World University Rankings 2016-2017: Standing still is not an option
- World University Rankings 2016-2017: Higher education's diverse mission
- World University Rankings 2016-2017: Keep the UK's doors open to students and scholars
- World University Rankings 2016-2017: Burgeoning rankings list has BRICS struggling to keep pace
- World University Rankings 2016-2017: The nine big challenges global higher education must confront over the next five years
- World University Rankings 2016-2017: Why it's hard to make global comparisons in





- Provide central university High Performance Computing (HPC) service
- Gateway to regional, national and international HPC activities
- Provide support, training, expertise and guidance on HPC
- First step into HPC for many users

Hardware

- Compute clusters
 - GPU resources

Arcus B	362 nodes	Haswell	5792 cores
Arcus A	108 nodes	Sandy Bridge/ Ivy Bridge	1728 cores
Arcus B GPU	30 GPUs	K40 / K80	

- Shared memory systems
 - Caribou: 1TB, Arcus-B: 2x 1.5TB, 6TB
- Novel architectures
 - Intel Phi (Knights Corner)
 - Power 8
- Storage: Panasas, GPFS



ARCUS (B) cluster

OCF & Lenovo – 94Tflops

- NeXtScale System M5
 - Nextscale N1200 6U chassis
- 340 Compute Nodes
 - 2 x Intel E5-2640v3 CPU
 - Intel TruScale, QDR Infiniband
 - 248 64GB Nodes
 - 88 x 128GB Nodes
 - 4 x 256GB Nodes
- *Of which 10 dual purpose GPU Nodes*
 - Each with 2 x Nvidia K40
- Storage (Panasas)
 - 2 additional Panasas AS14s shelves added to common HPC Storage service (now 6 shelves total)
 - Total ~ 400TB usable



lenovo FOR
THOSE
WHO DO.

OCF

panasas

NVIDIA.

intel

Software

- HPC software stack
 - OS: Linux (CentOS)
 - Compilers: GNU, Intel, Portland Group
 - Libraries
 - Applications
- Installation of software
- SLURM job scheduler



Research in Oxford

- Anthropology
- Computational Medicine
- Engineering
- Square Kilometre Array (SKA)
- Quantum Technology
- Zoology

A - Anthropology



to

Z - Zoology

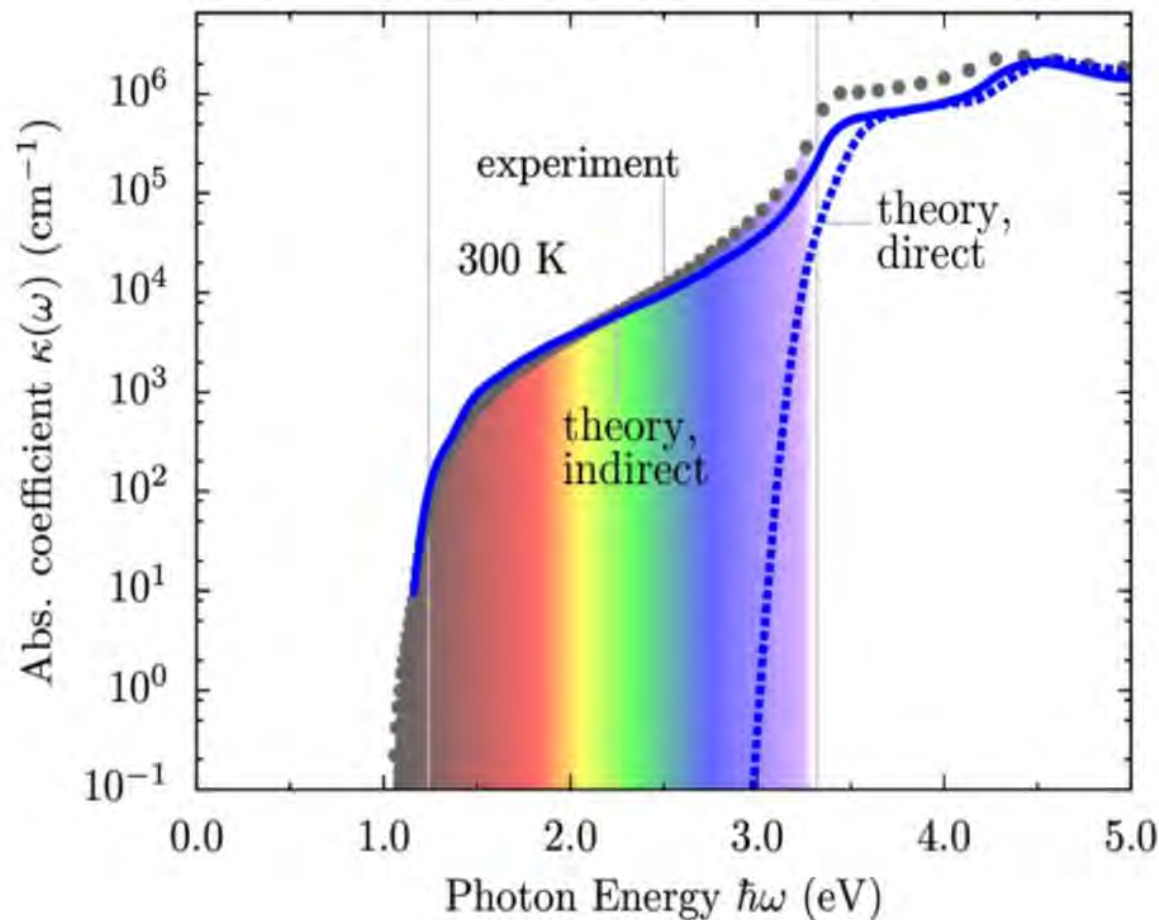


Case Studies

- Some examples of results generated through the use of ARC resources



Materials Science M. Zacharias, C. E. Patrick, and F. Giustino



We develop a first principles method for the calculation of optical spectra of semiconductors at finite temperatures ... Calculations (blue solid line) are based on a new approach developed in our group which describes interaction of light with quantum nuclear effects. ...

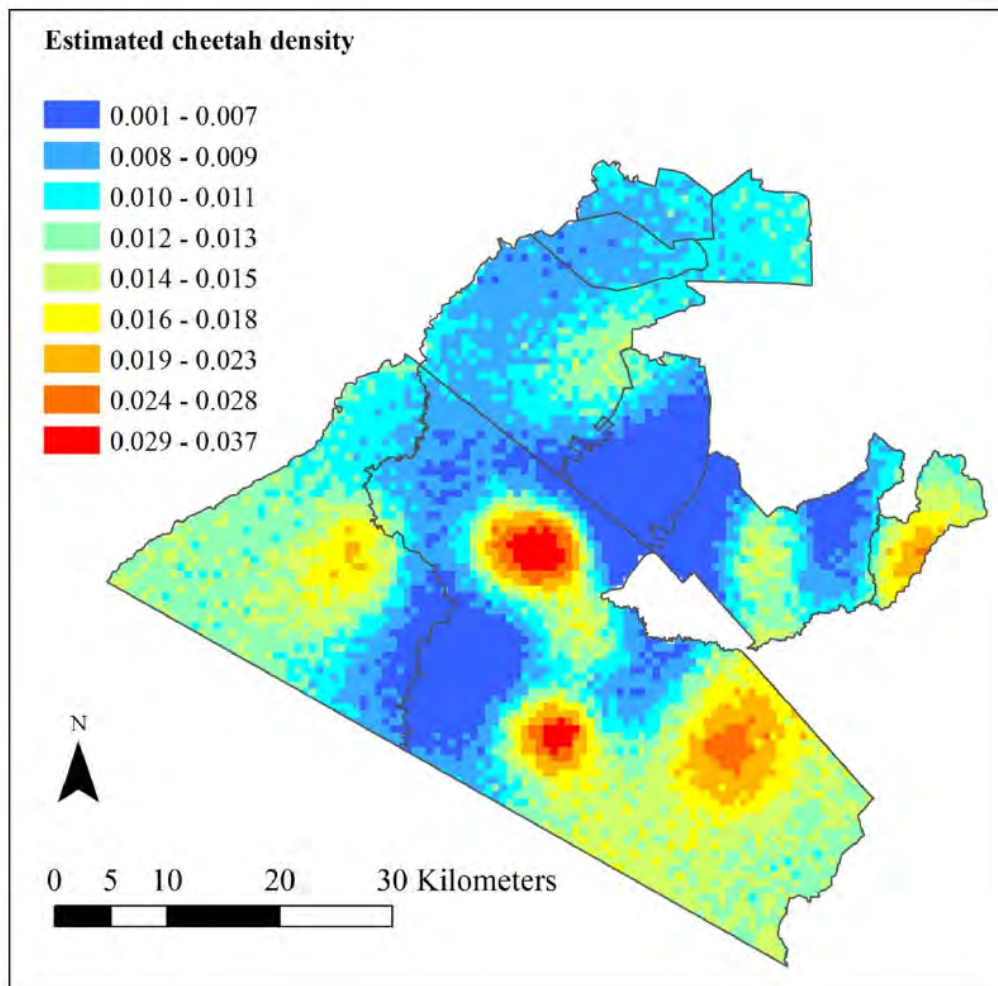
Here our method for the calculation of the absorption coefficient was obtained by sampling the nuclear motion using a stochastic approach and a simulation cell of 1024 atoms. The large simulation cell is necessary in order to accommodate all atoms vibrating with different phases and retain the periodicity.

Why?: e.g. LEDs, Solar Cells

Cheetah Conservation


F.Broekhuis, A.Gopalaswamy

Many ecological theories and species conservation programmes rely on accurate estimates of population density. Accurate density estimation, especially for species facing rapid declines, requires the application of rigorous field and analytical methods. However, obtaining accurate density estimates of carnivores can be challenging as carnivores naturally exist at relatively low densities and are often elusive and wide-ranging. In this study, we employ an unstructured spatial sampling field design along with a Bayesian sex-specific spatially explicit capture-recapture (SECR) analysis, to provide the first rigorous population density estimates of cheetahs (*Acinonyx jubatus*) in the Maasai Mara, Kenya



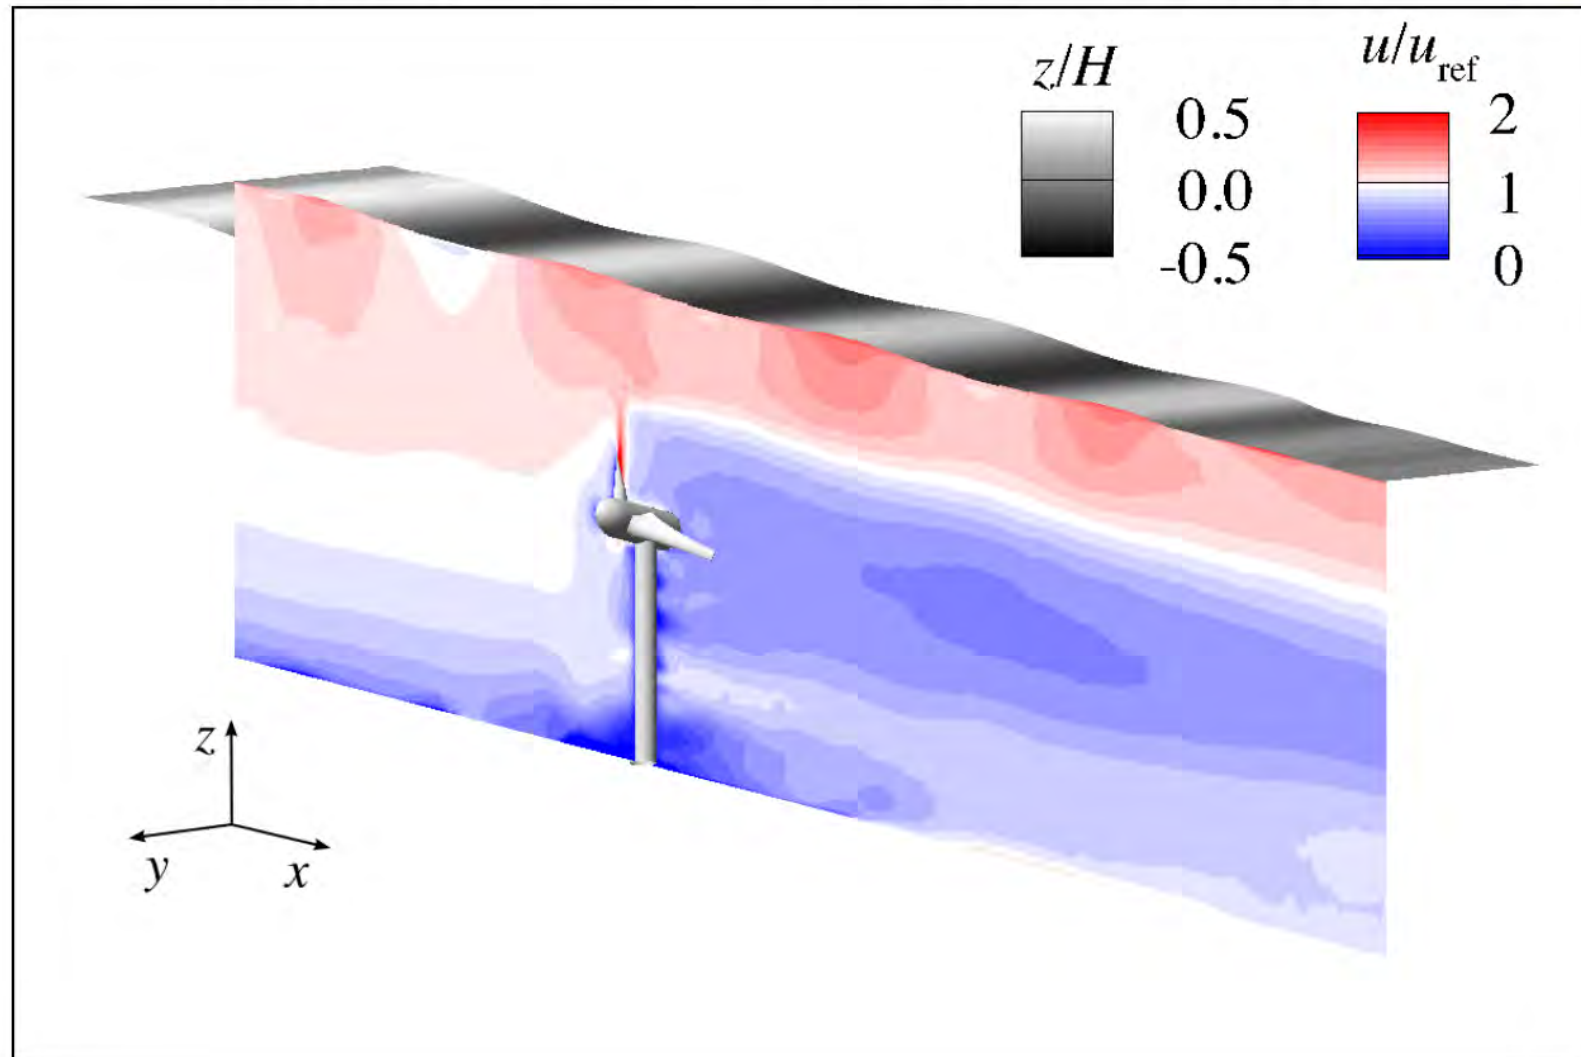
Chemistry: Molecular Design

– M. Peeks, H. Anderson



Computational chemistry has played an important role in both the design of new molecules and the understanding of spectroscopic results. Most recently, we synthesised a molecular nanoring composed of five porphyrins built on a molecular scaffold: a template. This and similar ring-shaped arrays of porphyrins resemble the light-harvesting systems found in the photosynthetic centres of plants and some bacteria. Density functional theory (DFT) calculations helped us to understand how well the template fits into the ring, and so the effectiveness of the nanoring assembly.

Engineering: Tidal Turbine Modelling – C.Fleming, R.Willden



Machine Learning: Kyle Zheng



Original image (hover to highlight segmented parts)




Semantic segmentation

Objects appearing in the image:



Biochemistry: Enzyme Simulations – N.Willems



My research mainly involves investigating how lipase enzymes interact with different surfaces using molecular dynamics simulations. ARC services provide substantial computational resources that enable this research to be performed in a high throughput manner, including substantial sampling of individual systems in order to make robust conclusions from the simulations. This has helped in gaining a better understanding of how interfaces of different surface properties, such as dynamic lipid bilayers and static, planar graphene surfaces, affect how enzymes bind to and interact with these interfaces. Our results include identifying different binding modes in function of the interface in question, providing insight into factors that generate enzyme-surface interactions. This would not be possible with ARC as a free computational resource for the University.

PENGUINS!: G.Clucas



To infer past population dynamics from contemporary gene sequences we rely on Bayesian coalescent inference with the Extended Bayesian Skyline model (implemented in BEAST2.0). This uses Markov Chain Monte Carlo (MCMC) sampling procedures to estimate a posterior distribution of effective population sizes through time, given a sample of gene sequences and a specified nucleotide-substitution model. This method co-estimates the evolutionary rate, substitution model parameters, phylogeny and ancestral population dynamics for multiple partitions (genes) in one analysis and so it is necessary to run the MCMC for many millions of generations and to do multiple runs to check for consistent convergence on the same parameter values. This would not have been possible without the ARC facilities as run-times would have been prohibitive. The results from this were published in the journal *Global Change Biology* and received quite a bit of media coverage as we found that (quite surprisingly) **only three small populations of emperor penguins survived the last ice age**, and so conditions back then appeared too harsh for large numbers of them, despite their adaptations to an ice environment.

HPC in Oxford

- Other groups, departments doing HPC
 - HPC Special Interest Group in Oxford
- Co-investment
 - Almost a third of our current hardware has been purchased by co-investment



Data Centre Space

- Power density
 - Half populated racks
- Options?
 - Move to a new data centre
 - Within Oxford?
 - External DC?
 - Upgrade the current DC
 - Improve A/C, contained aisle?
 - Water cooled options?



Questions?

