Numerical algorithms for high-performance computational science

DISCUSSION MEETING

Organised by Professor Nicholas Higham FRS, Laura Grigori and Professor Jack Dongarra.

8 – 9 April 2019

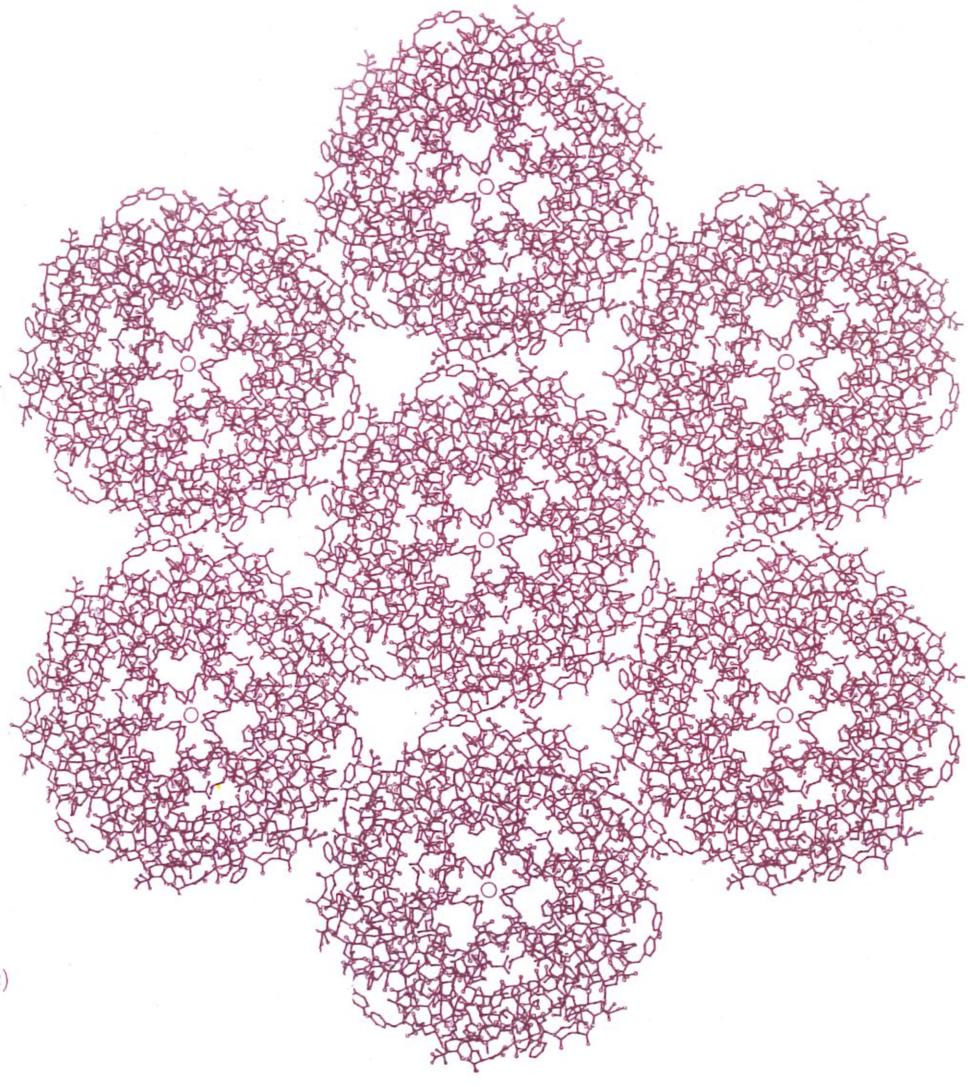

(a)

THE ROYAL SOCIETY

This meeting is part of the Royal Society scientific programme – connecting scientists from around the world in discussions which influence their field and inspire future research opportunities.

Cover image

Molecular diagram captioned 'The atoms in 42 2Zn insulin molecules (excluding hydrogen) projected along the crystal threefold axis'. The image shows the six insulin molecules in hexameter form round two zinc ions. The whole forms a rhombohedral unit cell of crystal.

Figure 2.1 (a) from plate 1 in the paper *The structure of 2Zn pig insulin crystals at 1.5 Å [Angstrom] resolution*, by Edward N. Baker, Thomas L. Blundell FRS et al [including Dorothy Crowfoot Hodgkin FRS], *Philosophical Transactions of the Royal Society B* 319 (1988) pp.369-456.

dimensions up to 1000 with very high precision, eventually allowing physicists to conjecture analytic formulae. Monte Carlo methods are not up to the challenge, while the tensor interpolation algorithm computes integrals to 100 decimal digits.

20. Cellular approach and accurate large scale environment simulation

Thao Thu My Truong[1,2], Vincent Rodin[2] and Bernard Pottier[2]

[1]Faculty of Information Technology, Vinh Long University of Technology Education (VLUTE), Vietnam
[2]Faculty of Sciences/Lab-STICC, Université de Bretangne Occidentale(UBO), France

Climate change impact on environment aspects: sea level, weather patterns, ecosystems. Modelling and simulation allow to monitor changes and mitigate consequences together with sensors, radars, and radio systems. One approach for simulation is to describe the environment into cell systems having geographic coordinates. Data from different sources can be provided to cell logic. Following the Cellular Automata paradigm, cells communicate by links representing physical dependencies. Examples are rain flooding, long range radio, insect swarms and city noise. The simulation model is designed from maps and translated into parallel programs for process systems (1k cells) or accelerators (100k cells). Results were obtained for real applications:

- Distributed swarm observatory as a network of automatic light traps in Mekong Delta, Vietnam.
- Tool to compute physical propagation on complex terrain, for radio, and rain.

The present work is to optimise accuracy and performance system in complex terrains. A test application is Harmful Algae Blooms monitoring. Processing is achieved by distributing data slices on processors maintaining margin dependencies along the simulation steps, removing interesting cells. The authors expect to simulate HABs on sensible parts of the Great Lakes. The transition rules will be compared to a theoretical model, and an original sensor capabilities.

21. LUTNet: Rethinking Inference in FPGA Soft Logic

Erwei Wang, James J Davis, Peter YK Cheung and George A Constantinides

Imperial College London, UK

Research has shown that deep neural networks contain significant redundancy, and that high classification accuracies can be achieved even when weights and activations are quantised down to binary values. FPGAs are devices featuring large quantities of programmable Boolean lookup tables (LUTs) interconnected by programmable routing, providing the potential to build customised architectures for a given application. Recently, FPGAs are gaining traction as an implementation medium for deep neural networks, especially those featuring such extreme quantisation. Network binarisation on FPGAs greatly increases area efficiency by replacing resource-hungry multipliers with lightweight XNOR gates. However, an FPGA's fundamental building block, the K-LUT, is capable of implementing far more than an XNOR: it can perform any K-input Boolean operation. Inspired by this observation, the authors propose LUTNet, an end-to-end hardware-software framework for the construction of area-efficient FPGA-based neural network accelerators using the native LUTs as inference operators. They demonstrate that the exploitation of LUT flexibility allows for far heavier pruning than possible in prior works, resulting in significant area savings while achieving comparable accuracy. Against the state-of-the-art binarised neural network implementation, they achieve twice the area efficiency for several standard network models when inferencing popular datasets.