HPC Applications Research at (or around) The Australian National University

Peter Strazdins
Computer Systems Group,
Research School of Computer Science,
The Australian National University

Intel Corporation
Santa Clara, 08 May 2013

(slides available from http://cs.anu.edu.au/~Peter.Strazdins/seminars)

THE AUSTRALIAN NATIONAL UNIVERSITY

1 Overview

- at the Research School of Computer Science:
  - MetUM global atmospheric model
  - Chaste cardiac simulation
  - ANUGA tsunami propagation
- at CSIRO Advanced Scientific Computing:
  - Self-Orienting Maps
- at the (Australian) NCI National Facility:
  - Landsat image processing
  - accelerating RAID in the Linux kernel
  - accelerating the OpenStack environment
- conclusions

2 The UK MetOffice Unified Model (MetUM)

- the Met Office Unified Model (MetUM, or just UM) is a (global) atmospheric model developed by the UK Met Office from early '90s
- extensively used in Australia (and elsewhere)
  - BoM: weather predictions using N320L70 ($640 \times 481 \times 70$) grid
  - ACCESS project: climate simulations, typically use a N216L85 grids
  - both wish to double resolution to increase 'skill' but is $\leq 8\times$ the work!

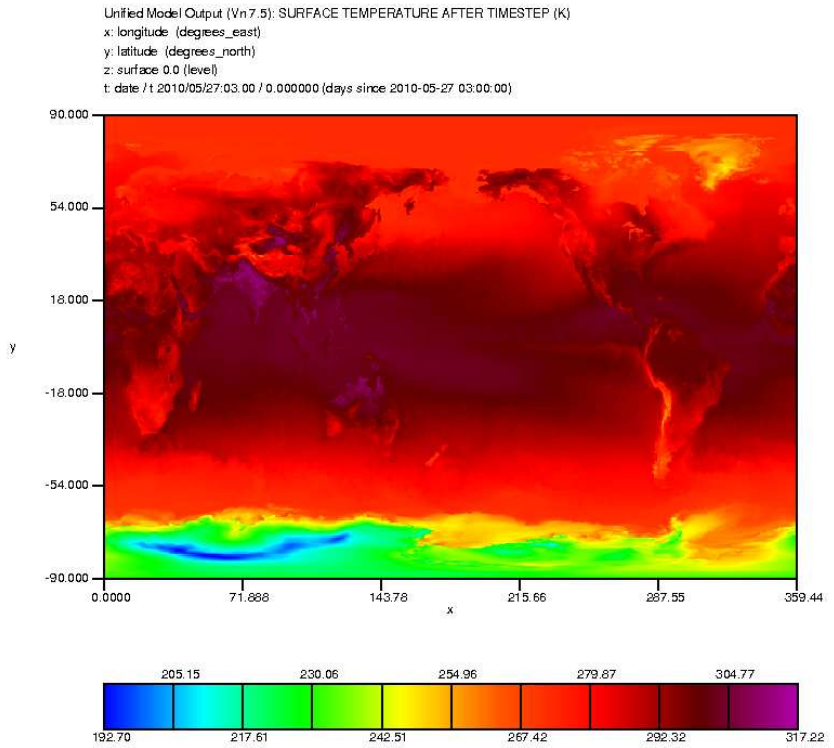

3 Unified Model Code Structure

- the software:
  - main input file is a 'dump' of initial atmospheric state (1.5GB for N320L70)
  - 'namelist' files for ≈ 1000 run-time settable parameters
  - in operational runs, periodically records statistics via the STASH sub-system
  - partition evenly the EW & NS dimensions of the atmosphere grid on a $P \times Q$ (MPI) process grid
  - codes in Fortran-90 (mostly F77; ≈ 900 KLOC) with cpp (include common blocks, commonly used parameter sub-lists, etc)
  - dominated by Helmholtz P-T solver (GCR on a tridiagonal linear system, coded in F90)
- work done: scalability analysis with in-depth profiling up to 2K cores
- todo: accelerating the computationally intense components

4 The Chaste Cardiac Simulation Project

- Chaste: software modelling the electro-mechanical properties of the heart (large system of C++ code, many dependencies)
  - required resolution necessitates parallelization via MPI
  - most computationally-intensive part is solution of a large sparse linear system once per timestep
  - workload uses a high resolution rabbit heart (Oxford University) (2 $\times$ 1 GB files – 4 million nodes, 24 million elements)
  - work done: in depth scalability analysis & optimize for 2K cores
  - current work (Ashley Valent): acceleration in CUDA & OpenACC (mainly at the PETSc level)

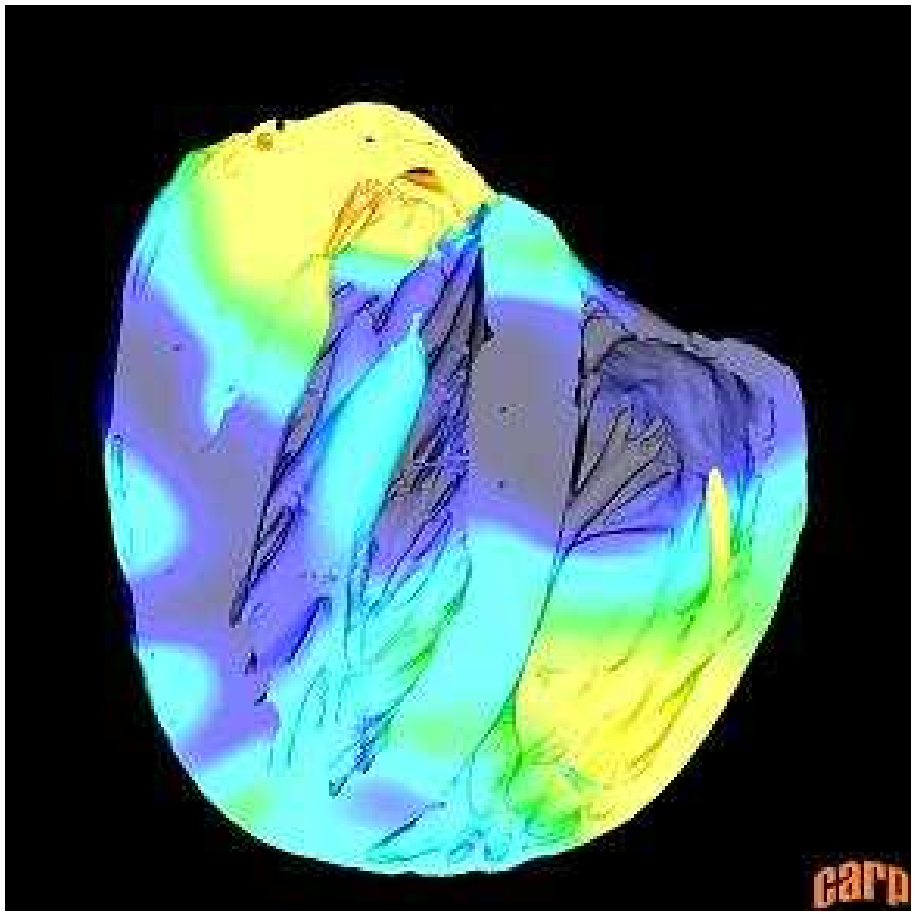

5 ANUGA Tsunami Propagation and Inundation Modelling

- website: ANUGA; open source: Python, C and MPI
- shallow water wave equation, takes into account friction & bed elevation
  - 2D triangles of variable size according to topography and interest
  - time step determined by triangle size and wave speed
- simulation on 40M cell Tohoku tsunami: super-linear speedup to 512 cores on Kei computer

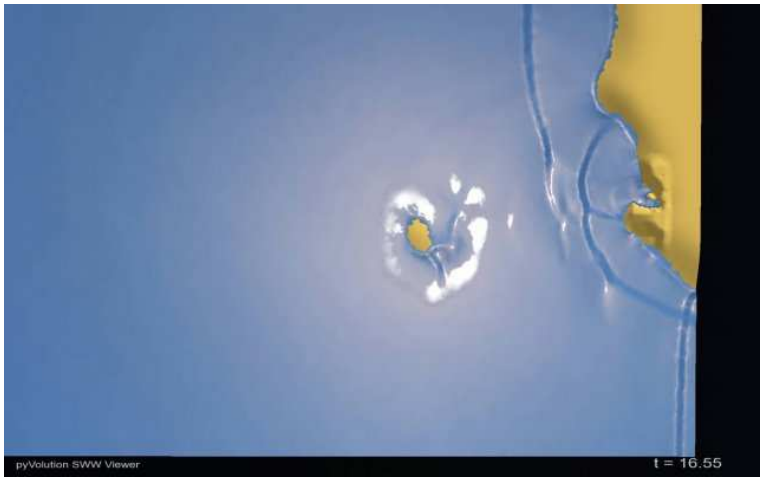

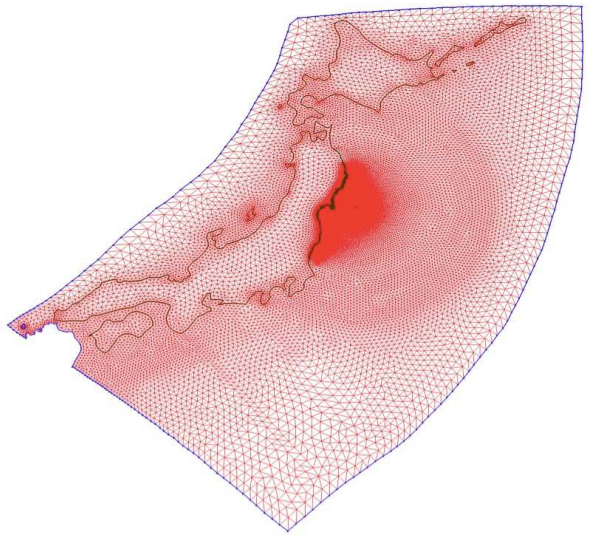

- current work (Zhe Weng): acceleration in CUDA and OpenHMPP

6 SOMA: Self-organizing Maps, Accelerated (CSIRO)

- recent work by James Barker, Advanced Scientific Computing, CSIRO
- SOM: 'train' grid points in a fixed topology toward a data set
- initialize weights on points using 1st 2 principal components
- then iteratively train the weights ($O(n^2)$)

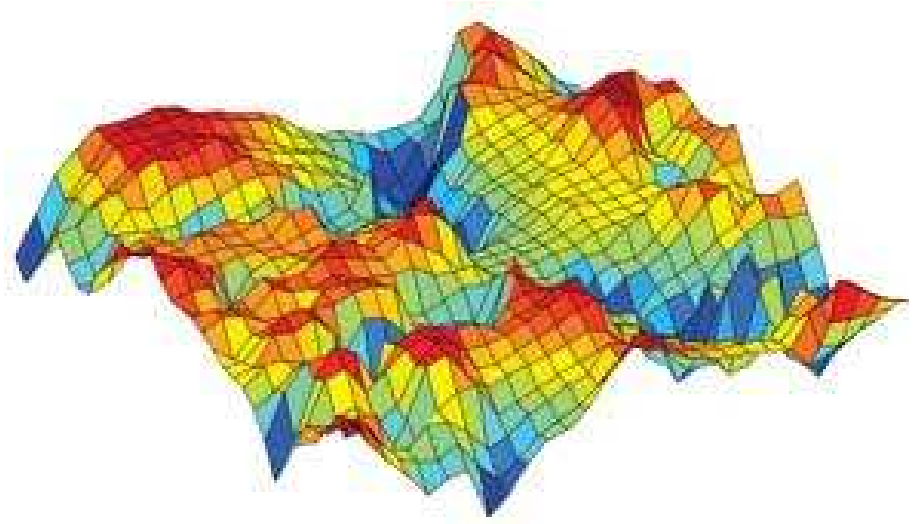

- simple 30-line loop, easily parallelized via OpenMP
- after loop tiling (+ cache alignment) and manual vectorization, got a speedup of 110 with 244 threads on a 61 core Phi
- loop is now 800 lines!
- this work was recently presented in a lecture in our Parallel Systems course

7 Applications & Middleware at the NCI NF

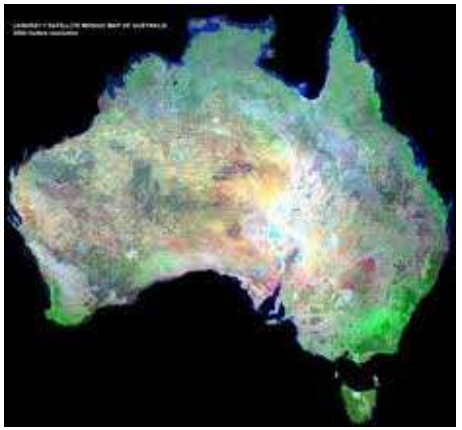

- contact: Joseph Antony, NCI National Facility
- Landsat image processing for the entire Australian continent:
  - NCI has a project with GA and Lockheed-Martin in this area
  - bottlenecks are CPU and IO constraints
- storage applications:
  - accelerating (GPGPU) RAID in the Linux kernel
  - longer term: ZFS backend for Lustre, accelerating RAIDZ
  - is it possible to create Lustre/storage building blocks using Phis + Atom + LSI SAS?
    - similar to Intel's storage hardware group's reference building blocks
- acceleration of OpenStack and HPC Clouds
  - issues with device virtualization under KVM/Xen

8 Conclusions

- ANU's School of Computer Science is closely linked with CSIRO Advanced Scientific Computing and NCI National Facility
- acceleration of many HPC applications are of interest for us. Also includes:
  - machine-learning apps, fast multipole method, Gaussian chemistry kernels