Enhancing eResearch productivity with NeSI's consultancy service

Alexander Pletzer (NeSI/NIWA), Chris Scott (NeSI/U Auckland),
Wolfgang Hayek (NIWA/NeSI), Kim Frew (NeSI/U Auckland), Georgina Rae (NeSI/U Auckland)
and Jana Makar (NeSI/U Auckland)
(alexander.pletzer@nesi.org.nz)
eResearch NZ 2020 – Dunedin 12-14 February 2020

New Zealand eScience Infrastructure

Consultancy

- Analysis, debug and optimization of user applications

NeSI
New Zealand eScience Infrastructure

Support

- Expert knowledge in multiple domains

NeSI
New Zealand eScience Infrastructure

Training

- Software Carpentry / Data Carpentry
- Intro & advanced HPC training

Data transfer

- high speed data input/output Partnership with Globus (global data management platform)

Hardware and software for compute and analysis

- ~700 compute nodes
- hundreds of software packages

NeSI's work horses

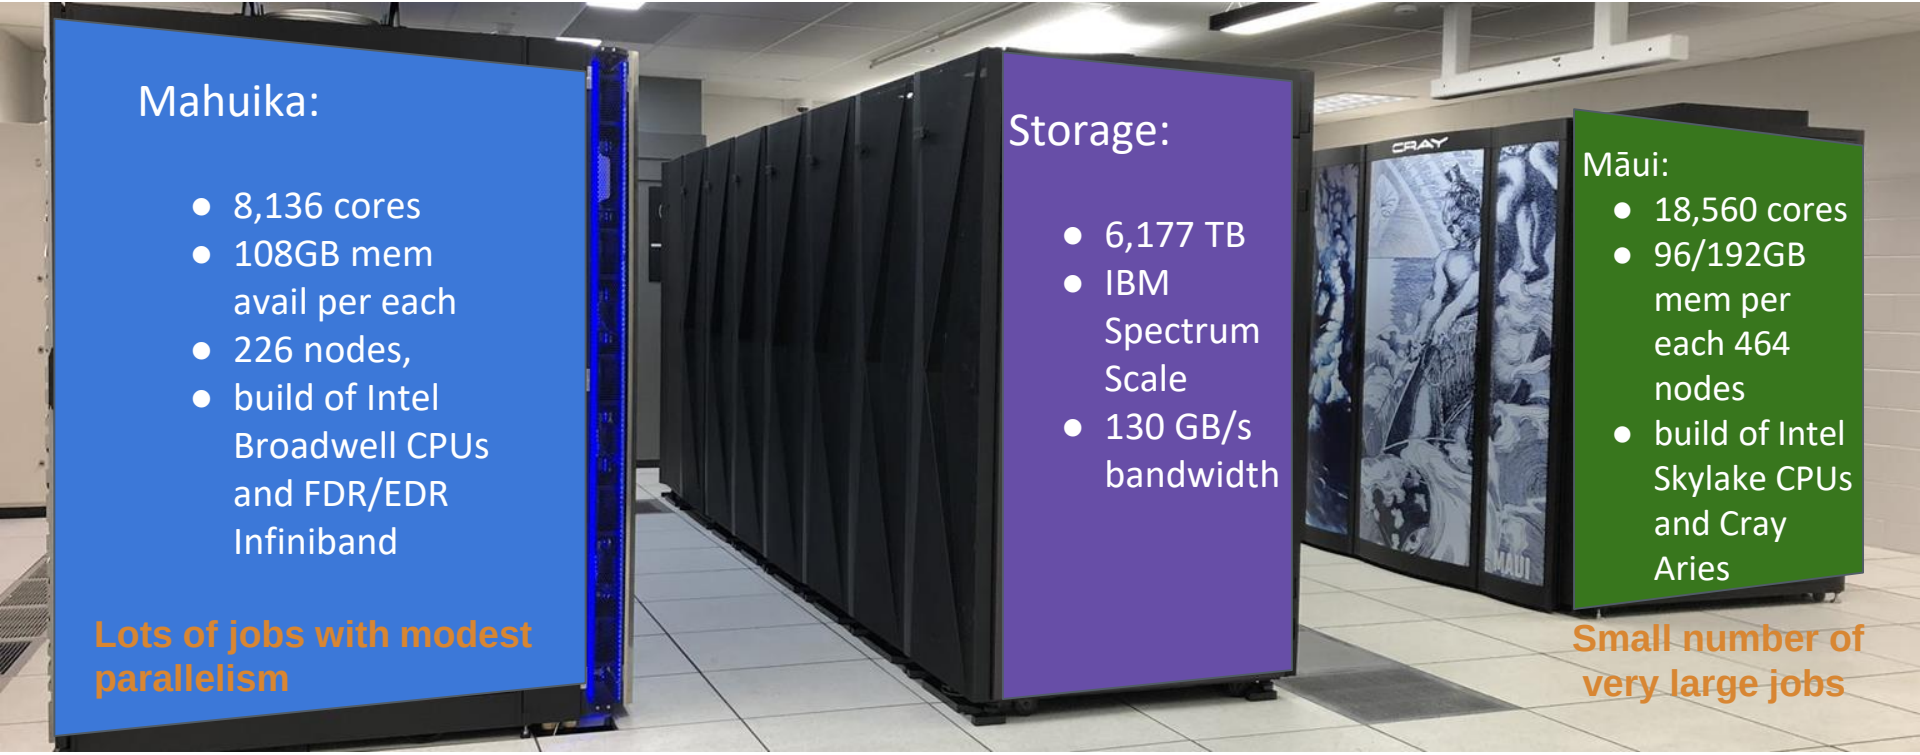

Robin Bensley
Business Operations Manager,
University of Auckland

Blair Bethwaite
Solutions Manager,
University of Auckland

Thomas Berger
Product Manager,
University of Auckland

Fabrice Cantos
HPC Operations Manager,
NIWA

Laura Casimiro
Operations Coordinator,
University of Auckland

Brian Flaherty
Data Services Product Manager,
University of Auckland

Kim Frew
Science Engagement Manager,
University of Auckland

Megan Guidry
Research Communities Advisor,
University of Auckland

Greg Hall
Systems Engineer,
University of Auckland

Yuriy Halytskyy
Systems Engineer,
University of Auckland

Wolfgang Hayek
Scientific Programmer,
NIWA

Matt Healey
Application Support Specialist,
University of Otago

Aaron Hicks
Systems Engineer,
NIWA

Jose Higino
Systems Engineer,
NIWA

Jun Huh
Business Innovation
and Growth Manager,
University of Auckland

Nick Jones
Director,
University of Auckland

Marko Laban
Software Product
Engineering Lead,
University of Auckland

Nancy Lin
Data Analyst,
University of Auckland

Nooriyah Lohani
Research Communities Advisor,
University of Auckland

Jana Makar
Communications Manager,
University of Auckland

Peter Maxwell
Application Support Specialist,
University of Auckland

Alexander Pletzer
Scientific Programmer,
NIWA

Nitharsan Puwanendran
Analyst Programmer,
University of Auckland

Georgina Rae
Engagement Manager,
University of Auckland

Kumaresh Rajalingam
Analyst Programmer,
University of Auckland

Ben Roberts
Application Support Specialist,
Manaaki Whenua –
Landcare Research

Albert Savary
Application Support Specialist,
University of Otago

Chris Scott
Scientific Programmer,
University of Auckland

Dinindu Senanayake
Genomics Support Specialist,
University of Auckland

Anthony Shaw
Application Support Analyst,
University of Auckland

Nick Spencer
Site Manager
Manaaki Whenua –
Landcare Research

Callum Walley
Application Support Analyst,
University of Auckland

Damian Wheeler
Site Manager,
University of Otago

Jeff Zais
Senior Science Advisor &
Platforms Architect,
NIWA

Lifting the computational capabilities of New Zealand researchers

Why NeSI's consultancy programme

eResearch productivity not all about hardware:

- **Tools:** profiling, visualization, …
- **Algorithms**. Complexity is a defining factor
- **Implementation**. Understand how computer works
- **Compilers**. Some produce more efficient code
- **Environment settings**. How you distribute threads across CPUs can impact performance

Aim is to enhance eResearch output and ensure that NeSI's platforms are used in the most efficient way

Understanding the microbial world of marine

- The problem

  Alexis Marshall needs to analyse and reconstruct 1.4 billion nucleotide fragments obtained from ocean sediment, to identify their species and gain a better understanding of how microbial life cycles nutrients but wasn't sure how to take full advantage of the parallel power of Mahuika

- The outcome

  NeSI modified the workflow to spread the work across the cluster, running in parallel across many nodes and **reducing run times from weeks to days**.

Simulating a mega-thrust earthquake in New Zealand

- The problem

  Yoshihiro Kaneko and Bryant Chow are developing a better crust deformation model for the Hikurangi fault that can reproduce the full spectrum of earthquake slip behaviours

- The outcome

  Leveraging vectorization and OpenMP on top of MPI saves 30% of execution time. Will potentially **save 100,000s of core hours**

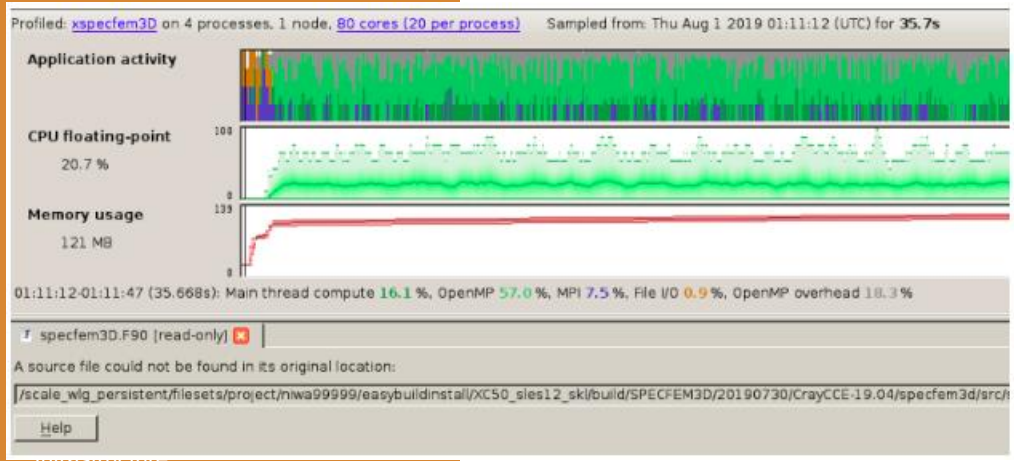

SPECFEM3D profile

- 21% spent in floating point instructions
- 57% in OpenMP parallel loops

Identifying whales by their songs

- The problem

  Estimate the number of marine mammals in NZ waters

- The outcome

  Giacomo Giorli collected hundreds of underwater spectrogram recordings. With NeSI's help, he created a neural network able to **classify > 95 percent accuracy**

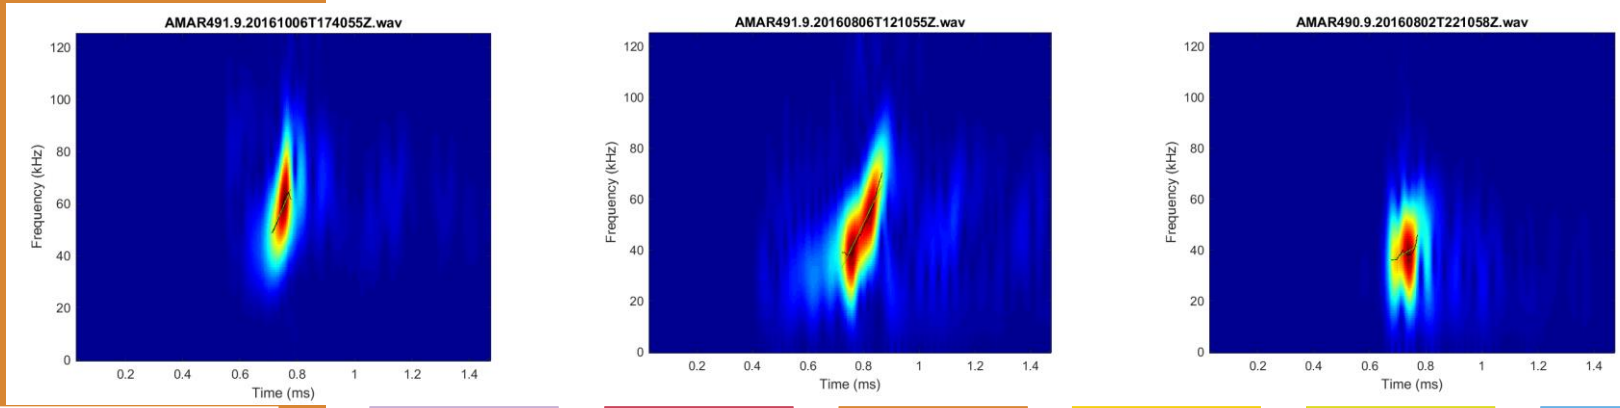

Predicting floods and storm surges

- The problem

Cyprien Bosserelle is developing a fast code for modelling inundation from tsunamis, rivers, … The model needed an efficient method for producing output without slowing down

- The outcome

Using the ParaView Catalyst library for in-situ visualisation, the code can now produce output while it is running. **Users can connect to the simulation for real-time output analysis**

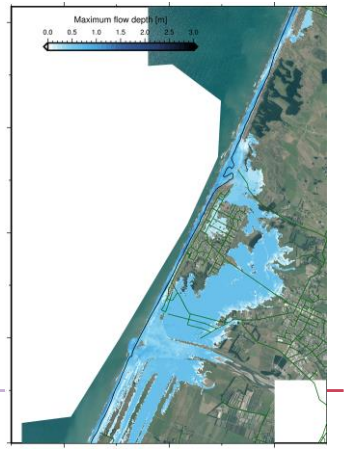

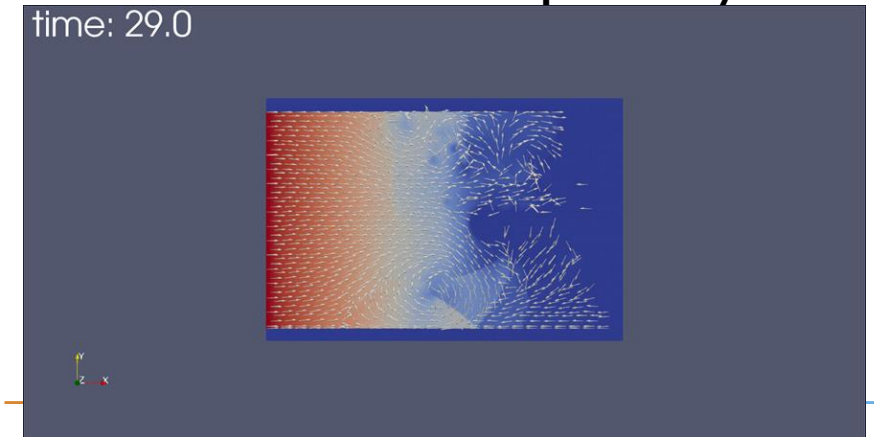

Simulating fluvial erosion and transport of sediments

- The problem

  The GRATE river sediment transport program, written by Jon Tunnicliffe, needed updating to make it more efficient and adaptable. This code has many applications, including modelling the effects of river gravel extraction, dam construction and post-earthquake impacts on rivers (one river near Kaikoura has over 30 metres of sediment piled up within the confining valley).

- The solution

  NeSI created a version of GRATE that can run on its cluster, **enabling large parameter sweeps to run in much shorter times**. We also worked on the code structure: adding tests, **continuous integration**, automatic build and deploy of releases, making GRATE much more robust and accessible to others.

Lessons learned

We found

- Large opportunities often exist to improve the performance of workflows and programs on NeSI platforms
- Good programming practice helps raise productivity of research and facilitate collaboration (source version control, testing, …). It's not only about floating point performance
- Scientific programming is evolving rapidly with new tools, programming languages and techniques being developed. Consultancies service help us understand what users want and adopt. Example: Julia programming language
- Many scientists prefer to focus on science rather than programming. Natural division of labour between researchers and NeSI
- Additional pairs of eyes improve quality of software
- Consultancies are a window allowing NeSI to see how researchers use our platforms

Optimisation pays off

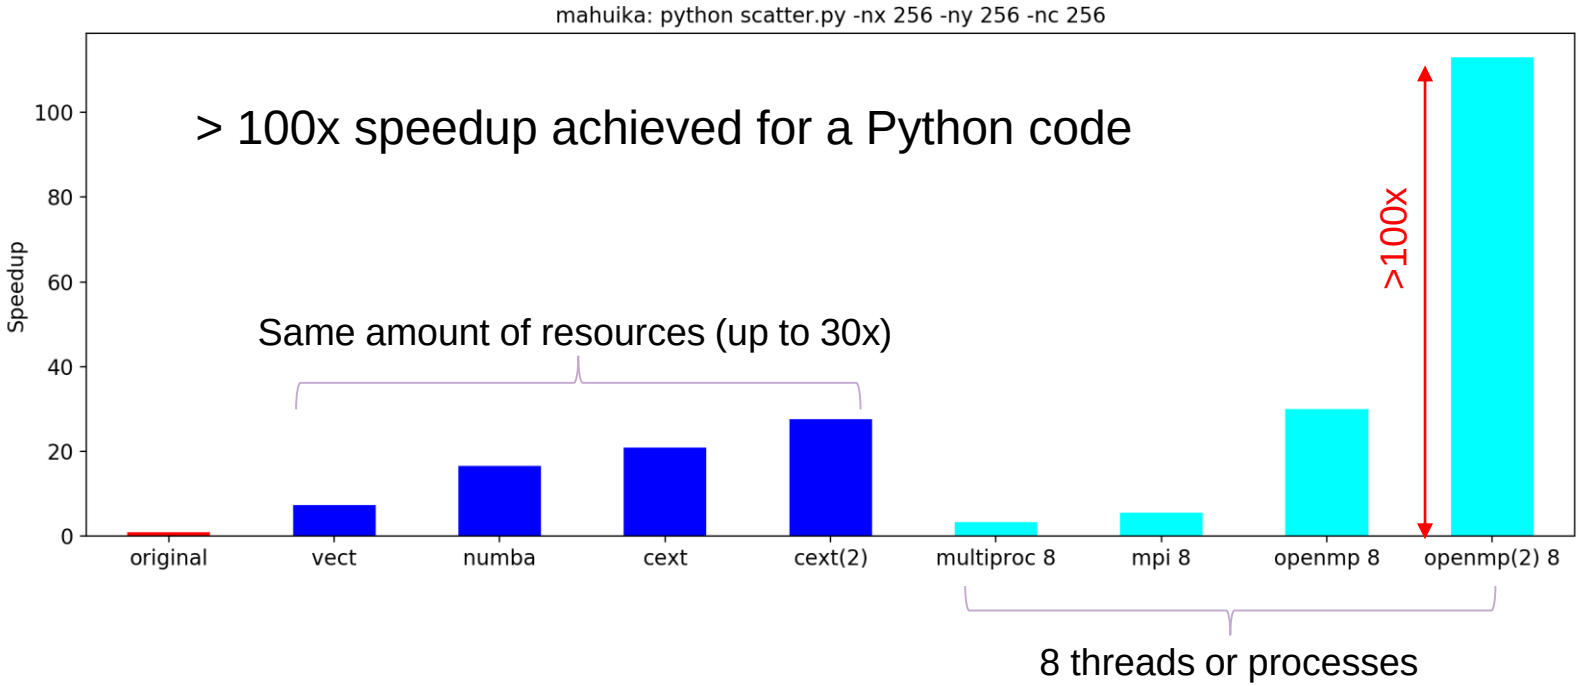

Would your research benefit from a consultancy? contact support@nesi.org.nz

Save the Date:

Science Coding Conference 2020
9 – 11 September 2020
Auckland, NZ

Call for Submissions open soon! Watch http://sciencecodingconference.nz for details

More from NeSI today at eResearch

Nooriyah Lohani

Research Software Engineering Community update and next steps in New Zealand

Glenroy Auditorium Friday 13:30

Thank you