A Grid-Based Facility for Large-Scale Cross-Correlation of Continuous Seismic Data

John Townend[1], Yannik Behr[1], Kevin M. Buckley[2], Martha Savage[1], John Hine[2]

Victoria University of Wellington
New Zealand

Kevin.Buckley@ecs.vuw.ac.nz

http://www.ecs.vuw.ac.nz/~kevin/Conferences/eRA09/foils/

Outline

■ Inside the boundary:

Seismic Cross-Correlation and Ambient Noise Tomography

■ Outside the boundary:

Workflows and Computational Grids

■ So, "No Boundaries" then ?

Inside the boundary: Existing Science 1

Work described here comes from a project funded by the

- NZ Earthquake Commission's Research Foundation (EQC08/566)

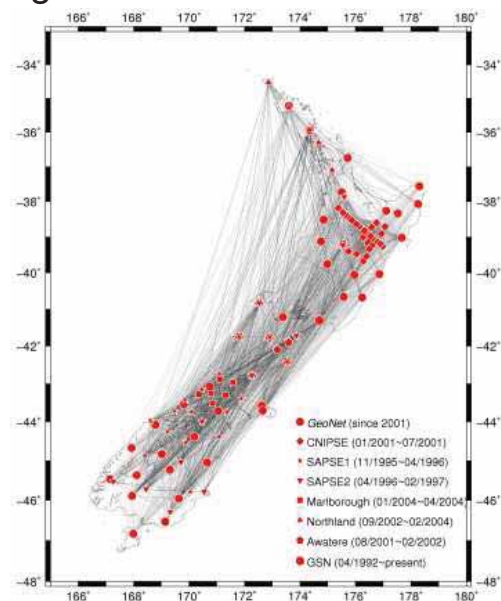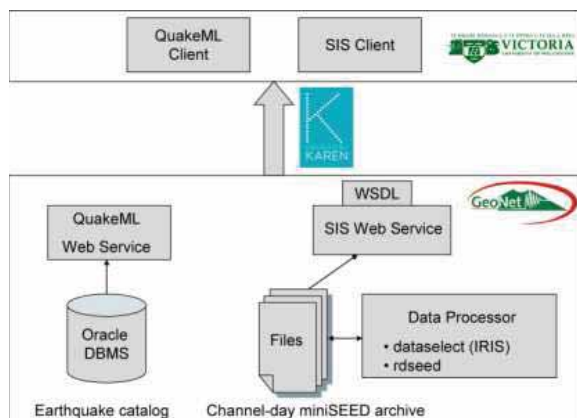and "builds" upon two talks given at eResearch Australasia 2008

- George Slim's (MoRST) overview of where NZ eResearch was heading
- Paul Grimwood's (GNS) Seismographic Information Service project

Also a second strand of work using the same underlying infrastructure
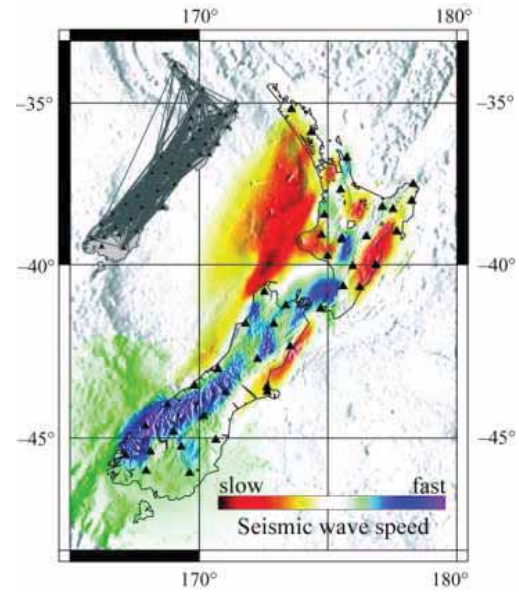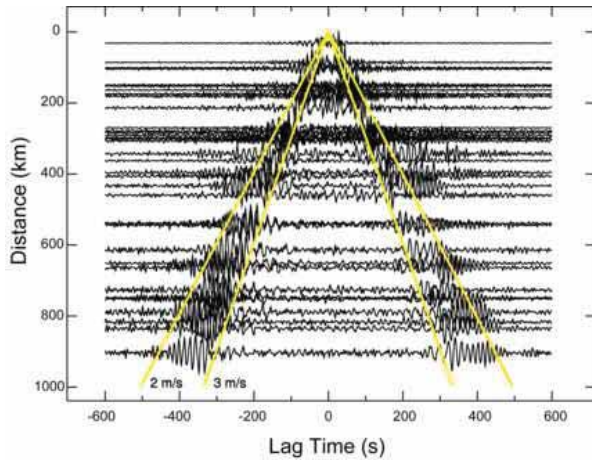that provided a teaching module and some Google Earth/Maps outputs

---

Inside the boundary: Existing Science 2

Inside the boundary: Existing Science 3

Inside the boundary: Existing Science 4

Outside the boundary: Workflows

Workflows provide a mechanism for knowledge transfer about processes within an automated representation of the process as a whole.

Two seemingly opposite drivers at work:

Hiding the raw computation in the sub-processes from the end user, eg
- Interact with a specified web service via a WSDL scraper

Making the sub-processes explicit to the end user
- End user might specify a different web service
- End user might not use a web service at all

Workflow tool-kits come with a lot of pre-defined process units for "doing things"

Does your domain science, let alone your sub-processes, have units in the workflow tool-kit you want to use, that let you do what you want?

Taverna's Bio Roots

eRA09: Kevin Buckley: A Grid-Based Facility for Large-Scale Cross-Correlation of Continuous Seismic Data

Kepler WSDL

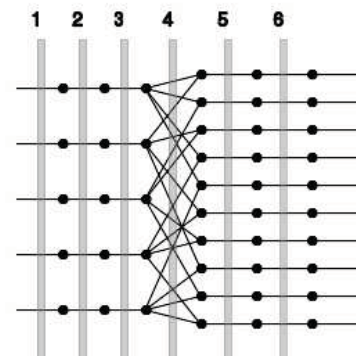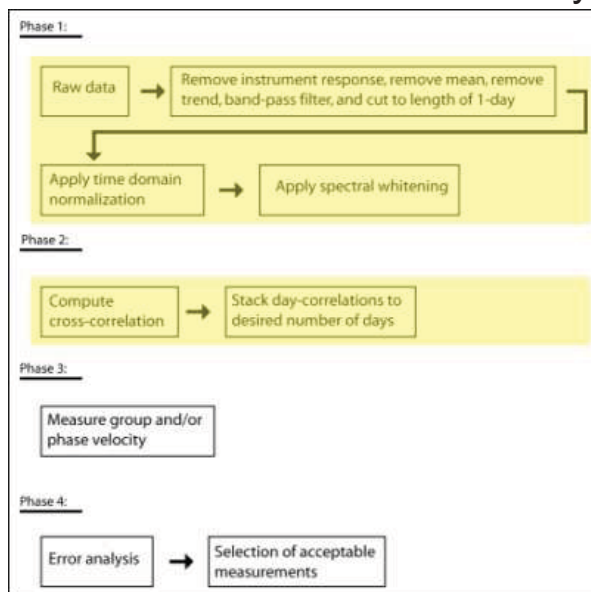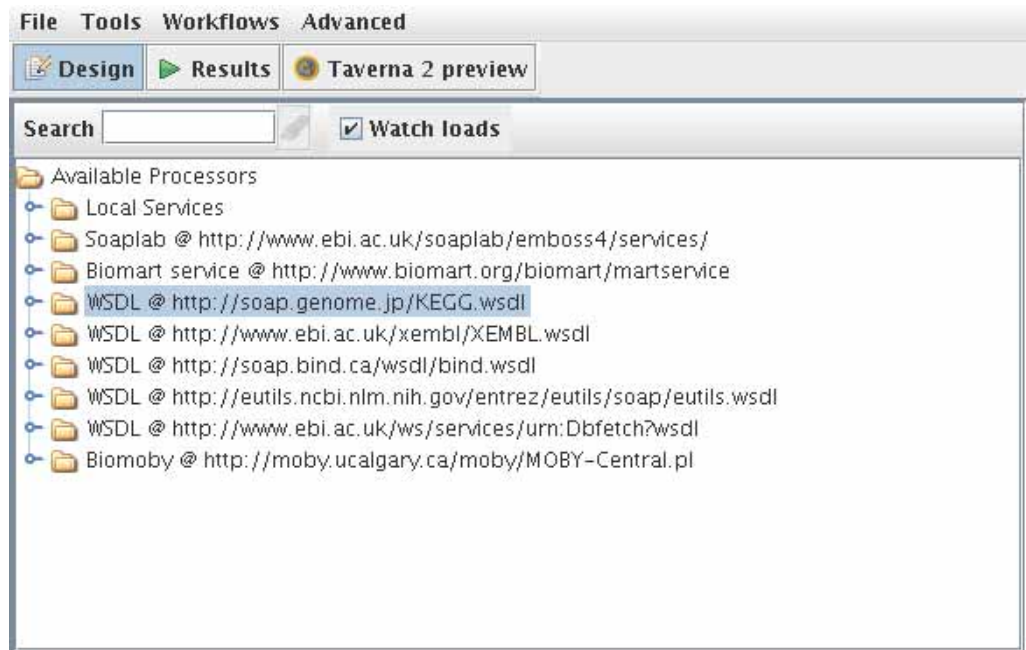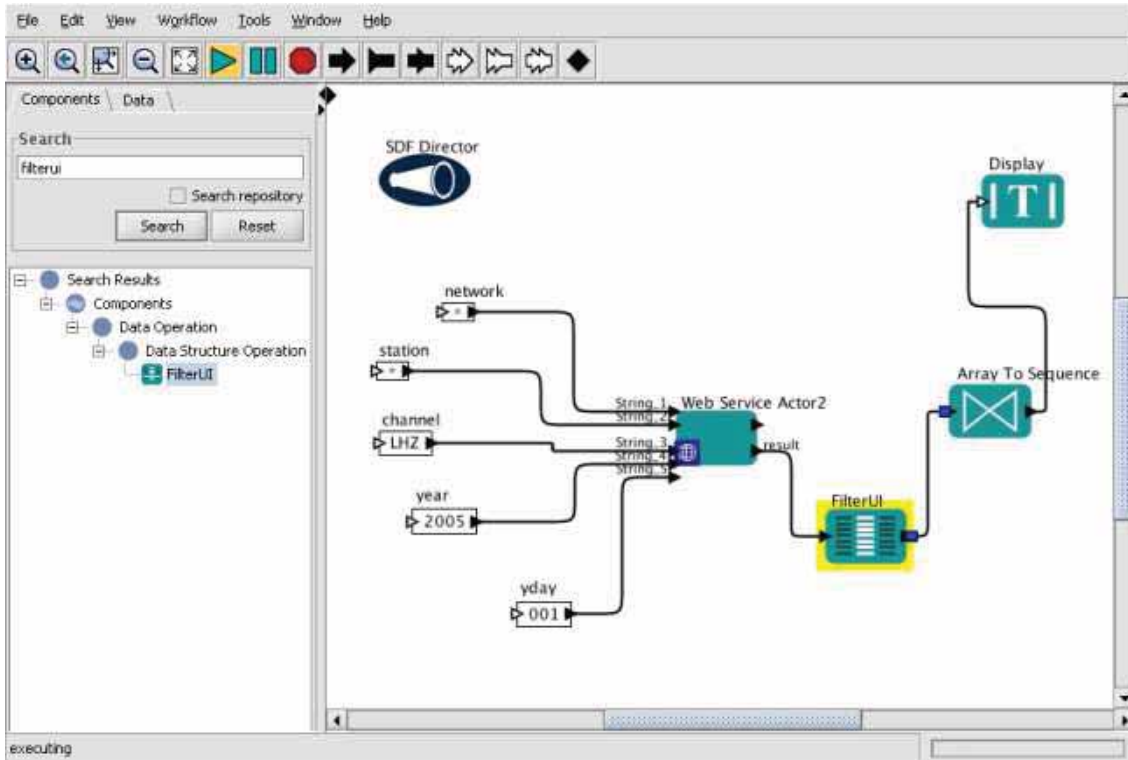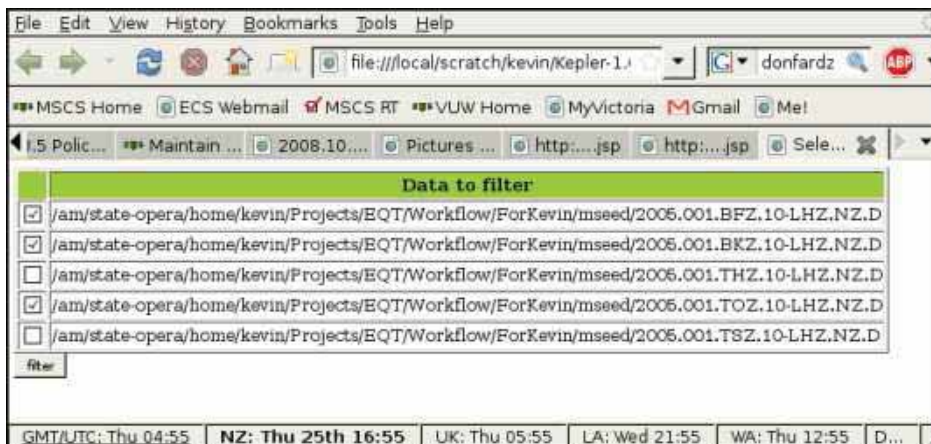eRA09: Kevin Buckley: A Grid-Based Facility for Large-Scale Cross-Correlation of Continuous Seismic Data

Kepler UI Filter

eRA09: Kevin Buckley: A Grid-Based Facility for Large-Scale Cross-Correlation of Continuous Seismic Data
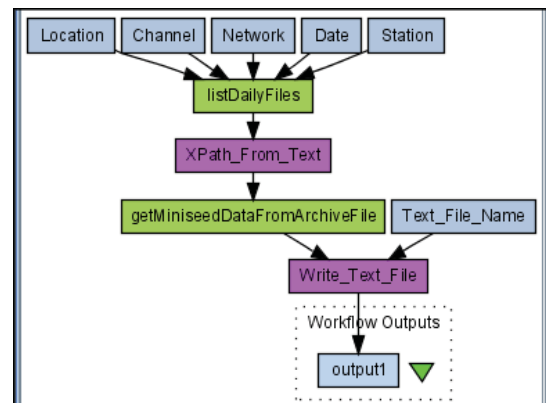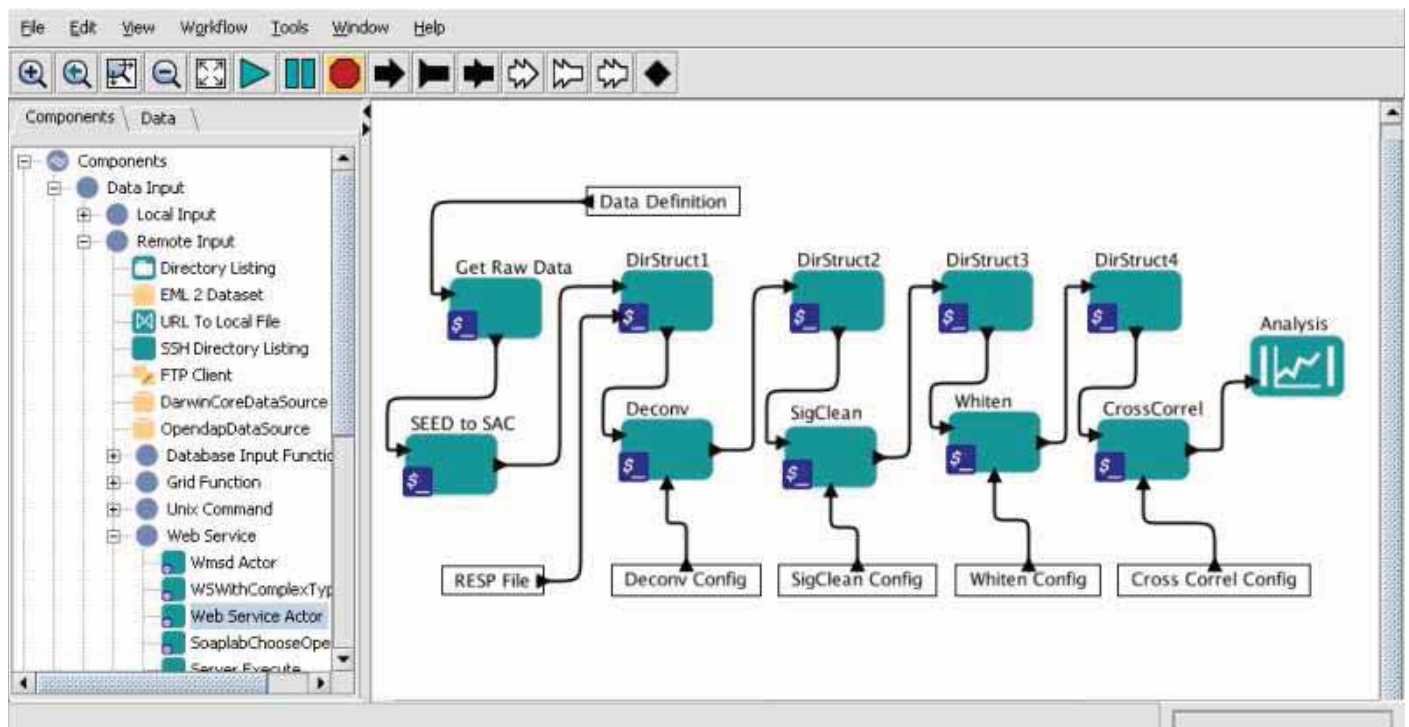
Taverna WSDL

---

eRA09: Kevin Buckley: A Grid-Based Facility for Large-Scale Cross-Correlation of Continuous Seismic Data

Taverna WSDL 2

Taverna WSDL 3

Outside the boundary: Workflow representation

Removing the boundary: Workflows

Once raw data is converted, stage inputs and outputs are the same format - simple pipeline ?

Compare that workflow representation, itself a mere schematic, with the schematics of what the domain scientists think of as what's happening

A lot more stuff (local/domain knowledge) going on, eg
- overcoming directory requirement impedences between the components
- handling errors at various stages

and that "stuff" becomes visible in the workflow tools.

Tempting to say "keep things hidden" - but it's not a package, it's a workflow

Other benefits to being explict

Outside the boundary: Grids

Underlying expectation (and maybe even what's being promised?) from "the grid"

"Hand something over and just have it run faster"

Different people have different ideas and expectations about grids - hardly surprising

Different grids: Compute, Data and Collaboration

Grids vs Clusters?   The Grid as Utility?   The Grid /as a/in the/ Cloud ?

What about "Package Grids" ? Do users even see the grid they are using?

Headline figures can hide the realities:
- 4096 processors at UoC: ~1900 at VUW: 512Mb vs 1Gb (or 2Gb)
- More/Faster processors => smaller run time - did we mention the queuing time?

Now remove the boundary

Our grids
- VUW actually has two totally seperate grids
- Neither are dedicated resources: cycle-stealing grids
- So neither really allow for message-passing approaches

Our science: where were the boundaries for us now?
- Data - constant access to raw data vs storing raw/modified data
- Data - three distinct processes: only need the whole at the end
- Compute - discretisation gives us resilience to failure
- Collaboration - not all users happy at the command line

No boundaries: Grid submission

After all the knowledge transfer, we came up with something as simple as this:
(Stage 2 for 30 stations over three months, rather, 90 days: 90 * 435 x-Crlns)

```
qsub -t 1-39150 \
   -q all.q \
   -wd /vol/grid/sgeusers/emily \
   -M Emily.Alias@vuw.ac.nz \
   -m e \
  /vol/grid/pkg/SGEES/gbfcccsd/bin/stage2-args-001.sh \
   /local/tmp/emily \
   /vol/scratch/emily/EQC/data/SAC \
   /vol/scratch/notemily/EQC/data/miniSEED/2008-000-filelist.txt \
   /vol/scratch/emily/EQC/data/SAC \
   5 100 90
```
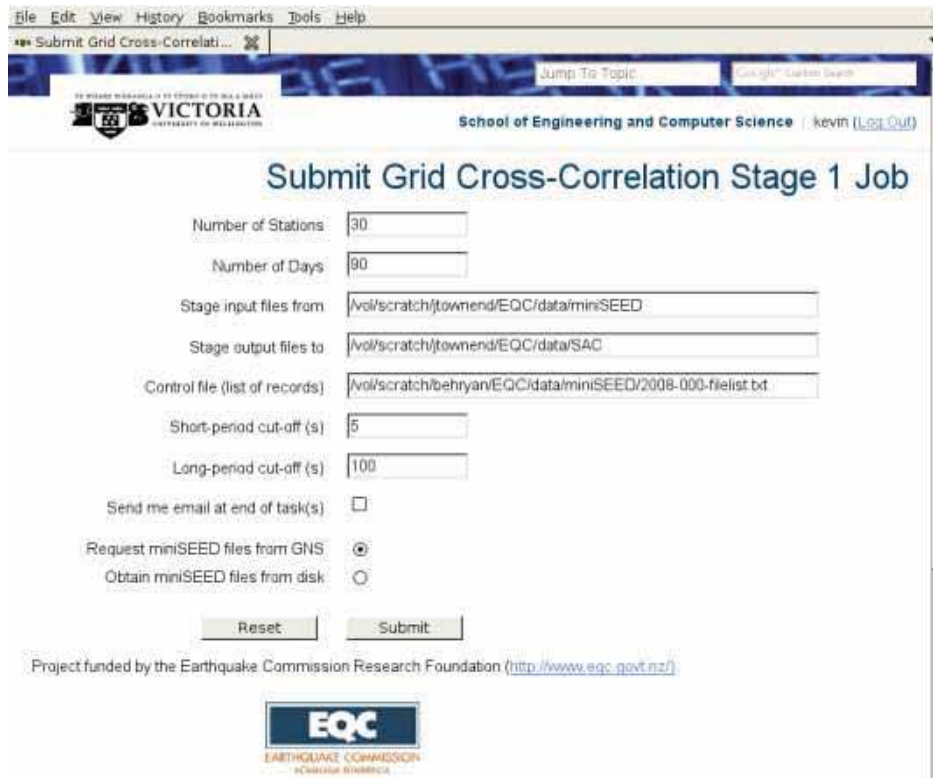
The DQS/SGE aware will notice there's an implied step size (of 1) there

No Boundaries: Grid Submission - Web

Beyond the Boundary

One job was run whilst Yannik was back in Germany - no boundaries indeed !

Reductions of up to 87% (91% and 51%, stage 1 and 2 resp.) in processing times were seen
  (5 stations over 99 days, 1Hz sampling data, 170Mb of disk)

```
(Time in mins)   Desktop   Grid
Stage1 time        34.2     2.9
Stage2 time         5.2     2.2
Total  time        39.4     5.1
```

Managed to get a concurrent usage of 180 (out of ~230) otherwise idle processors
  (10 stations over 360 days, 100Hz sampling data, 114Gb of disk)

```
(Time in hrs)    Desktop   Grid
Stage1 time        ~120      ~3
```

Summary: Seismic Processing

- Moved a desktop process onto the "grid"
- Built in resilience and flexibility across generic grid resources
- Maximised use of local grid resources
- Hidden "the grid" from the user
- Questions raised informed future domain science approaches

As to the future:

- Investigate other approaches upon dedicated grid resources
- Actual implementation of processes in a workflow tool

Summary: Here be boundaries
Domain Science

- Processes operating in a simple, understandable, pipeline
- Become wrapped, for operational simplicity, into one process

Workflow Boundary

- Individual processes become hidden - (package or workflow?)
- Additional operational artefacts not externally visible

Grid Computing Boundary

- No visible loops to parallelise
- Data placement/distribution is hard

Avoiding Boundaries

- Clearly identifiable processes with simple (maybe NO) interfaces between them
- Explicit loops over monolithic processing

And finally ... who is supposed to do this "no-boundary" science ?

Domain scientists ?

- Might need to brush up on their BEPL, SCUFL, WSDL, MPI and DRMAA

Workflow/Grid computer scientists ?

- Might need to read up on SEED, SAC, QuakeML

Phew !

Might still be a few jobs for people who can straddle the boundaries

---

Colophon

Slides for this talk have been created using MagicPoint

http://member.wide.ad.jp/wg/mgp/

An X11 based presentation tool which has the slide sources in plain text
and which also provides for creation of an HTML slide-set.