Lessons Learned from Running Large-Scale Simulations on a Volunteer Platform

Carl Christensen carlc@atm.ox.ac.uk carlgt1@yahoo.com
Chief Software Architect, Quake-Catcher Network
School of Earth Sciences, Stanford University

Senior Consultant, climateprediction.net
Department of Physics, Oxford University

STANFORD UNIVERSITY

Volunteer Computing Challenges

- No control over quality or amount of resources
- Great for brute-force or 'embarassingly parallel' – but can't access entire CPUs & memory parallel a la supercomputer (i.e not appropriate for all problem, can now run MPI jobs locally however), current theoretical max 8 core, 8GB
- Can be risky as far as numbers needed (i.e. "if you build it, they may not come!" – or maybe too many come?) – may need to look into PR/ media; address user attrition issues etc
- Interaction with participants can scare off more socially awkward academics and staff (successful projects have "hams" :-) – the "social aspects" take time & resources

Volunteer Computing Benefits

- Enormous potential for “free” resources from computing enthusiasts and corporate partners (at the “top supercomputer” levels and above, esp with GPU applications e.g. NVidia CUDA)
- Great for public outreach, education, publicity (QCN & CPDN in schools, BBC programmes etc)
- Relatively low cost as compared with paying for supercomputer resources (3-5 computing staff); fraction of a percent compared to supercomputer time
- Can be applied to a wide range of platforms (gaming consoles, graphics cards, all CPU O/S’, 32 & 64-bit) and problems

Case Study 1 – climate*prediction*.net

- ~12 years ago ("Do It Yourself Climate Prediction") pondering the feasibility of large ensemble climate modelling & the SETI@home paradigm (what is now called "volunteer computing")
- ~8 years ago (12th September '03) – CPDN "Classic" (pre-BOINC) launched at the Science Museum in London
- ~7 years ago (August '04) – CPDN merged with the BOINC project and launched our first app…

climate*prediction*.net Challenges

- Climate models (ESM's, AOGCM's etc) are very large, complex systems developed by physicists sometimes over decades (& proprietary in case of UKMO)
- ~1 million lines of Fortran code (HadCM3 - 550 files, 40MB text source code), 2-2.5e6 w/ regional model
- Little documentation (the science is often well documented but not the software and design of the system per se)

climate*prediction*.net Challenges (cont'd)

- Also "utility" code written by various scientists & students over the years (outside of model code, 220 files, 12MB source, 250K lines); often workable but hard to implement on a cross-platform PC project (i.e. python, shell scripts)
- Meant to be run on supercomputers, primarily 64-bit – not designed (or indeed envisioned) to be run on anything other than a supercomputer or at the very least, a Linux cluster

climate*prediction*.net Overview

- Proprietary, licenced by UK MetOffice, distribute executable/binary form only
- Resolution originally used: 2.75x3.75 degrees (73 lat x 96 long), now N144 (HadAM3) also (about 1 degree square) and regional (30km)
- Typically run on a supercomputer (i.e. Cray T3E) or 8-node Linux cluster (minimum)
- Ported to a single-processor, 32-bit Linux box, Original: Windows only, now also Mac OS X, Linux; also doing 64-bit (HadGAM) now (25% of all active users, about 13K machines, are 64-bit)
- Intel Fortran Win & Linux, IBM XLF for Mac, now all Intel compilers (Mac, Linux, Win), also looking into GNU Fortran (new versions can compile the UM)
- As of today (3/2011) about 100 million model-years completed in 1.5 million simulations (since 9/2003)

BBC

climateprediction.net

BBC Climate Change Experiment

- 160 year coupled-model (HadCM3) run
- Promoted as part of the BBC “Climate Chaos” season of programmes & documentaries in ’06
- “Meltdown” documentary featured CPDN and launched the experiment in February of ‘06
- Shown on BBC4 , had two million viewers, so was also shown on BBC 1 (promos during “Eastenders”!)
- Results show featuring David Attenborough broadcast in January of ’07
- Nominated for a BAFTA in “computer interactive” category (lost to Terry Pratchett’s Hogfather)
- http://www.bbc.co.uk/sn/climateexperiment/

climate*prediction*.net Users Worldwide

>300,000 users total (typical 50K at once ~50TF)

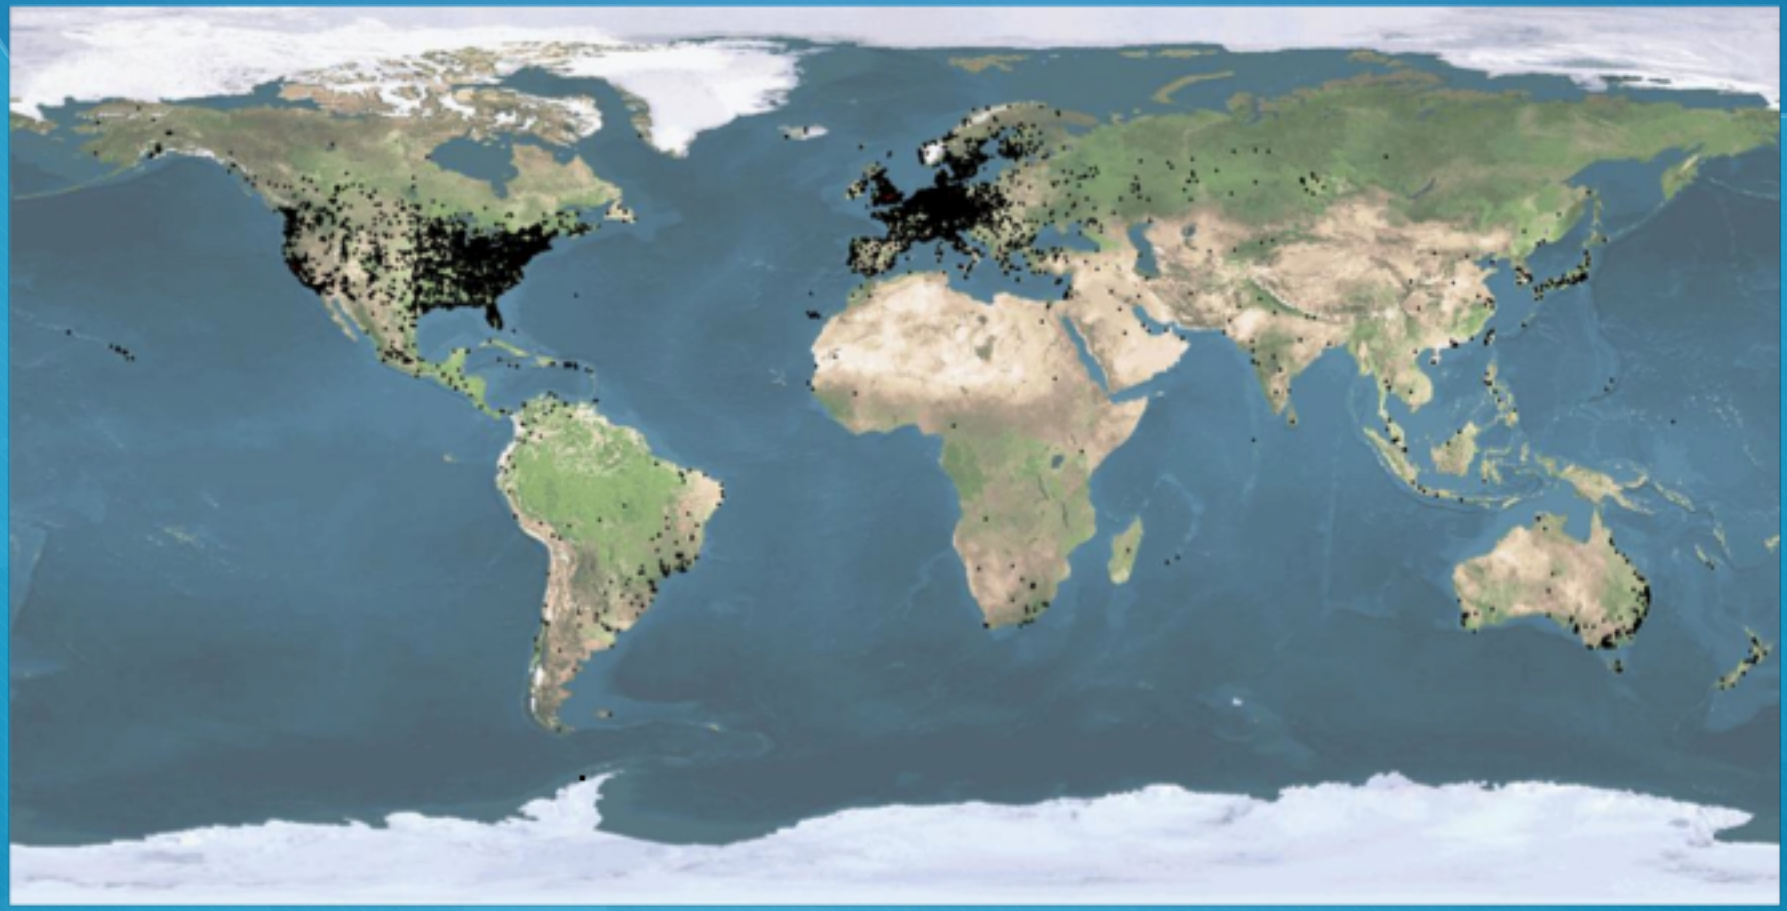

hadcm3l version 5.15 [workunit: hadcm3ln_2bjz_00447461]

BBC

Climate Change Experiment

Use keyboard keys to change view

Use CTRL + key when in screensaver mode
T - Temperature
R - Rain & snow
P - Pressure
C - Clouds
S - Stop/Start rotation
G - Show/Hide grid
H - Help & more options

Current view:
Clouds & surface

This computer model of Planet Earth simulates the atmosphere & ocean on a 3-D global grid. Gridded view shows the model grid scale. Switch between cloud (C) and rain (R) views to observe your model's weather.

This globe shows your climate model running

Model date and time: 26/01/1921 14:00

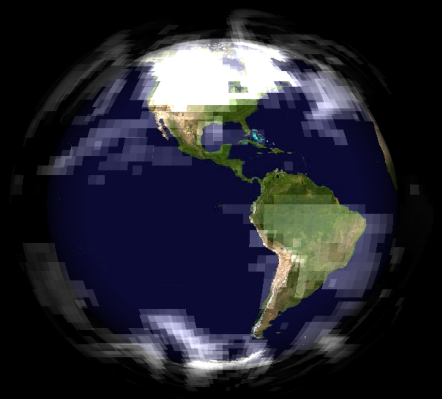

Thanks for taking part! Modelling the first few years is extremely useful for us, so do ke

bbc.co.uk/climatechange

created by climateprediction.net

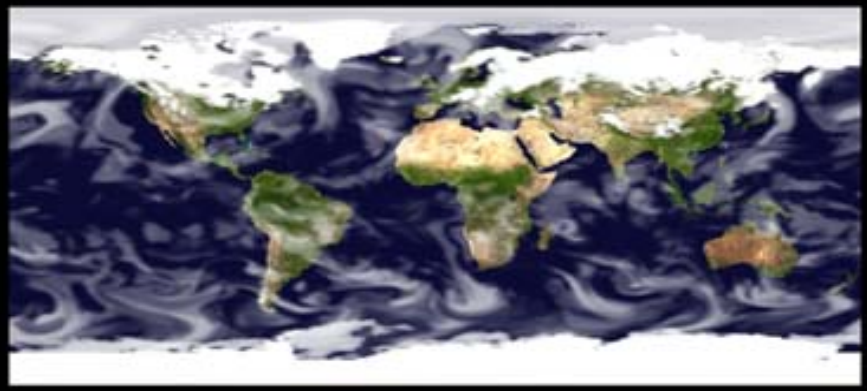

hadcm3l version 5.15 [workunit: hadcm3ln_2bjz_00447461]

BBC

Climate Change Experiment

Use keyboard keys to change view

Use CTRL + key when in screensaver mode
T - Temperature
R - Rain & snow
P - Pressure
C - Clouds
S - Stop/Start rotation
G - Show/Hide grid
H - Help & more options

This computer is contributing to the biggest climate change experiment in the world

To find out more about how you are helping climate research, visit the project website:

This globe shows your climate model running

Model date and time: 27/01/1921 04:00

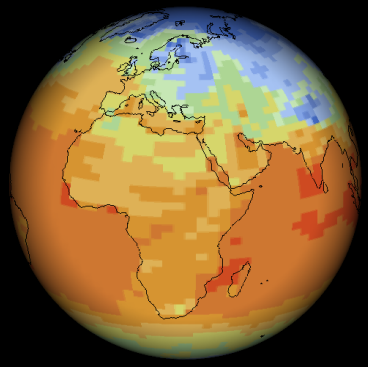

bbc.co.uk/climatechange

created by climateprediction.net

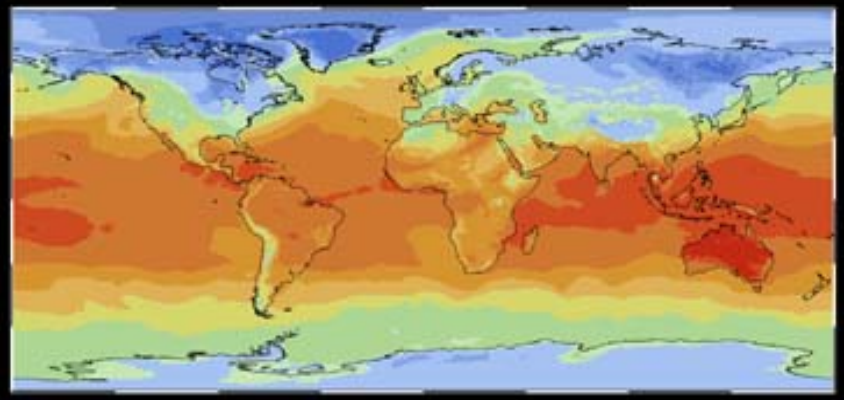

start

climate*prediction*.net
for Educational Outreach

Students at Gosford Hill School, Oxon, viewing their CPDN model

- CPDN has public education via the website, media, and schools as an important facet of the project
- Website has much information on climate change and related topics to the CPDN program.
- Open University (UK) offered a short course (S199) utilizing the climateprediction.net experiment (MS Windows client)
- Students hosted a debate on climate change issues, compared and contrasted their results, etc.

climate*prediction*.net Summary

- A good example of a high-CPU requirement volunteer computing project
- Ongoing: new experiments with MetOffice Hadley Centre models (slab model, coupled model, atmos only, regional models, newer versions i.e. HadGam/HadGem), 64-bit
- Interested in collaboration with other modelling groups to get other models on CPDN (ECHAM5?, CCSM?, WRF?)
- Thanks to Myles Allen, Tolu Aina, Milo Thurston, Hiro & Kuniko Yamazaki, Sue Rosier & David Frame

BOINC Modifications – climate*prediction*.net

- "trickle" system carried over from "old/original cpdn" and implemented by DA into BOINC (for all projects)
- This enables updates of the run daily (as run can take a month or sometimes 6-12 months!)
- Trickle uploads credit stats as well as scientific diagnostic data from the model run (which provides user feedback of run as well as extra data for the scientists)
- Also uploads of results during a run (also implemented in BOINC overall)
- Customized website for delivery of results to participants, as well as a portal for scientists
- Customized credit system based on trickles – outside of BOINC system – probably wasn't worth it

Case Study 2- The Quake-Catcher Network

Prof. Jesse Lawrence (Stanford); Prof. Elizabeth Cochran (UC-Riverside)

- Based at Stanford with collaborators in UC-Riverside and international collaboration starting with Chile and New Zealand
- "Opposite" the usual volunteer computing projects with high computing requirements (low CPU but larger & faster network bandwidth requirements)
- Sensors report seismic events ("triggers") over the Internet to our servers via internal (laptop) or external (USB) sensors
- http://qcn.stanford.edu

QCN: No-Cost Network

- Most Modern Laptops Have Sensors:
- HP, Apple, ThinkPad, Acer, Lenovo
- Very noisy data, not coupled to ground
- Sensors move location – track online
- Drawback: info on sensor detection and usage tough to get from manufacturers: only Apple Mac & ThinkPad laptops supported by QCN

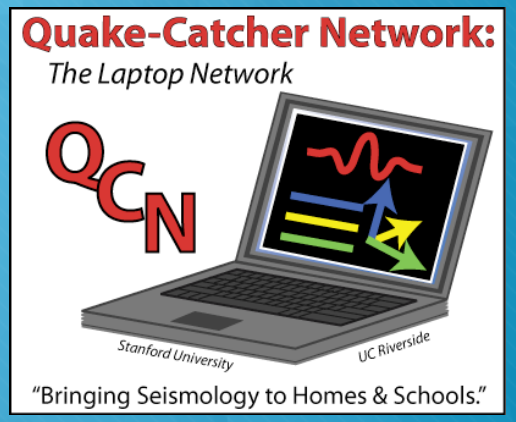

QCN: Low-Cost Network

- Desktops with connected USB sensors
- Cost: $35-$150 per sensor
- In noisy environments (homes & businesses)
- Over time hardware is getting cheaper, sensitivity/features are increasing
- 5000 purchased for distribution (Paypal purchases etc)
- “new & improved” sensors coming i.e. 16 & 24 bit

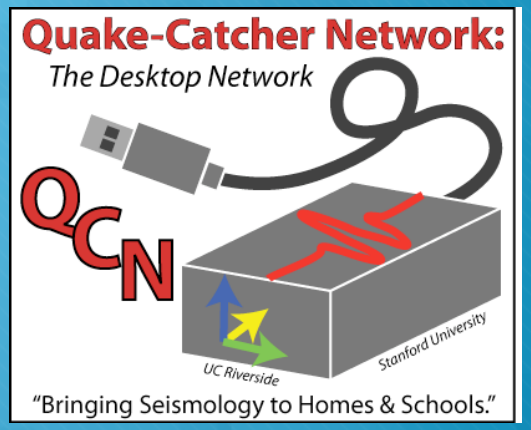

QCN Challenges

- **Better Sensors:** More sensitivity = better science
- ***Location:*** always changing – track IP addresses on our website, user inputs their location lat/lng
- ***Noise:*** typing, bouncing on laps, slamming doors, running kids, …
- ***Timing:*** no GPS clock – synchronized to our servers via ntp (network time protocol)
- Funding: as with may BOINC projects “pushing th e envelope”, it can be challenging to get funding, thankfully we just received an NSF grant!

Location

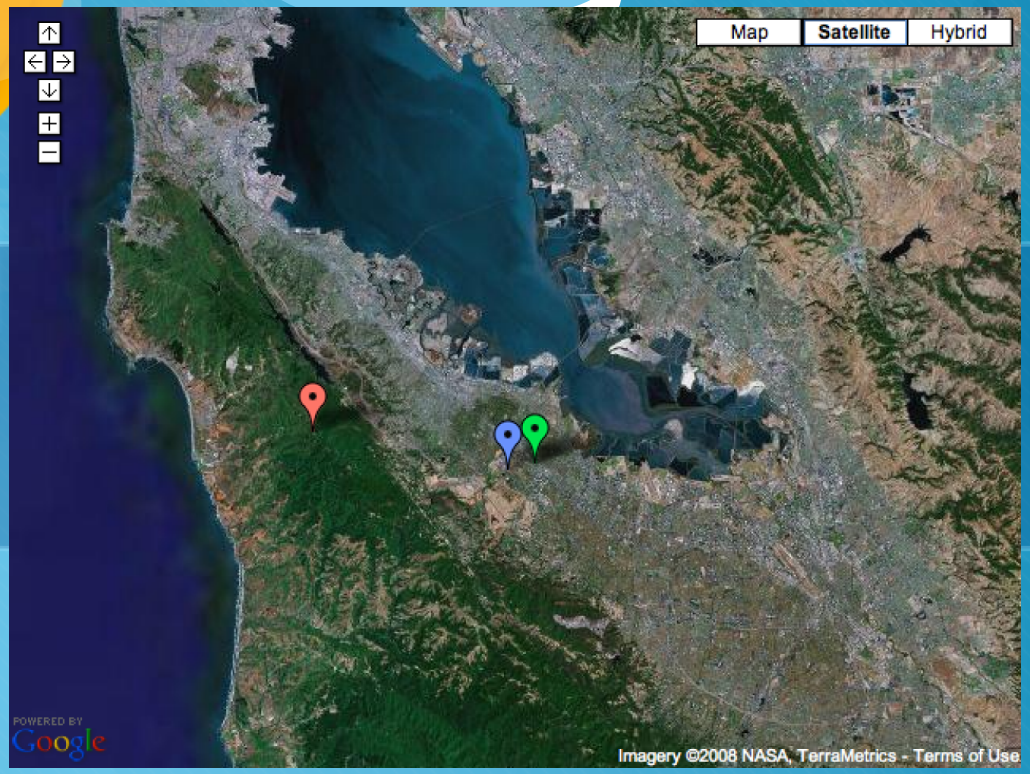

3-step Location System:

- Estimate location based on last known router (geoip)Often accurate within several kilometers
- Participants provide their "favorite five locations" using a Google Maps Interface.
- linked to IP or set a default location/address

Future: a "Where Were You?" website

Select	Location Name (optional)	Latitude	Longitude	Net (IP) Addr	Set Net Addr	Clear Net Addr
	Home	34.0971731803043	-117.72793114185	76.170.119	Set Current	Clear
	Work	33.9745572764349	-117.32615232467	138.23.128	Set Current	Clear
					Set Current	Clear
					Set Current	Clear
					Set Current	Clear
	Update info					

Italy 04/2009 – 3 sensors

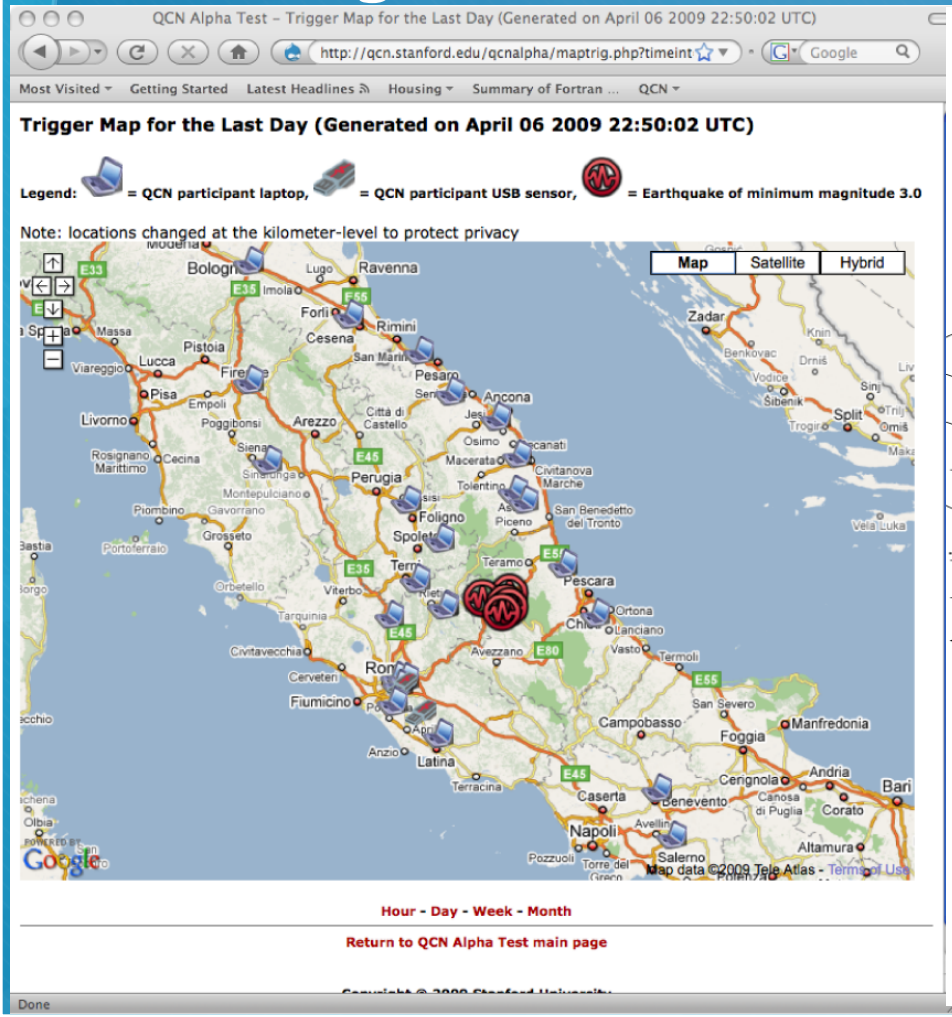

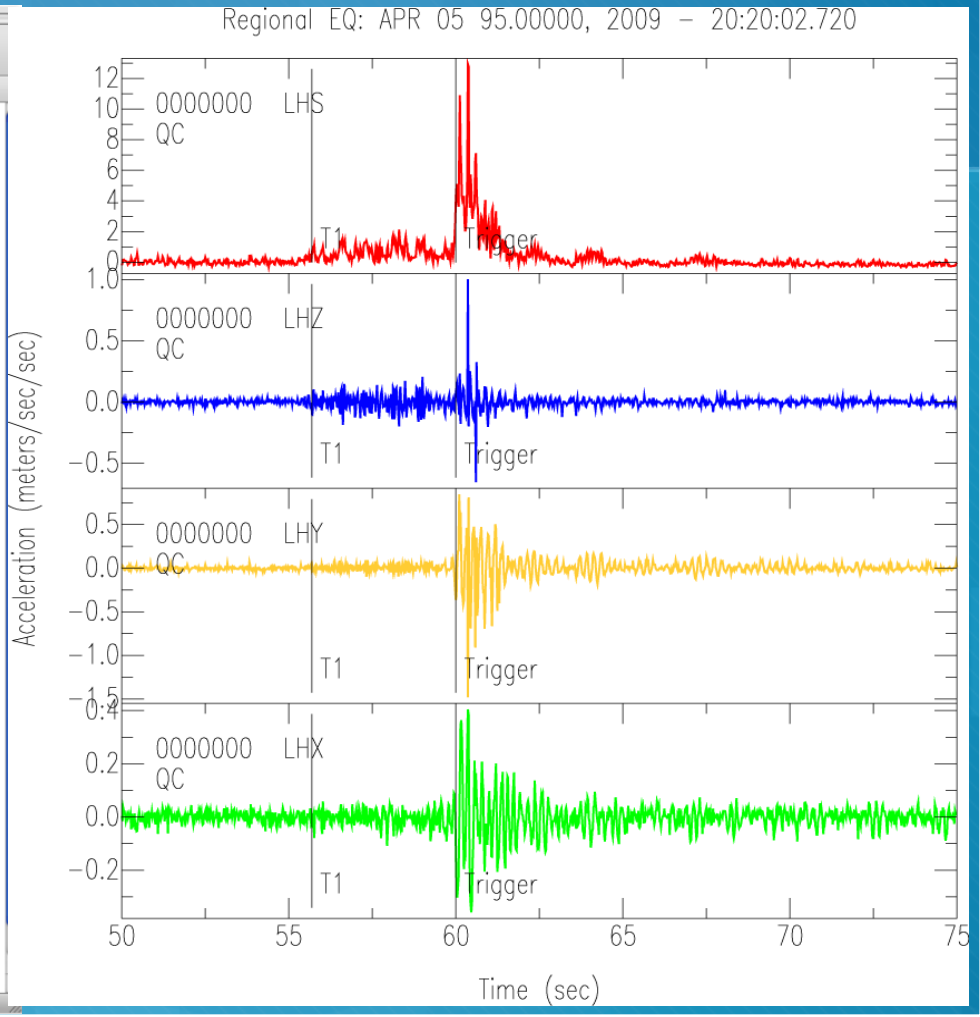

Japan 03/10-14/2011

each circle at least M4 quake

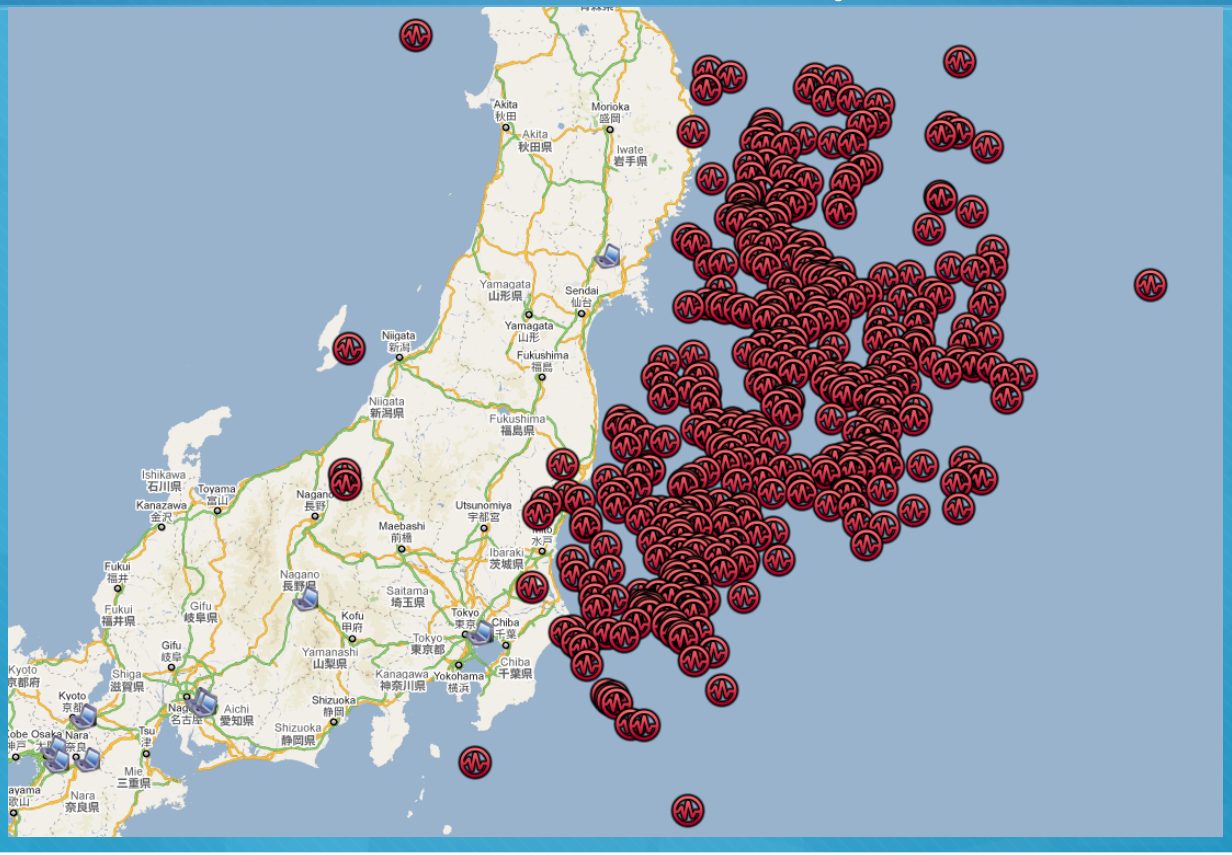

M9 – 3/11/2011 – Japan

Lenovo Laptop in Kurihara (simple 2d, 8-bit sensor) 120km away

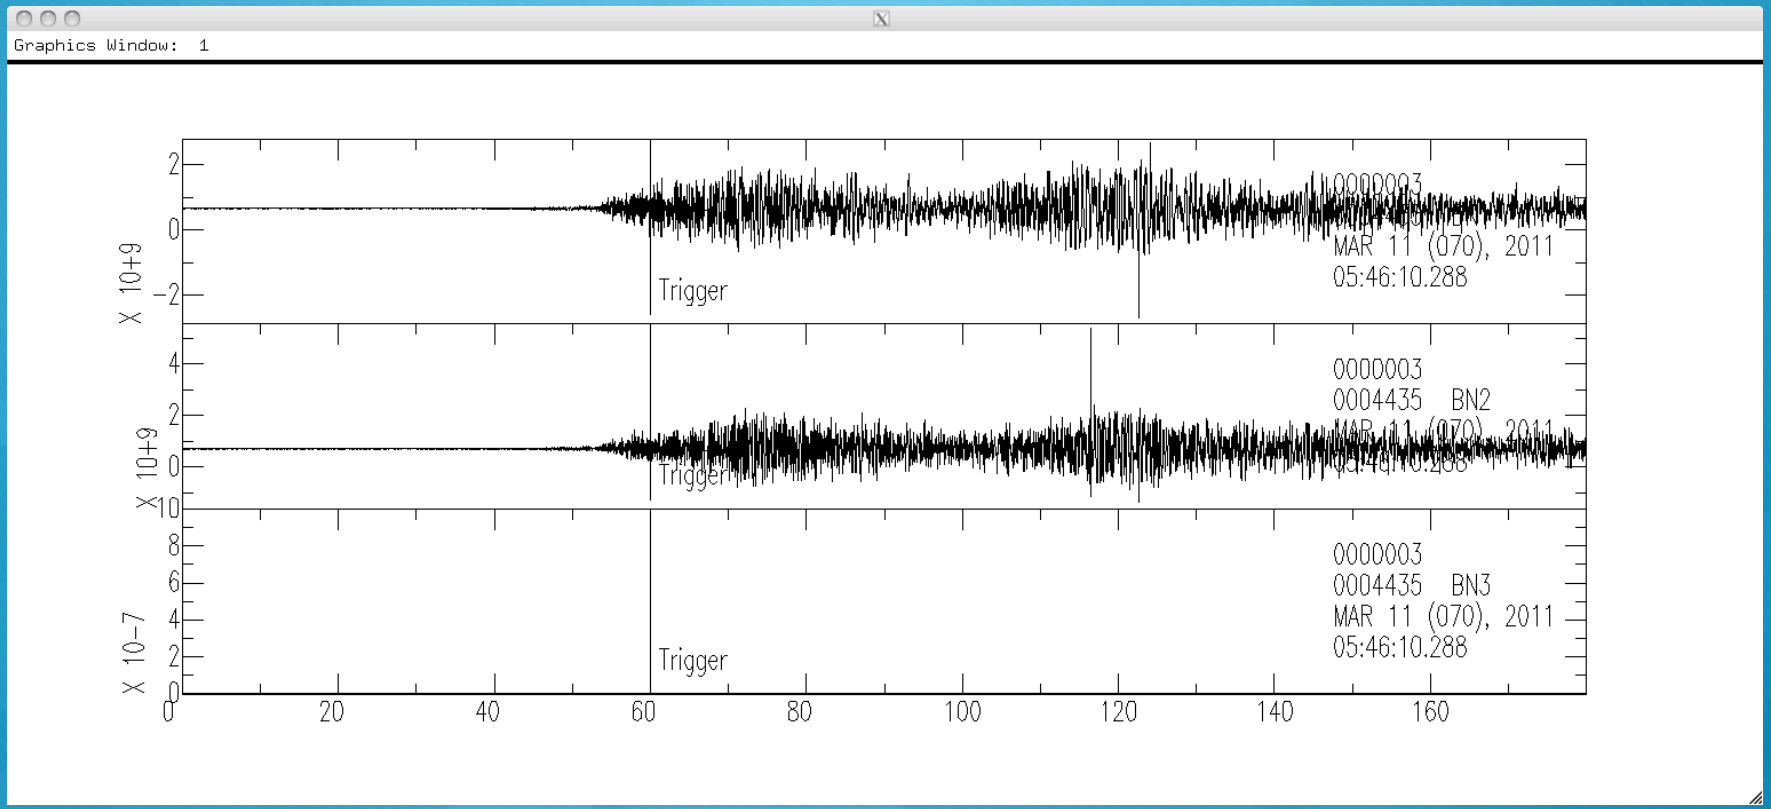

Educational Outreach

- What we provide:
  - Classroom Demo software - QCNLive
  - Seismology related in-class activities.
  - Classroom USB Sensor.
  - Classroom BOINC Software.
  - Porting to Google Android and iPhone (i.e. QCNLive on mobile devices with optional BOINC sreever support)

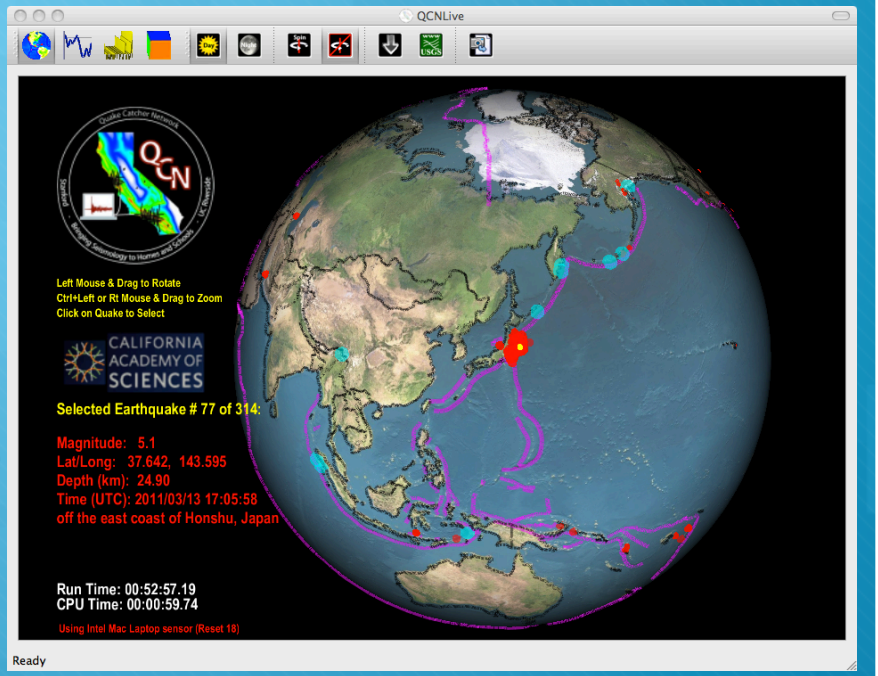

BOINC Modifications – Quake-Catcher Network

- These were more extensive, and not “standard” in BOINC due to the different nature of QCN and in fact a whole talk in itself
- **ntpdate** usage in the client – to get a reasonably accurate time (within 1/100th second at least)
- Trickles still used – now they deliver run-time stats as well as trigger information for a seismic event

BOINC Modifications for QCN (continued)

- Also using “intermediate uploads” feature of BOINC as a mechanism to upload complete SAC files surrounding a potential seismic event, or “continual” monitoring
- Substantial updates to the BOINC scheduler to handle the incoming trickle/trigger information, do the lat/lng lookup by IP address, pipe to event detection algorithms etc (bypassing database for speed)
- Still have a customized credit system based on trickles due to the non-CPU-intensive nature of the project

Other Considerations

- Minimum job 1 day; maximum 6 month (if you provide feedback and credits via trickles)
- So consider the job sizes you want to send out, i.e. better to split up big jobs temporally and/or spatially if possible
- Consider download/upload sizes (subset data, compression either lossy or lossless etc), 1GB per day probably too much
- Don’t underestimate user feedback/interaction to keep participants interested and involved

MPI Programs in BOINC

- http://boinc.berkeley.edu/trac/wiki/MpiApps
- Possible using mpich2 library from ANL
- Link your app against BOINC & mpich2 on Linux/Mac/Windows
- Simplest – use BOINC wrapper with appropriate “job description file” (job.xml)
- May need some processing of namelists etc to setup for the number of procs BOINC says is available to use
- We’ll explore in the “hackfest” with SPECFEM3D

Conclusion

- Volunteer computing with BOINC is a proven resource for large earth science projects whether low or high-CPU requirements
- BOINC has a lot of “stock” functionality, but if you need customization either at a low-level (scheduler) or web, it’s feasible.
- Exciting new avenues of approach – 64-bit, GPU, gaming consoles, multicore/multithreaded/multiprocess/MPI, mobile devices
- I would love to hear of any new ideas, suggestions, or proposals you may have, please email:

  carlc@atm.ox.ac.uk or carlgt1@yahoo.com