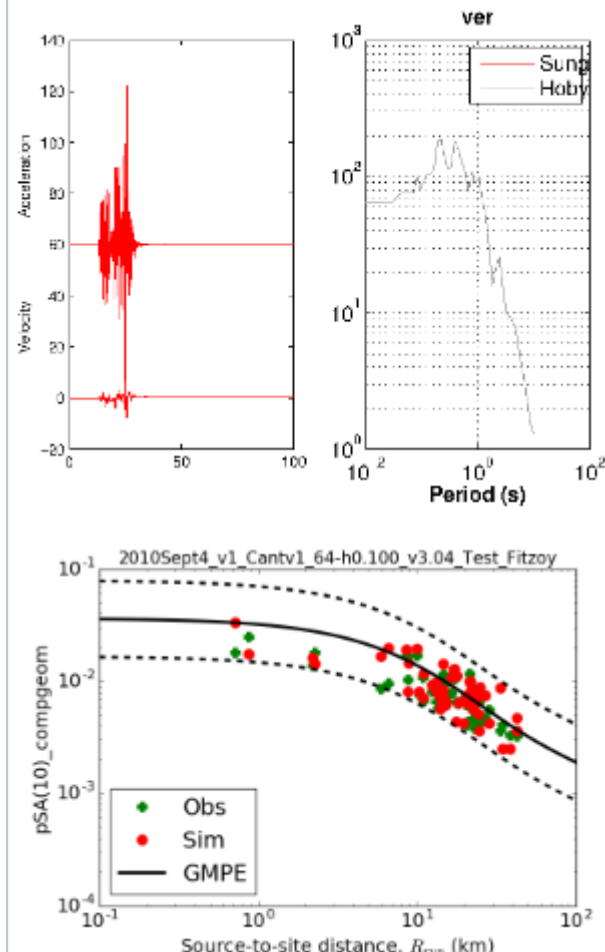


Verification of workflow porting from Foster to Fitzroy

Background/Inspiration or Problem Statement
QuakeCoRE has used NeSI's BlueGene/P through UC HPC collaborator allocation and NeSI research allocation (nesi00213) to run HPC workflow (Emod3d) based on Graves & Pitarka hybrid ground motion simulation methodology since 2013.
Due to NeSI's decommission of BlueGene/P and future upgrade of their hardware, it was suggested that QuakeCoRE migrate to Fitzroy, a POWER6 cluster installed at NIWA.
Project Members
Sung Bae, Alex Pletzer, Wolfgang Hayek (NeSI)
Hoby Rafzfindrakoto, Richard Clare and Brendon Bradley (QuakeCoRE)
Description (Objectives / Outcomes)
<ol style="list-style-type: none">Port all the software components to Fitzroy to produce output that is consistent with the output from BlueGene/PCreate a workflow that utilizes Fitzroy and other resources efficiently to match or surpass the user experience with BlueGene/P
Tasks
<ul style="list-style-type: none">Create a GIT repository for code and test inputs and implement a structured environment for effective collaborationBuild Emod3d code on Fitzroy and verify - rerun small/large input data and compare with the BlueGene/P output.Produce a few sets of verified sample outputs for benchmark purposesMeet the performance target (3-4 times of BG/P speed for the same number of cores used)Complete the workflow:<ul style="list-style-type: none">Port the current solution for HF Sims and BB SimsPort the current solution for plotting the output from Emod3d (current solution still under development (90% complete) written in Python)Port the current solution for creating an animation movie to Fitzroy (currently relying on GMT and QuickTime to stitch frames)Profile the code to identify performance bottlenecksWrite a verification codeTrain PI and the team to run the simulation on Fitzroy
Schedule
The project has completed within the scheduled timeframe. <ul style="list-style-type: none">Code ported and running on Fitzroy with confirmed outputs - 15 MarProfiling/Performance target met - 31 MayWorkflow tools ported/completed (plotting,animations etc) - 31 MayFinal tuning, PI training, acceptance test - 30 Jun
Verification of Computation
Strategy
BlueGene/P and POWER6 share the same CPU architecture (IBM POWER-based), but have different OS and different development toolchains including C and Fortran compilers. Even if all the software components are successfully built, the outputs from POWER6 may differ from that from BlueGene/P for one of the following reasons. <ol style="list-style-type: none">Floating-point computation: Certain errors inherent to floating-point computation were expected. For deterministic computation, the level of precision should be contained (The target precision level was not set at the time of project start).Random number: Different machine produces a different sequence of random numbers unless a customized random number generator is used. Some portion of the workflow is based on stochastic simulation (eg. High frequency simulation). The output should be analysed and verified by domain experts.Non-standard compliant code: Different compilers interpret the syntax of code differently if the code is written with ambiguity (ie. not sticking to the language standard). This often causes different behaviour of the code.Data format: During the complicated workflow, the output from one step may serve as the input for the next step. If such output is in ASCII format in different style (eg. 0.012345 vs 1.234E-02), and the next step assumes one particular style, the code will need to be modified.Aggressive compiler optimization: If a compiler level optimization is applied too aggressively, the behaviour of code can change (eg. the order of computation), resulting in a different output. Due to the reasons outlined above, and unknown behaviour of the software ported to POWER6, automated testing appeared to be difficult. NeSI and QuakeCoRE agreed to verify the computation output by domain experts' visual inspection utilizing plotting/analysis tools developed in-house. QuakeCoRE supplied two input sets of Dec26 and Sep11 events.
Result
We have concluded that the difference between the output from POWER6 and that from BlueGene/P is only caused by the first reasons listed above. All test cases passed.

Unit	Test Case	Verdict	Tested by	Notes
emod3d+merge_ts+plot_ts	Compare GMT results of time slices	Pass	Sung (20/04/16)	blocked URL
winbin_aio	Compare intensity measurements/waveforms	Pass	Hoby/Sung	Observation (POWER6) and Simulation (BlueGene/P) matched
hfsims-stats	<p>Compare the acceleration spectra at the same station - "minimal" difference expected</p> <p>Compare 10 results with different random seeds.</p>	Pass	<p>Hoby (26/04/16 - used MATLAB code)</p> <p>Sung/Hoby/Brendon(02/05/16)</p>	 <p>BB commented 'the results are 'identical' given the inherent randomness resulting from the random seeding.'</p>
match_seismo	<p>Compare (1) vs (2) and (1) vs (3)</p> <p>(1) Running on BGP with LP/HF from BGP (baseline)</p> <p>(2) Running on Fitzroy with LP/HF from BGP</p> <p>(3) Running on Fitzroy with LP/HF from Fitzroy</p>	Pass	Sung/Richard/Hoby/Brendon (10/05/16)	<p>A. Comparison (1) vs (2): The quality of "just" BB on Fitzroy : https://www.dropbox.com/sh/1vdzmv0gxyza27/AAAYH8Y7rWIBU82ogV49shlta?dl=0</p> <p>B. Comparison (1) vs (3) : The quality of overall computation on Fitzroy: https://www.dropbox.com/sh/b8kwige77vzeepc/AACKNxGsWlmXVh26z73bmlda?dl=0</p>