

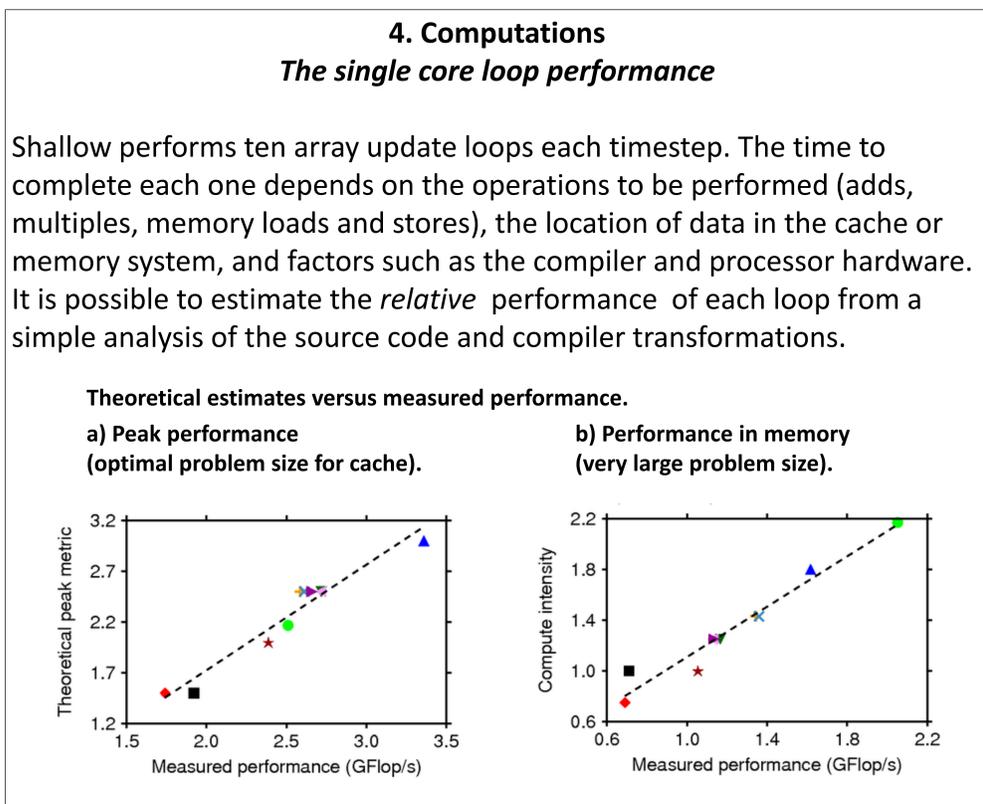
Modelling the Performance of a Shallow Water Code

Annette Osprey¹, Graham Riley², Bryan Lawrence¹ and Muniyappa Manjunathaiah³

a.osprey@reading.ac.uk

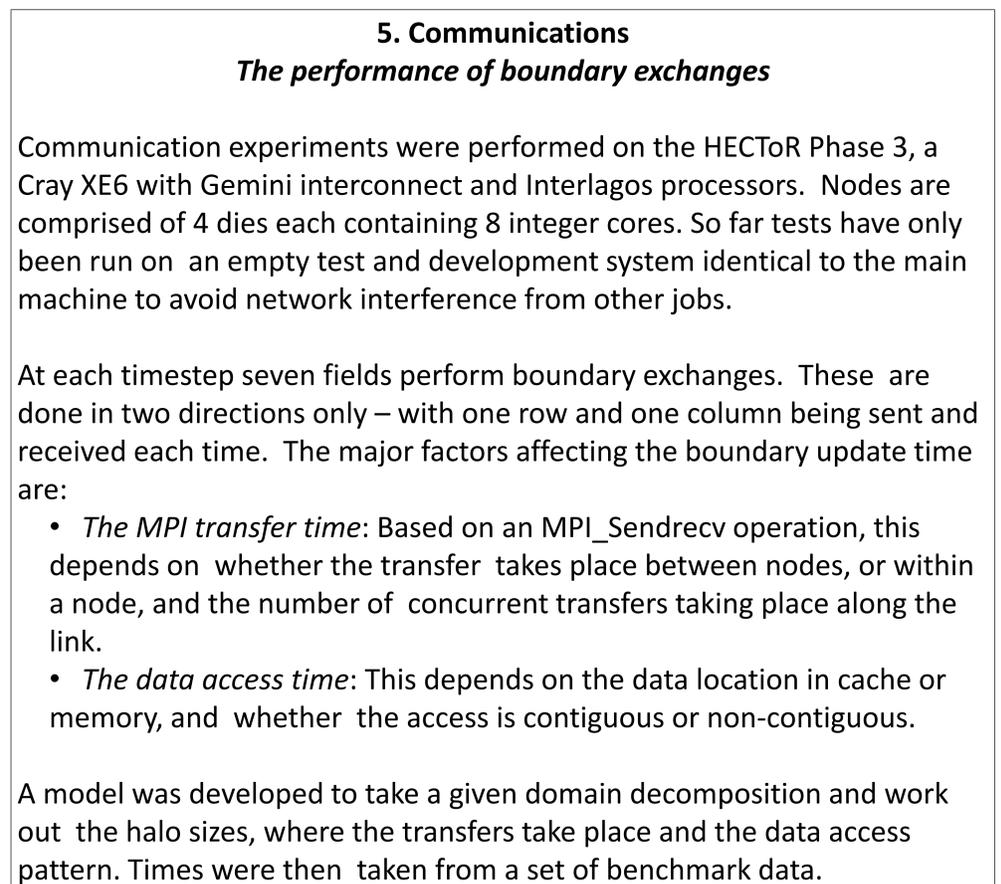
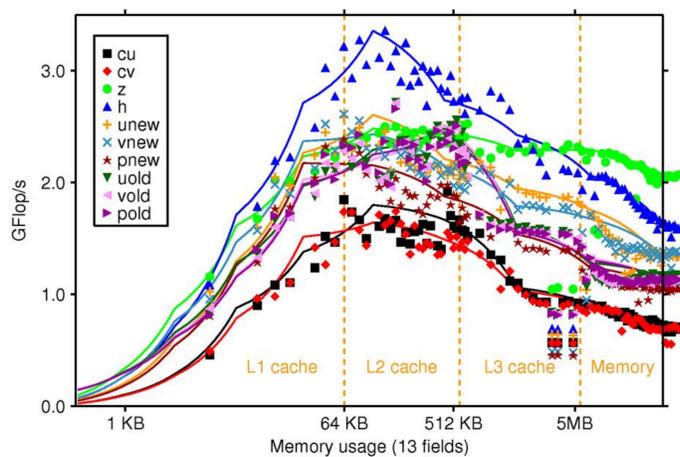
¹ Department of Meteorology, University of Reading, ² School of Computer Science, University of Manchester
³ Department of Computer Science, University of Reading

1. Performance modelling	2. Shallow water model	3. A performance model of shallow
<p>A performance model is an equation for the time to run an application on a given machine, based on software and hardware characteristics.</p> <p>Performance models can be used to:</p> <ul style="list-style-type: none"> • <i>understand</i> current behaviour; • <i>make predictions</i> about future behaviour; <p>and</p> <ul style="list-style-type: none"> • <i>inform</i> design choices when developing or updating code, for example in Gung-Ho! 	<p>The application studied here is a simple program <i>shallow</i>, based on the NCAR shallow water model: http://www.cisl.ucar.edu/docs/hpc_modeling/</p> <p>Shallow mimics key parts of more complex climate model applications such as the UK Met Office Unified Model, particularly:</p> <ul style="list-style-type: none"> • time-step iterations; • loop-based calculations to update array values based on other arrays; • exchanges of boundary data between cores. 	<p>Our performance model takes the form:</p> $T_{total} = T_{comp} + T_{comm}$ <p>where:</p> <ul style="list-style-type: none"> • T_{comp} is the time spent in computational loops updating array values, and • T_{comm} is the time spent performing MPI data exchanges between cores. <p>The times for each part are derived from code analysis and benchmarking.</p>

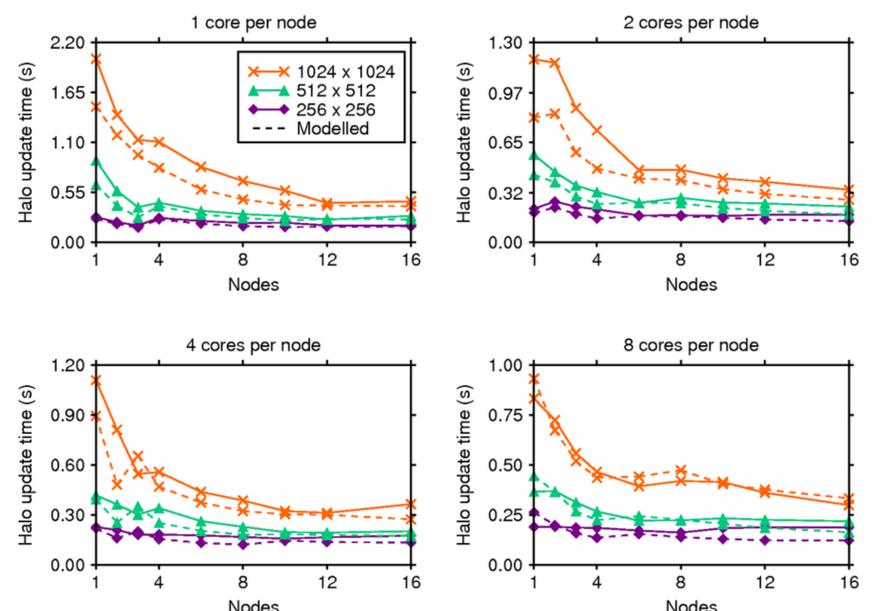


Experiments were performed on HECToR Phase2b with the PGI compiler. The Phase2b was a Cray XE6 comprising Magny-Cours processors with 16 KB of L1 cache, 512 KB of L2 and 5 MB of L3.

To derive *actual* performance values requires some level of machine benchmarking. In this case the entire loops are measured for a series of exponentially increasing problem sizes. The predicted performance for any problem size is then based on a linear interpolation from this data.



Predicted time to complete halo exchanges (dashed lines) and measured times (solid lines) for different problem sizes, cores per node and total number of cores.



6. Next steps	Future work
<h4>Finishing the performance model</h4> <ul style="list-style-type: none"> • Include multi-core loop performance with shared cache effects. • Update computation model to Phase 3 with Interlagos processor. • Put computation and communication models together. 	<ul style="list-style-type: none"> • Communications benchmark on the main HECToR system for some measure of run time variability. • Evaluate model process on a different machine.