EMS Annual Meeting Abstracts
Vol. 15, EMS2018-295, 2018

To InfiniBand, and beyond? Cloud computing for the WRF model

Richard Bassett, Gordon Blair, Faiza Samreen, William Simm, and Paul Young
Lancaster University, United Kingdom (r.bassett@lancaster.ac.uk)

WRF's intensive computing requirements means the deployment of the model is realistically limited to High-Performance Computing (HPC) environments. Access to WRF is further restricted by the steep learning curve, particularly for model installation. We envisage these barriers to WRF, and similar environmental models, can be overcome by the use of cloud computing technologies and software engineering tools and techniques. The potential for theoretically infinite resources, combined with the abstraction of system and model complexities, will also help expand the WRF user community to smaller research organisations or citizen scientists without access to HPCs and enable new ways of working, for example deploying massive parallel simulations.

In our research, WRF was installed on Microsoft Azure's cloud service through a series of automated scripts. WRF's performance was tested using cases of increasing domain size on low and high-cost cloud CPUs, and on a local HPC. Simulations using the low-cost cloud CPUs were slower than the HPC, caused by latency issues (i.e. cloud CPUs potentially split over multiple locations). To overcome this, high-cost CPUs with a dedicated InfiniBand (112Gb bandwidth) were tested and similar model performance to the HPC was found. This demonstrates the suitability of cloud services as a replacement for traditional modelling environments. Although clouds charge for CPU time, this is small in comparison to the cost and overheads of a HPC and there is no queue time.

Beyond the cloud, we are creating a model of the WRF system using a "Domain Specific Language" (DSL), a computing language written specifically for a given situation. Users often encounter challenges installing the model followed by a large number of repeated tasks (e.g. data transfer). A WRF-DSL can orchestrate all aspects of running the WRF model, by encoding these processes, and allow simulations to be deployed to the cloud with only a limited knowledge of the computing environment. Furthermore, the WRF-DSLs can enable connection to cloud services such as data storage, apply software engineering tools such as machine learning. We aim to encapsulate these processes in an open-source visual web-cloud interface, e.g. Jupyter notebooks (http://jupyter.org/), noted for their collaborative and open working facilities. The use of these tools will also allow users to contribute model configurations (and errors) to a common, semantically enhanced namelist bank.

Through the "Models in the Cloud" project (www.ensembleprojects.org/projects/models-in-the-cloud/; EPSRC grant: EP/N027736/1) we advocate a combination of a new WRF-DSL coupled with cloud computing to create a new, enriched modelling environment for WRF. The advantages of this approach include: (i) abstraction from computer system and model complexity, (ii) access to scalable and tailored computing resources, (iii) model democratisation, i.e. access to computing resource without the need for a HPC, and (iv) improve project collaboration.