

Optimisation of Data Movement in Complex Workflows

Friday, 28 September 2018 11:00 (30 minutes)

Since its formation in 2015, the Cray EMEA Research Lab has taken a keen interest in new approaches to optimise data location, distributed task computation and HPC I/O. To fully address HPC I/O optimization we must move beyond configuration, application and library optimisation to looking at the whole workflow and to optimise data movement required by dependencies in the workflow. Our **Octopus** project will deliver a high-level workflow description, scheduling, and execution framework, built around data and memory hierarchy awareness. We believe that being able to schedule applications while being able to reason about their data production, data locality, and data consumption, and accounting for the cost of moving data between tiers of the memory hierarchy will enable efficient execution of coupled applications in complex workflows. Our use-cases cover: in-situ analysis and visualization; simulation with multiple consumers; coupled simulations with multiple concurrent analysis, external-data source and archiving dependencies; and even sample distribution and shuffling in high-throughput machine learning tasks. Octopus is an ambitious project that will provide optimisations impossible if we constrain ourselves to the individual application. We have an outline design and will develop Octopus collaboratively with partners of recently funded EU projects, ECMWF is one of these partners.

We have developed an enabling library-based 'Data Junction' to transport (and redistribute) data encapsulated in Octopus objects between applications, this already supports various transports (including MPI, Dataspaces, Ceph RADOS, libfabric and POSIX file).

The talk will describe Octopus, work completed so far on data transport and how this will continue via collaborative EU H2020 projects MAESTRO and EPIGRAM-HS. We are interested in how the Octopus framework can be applied to a variety of NWP and climate workflows.

Affiliation

Cray EMEA Research Lab

Primary authors: TIM, Dykes (Cray UK Ltd); FOYER, Clement (Cray UK Ltd); HAUS, Utz-Uwe (Cray EMEA Research Lab); RICHARDSON, Harvey (Cray UK Ltd); SIVALINGAM, Karthee (Cray EMEA Research Lab); TATE, Adrian (Cray EMEA Research Lab)

Presenter: RICHARDSON, Harvey (Cray UK Ltd)

Track Classification: 18th Workshop on high performance computing in meteorology