nextgenio newsletter

Welcome to the fourth issue of the NEXTGenIO newsletter!

NEXTGenIO at SC 2017

PARTNER BOOTHS

intel	TECHNISCHE UNIVERSITÄT DRESDEN	BSC
1203, 1301, 1319	1881	1975
allinea Now part of ARM	epcc	FUJITSU
610	201	1125

EVENTS

The best place to come and talk to us at SC is at any of our partners' booths, but you will be able to find us at different events throughout the conference as well.

RESEARCH POSTERS

Our Partners at TUD are presenting two research posters during the conference. The poster presentations are on November 14th, 15th and 16th in the Four Seasons Ballroom.

Analyzing Multi-layer I/O Behaviour of HPC Applications - Ronny Tschüter, Christian Herold, Bert Wesarg, Matthias Weber

Measuring I/O Behaviour on Upcoming Systems with NVRAM - Christian Herold, Ronny Tschüter

TALKS

Emmanouil Farsarakis (EPCC) will be giving two talks at the conference.

NEXTGenIO: Storage Class Memory and Job Scheduling - PRACE booth (2171), 14th November, 3pm

NEXTGenIO - PDSW-DISCS 2017 workshop, 13th November, 9am - 6pm

NEXTGenIO Co-design Applications: MONC, Halvade and OpenFOAM

(Michele Weiland, EPCC)

This article presents a brief overview of three of the applications that are representative of the types of application that will run on the system we are designing, and that will be used to test and evaluate our prototype platform. Each of the applications has different demands on the memory and I/O subsystem of the high-performance computing and high-performance data analytics platform we are developing, highlighting the multi-disciplinary, multi-user nature of current and future research computing systems.

The Met Office NERC Cloud Model (MONC) is a very high resolution (~2 to 50 m), flexible, portable cloud modelling framework, which includes its own I/O server. Unlike most computational simulation applications, it dedicates specific processes to handling all diagnostic and I/O aspects during a given simulation run. Typically, one core per processor will run the I/O server and support the remaining cores that execute the cloud model. MONC asynchronously "fires and forgets" any raw prognostic data to the I/O server for processing (see Figure 1).

The NEXTGenIO prototype will allow the MONC I/O server to transfer data to NVRAM instead of disk. This data can then be transferred to the parallel file system while the computation and analysis continue within MONC. This will enable higher write frequencies, and therefore increased diagnostics, without affecting MONC's performance. In addition, the in-situ data analytics functionalities in MONC will be able to exploit the capacity of the NVRAM for increasingly complex diagnostic computations.

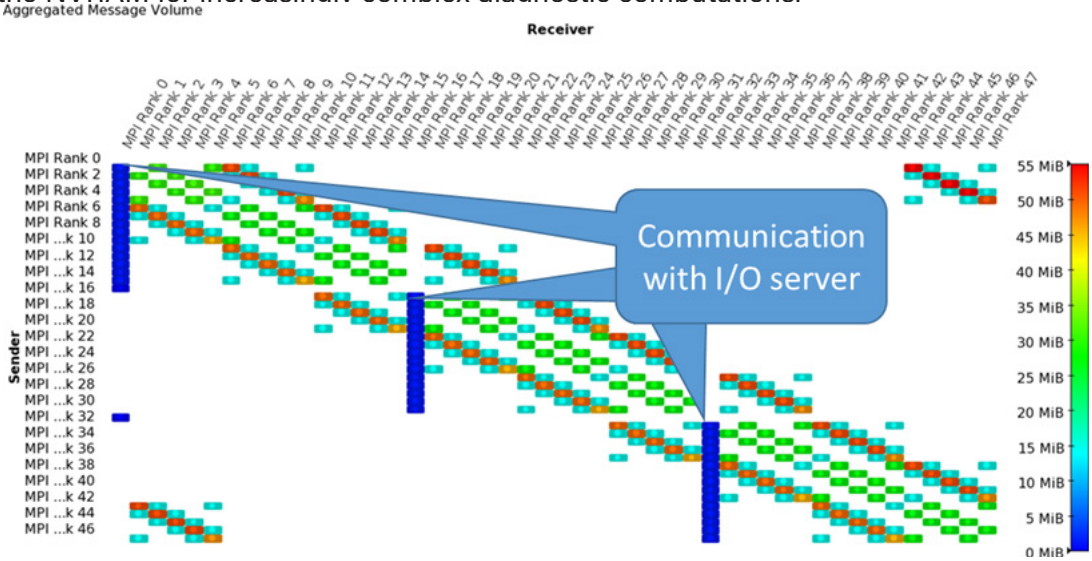

Figure 1: MONC I/O server communication.

Halvade is a MapReduce implementation used for genome sequencing. Due to rapid progress in the field of next-generation sequencing platforms, the speed of DNA sequencing has increased considerably: modern Illumina systems, for example, can generate several hundreds of GigaBases per run, which result in several hundreds of gigabytes of raw sequence data to be processed. A typical Halvade analysis consists of 5 steps, see Figure 2.

With the expected significant improvements in I/O performance and the opportunity to keep data passed between the different steps in memory, in NVRAM, NEXTGenIO can make a major difference in the field of genome sequencing. The expected performance increases for storing data or passing it between steps will very likely "unmask" inefficiencies in the compute parts and open up the opportunity to create much more efficient genomics tools and workflows.

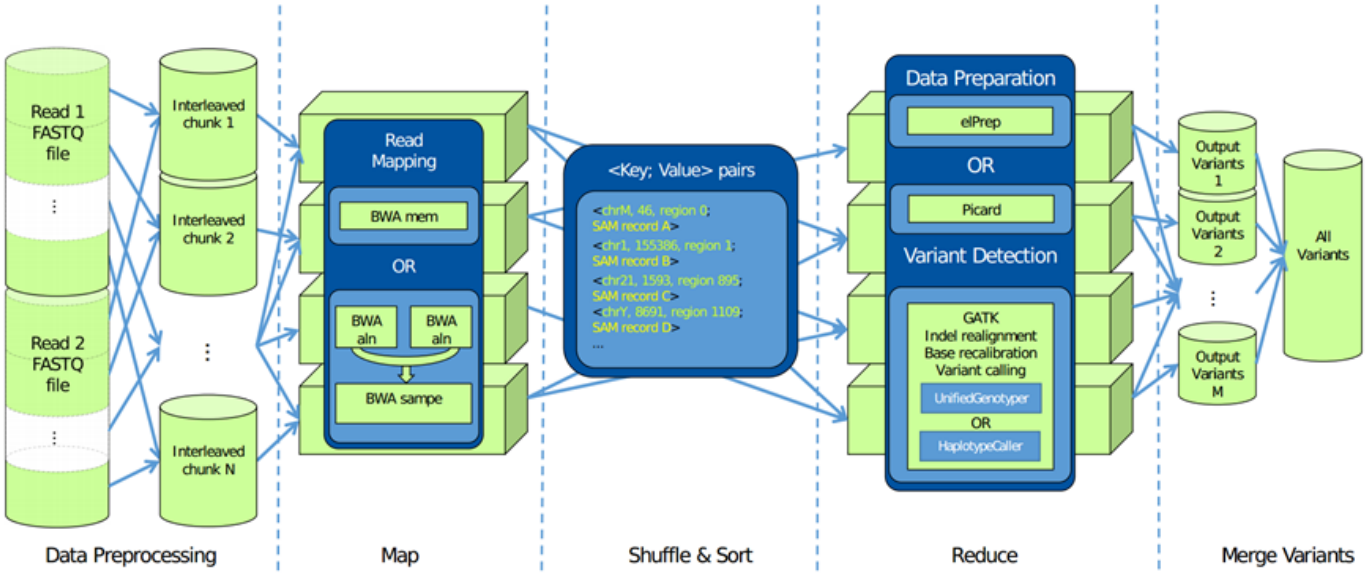

Figure 2: Overview of the Halvade framework showing the five principal execution steps. (Decap et al, 2015)

OpenFOAM is a free, open source CFD software with a large user base across many areas of engineering and science. OpenFOAM has an extensive range of features to solve complex fluid flows involving chemical reactions, turbulence and heat transfer, to acoustics, solid mechanics and electromagnetics. In NEXTGenIO, we are using a test case provided by Pipistrel, a leading designer and manufacturer of light aircraft based in Slovenia (see Figure 3).

When an OpenFOAM simulation is running, it outputs the state of the fields and the mesh at a given point in time. Ideally, output is produced at every time step, but the demands on the I/O subsystem can be extreme. The availability of NVRAM will allow OpenFOAM to write output more frequently, thereby enabling more detailed post-processing. Furthermore, it will also allow writing the uncompressed data directly at a faster rate.

On the NEXTGenIO platform, we will explore using NVRAM together with the job schedulers to create an OpenFOAM workflow that can pass data between the different steps of the workflow without the need to write to the parallel file system until the end of the simulation, when the final result data can be moved after the main computation has completed. This will enable visualisation and analysis to be undertaken alongside simulation, consuming data produced by the solver. This has the potential to greatly reduce the runtime for the whole simulation workflow required to setup, run, and analysis a simulation, saving both valuable compute time, and the real world time users need to produce scientific results.

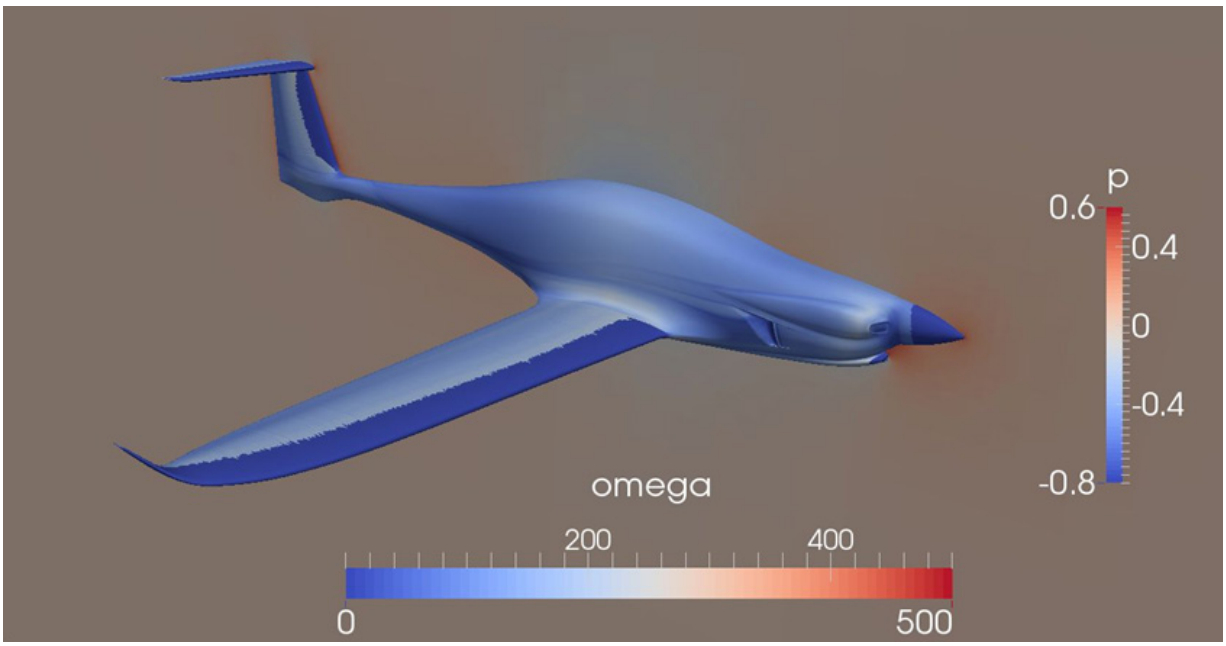

Figure 3: Pipistrel test case used with OpenFOAM

PyCOMPSs @ NEXTGenIO

(Javier Conejero, BSC)

- ***The Barcelona Supercomputing Center provided an updated version of PyCOMPSs to the NEXTGenIO project.***

- ***This update includes a new interface for the PyCOMPSs applications debugging and performance analysis with the tools provided by the NEXTGenIO partners.***

The Workflows and Distributed Computing team at the Barcelona Supercomputing Center provided an updated version of the programming environment PyCOMPSs to the NEXTGenIO project. PyCOMPSs is a task based programming model that enables the execution of sequential Python applications in large scale distributed computing platforms by automatically parallelizing them.

This version of PyCOMPSs, updates the latest release (version 2.1) with a new interface that enables the interaction between PyCOMPSs and the debugging and profiling tools provided by the NEXTGenIO partners, which are being improved with support for NVRAM. Therefore, this integration will enable users to retrieve hardware metrics, such as NVRAM usage, that can be later analysed in order to understand the applications behaviour, detect bottlenecks, and parallelization opportunities in the future NEXTGenIO prototype cluster.

In particular, this interface is capable of enabling the interaction between PyCOMPSs with ScoreP in order to obtain OTF2 traces that can be later analysed with Vampir, with Allinea MAP for obtaining profiling reports, and with Allinea DDT for debugging purposes.

PyCOMPSs is currently being enhanced for NVRAM support with the integration with dataClay (object storage) and the usage of the data scheduler (traditional storage).

For more information on COMPSs please visit: *http://www.bsc.es/compss*

Recent Publications

The project has recently published two papers on our work, both as part of the proceedings of two international conferences. These papers cover different aspects of the work being done as part of the project.

Simon D. Smart, Tiago Quintino, and Baudouin Raoult. 2017. ***A Scalable Object Store for Meteorological and Climate Data.*** In Proceedings of the Platform for Advanced Scientific Computing Conference (PASC '17). ACM, New York, NY, USA, Article 13, 8 pages. DOI: https://doi.org/10.1145/3093172.3093238

Harold E.B. Dennis, Adam S. Ward, Tyler Balson, Yuwei Li, Robert Henschel, Shawn Slavin, Stephen Simms, and Holger Brunst. 2017. ***High Performance Computing Enabled Simulation of the Food-Water-Energy System: Simulation of Intensively Managed Landscapes.*** In Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact (PEARC17). ACM, New York, NY, USA, Article 43, 10 pages. DOI: https://doi.org/10.1145/3093338.3093381

The project was also featured in the German magazine ***iX*** in September.

Bernhard Homölle. ***Einfach Direkt.*** iX Magazin für Professionelle Informationstechnik. September 2017. Print.