nextgenio newsletter

Welcome to NEXTGenIO!

Welcome to the first NEXTGenIO newsletter! NEXTGenIO is a project that was funded under the European Commission's Horizon 2020 program to design and develop hardware and software solutions to address the I/O bottleneck at the Exascale.

Twice a year, this newsletter is going to bring you a progress update on our work, plus some sneak previews of our research, and let you know what events you can meet us at in the coming months. We hope that you'll find the information in this newsletter interesting and useful – for further information, please visit www.nextgenio.eu!

The First Six Months of NEXTGenIO

NEXTGenIO partners are excited to be part of a project that will develop transformational technology for Exascale I/O. The first six months of NEXTGenIO have seen significant progress across all work packages.

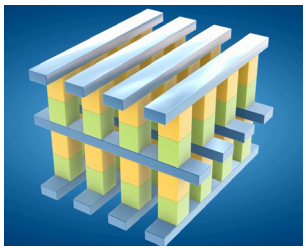

The NEXTGenIO project was officially launched on the 7th & 8th October 2015 with a kick-off meeting hosted by the project coordinator EPCC at the University of Edinburgh. All project partners were represented at this meeting: Intel, Fujitsu, Allinea, ECMWF, TU Dresden, Barcelona Supercomputing Centre, and Arctur.

NEXTGenIO is developing a hardware prototype that is based around Intel's 3D XPoint™ non-volatile memory technology. As such, the key focus of the first months of the project was around gathering the requirements for this prototype, both in terms of the hardware functionality and the software it will need to support. Once the initial set of requirements had been agreed and submitted as an early deliverable, the next steps involved defining the NEXTGenIO architecture – an ongoing process.
All project partners had to work closely together from day one in order to obtain a full set of requirements that would lead to a complete and workable architecture; it has been an interesting and (at times) challenging exercise in co-design and our meetings have not been short of topics for discussion.

NEXTGenIO has already been presented at a number of events in both Europe and the US, and it has piqued the interest of many HPC users. The I/O bottleneck is a limiting factor even for today's supercomputers and addressing it will create new opportunities for large-scale data intensive scientific simulations.

The project is working hard to design and develop an architecture that will deliver transformational I/O performance. The first six months have laid the groundwork and we are looking forward to a productive time ahead!

Report: What We've Been Up To

Since the project officially began in October, we have attended and presented at a number of events. This has allowed us to really engage with the community over our objectives and spread the message of NEXTGenIO to a wide audience. Some of the events we have attended include:

The project took part in the HPC Workshop hosted by EXDCI in Rome in September 2015. This was an opportunity for members of FET and CoE projects to meet and discuss the projects.

We visited SC15, Austin, Texas in November, with all the partners handing out NEXTGenIO flyers. Mark Parsons also gave a talk at the Intel booth.

Tiago Quintino from ECMWF gave a talk at the 96th Annual Meeting of the American Meteorological Society in New Orleans in January 2016. This talk was specifically about how the NEXTGenIO system will be used for weather forecasting.

Mark Parsons presented a talk on the project at the 4th Annual ENES Workshop on Climate and Weather, in conjunction with ESiWACE. This took place in Toulouse in April 2016.

NEXTGenIO also took part in the European HPC Summit Week in Prague, in May 2016.

Slides from all talks and presentations are available online at nextgenio.eu/publications.

Work Package 2: Prototype System Architecture

One of the first tasks in the project has been to capture the requirements we have on the system we are developing. We then use these to define the architectures for the different components in the prototype we will eventually build. We actually need to create 3 different architectures:

- a hardware architecture, which defines the hardware components we will use, how they will be connected together, and packaged up. This is a bit more complicated than just defining the compute nodes for our prototype, and has to cover any login nodes, management resources, networks, and I/O systems.

- a software architecture, one of the key contributions that the project will make is the definition of a complete software architecture for enabling the use of NVRAM in scientific applications. We will use existing software components that are already under development for NVRAM usage, and implement any pieces that aren't currently available but we need to effectively use this technology.

- a data architecture, providing users with an understanding of how data can move and when it can be located in our hardware and software architectures, enabling users and developers to understand how to get the best performance on the system for their applications.

The first part of the process of defining the architectures was to capture all the requirements (from the application, development, and system deployment, operation, and management perspectives) and the constraints (from the hardware and software we know we will be using) that are relevant to the system. This was done by defining a number of different usage scenarios; descriptions of ways we could imagine a HPC system with NVRAM technology being used, and then refining these into distinct use cases. The use cases were then examined to generate specific requirements, that we documented and refined.

Constraints were gathered from the hardware partners in the project (Fujitsu and Intel) by examining the documentation for the hardware components were we using, and using their experience of designing and creating systems like the one NEXTGenIO is aiming to produce.

Now we have the requirements, we can start the task of designing architectures that will satisfy those requirements! More of that process next time.

Extensive Performance Analysis for NEXTGenIO

Partner Profile

Vampir implements optimised event analysis algorithms and customisable displays which enable fast and interactive rendering of very complex performance monitoring data.

The product has been developed at the Center for Applied Mathematics of Research Center Jülich and the Center for High Performance Computing of the Technische Universität Dresden. The development is continued by Center for Information Services and High Performance Computing (ZIH) of the Technische Universität Dresden.

Score-P
Scalable performance measurement infrastructure for parallel codes

VAMPIR

TECHNISCHE UNIVERSITÄT DRESDEN

Vampir has been available as a commercial product since 1996, and has been enhanced over time within the scope of many research and development projects.

The Vampir Performance Analysis Tool

Vampir is a framework for performance analysis, which enables developers to quickly study program behaviour at a fine level of detail. In contrast to a profiler a fine level of detail is achieved by means of recording and processing extensive program traces. An important and unique feature of Vampir is its intuitive and interactive graphical representation of detailed performance event recordings over time (timelines) and as aggregated profiles.

NEXTGenIO aims to exploit vast resources including NVRAM in an Exascale system. Utilisation of Vampir in NEXTGenIO project can allow extensive searching and filtering capabilities to determine critical bottlenecks related to either computation, communication or I/O. In contrast to traditional profiling, one may drill-down into the fine-grained details to discover the cause of a problem(s).

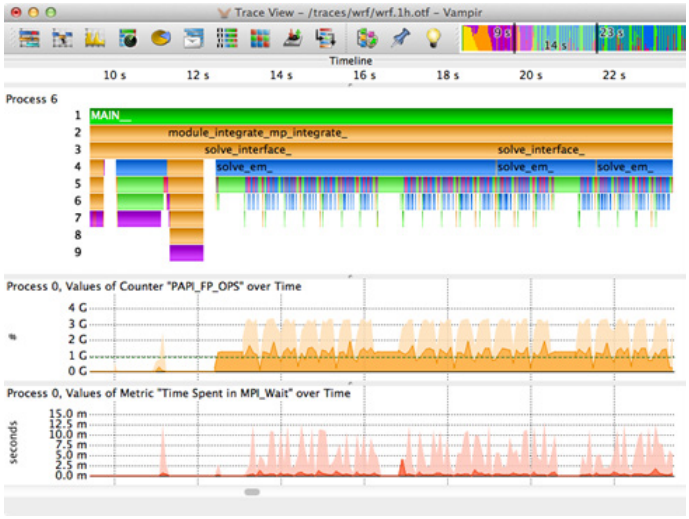

For NEXTGenIO, we aim to provide extensive performance monitoring for application development process with the following goals:

- Tracking of memory usage to support developing and tuning applications for NVRAM.
- Monitoring data transfers from/to NVRAM to determine memory extensive parts in application.
- Monitoring I/O operations and their behaviour to determine I/O bottlenecks for optimal resource utilisation.

This monitoring concept will be aggregated with the current state of the art performance analysis in order to support development of resource-friendly applications for Exascale systems.

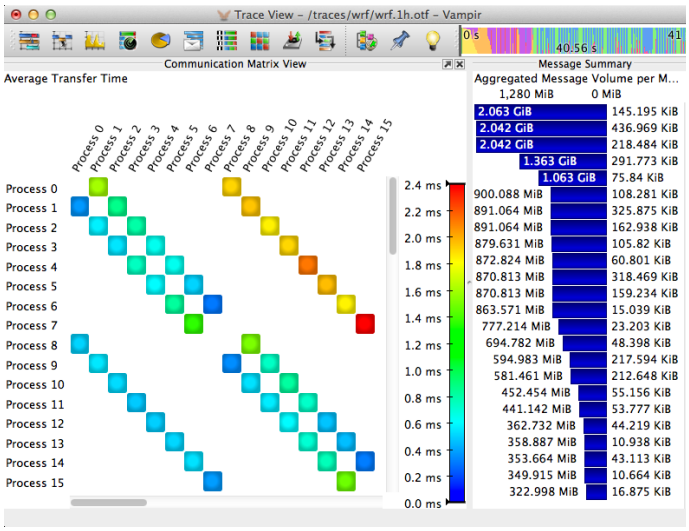

Dates for the Diary

ISC HPC 2016
Frankfurt, Germany
June 19 - 23

SUPERCOMPUTING 2016
Salt Lake City, Utah, USA
November 13 - 18

Workshops, Talks and Other Presentations

The project plans to stage a number of workshops, talks, BoFs, and other sessions at events throughout Europe and the wider HPC community in order to spread the project goals as widely as possible. These will continue throughout the life of the project.

At this time, we are pleased to announce the following sessions at ISC and SC 2016.

- A talk on NEXTGenIO at the ExaHYPE workshop at ISC

Thursday 23rd June 2016, 1pm - 6pm

- A BoF session at ISC 2016 and workshop at SC 16:
'Exascale I/O: Challenges, Innovations & Solutions'

Wednesday 22nd June 2016, 11.15am (ISC)

Friday, 18th November 2016, morning (SC)

In addition, there have been several talks that predate the publication of this newsletter.

More info regarding locations and registering to attend any events can be found on our website, nextgenio.eu .

Coming Up Next...

We are now heading into the second half of this first year in the project, and NEXTGenIO has proven to be every bit as exciting and challenging as we expected. The next six months will see considerable amounts of work to push the NEXTGenIO architecture into its implementation phase and progress the development of the supporting ecosystem. The project will also undergo its first technical review in July, giving us the opportunity to take stock, present our work and receive constructive feedback.

In the next six months, NEXTGenIO will also be represented at a number of events, such as the European HPC Summit, ISC High Performance and SC - if you are attending any of those occasions, please keep an eye out for our presentations and come speak to us. We are planning on publishing the 2nd edition of this newsletter in time for SC16; that issue will take a look back at the first complete year of NEXTGenIO. But until then, there's a lot of work still do be done!