

Workshop on Blue Gene Active Storage

The term active storage refers to computer architectures which comprise a storage sub-system that integrates significant processing power. The Jülich Supercomputing Centre plans in collaboration with IBM to realize such a concept for its Blue Gene/Q system JUQUEEN. As part of this project a workshop [1] was organized in January to bring computer architects and application developers together.

The architecture of I/O sub-systems is a particular challenge when progressing towards exascale machines. Already today it is difficult to maintain a reasonable balance between compute performance and performance of the I/O sub-system. In practice, this gap is widening and systems are moving away from Amdahl's rule of thumb for a balanced performance ratio, namely a bit of I/O per second for each instruction per second. Additionally, traditional disk based storage systems do not perform particularly well in case a large number of random I/O requests are performed.

Technology

Active storage is not a new concept (see, e.g., [2]). It has the potential of mitigating the mentioned problems as data processing is moved closer to the data. This approach therefore is additionally promising as it helps to reduce energy consumption due to data transport. Performance in terms of bandwidth and in particular I/O access rates as well as energy efficiency can additionally be improved by using non-volatile memory technologies like flash memory.

Fitch et al. [3] analysed already in 2010 the vision of an active storage concept based on the (at that time) emerging Blue Gene/Q architecture and solid state storage devices. The latter are integrated inside the compute racks and form one part of what becomes a tiered storage system. While the internal storage due to the high costs of suitable flash memory, i.e. SLC NAND flash, is limited in terms of capacity it provides high bandwidth and – compared to disk technology – very high I/O access

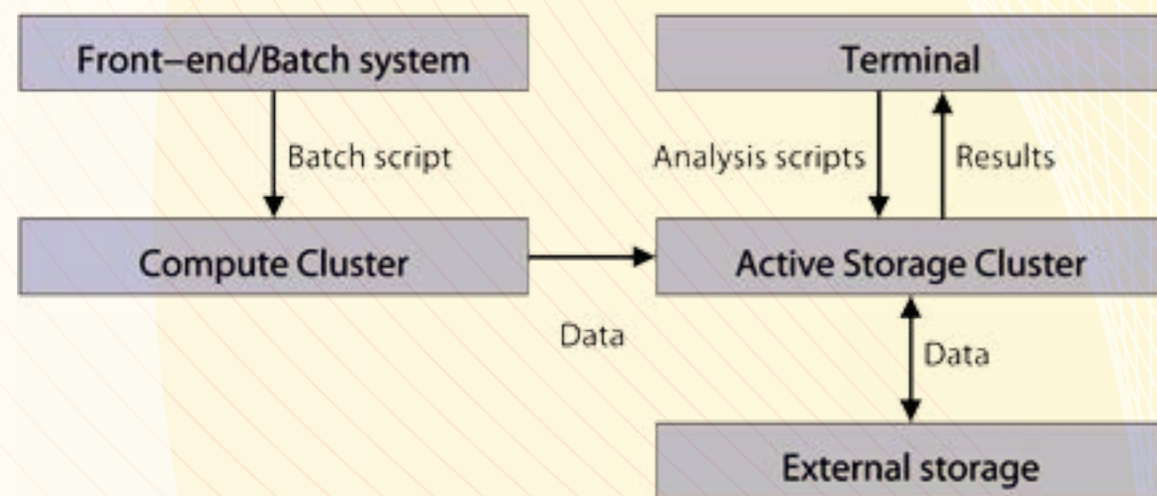


Figure 1. Work-flow of the neuronal network simulator NEST using active storage.

rates. The second tier, an external storage system based on traditional technologies, i.e. disk, continues to be available to provide large storage capacity.

Such a tiered storage architecture comprising active storage can be used in different ways. For instance, in case of applications where the amount of data generated is too large to be written to external storage systems the processing capabilities of the active storage enables data post-processing such that the remaining amount, which needs to be written to disk, is significantly reduced. Other use cases are out-of-core computations, where main memory capacity limitations are mitigated by temporary swapping data to storage, or multi-pass analysis, where multi-terabyte data-sets are randomly accessed many times.

Applications

Neuronal network simulators are an example for applications which plan to exploit BGAS for data post-processing. Markus Diesmann (Forschungszentrum Jülich) presented the vision of the developers of the simulator NEST (see Fig. 1). Fast storage being available will allow to not only write (and later analyse) information about spike events, but also about membrane potentials and synaptic weights. This information would otherwise be much too large to be written. The size of the simulated networks is mainly limited by the available memory. Since non-volatile memory could provide additional memory space, future node architectures comprising large amounts of such memory would en-

able simulation of realistic neural tissue models as shown by James Kozloski (IBM). Such an out-of-core computing ansatz could also be used by other applications. Stefan Blügel and Paul Baumeister (Forschungszentrum Jülich) considered this approach for calculations based on density functional theory.

Multi-pass analysis is a use case which occurs in genetic epidemiology or radio astronomy. Paolo Bientinesi (RWTH Aachen) analysed the computational requirements of genome association studies, which are used to examine common genetic variants of different individuals to identify variants associated with a trait. For different regions of the genome measured data have to be processed many times. In a related field, genomics, active storage allows to deal with the exploding amount of data generated by next-generation sequencing methods. David Carrera (Barcelona




Figure 2. Brain tissue containing two neurons needs to be volume decomposed for parallel simulations. Non-volatile memory can help to keep these volumes large. (© 2011 Kozloski and Wagner, Frontiers Media)

Supercomputing Center) presented his Parallel In-Memory Database, where the active storage is used to implement a key-value store.

Active storage concepts could also be utilized by climate science applications, as pointed out by Nathanael Hübbe (University of Hamburg), to implement lossless compression and thus reduce the growing amount of data written to and read from large capacity, external storage systems.

Managing increasing data volumes is also a challenge for research in astronomy and radio astronomy. David Champion (MPI for Radio-Astronomy) explained how to search for pulsars in petabytes of data generated by planned surveys of the universe with high time resolution. Astronomy is traditionally data driven, as pointed out by Alex Szalay (Johns Hopkins

University), who had been one of the architects of the archive of the Sloan Digital Sky Survey project. By making data accessible and by enabling any scientist to process this data, such archives turn into unique research instruments. Active storage concepts can help to bridge the gap between large capacity data services and HPC.

References

- [1] http://www.fz-juelich.de/ias/jsc/EN/Expertise/Services/Documentation/presentations/presentation-bgas_table.html
- [2] A. Acharya et al., "Active Disks: Programming Model, Algorithms and Evaluation," 5th conference on Computing frontiers, 2008.
- [3] B. Fitch et al., "Blue Gene Active Storage," HEC FSIO Workshop, 2010.

- Dirk Pleiter
- Marcus Richter

Jülich
Supercomputing
Centre



Figure 3. A galaxy-quasar combination found using the Sloan Digital Sky Survey (Credit: NASA, ESA/Hubble and F. Courbin (Ecole Polytechnique Federale de Lausanne, Switzerland)).

International Workshop on Cooperative Quantum Dynamics and Its Control (CQDC2012)

The "International Workshop on Cooperative Quantum Dynamics and Its Control (CQDC2012)" took place from 29 to 31 October 2012 at the rotunda of the Jülich Supercomputing Centre. The goal of the workshop was to discuss the possibility of essentially new quantum phenomena, and to discuss the observations and/or operations that would be required for their realization, detection, and understanding.

About 40 researchers from Germany, France, Switzerland, the Netherlands, Canada, the United States and Japan participated in the workshop. Recent developments in quantum annealing, open quantum systems, quantum computer hardware, quantum dynamics of spin systems, equilibration and thermalization of quantum systems, decoherence, entanglement and related topics were highlighted in talks and posters.

Quantum dynamics and novel quantum states in systems of many particles or spins, many of which have been beyond the reach of experimental realization, are now getting more and more realistic due to recent advancements in experimental techniques, such as the synthesis of molecular magnets, nano-engineering of quantum dots, time-resolved measurements with ultra-short pulses, and optical lattices of cold atoms. In particular, many attempts have been made for characterizing the wave function from the viewpoint of entanglement, which is often accomplished by the aid of supercomputers through the method of direct numeri-

cal solution of the evolution equation and/or through quantum Monte Carlo methods. In addition, new numerical methods, such as the tensor-network variational approximation, are making many previously hard problems now tractable.

The exploitation of quantum effects, requiring coherent control of dissipative dynamics and entanglement control, is expected to have profound implications for future emerging information technologies. Managing and designing complex quantum systems with specified behaviour for quantum information processing requires a deep understanding of the cooperative behaviour of their components. Unraveling this behaviour necessitates an intensive collaboration between theoreticians and experimenters. The workshop successfully presented an overview of the current research on various topics very closely and less closely related to quantum information processing.

- Kristel Michielsen

Jülich
Supercomputing
Centre

