The Square Kilometre Array Computing Challenges and the DOME Research Approach: Intermediate Results

A.J. Boonstra[1] and T. Engbersen[2]

[1]ASTRON, PO Box 2, 7990 AA Dwingeloo, The Netherlands, email: boonstra@astron.nl

[2]IBM IBM Research GmbH, Saeumerstrasse 4, CH 8803 Rueschlikon, Switzerland, email: apj@zurich.ibm.com

Abstract

The Square Kilometre Array (SKA) radio telescope is aimed to be much more sensitive than any radio telescope in existence today. As the compute challenges for the full SKA lie in the exa-scale range, research is needed to make the SKA affordable and feasible. The DOME project, a collaboration between ASTRON and IBM, addresses the compute challenges from three perspectives: sustainability, nano-photonics and extreme streaming. This paper describes the intermediate results of the DOME project.

1. Introduction

The Square Kilometre Array (SKA) is a radio telescope concept promising better sensitivities and larger survey speeds than any radio telescope in existence today. The SKA will be deployed in two phases. Construction of the first phase, SKA1 [1], is targeted to start in 2018 [2], one year after the planned critical design review of the SKA1. Detailed design of the SKA2 is expected to start in 2018, and early science results of the SKA are expected from 2020. Due to the wide bandwidths required, large field-of-view, high sensitivities, and high spatial resolutions, the computing challenges for the SKA lie in the exa-scale range. With current-day technologies the SKA2 would require well over 1 GW electric power to operate it. Even taking into account Moore's law, there probably will remain a gap between the available affordable compute resources and the required compute power if current designs are scaled-up naively. In order to address this challenge, ASTRON and IBM started the DOME project [3] early 2012. This project addresses the SKA compute challenge from three perspectives: (1) Sustainable Computing addresses technologies to radically reduce the power needed to do computationally intensive work on extremely large amounts of data, (2) Nano-Photonics addresses technologies needed to drastically reduce the power of data transport over longer distances and inside computing machines, and (3) Data & Streaming addresses technologies to process data on-the-fly and store data at a high efficiency for later use. This paper presents intermediate DOME project results in these three research areas.

2. Facing the SKA computing challenge in DOME

2.1 Algorithms and Machines

The three areas mentioned in the introduction are connected via a central DOME sub-project "Algorithms and Machines" which is the coordinating architectural activity in DOME. The purpose of this work is to develop a design methodology, based on mathematical principles addressing fundamental equations that govern the system. This methodology will allow to reason on system optimizations and parameter space. Initially tested on LOFAR [4] in a retrospective analysis, this approach has been applied to phased-array synthesis telescopes [5] and to the SKA1 [6]. The analysis results using an end-to-end compute model of the SKA [6] based on the SKA1 baseline design [1] results in 4.2 exa operations per second (Eops) for the most demanding operational mode of one of the three envisioned SKA1 concepts: the imaging processing (gridding to allow a two dimensional FFT, and visibility prediction) for the "Mid-Array".

Extrapolation of the current trend in top-500 (green) supercomputers shows an estimated 500,000 Mflops/W in 2020, as shown in Figure 1. This is, also for the "green" list, not sufficient for the project SKA1 timescales for the most demanding mode. This "gap" can be reduced by relaxing part of the requirements, or by applying alternative approached such as compressed sensing, imprecise computing, or novel algorithmic approaches. As the bulk of the compute resources is spent in the gridding and predict/imaging step, alternative approaches could lead to huge gains in processing requirements. These approaches are studied in DOME as of early 2014.

978-1-4673-5225-3/14/$31.00 ©2014 IEEE

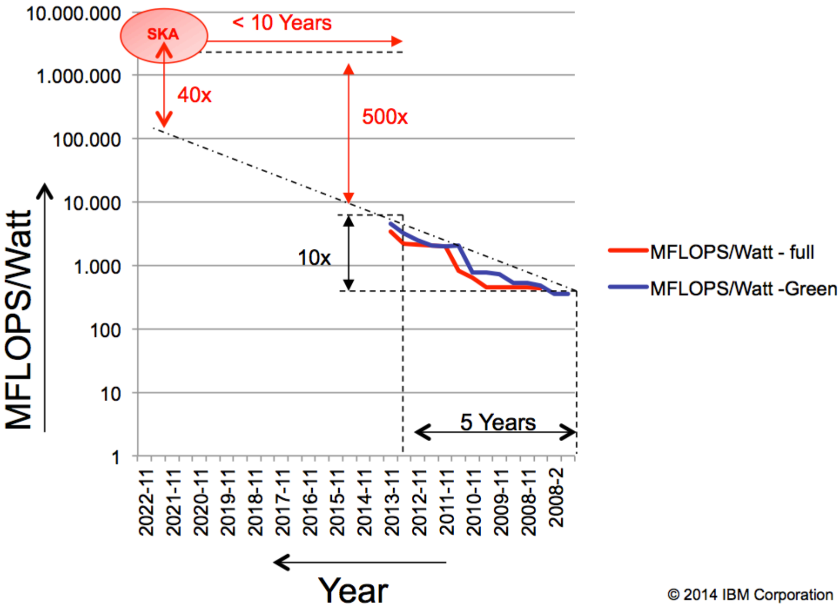

Figure 1 Graphs showing the development over time of the power required by the top-500 supercomputers, both the full list and the "Green" list.

2.2 Sustainable Computing

"Sustainable Supercomputing", the reduction of energy consumption in large computer systems and data centres is crucial for the SKA, but it is also a challenge in a much broader sense, as major Internet data centres and search engines require more and more computing power. There is a great economic importance to have the latest technology and associated knowledge available when deploying compute and datacentres. The largest fraction of power consumption in computing systems is spent on data transport rather than on the compute (multiplication, addition) itself. Therefore, miniaturisation, not just within chips but also in systems is part of achieving sustainable computing facilities. In this context, the DOME team has designed and built a prototype micro-server [7], a 64 bit high-density capability demonstrator designed for running server class operation systems. The system is designed to use hot water cooling, aimed at 40% reduction in power consumption compared to air-cooled systems. Water cooling also leads less thermal cycling which improves reliability. Figure 2 shows a photo of the first running prototype of a micro-server. The micro-server has a DIMM memory module form factor.

Figure 2 Micro-server board with Freescale P5020 dual-core chip mounted onto a test board

Based on initial insights provided by using the model that was developed for the LOFAR retrospective analysis, as described in section 2.1, a modular approach for designing the initial processing steps of a SKA correlator was chosen. By means of three integrated circuits [8], which can be connected with each other in various configurations, about 70 % of the desired SKA signal processing functions can be realized. The power consumption is estimated to be a factor three smaller than the approach that has been used in LOFAR. This means that the integrated circuit option clearly stays well within the SKA power requirements.

However, in the SKA signal processing stream there are still several places at which very complicated calculations are required. With such large amounts of data as the SKA generates and needs to process, it is important to use the correct computer architectures, whether it is based on modern Graphics Processor (GPU) or special hardware (accelerators). The DOME project is currently analysing and measuring the compute efficiency and power efficiency of many different radio astronomy algorithms mapped onto GPUs, CPUs, DSPs, and other accelerators. Results are expected later this year, and it will be input to the SKA designs.

2.3 Nano-Photonics

Using optical connection technology", Nano-photonics", exponentially more data can be transferred as compared to conventional electrical connections. In addition, the energy consumption is much lower. The DOME project aims at adapting both existing and future optical technologies, and evaluates applicability of these technologies for the SKA. Based on measurements made with Avago optical connections ASTRON decided to use this technology in their next generations of data processing systems.

As the requirements for the design of the SKA have slowly become more concrete, the DOME project looked into the use of optical data transfer. Given the enormous sensitivity of the SKA antennas for noise, in particular also noise caused by human activity (electronics), a study was initiated to analyse the use of analogue antenna signal transport through optical fibres. This would make it possible to convert the analogue signals into digital signals in a shielded location (Faraday cage), at least at some 100 meters distance from the antennas arrays. The first theoretical results show that this approach is just at the limit of what can be achieved. Therefore, we decided to build a demonstrator in order to get more accurate data and insights in the issues.

As indicated in section 2.1, the SKA compute challenge could also be tackled by investigating compression techniques and alternative algorithmic approaches. For example, at some locations in the signal processing stream imprecise computing techniques may be applied, and there may be algorithmic alternatives for the computationally expensive gridding algorithm. These studies potentially may have a big impact on the SKA designs. This line of research has just started within the DOME project; the first discussions between ASTRON and IBM show that this is an extremely complex area but with high potential gain.

2.4 Data and Streaming

More and more devices and sensors are permanently "on-line", producing so much "Extreme Streaming" data that not all raw data can be stored. The continuous process of this flow of information requires innovative data reduction hardware and software, not only for scientific instruments, but also for safety applications, weather forecasting, logistics guidance, prognostication, etcetera. In order to handle the large streaming data volumes, meaning must be derived from the data streams, and data reduction must be applied. In processing chains, often a multi-tiered storage approach is needed with fast-access memory in the first part of the processing chain, and slower bulk storage systems at the end of the chain.

In radio astronomy, after correlation and calibration, data are stored for subsequent (re-)calibration and astronomical processing and feature extraction. As the datasets are becoming extremely large, optimal multi-tiered storage systems are required. The LOFAR telescope for example [4] already maintains an archive in excess of 20 peta byte. Different storage media have different access times and cost: tape has extremely low cost (and power requirements) but significant access times, solid state devices have very low access times but are rather costly, and the properties of disks lie between those of tape and solid state media. Other promising technologies include novel phase change memories.

The DOME project designed and constructed a multi-tier system model for simulating a full range of storage technologies, the "memory-tiers": from high-speed memory technologies over hard disk to magnetic tape technology. This simulation system derives the optimal mix of these technologies for given costs, memory sizes and memory latency requirements.

Finally, IBM extended the distributed hard-disk data system GPFS (General Parallel File System) with LTFS (Linear Tape File System) based tape storage. Project DOME, TARGET and the University of Groningen have defined a test-system, aimed to assess the usability of GPFS and LTFS for DOME and SKA. This system will be put into operation in February 2014.

3. Conclusion

The work done in the DOME project shows that significant processing improvements can be made by applying carefully designed signal processing, computing, and storage modelling techniques. An important cost driver for the SKA is power consumption. Careful design, applying clever scaling approaches, and Moore's law (with its changing rate) will help ensuring that the full SKA can be built for affordable cost. However, there still is a gap concerning required power efficiency, which must be solved by applying novel algorithmic approaches.

4. Acknowledgements

DOME is a Public Partnership between IBM and ASTRON, the Netherlands Institute for Radio Astronomy. ASTRON is an institute of the Netherlands Organisation for Scientific Research, NWO. Project DOME is supported by the Dutch Ministry of Economic Affairs, and by the Province of Drenthe. The presented results are based on work conducted by the DOME teams located at ASTRON and at the IBM Research Laboratory in Zurich.

5. References

1. P.E. Dewdney, W. Turner, R. Millenaar, R. McCool, J. Lazio, and T.J. Cornwell, SKA1 System Baseline Design, SKA Technical Report SKA-TEL-SKO-DD-001, March 2013, see also www.skatelescope.org
2. SKA Office, The Project Timeline, Url: www. https://www.skatelescope.org/project/projecttimeline/
3. P. Chris Broekema, et al, DOME: Towards the ASTRON & IBM Center for ExaScale Technology, Proceedings of the 2012 workshop on High-Performance Computing for Astronomy (AstroHPC'12), Delft, the Netherlands, June, 2012
4. M. van Haarlem et.al., "LOFAR: The LOw-Frequency ARray" Astron.& Astroph., arXiv:1305.35550v2, Vol. 556, Aug 2013, 53pages, DOI:10.1051/0004-6361/201220873, see also: www.LOFAR.org
5. A. Anghel, R. Jongerius, G. Dittmann, J. Weiss, and R. Luijten, Holistic Power Analysis of Implementations Alternatives for a Very large Scale Synthesis Array with Phased Array Stations, IEEE ICASSP 2014, accepted for publication
6. R. Jongerius, S. Wijnholds, R. Nijboer, and H. Corporaal, "End-to-end compute model of the Square Kilometre Array", 2014, IEEE Computer, accepted for publication.
7. R. Luijten, "IBM high-density micro-server demonstration platform leveraging PPC, Linux and hot2water cooling", ASIA Power Architecture Conference, Shanghai, China, October 2012.
8. M. Schmatz, et al, "Scalable, Efficient ASICs for the Square Kilometre Array: from A/D Conversion to Central Correlation", ICASSP 2014, Florence, Italy, May 2014, accepted for publication