

TRYTON SUPERCOMPUTER CAPABILITIES FOR ANALYSIS OF MASSIVE DATA STREAMS

Henryk Krawczyk, Prof.

Michał Nykiel, Ms. C.

Jerzy Proficz, Ph. D.

Gdańsk University of Technology, Poland

ABSTRACT

The recently deployed supercomputer Tryton, located in the Academic Computer Center of Gdansk University of Technology, provides great means for massive parallel processing. Moreover, the status of the Center as one of the main network nodes in the PIONIER network enables the fast and reliable transfer of data produced by miscellaneous devices scattered in the area of the whole country. The typical examples of such data are streams containing radio-telescope and satellite observations. Their analysis, especially with real-time constraints, can be challenging and requires the usage of dedicated software components. We propose a solution for such parallel analysis using the supercomputer, supervised by the KASKADA platform, which with the conjunction with immerse 3D visualization techniques can be used to solve problems such as pulsar detection and chronometric or oil-spill simulation on the sea surface.

Keywords: supercomputer, data streams, distributed systems, radio-telescope, satellite

INTRODUCTION

The emerging need to analyze radio-telescope data means that scientists are required to use the most sophisticated HPC equipment in their research. The amount of data produced by such devices can only be processed by extremely efficient computers, usually large computer clusters. Thus, HPC solutions seem to be the most appropriate for radio-telescope observation data analysis. Because of the (soft) real-time constraints, dedicated software solutions need to be used.

Moreover, due to the possible interference with the space observations, such powerful computation equipment cannot be located next to the radio telescope, so fast and reliable data transfer is yet another problem.

Currently used at Nicolaus Copernicus University in Torun, radio telescope RT-4 has a 32 m diameter and a number of receivers for various frequency bands, however the proximity to urban areas decreases its usability. Thus, a new planned radio-telescope RT-90+ is going to have a 90 m (or even larger) diameter and will be located in Tuchola Forest, where the interferences are significantly lower (even 1 to 1000 ratio). Such a sophisticated device gives a great research opportunity, but it also introduces new challenges. One of the most important ones is the amount of data gathered during the observations. The data streams produced by the radio-telescope need to be transferred, decoded and analyzed by a powerful computer system, and due to the usual interference produced by digital devices, located at some distance from the radio-telescope.

Considering the above constraints, we perceive the Academic Computer Center of Gdansk University of Technology to be an excellent choice for the gathered data analysis. It is the nearest data center to the proposed radio-telescope's location and the connection to the PIONIER network enables fast data transfer. The Center is also equipped with the most powerful computer in Poland (1.37 PFLOPS) and finally, thanks to the Center of Excellence in Scientific Application Development Infrastructure (CD NIWA), it provides a suitable base for proper software solutions designated for the described data analysis.

The rest of the paper provides the information about the potential of the Academic Computer Center including the Tryton supercomputer. Two following sections present more information about the processed data, the next one describes the supercomputer including its internal characteristics and networking capabilities. Section 5 presents the KASKADA platform as a universal solution for data stream processing, section 6 describes immerse 3D visualization, and finally in the last section, some conclusions and future works are provided.

RADIO-TELESCOPE DATA STREAM

One of the most popular techniques in the field of radio astronomy is called astronomical interferometry. It involves using multiple connected radio-telescopes and superimposing

the signal waves to significantly increase resolution of the observations. The maximum resolution is determined by a baseline, i.e. the distance between the furthest radio-telescopes. The most known examples of using this method are the Very Large Array (VLA) observatory in USA with a maximum baseline of 36 km and Multi-Element Radio Linked Interferometer Network (MERLIN) in England with the longest baseline of 217 km.

It is possible to use interferometry with even more distant radio-telescopes, for example located in various countries across Europe. This method is known as Very Long Baseline Interferometry (VLBI) and has been used since the 1960s [9]. Until recently, sending data over such distances was possible only by physically transporting magnetic tapes or hard drive arrays to a central location. However, with modern computer networks based on fiber-optic cables it is possible to transfer large amounts of data in real-time, significantly speeding up and simplifying the observation process.

In order to transmit data from a radio-telescope through the Ethernet, it must be first converted from an analog to a digital signal. The Digital Baseband Converter (DBBC) performs this transformation with sampling frequency up to 1 GHz. Digital samples are then transmitted to a VLBI data system such as Mark 5. The data system bundles samples into data frames and aggregates them into streams and threads which can be sent over the Ethernet using TCP or UDT protocol. This sequence is shown in Figure 1.




Fig. 1. Data stream flow from a radio-telescope to a computer cluster

Using a supercomputer for data processing is beneficial in many ways. First, the data stream can be analyzed in real-time and the output is almost immediately visible for the observer. Besides the obvious advantage of obtaining results of the observation faster, it allows the researchers to, for example, recalibrate the radio-telescope if some interesting event is detected. Additionally, the supercomputer could be placed at a remote location, many kilometers from the observation site, which results in reduction of radio noise produced by CPUs and other computer equipment.

A typical single e-VLBI data stream contains multiple threads and has a total bandwidth ranging from 1 Gbit/s (Mark 5B) to 16 Gbit/s (Mark 6). For example, the RT-30 radio-telescope, which is used by Torun Centre for Astronomy at Nicolaus Copernicus University, produces a data stream with 2 polarizations, 4 frequency bands, each with 2 sub-bands, sampled at 32 MHz and using 2 bit digitization. As mentioned before, the interferometry technique requires data from multiple streams to achieve high resolution of the observation, which results in even higher data throughput.

Perhaps one of the greatest challenges would be processing and analyzing the data acquired by RT-90+ radio-telescope, which will be built in Tuchola Forest in Poland. A radio-telescope with a diameter of over 90m and state-of-the-art equipment could produce even several Tbit/s of observation data samples.

SATELLITE DATA STREAM PROCESSING

Satellite images are commonly used in Geographical Information Systems (GIS) to support decision-making in urban development and infrastructure protection. In general, it is realized by integrating numerical models of hazard scenario simulations with existing information from geodatabases and earth observation data streams from satellites.

SafeCity GIS [6] is an innovative system that allows monitoring threats and analyzing risks for the municipal areas developed by the Geoinformatics Department of Gdansk University of Technology. The system can integrate data from multiple sources, including data from a 1.5 m-wide HRPT/MetOp-A/B satellite ground station and information from Service Support Environment (SSE), and then perform various hazard simulations like blast and chemical attacks, floods or oil-spill.




Fig. 2. Architecture of a GIS system integrated with computing cluster

Such simulations often require significant computing power, especially when risks must be analyzed in real-time. For example, oil-spill simulation requires parallel modeling of sea currents and winds, analyzing diffusion of oil spill droplets, processing vaporization and dispersion, and forecasting the size and shape of oil-spill for several hours ahead. The GIS system that handles these types of simulations must offload computation to a high performance computing cluster. Figure 2 demonstrates this concept.

SafeCity GIS system includes two client applications: the complete web client and mobile client. Due to the nature of mobile devices and their restrictions, the mobile client does not offer all functionalities that are available in the complete web client: for example launching a hazard scenario simulation is not possible on a smartphone or tablet. These limitations could be bypassed by offloading some tasks from the mobile device to the cloud [8]. However, the Internet connection required for communication with the cloud services may not be always available and developers have to use adaptive application architecture to overcome this problem, i.e. enable or disable features based on the context of a mobile device, or move the execution between the mobile device and the cloud dynamically.

One of the solutions to this problem is the Mobile Offloading Framework (MOFF) [3] developed as a part of the Center of Excellence in Scientific Application Development Infrastructure “NIWA” project. It provides mechanisms for optimizing the execution of mobile applications in run-time, based on various parameters, such as network availability, network speed, CPU power or battery level. The framework is constantly monitoring parameters of the device and execution time of mobile application components to optimize user experience. The process of offloading execution to the cloud is transparent for the user, as well as for the application developer. The concept of integrating the mobile devices with the cloud is presented in Figure 3.




Fig. 3. The concept of Mobile Offloading Framework

TRYTON SUPERCOMPUTER DESCRIPTION

Tryton supercomputer [10] is the first computer cluster in Poland which exceeded 1 PFLOPS computing power. It is located in the Academic Computer Centre in Gdansk University of Technology and was built to support the newly created Center of Excellence in Scientific Application Development Infrastructure “NIWA”. Its primary purpose

is to provide computing capabilities for the scientific community, including radio-telescope and satellite data streams processing.

Every node of the cluster contains two processors (Intel Xeon Processor E5 v3, 2.3GHz, Haswell), with 12 physical cores (24 logical ones, due to Hyperthreading technology) and 128-256GB RAM memory. Some of the nodes have additional accelerators, such as nVidia Tesla, Intel Xeon Phi or AMD FirePro.

In total the supercomputer consists of 40 racks with 1483 servers (nodes), 2966 processors, 35592 compute cores, 48 accelerators, 202 TB RAM memory. It uses fast FDR 56Gb InfiniBand in fat tree topology and 1Gb Ethernet for networking. The total computing power is 1.37 PFLOPS. The cluster weighs over 20 metric tons.

The computer center also has extensive storage capabilities. Within a few projects, such as PL-Grid, Platon (U4) or Mayday Euro 2012, storage servers can store up to 5PB of data, using modern hard disk and tape archives. All this storage space is easily accessible for the applications and services working on the supercomputer, via InfiniBand or other optic fiber connections.

The supercomputer is placed in one of the most important nodes of the wide area network connecting the Gdansk metropolis and the whole country. The MAN (metropolitan area network) is based on 10 Gigabit Ethernet & ATM technologies. It creates a ring structure connecting academic nodes in Tricity, consisting of 260 km of the fibers in total.




Fig. 4 The PIONIER network structure. Source: <http://www.pionier.net.pl>

The Internet outside connection is facilitated by the PIONIER network (the computer center is one of the shareholders), providing fast access to the current radio-telescope in Torun, and it can be easily extended to the possible location of RT-90+ in Tuchola Forest. The foreign connection is provided by the access to Geant and CERN (through Amsterdam) networks, see Fig. 4.

KASKADA PLATFORM – FRAMEWORK FOR DATA STREAM PROCESSING

The KASKADA platform is a solution dedicated to processing of continuous streams of data. Initially, it was designed for video and audio processing, however with time it was extended for other data types, like endoscopy examination results or text documents. Moreover, the platform provides the support for high numbers of processed streams and is capable of doing so satisfying (soft) real-time constraints.

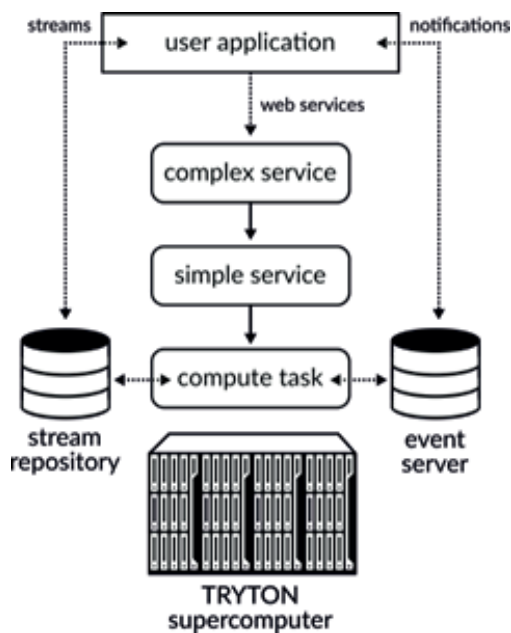


Fig. 5. KASKADA platform architecture

Fig. 5 presents the layered architecture of the KASKADA platform. It is based on the computer cluster equipped with multiple compute nodes connected by a fast network (see section 3 for more details). The second layer uses compute tasks to perform processing dependent on the implemented stream processing algorithm. The next layer provides the support for web services, the simple and the complex ones, respectively. Finally the top layer: user applications enable direct interactions using the proper GUI mechanisms.

The additional components: stream repository and event server are used to support the data stream online computations. The repository uses the mass storage for archiving and managing the incoming streams to facilitate the compute and to provide the data to the user through the user application. The event server provides the means to store and transfer the messages about the events detected during the online processing, and deliver them to the user with the various protocols. See [2] for more details about the architecture and task/service management in general, and [4],[7] specifically for the KASKADA platform.

It is important to emphasize the support for service scenarios implemented in the complex service layer. They enable usage of the pipeline processing including parallel execution of the underlying algorithms. This approach is

crucial for processing of massive data streams because they must be distributed between multiple computing nodes to satisfy real-time constraints. When the user chooses to start some data processing scenario, the application sends a request to the KASKADA platform. The corresponding complex service is selected from the repository, required simple services are assembled, and one or more compute tasks are executed for every simple service. These tasks communicate directly with the stream repository to send or receive data streams and, with the event server, to send notifications to the user application.

Fig. 6 presents an example of such a scenario described in MSP-ML (Multimedia Stream Processing Modeling Language) [1]. The first service (Decoder) is responsible for decoding the data, the second one (SpectralDensity) is used for converting the data from a time to a frequency domain, the two others analyse the pulsar timing (PulsarTiming) and produce the data used for plotting the spectrum function and time series data for the user (GraphPlot). The pulsar timing is only one example of a service scenario that could be used for analyzing data streams from a radio-telescope. Other scenarios that could be easily executed on the KASKADA platform are pulsar discovery, polarimetry or detecting unknown astronomical objects.




Fig. 6. An example of a service scenario in the KASKADA platform

On the other hand, the KASKADA platform can be also used for non-stream related processing, including such services like simulations, numerical computations or implementation of other scientific algorithms. The platform supports the master-slave processing model (see Fig. 7), where the problem can be split into parts in the master task, and then sent to the slaves; who can be placed on different cluster nodes, increasing the speedup of computations.




Fig. 7. Master-slave processing model

IMMERSE 3D VISUALIZATION LABORATORY

The KASKADA platform is also connected with the Immerse 3D Visualization Laboratory [5] that could be used to visualize the results of data processing. The aim of I3DVL is to ensure the highest possible degree of immersion together with the least amount of equipment worn by the user. To achieve this, the laboratory is equipped with a Spherical Walk Simulator placed in a cubic CAVE (Cave Automatic Virtual Environment). Visualization of the laboratory is shown in Figure 8.



Fig. 8. Visualization of the rotary sphere inside the CAVE. Source: [5]

The Spherical Walk Simulator is a transparent sphere rotating on rollers. Therefore, the user can move freely in any direction without actually changing location. Transparent panels allow the user to see the CAVE screen-walls through the sphere. The image is rear-projected on all six faces of the cube structure using stereoscopic projectors. Every wall displays an image with a resolution of 1920x1920 pixels and stereoscopy is achieved either, using an active shutter, or spectrum channels separation glasses.

The I3DVL uses 12 projectors which are driven by a visualization cluster consisting of 12 computing nodes with Nvidia GPUs, connected in an all-to-all communication model. The visualization cluster is controlled by two management nodes, which are responsible for analyzing the context of the simulation, i.e. user position, speed and direction of movement. Architecture of the I3DVL is presented in Figure 9.

This computation model is suitable for visualization of authentic scenarios modeled by maritime applications, including oil-spill simulation, algae bloom monitoring or physical simulation of a ship in rough seas. The results of data processing from the radio-telescope and satellites could be easily transferred from the KASKADA platform to I3DVL. 3D models of planets, stars, galaxies and other astronomical objects could be generated from the observation data and used to better study and analyze results of data processing. The concept of this solution is presented in Figure 10.




Fig. 9. Model of the I3DVL




Fig. 10. Immerse 3D Visualization Laboratory

CONCLUSION

The paper presented the potential of the ACC in GUT for radio-telescope and satellite data processing. Three different computational solutions are suggested (The KASKADA pipeline, master/slave, and I3DVL). The hardware capabilities of the Tryton supercomputer (the highest computational power in Poland), as well as the software provided in the center (KASKADA platform) enables dependable and fast adaptation to the current computational requirements following the different radio-telescope and satellite analysis scenarios, and are ready for the challenges of the new, planned, equipment (e.g. RT-90+).

Although the platform is ready to use for any data stream processing applications devoted to sea and space analysis, we plan to introduce additional features especially designed for the radio-telescope data, such as dedicated input data archiving, or output stream generation. Also, additional effort is required to provide the proper GUI interface and frontend application, first of all for CAVE presentations.

ACKNOWLEDGEMENT

This work was carried out as a part of the Center of Excellence in Scientific Application Development Infrastructure "NIWA" project, Operational Program Innovative Economy 2007–2013, Priority 2 "Infrastructure area R&D".

REFERENCES

1. Bobkowska, A., Nykiel, M., Proficz, J., A Model-Driven Solution for Development of Multimedia Stream Processing Applications, Proceedings of the 13th International Conference on Software Engineering, Parallel and Distributed Systems, 2014
2. Hennessy, J. L., Patterson, D. A., Computer architecture: a quantitative approach, Elsevier, 2012
3. Krawczyk, H., Nykiel, M., Proficz, J., Mobile Offloading Framework: Solution for Optimization of Mobile Applications using Cloud Computing, Communications in Computer and Information Science, 2015
4. Krawczyk, H., Proficz, J., Real-time multimedia stream data processing in a supercomputer environment, Interactive Multimedia, InTech, 2012
5. Lebiedź, J., Mazikowski, A., Innovative Solutions for Immersive 3D Visualization Laboratory, WSCG2014 Conference on Computer Graphics, Visualization and Computer Vision, 2014
6. Moszyński, M., Kulawiak, M., Chybicki, A., Bruniecki, K., Bieliński, T., Łubniewski, Z., Stepnowski, A., Innovative Web-based Geographic Information System for Municipal Areas and Coastal Zone Security and Threat Monitoring Using EO Satellite Data, Marine Geodesy, 2014
7. Proficz, J., Krawczyk, H., Task allocation and scalability evaluation for real-time multimedia processing in a cluster environment, proceeding of 6th International Conference on Pattern Recognition and Machine Intelligence, Warsaw, Poland, 2015 (accepted)
8. Rittinghouse, J.W., Ransome, J.F., Cloud Computing: Implementation, Management, and Security, CRC Press, Taylor and Francis Group, 2010
9. Thompson, A., Moran, J., Swenson Jr, G., Interferometry and synthesis in radio astronomy, John Wiley & Sons, ch. 9, pp. 304-325, 2008
10. Tryton Supercomputer, Academic Computer Centre in Gdansk, <http://task.gda.pl/kdm/sprzet/tryton/>

CONTACT WITH AUTHOR

Henryk Krawczyk
Michał Nykiel

Faculty of Electronics, Telecommunications
and Informatics,
Gdańsk University of Technology,
G. Narutowicza 11/12 Str.
80-233 Gdansk
POLAND

hkrawk@eti.pg.gda.pl
mnykiel@task.gda.pl

Jerzy Proficz

Academic Computer Centre in Gdansk
Gdańsk University of Technology
ul. G. Narutowicza 11/12 Str.
80-233 Gdansk
POLAND

jerp@task.gda.pl