HIGH-PERFORMANCE COMPUTING FOR SOIL MOISTURE ESTIMATION

S. Elefante[1], W. Wagner[1], C. Briese[2], S. Cao[1], V. Naeimi[1]

[1] Department of Geodesy and Geoinformation, Vienna University of Technology, Vienna, Austria
[2] EODC Earth Observation Data Centre for Water Resources Monitoring, Vienna, c/o University of Technology, Vienna, Austria

ABSTRACT

The aim of this work is to study the performance of a Synthetic Aperture Radar (SAR) processing chain for soil moisture retrieval using a supercomputer (Vienna Scientific Cluster 3) and the capabilities of the Earth Observation Data Centre (EODC). Different processing tests were designed and performed using three benchmark data sets: the whole ENVISAT ASAR Global Mode archive, part of the ASAR Wide Swath Mode archive and one large Sentinel-1 data set. The experiments showed the feasibility of processing large amount of earth observation data using the services offered by the EODC supercomputing environment.

Index Terms— Big data, parallel computing, HPC, high performance computing, soil moisture

1. INTRODUCTION

The availability of large amounts of Synthetic Aperture Radar (SAR) data stemming from the completed ENVISAT satellite mission and from the Sentinel-1 satellites constellation represents a major opportunity for the worldwide scientific community. Earth observation satellite data are widely employed for many operational and scientific applications. One of the chief applications of Sentinel-1 over land is to monitor surface soil moisture that is regarded as a key parameter for flood forecast and numerical weather prediction systems. [1] [2]. However, the processing of such a massive data amount can be executed only by following the well-established paradigm of distributed computing. The processing of the whole ENVISAT archive and of the Sentinel-1 data in a time frame that allows researchers to make scientific conclusions within a reasonable time span can occur only by using computing environments that fully support parallel processing. The high performance computing (HPC) system represented by the Vienna Scientific Cluster 3 (VSC-3) [3] suits to this aim well. The VSC-3 is a large distributed platform installed in summer 2014 in Vienna. Due to the massive data volume, the accessing and downloading of the raw data is also not a trivial issue. To effectively deal with these challenges, the paradigm is to bring the processing algorithms close to the data as well as to the computing resources. The Earth Observation Data Center (EODC) [4] provides an effective answer to this paradigm as it is a private-public partnership aiming at delivering a collaborative cloud infrastructure for archiving, processing, and distributing EO data. Through multi-national partnerships from science, the public and private sectors, users can get direct access to Sentinel data storage and running data-intensive geoscientific models. One of the services offered by EODC is the capability of performing data processing through VSC-3. EODC environment is, therefore, a comprehensive infrastructure in which users have the possibility to directly access EO data, process them with their own algorithms and retrieve the final products.

The aim of this work is to assess the possibility of processing Big Data for remote sensing applications with a particular attention to the TU Wien soil moisture retrieval processing chain. The results of five case studies, which were performed by using the EODC infrastructure and the VSC-3 high performance computing platform, are shown.

3. TU WIEN SAR TOOLBOX AND PRODUCTS

The TU Wien SAR Geophysical Toolbox (SGRT) is a software package developed to process Envisat and Sentinel-1 Synthetic Aperture Radar (SAR) data. The SGRT is written in the Python programming language and includes some external software modules, in particular ESA's Sentinel-1 toolbox (S1TBX) for SAR data geocoding, radiometric corrections and calibration [5]. The SGRT program is split into three main processing components to ease workflows traceability and reproducibility: pre-processing, model parameter extraction, and data production. During the pre-processing step, the SAR data are first calibrated and georeferenced using S1TBX. After some quality checks and data conversions and corrections, the geocoded SAR scenes in orbital image format are resampled to the TU Wien predefined fixed, planar grid called *Equi7 Grid* that consists of seven continental grids in Azimuthal Equidistant projection. The TU Wien Equi7 Grid is designed to minimize the oversampling rate of the high resolution satellite data globally, while keeping its structure simple [6]. The output of the pre-processing step are time series of the

doi: 10.2788/854791

terrain corrected and georeferenced SAR data resampled to the respective continental Equi7 subgrid.
The SGRT is aimed to be easily adapted for new products and workflows. As a result, a wide range of algorithms can be implemented within the SGRT framework for generating various products. One of the main SAR derived products developed by TU Wien is the Surface Soil Moisture (SSM) product.

3. VIENNA SCIENTIFIC CLUSTER

The Vienna Scientific Cluster-3 (VSC-3) is an advanced HPC system and consists of 2020 nodes, each equipped with 2 processors (Intel® Xeon® Processor E5-2650 v2 from the Ivy Bridge-EP family) and internally connected with an Intel QDR-80 dual-link high-speed InfiniBand fabric. The Scientific Linux release 6.6 is installed on each node as operative system. The Simple Linux Utility for Resource Management (SLURM), that is an open source, fault-tolerant, and highly scalable cluster management and job scheduling system for large and small Linux clusters, is installed as middleware [7]. SLURM organizes the access to the computing nodes through the management of a queue of pending work and provides a framework for executing and monitoring jobs. The parallel cluster file system BeeGFS (formerly FhGFS) [8] is installed on the VSC-3 distributed volume which is constituted by 360 spinning disks connected through around 160 Gb/sec bandwidth (evaluated experimentally). The satellite data are stored in 93 NFS disks (83 x 4T, 10 x 8T) of the EODC data archive. They are connected through high-speed InfiniBand (40 Gb/sec) to the VSC-3 system.

4. CASE STUDIES

Five case studies were executed to evaluate the VSC-3 performance in processing SAR data sets of different size and spatial resolution of the images. Table 1 indicates specifications and processing results of the five case studies. In all case studies, the SGRT pre-processing workflow was selected for the evaluation which is the most time-consuming part of the SAR level-1 data processing. The workflow comprises calibration, orbit corrections, geo-referencing including terrain correction using a digital elevation model (DEM), image-edges' noise removal, data conversion and encoding, and resampling to Equi7 Grid. Three representative sample data were considered as benchmark: Envisat ASAR Global Mode (ASAR GM), ASAR Wide Swath (ASAR WS) and Sentinel-1 IW GRDH. The SGRT preprocessing is utilizing S1TBX (version 1.1.1) for georeferencing and calibration which has multi-core parallel processing capability. The SAR data are acquired in form of image slices and can be independently processed as the elaboration of one image does not require any information from the other image slices processing. This simplifies parallelization of the processing that can be implemented through the run of many independent jobs. However, the large number of data files and considerable different data file sizes from a few MBs in case of ASAR GM (in case-1) to 1-2 GBs of Sentinel-1 IW GRDH (in case-3 and case-4) give rise to different challenges when they are processed on any HPC system.
The full archive of the Envisat ASAR GM data, 189,621 relatively small files (1 to 73 MB) with a total size of about 1.6 TB and a part of the ASAR WS archive including images acquired over mid-latitudes of the Earth), 31,199 files (12 – 692 MB) with a total size of about 5.4 TB were successfully processed during case-1, case-2 and case-3 (Table 1). In all these three cases, the large number of relatively small files (more than a million in the case of ASA GM) made it burdensome to save the output directly in the BeeGFS distributed volume. The storage device was not capable to simultaneously perform the I/O operations that were needed to fulfill all the computing nodes requests. All attempts for storing the output data directly on distributed volume failed due to fatal stucks of the nodes. Thus, the images were divided in small groups of 8 and 2 images per job for ASAR GM and WS, respectively. Such a job file was, afterwards, sent to a single computing node, the output was temporary cached in the memory available at the node and only at the end of the processing the output was copied on the persistent BeeGFS storage volume. The selected number of images, that was assigned to each group (8 and 2), was a constraint imposed by the available RAM on each node. In total, an array of 23,703 and 15,600 independent jobs were submitted to the SLURM middleware respectively for ASAR GM and ASAR WS datasets processing.
As shown in Table 1, the same dataset was processed in case-2 and case-3 as these two tests were ad-hoc designed to further identify the processing chain bottleneck. The number of nodes that were provided by the SLURM middleware during the overall ASAR WS processing is shown in Figures 1 and 2 for case-2 and case-3, respectively. A maximum of approximately 400 and 600 nodes was assigned during the two processing. Such a number guarantees a significant data throughput between storage device and nodes. For instance, the processing of case-1 started in the night of the 4th December 2015 and ended early in the morning of the 9th December 2015. The peak of assigned nodes was reached during the night between the 8th and 9th of December. The average processing time was 5.65 and 2.39 MB/sec for case-2 and case-3 respectively. Figures 3 and 4 show the elapsed times with respect the data file size; while in case-3 the elapsed times have almost a processing linear trend, case-2 is characterized by a large dispersion meaning that the system got stuck quite a few times. The difference in the two case studies is that, although both tests cached the output, only the case-3 cached also the input images. We hypothesize that the reason for this is that the SGRT and S1TBX keep

continuously reading and writing data on disk. Such a programming strategy can cause either bandwidth saturation or storage volumes failing when many nodes are simultaneously active. Copying the input image files to the node-cash could solve the issue as during the processing, the node has a direct and fast access to the cache without incurring the risk of saturating the storage volume and bandwidth. However, further studies are needed to confirm that hypothesis and to exactly identify the reason of such a different behavior.

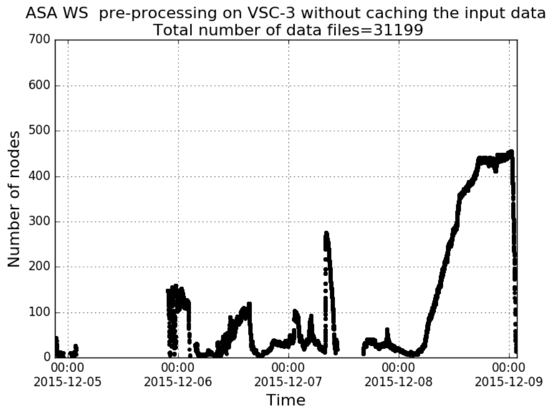

Figure 1. Nodes assignment during case study 2

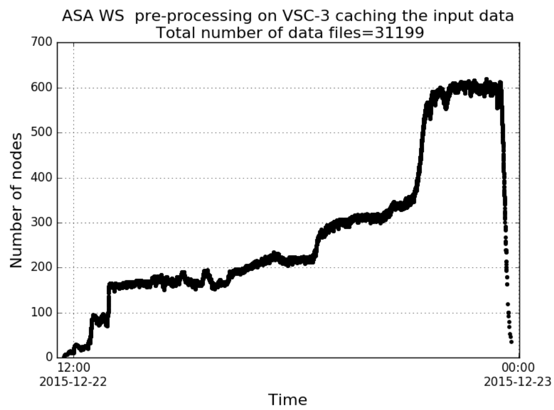

Figure 2. Nodes assignment during case study 3

The case studies 4 and 5 were designed to identify the optimal S1TBX configuration in terms of running cores within a node. The same dataset of Sentinel-1 data consisting of 1,075 data slices of Europe were considered. The input image was immediately cashed in the node to ensure a fast access to the data to SGRT and S1TBX. The output directory was also set to be the cash, and only at the end of the processing the output data were copied in the persistent parallel disk. Due to the size of Sentinel-1 data files and the limited RAM available at each node, only one image was processed at each single nodes. A total of 1,075 independent jobs, equal to the number of images, were therefore submitted. The queue system provided by the SLURM middleware has made a maximum of nearly 396 and 182 nodes simultaneously running. Table 1 shows the average processing times, for case-4 is 2.69 MB/sec in contrast of 2.83 MB/sec for case-5. Case studies 4 and 5 showed that increasing the number of cores used by S1TBX did not improve the performance.

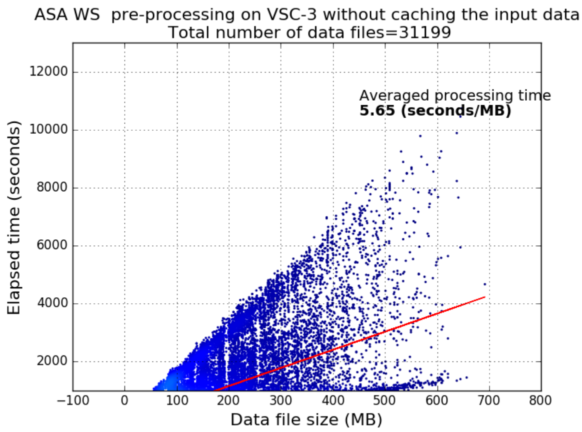

Figure 3. Average processing time for case study 2

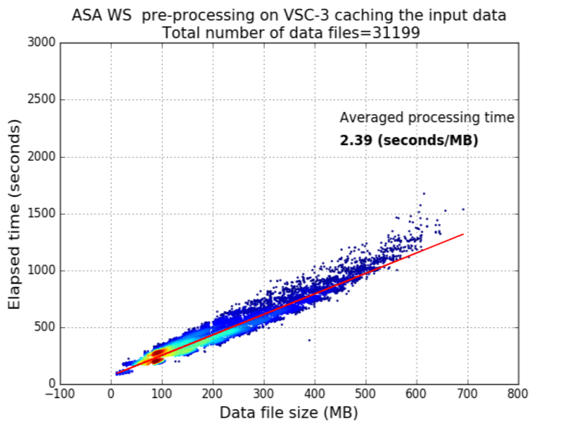

Figure 4. Averaged processing time for case study 3

5. CONCLUSIONS

This paper showed the feasibility of processing large EO data sets at the EODC high-performance platform. Different tests involving the processing of the whole ENVISAT ASAR GM archive, part of the ASAR WS archive and one large Sentinel-1 data sets were performed. The experiments showed that EODC is a comprehensive infrastructure in which users can bring their own algorithm close to both EO data and use the powerful computational resources offered by VSC. Distributed computation can be, therefore, the tool to open

doi: 10.2788/854791

novel roadmaps for addressing new research questions and challenging operational applications.

6. ACKNOLEDGEMENTS

This study was supported by the Austrian Research Promotion Agency (FFG) through “P4EODC - Preparing for the Initial Services of the EODC” (project number 36786), “WetMon - Enabling an operational Sentinel-1 wetland monitoring service for the EODC” (project number 848001) and by the EUMETSAT through “Satellite Application Facility on Support to Hydrology and Water Management (H-SAF)”. The authors would like to express their gratitude to B. Bauer-Marschallinger, A. Dostalova, S. Hahn, S. Hasenauer and to H. Thüminger for their valued and continued support.

7. REFERENCES

[1] Hornacek, M., Wagner, W., Sabel, D., Truong, H.L., Snoeij, P., Hahman, T., Diedrich, E. and Doubkova, M., “Potential for High Resolution Systematic Global Surface Soil Moisture Retrieval via Change Detection Using Sentinel-1”, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2012, 5 (4), 1303-1311.
[2] Pathe, C., Wagner, W., Sabel, D., Doubkova, M. and Basara, J.B., “Using ENVISAT ASAR Global Mode data for surface soil moisture retrieval over Oklahoma, USA”, *IEEE Transactions on Geoscience and Remote Sensing*, 2009, 47 (2), 468-480.
[3] Vienna Scientific Cluster 3. Available from <http://vsc.ac.at/systems/vsc-3/>
[4] Briese, C., Wagner, W., Boresch, A., Federspiel, C., Aspetsberger, M., Hasenauer, S., Mücke, W., Hoffmann, C., Wotawa, G., “Challenges in the exploitation of big earth observation data”, *Proceeding of the 2014 conference on Big Data from Space*, Frascati (IT)
[5] https://sentinel.esa.int/web/sentinel/toolboxes/sentinel-1
[6] Bauer-Marschallinger, B., Sabel, D., Wagner, W., 2014, "Optimisation of Global Grids for High-Resolution Remote Sensing Data", Computers & Geosciences, ISSN 0098-3004, Available online 11 July 2014.
[7] SLURM. Available from <http://slurm.schedmd.com/>
[8] BeeGFS. Available from <http://www.beegfs.com/content/>

Case study	case-1	case-2	case-3	case-4	case-5
SAR data product mode	ASAR GM	ASAR WS	ASAR WS	S-1 IW GRDH	S-1 IW GRDH
Spatial resolution	1 km	150 m	150 m	20 m	20 m
Grid spacing	500 m	75 m	75 m	10 m	10 m
Total number of data files	189,621	31,199	31,199	1,075	1,075
Number of images for job / Total Number of jobs	8 / 23,703	2 / 15,600	2 / 15,600	1 / 1,075	1 / 1,075
Input data file size range	1 - 73 MB	12 - 692 MB	12 - 692 MB	0.8 – 1.7 GB	0.8 –1.7 GB
Total input data files size	1.579 TB	5.401 TB	5.401 TB	1.2 TB	1.2 TB
Max. number of simultaneous running nodes	417	454	612	396	182
Number of cores used by Sentinel-1 Toolbox	4	8	8	8	16
Input data cashing on node	False	False	True	True	True
Output data cashing on node	True	True	True	True	True
Processing time range (seconds/MB)	0.8 - 90	1.1 – 22.1	1.1 – 3.1	1.5 – 5.6	1.5 – 5.8
Averaged processing time (seconds/MB)	9.18	5.65	2.39	2.69	2.83
Elapsed time including SLURM queueing	≈ 3.5 days	≈ 4 days	≈ 8 hours	≈ 3.5 hours	≈ 15 hours

Table 1. Specification and results of the five different case studies carried out on VSC-3.