SCALING THE PIPE: NASA EOS TERRA DATA SYSTEMS AT 10

Robert E. Wolfe[1,] and Hampapuram (Rama) K. Ramapriyan[2]*

1. NASA Goddard Space Flight Center (GSFC), Code 614.5, Greenbelt Rd., Greenbelt, MD, USA
* Corresponding author: phone: 301-614-5508, e-mail: robert.e.wolfe@nasa.gov
2. NASA GSFC Code 423, Greenbelt Rd., Greenbelt, MD, USA

ABSTRACT

Standard products from the five sensors on NASA's Earth Observing System's (EOS) Terra satellite are being used world-wide for earth science research and applications. This paper describes the evolution of the Terra data systems over the last decade in which the distributed systems that produce, archive and distribute high quality Terra data products were scaled by two orders of magnitude.

Index Terms— EOS, Terra, science data system, MODIS, CERES, ASTER, MISR, MOPITT

1. INTRODUCTION

In the late 1990s when the initial version of NASA's Earth Observing System's (EOS) distributed Terra data systems was being developed, a research data system that produced such a large number and volume of research products was unprecedented. The Terra satellite carries five instruments: Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER), Clouds and Earth's Radiant Energy System (CERES), Multi-angle Imaging SpectroRadiometer (MISR), MODerate Resolution Imaging Spectrometer (MODIS) and Measurement Of Pollution In The Troposphere (MOPITT). Over 70 calibration and geophysical science algorithms with complex inter-dependencies had to be integrated and tested before the first on-orbit Terra data from these instruments were available. Two of the instruments involved the additional complexity of collaborating with international partners: Japan for ASTER and Canada for MOPITT.

At the time of the Terra launch in December 2009, the Terra data system was the most complex component of the distributed EOS Data and Information System (EOSDIS) and involved a distributed set of production, archive and distribution facilities. Once the Terra instruments data became available in early 2000 the data systems helped the instrument science teams to continuously improve the science algorithms. Multiple reprocessing campaigns have continuously improved the algorithms giving the community stable high quality validated earth science products. Over the last decade, the data systems needed to be scaled to keep up to support these activities. For example, the data archive system went from primarily a tape-based archive to an on-line multi-petabyte disk archive that greatly improved user data access and sped up reprocessing activities.

Several factors were key to the success of the initial Terra data systems and evolution of the systems over the last decade. The most important were strong leadership of the Terra project, EOSDIS, and instruments' science team leaders. This leadership helped coordinate the large NASA science teams and keep the data system development focused on the science goals and objectives. Second was an evolvable and scalable set of data systems that were developed through close interaction with science teams. These data systems evolved over time and grew by two orders of magnitude in terms of processing and storage to allow for the daily (forward) processing, science testing and reprocessing rates needed to meet of the science teams' and the community's expectations. Finally, active applications and outreach activities facilitated a rich and varied set of Terra products that is widely used by the global community for near real-time and regional research and applications. The feedback from this community was also invaluable in the continual improvement of the standard algorithms as well as the data system capabilities.

Part of this data system evolution occurred before the launch of Terra. In an effort to make the data system more distributed and reduce its components to more manageable "chunks", the generation of standard products was moved, in most cases, from the EOSDIS Core System (ECS) to Science Investigator-led Processing Systems (SIPSs) developed and operated by the respective instrument teams. The approach to the development of ECS was also modified to result in more frequent releases of its Science Data Processing Segment based on priorities expressed by the science community. These steps led to the successful completion of all subsystems needed to support Landsat-7 (launched in April 1999) and Terra. Given the experience in getting ready for Landsat-7 and Terra, especially the multiple end-to-end tests (dubbed Mission Operations and Science System or MOSS tests), the overall readiness of the data systems for the Aqua, ICESat and Aura missions was

better and so the initial production data flows went much more smoothly [1].

In the next sections we will talk about the growth of the three main components of the EOS data systems: production, archive and distribution. This is followed by a discussion of the recent evolution in the data systems.

1. PRODUCTION

Two areas of processing need to be addressed in the EOS mission science data system, forward processing and reprocessing. In this discussion, we use the MODIS as an example of production system growth since MODIS produces the majority of the Terra products [2].

Forward production needs to keep up with the data stream from the instrument. Typically, Terra science products are produced within a few days of acquisition. To maintain this production rate, the forward processing system needs to be able to process a data-day worth of instrument data in a calendar day (we refer to this rate as 1X). To accomplish this, and to be able to catch up when issues occur because of anomalies in the instrument, data stream or production hardware, a system capable of a production rate close to 2X is needed. At launch, because of limited resources and the late-1990s technologies, the initial processing rate for the Terra missions was very close to 1X and so the catch-up capability was limited.

In addition to the standard forward processing, a near real-time (NRT) processing capability may be needed for instruments to support operational and application users. These systems may also be used for education and outreach. The NRT systems need to be designed to provide products specific to the application users and because of the low latency requirements of these users, may not produce the best science quality products. For Terra, one of the key systems in this area is the MODIS Rapid Response system which has played a key role in the development of many applications and has been widely recognized for its contributions to the Terra mission's public outreach. Because of the low-latency and NRT requirements, a production rate similar to the forward production system is needed, though typically for a more limited set of products.

Over the Terra mission lifetime, better calibration and characterization of the instruments have been performed and new and improved algorithms have been developed. This means that to extract the maximum value from the investment in the Terra research mission, multiple reprocessing campaigns are needed. For MODIS, reprocessing is also driven by the complex interdependencies between the algorithms. Terra, like many of the NASA missions, has lasted well beyond its design lifetime of 5 years, and is expected to survive long enough to produce a total of more than 15 years of high quality data. These reprocessing activities are typically driven by the science team and community. The reprocessing cycle involves three phases: development of the algorithm improvements, testing the improvements and then the actual reprocessing.

A major challenge for the EOS data systems is scaling the data production system to meet these reprocessing needs. For example, the MODIS team is in the process of preparing for the fourth major reprocessing for which the production phase is schedule to begin later this year, about 11 years after launch. For MODIS, the average time between reprocessing is about 3 years with the first one taking place after the algorithms were stabilized after launch. With a three year cycle, the actual reprocessing needs to be done in less than a year so as much time as possible is available for the important algorithm development and testing phases. This means that the needed reprocessing resources are more than 1X per mission-year. For example for MODIS's current reprocessing campaign of almost 11 mission-years of MODIS/Terra data, the reprocessing capacity is 15X (a conservative estimate that is based on the ingest capabilities at the archives).

The growth in data system to meet these reprocessing must scale with number of mission-years of data. So at 10 years, the data system must have ten times the reprocessing capacity as the system a year after launch. In addition, for reliability and other reasons, computer technology should be refreshed (older computers and disk retired) on a three to five year basis. These reprocessing resources are also used to perform extensive science tests that must be performed prior to reprocessing to ensure high quality science products are produced. For Terra, a steady progression of technology has been infused over the years, going from large single computers running proprietary operating systems and storing intermediate data on tape to clusters of inexpensive commodity compute servers running Linux and storing all data products on disk. Another key to scaling to the higher reprocessing rates has been to move the Level 0 data (raw instrument data) from tape to disk. The Terra data systems have also benefited from Moore's law [3]. By buying the latest technology in the period leading up to the reprocessing production phase, the overall cost of procuring the needed hardware is minimized. So, over the 10 years since launch as the system capacity grew by an order of magnitude, the overall yearly hardware procurement has been constant while keeping up-to-date with the most recent computer technologies.

2. ARCHIVE

At launch, Terra archive data was held in robotic tape archives. In early 2002, the total EOSDIS archive first exceeded 1 petabyte in size. To improve the distribution of data to users, we started migrating some of the data of high interest to the community onto data pools, which were disk caches of the order a few tens of terabytes, considered "large" at the time [4]. Today the EOSDIS archive is close

to 5 petabytes in size, and most of the data are held on on-line disks. The disk capacity has become very affordable over time, which has helped this migration. The advantages of on-line storage is improved access with reduced latency, easier maintenance of the archive and the ability to develop and provide many on-line services such as subsetting, reprojection, visualization, and data fusion.

3. DISTRIBUTION

One way to illustrate the growth in the EOS data system over time is to look at the amount of data that was distributed to the public in the early mission and more recently (Table 1). Over 10 years, there has been a steady increase in the volume and number of file distributed (Figures 1 and 2). Over the last seven years, the number of files distributed has increased an average of 37% per year and the volume by 21% per year. Since the mission start through December 2009, 230 million files totaling 5,498 TB have been distributed to the earth science community.

Table 1. Average daily distribution of Terra data in the first and tenth year of the mission.

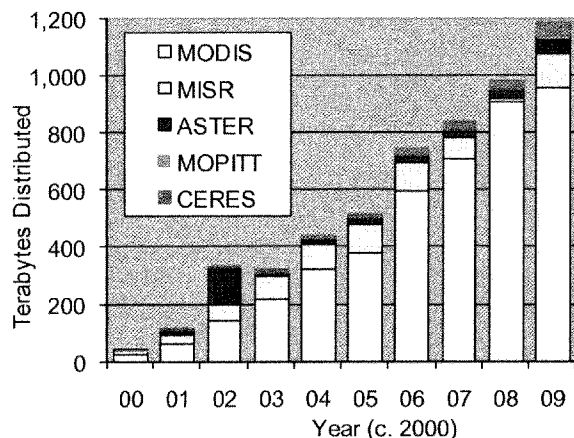|	2000	2009	Ratio
Files distributed (1,000/day)	3.0	170.4	67.6
Volume (TB/day)	0.10	3.25	31.6

Figure 1. Terra data volume distributed per year from 2000 to 2009.

In addition to the increases in the absolute numbers of files and terabytes distributed to the user community, the nature of distribution has changed over time as well. There are several methods for searching for (or discovering) the data of interest to a user [5]. These include the Warehouse Inventory Search Tool (WIST) for cross-data-center searches and tailored search clients for different disciplines provided by each of the data centers. In the early days of EOSDIS, a significant amount of distribution to users was

through media. Today, almost all distribution occurs on-line. Also, the number of on-line services has increased. Many of the MODIS products can be produced on demand and users can request post-processing of the data products as they order them. Post-processing capabilities include reprojection, subsetting, mosaicing, gap-filling, etc. The MODIS products and these post-processing capabilities can be ordered through the standard web interface as well as through a set of Web Services that allows for external client based access and automated (scripted) access.
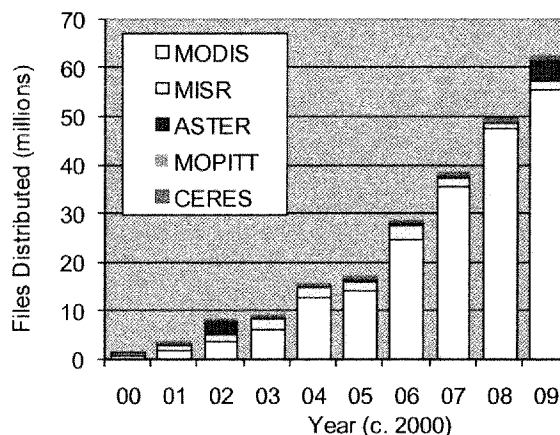

Figure 2. Terra data files distributed per year from 2000 to 2009.

4. EVOLUTION

The evolution of the data system has continued over the last few years. During 2005, a focused study was conducted by a NASA Head Quarters sponsored external Study Team in collaboration with a Technical Team consisting of members involved in the development and operation of EOSDIS. A vision of 2015 was developed and an implementation plan for the first step towards the vision was prepared [6]. The implementation was completed during 2006-2008 [7]. From the point of view of Terra, the significant changes that occurred as a part of this step were:

- Simplification of ECS at the Atmospheric Science Data Center (NASA Langley Research Center), Land Processes DAAC (USGS EROS), and the National Snow and Ice Data Center;

- Transition of responsibility for MODIS Level 1 processing as well as archiving and distribution of level 1 products and atmospheric products from the Goddard Earth Science DAAC to the MODIS Data Processing System (MODAPS)/Level 1 and Atmosphere Archiving and Distribution System (LAADS);

- Movement of most of the data into on-line archives (with tape back-up) to provide improved access and on-line services upon request;

- Improvements in operational capabilities of the EOS Clearing House (ECHO) and the Warehouse Inventory Search Tool (WIST) as a search and order client [8].

As indicated above, many of the lessons learned from Terra were applied to facilitate getting ready for the follow-on EOS missions. It is also expected that NASA's upcoming Earth Science Decadal Survey missions will benefit from the Terra and EOSDIS experience.

5. CONCLUSION AND THE FUTURE

There has been a significant increase in the capacity and capabilities of the data system supporting the Terra mission in the ten years of its operation. Production, archiving and distribution have all improved significantly, the most notable trend being the increase in distribution to the external user community. This has been achieved despite reductions in budget due to improvements in technology as well proactive re-architecting and evolution of the data system. In fact, the cost of operating EOSDIS has been reduced by about 30% since 2009 as a result of the recent evolution activities. Most of the data are on-line and hence are more easily accessible to users. Near real-time capabilities are being provided to support applications requiring the data within a few hours of acquisition.

The Terra data systems need to remain agile to serve the community as technology and users' expectations change. The goal is to make Terra data more useful for the scientific and broader user community and to make scientific collaboration easier. Some of the emergent technologies on the horizon that need to be considered are data fusion, cloud computing and access from mobile devices. The data processing teams are evaluating the possible use of cloud computing because of the potential for possible savings in terms of cost and schedule. Some issues that need to be addressed before adopting this technology are data stewardship, I/O bandwidth, scientific reproducibility and governance and the cost of computing and storage in the cloud.

As the Terra spacecraft and its instruments age, an important consideration is planning for ensuring that all the ancillary data that are needed to be preserved along with the data products are captured from the currently distributed sources (e.g., instrument teams) and placed at the appropriate EOSDIS Data Centers.

7. REFERENCES

[1] H. K. Ramapriyan, "EOS Data and Information System (EOSDIS): Where We Were and Where We Are", (2-part series), *Earth Observer*, vol. 21, no. 4 and 5, 2009.

[2] R. E. Wolfe, B. L. Ridgway, F. S. Patt and E. J. Masuoka, "MODIS Science Algorithms and Data Systems Lessons Learned", *IGARSS*, Cape Town, South Africa, July 2009.

[3] G. E. Moore, "Cramming more components onto integrated circuits", *Electronics*, vol. 38, no. 8, 1965.

[4] J. M. Moore and D. Lowe. "Providing rapid access to EOS data via Data Pools", *Proceedings of the 47th Annual Meeting of SPIE* paper 4814-56, July 2002.

[5] Ramapriyan H. K., R. G. Pfister and B. E. Weinstein, "An Overview of the EOS data Dissemination Systems", Chapter in *Land Remote Sensing and Global Environmental Change: NASA's Earth Observing System and the Science of ASTER and MODIS*, Springer, 2010 (*in press*).

[6] M. A. Esfandiari, H. K. Ramapriyan, J. Behnke and E. J. Sofinowski, "Earth Observing System (EOS) Data and Information System (EOSDIS) – Evolution Update and Future", *IGARSS*, Barcelona, Spain, 2007.

[7] Ramapriyan, H. K., J. Behnke, E. J. Sofinowski, D. R. Lowe and M. A. Esfandiari, "Evolution of the Earth Observing System (EOS) Data and Information System (EOSDIS)" - Chapter 5 in *Standards-Based Data and Information Systems for Earth Observations*, Springer-Verlag Berlin Heidelberg, 2009.

[8] A. E. Mitchell, H. K. Ramapriyan and D. R. Lowe, "Evolution of Web Services in EOSDIS Search and Order Metadata Registry (ECHO)", *International Geosci. and Remote Sens. Symposium (IGARSS)*, Cape Town, South Africa, July 2009.