

Article development led by **acmqueue**
queue.acm.org

Astronomers are collecting more data than ever. What practices can keep them ahead of the flood?

BY G. BRUCE BERRIMAN AND STEVEN L. GROOM

How Will Astronomy Archives Survive the Data Tsunami?

ASTRONOMY IS ALREADY awash with data: currently 1PB (petabyte) of public data is electronically accessible, and this volume is growing at 0.5PB per year. The availability of this data has already transformed research in astronomy, and the Space Telescope Science Institute (STScI) now reports that more papers are published with archived data sets than with newly acquired data.¹⁸

This growth in data size and anticipated usage will accelerate in the coming few years as new projects

such as the Large Synoptic Survey Telescope (LSST), Atacama Large Millimeter Array (ALMA), and Square Kilometer Array (SKA) move into operation. These new projects will use much larger arrays of telescopes and detectors or much higher data acquisition rates than are now used. Projections indicate that by 2020, more than 60PB of archived data will be accessible to astronomers.¹⁰

The data tsunami is already affecting the performance of astronomy archives and data centers. One example is the NASA Infrared Processing and Analysis Center (IPAC) Infrared Science Archive (IRSA), which archives and serves data sets from NASA's infrared missions. It is going through a period of exceptional growth in its science holdings, as shown in Figure 1, because it is assuming responsibility for the curation of data sets released by the Spitzer Space Telescope and Wide-field Infrared Survey Explorer (WISE) mission.

The volume of these two data sets alone exceeds the total volume of the 35-plus missions and projects already archived. The availability of the data, together with rapid growth in program-based queries, has driven up usage of the archive, as shown by the annual growth in downloaded data volume and queries in Figure 2. Usage is expected to accelerate as new data sets are released through the archive, yet the response times to queries have already suffered, primarily because of a growth in requests for large volumes of data.

The degradation in performance cannot be corrected simply by adding infrastructure as usage increases, as is common in commercial enterprises, because astronomy archives generally operate on limited budgets that are fixed for several years. Without intervention, the current data-access and computing model used in astronomy, in which data downloaded from archives is analyzed on local machines, will break down rapidly. The very scale of data sets such as those just described will transform the design and operation of archives as places



that not only make data accessible to users, but also support in situ processing of these data with the end users' software: network bandwidth limitations prevent transfer of data on this scale, and users' desktops in any case generally lack the power to process PB-scale data.

Moreover, data discovery, access, and processing are likely to be distributed across several archives, given that the maximum science return will involve federation of data from several archives, usually over a broad wavelength range, and in some cases will involve confrontation with large and complex simulations. Managing the impact of PB-scale data sets on archives and the community was recognized as an important infrastructure issue in the report of the 2010 Decadal Survey of Astronomy and

Astrophysics,⁶ commissioned by the National Academy of Sciences to recommend national priorities in astronomy for the coming decade.

Figure 3 illustrates the impact of the growth of archive holdings. As holdings grow, so does the demand for data, for more sophisticated types of queries, and for new areas of support, such as analysis of massive new data sets to understand how astronomical objects vary with time, described in the 2010 Decadal Survey as the "last frontier in astronomy." Thus, growth in holdings drives up storage costs, as well as compute and database costs, and the archive must bear all of these costs. Given that archives are likely to operate on shoe-string budgets for the foreseeable future, the rest of this article looks at

strategies and techniques for managing the data tsunami.

At the Innovations in Data-intensive Astronomy workshop earlier this year (Green Bank, WV, May 2011¹⁵) participants recognized that the problems of managing and serving massive data sets will require a community effort and partnerships with national cyber-infrastructure programs. The solutions will require rigorous investigation of emerging technologies and innovative approaches to discovering and serving, especially as archives are likely to continue to operate on limited budgets. How can archives develop new and efficient ways of discovering data? When should, for example, an archive adopt technologies such as graphical processing units (GPUs) or cloud computing? What kinds of

Figure 1. Growth in the scientific data holdings of IRSA, projected to 2014. The graphic calls out the dramatic impact of the Spitzer and WISE missions on the volume of the archive's science data holdings.

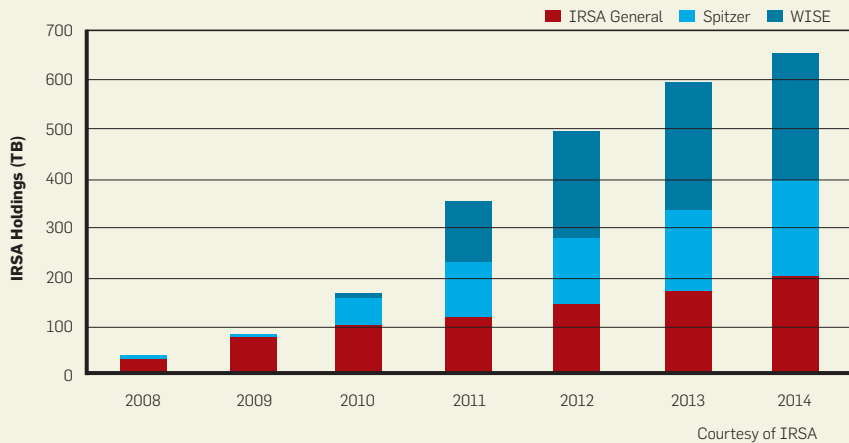
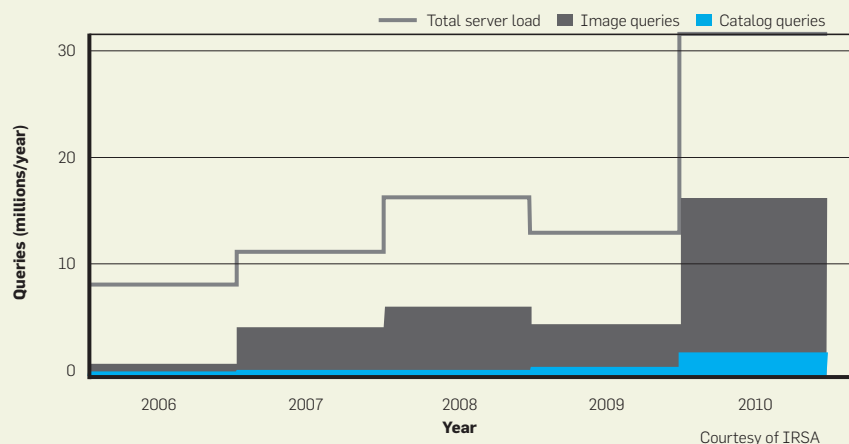
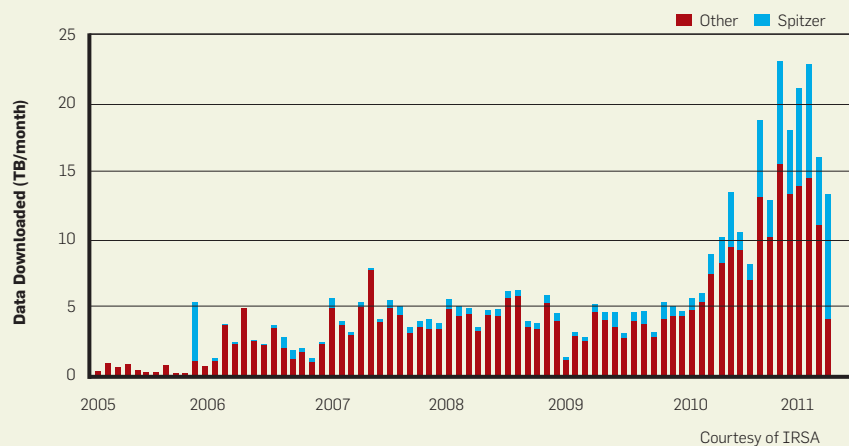


Figure 2. Growth in usage of IRSA from 2005 until the beginning of 2011. WISE data was not available until spring 2011.



technologies are needed to manage distribution of data time, compute-intensive data-access jobs, and end-user processing jobs?

This article emphasizes those issues we believe must be addressed by archives to support their end users in the coming decade, as well as those issues that affect end users in their interactions with archives.

Innovations in Serving and Discovering Data

The discipline of astronomy needs new data-discovery techniques that respond to the anticipated growth in the size of data sets and that support efficient discovery of large data sets across distributed archives. These techniques must aim to offer data discovery and access across PB-sized data sets (for example, discovering images over many wavelengths over a large swath of the sky such as the Galactic Plane) while preventing excessive loads on servers.

The Virtual Astronomical Observatory (VAO),¹⁹ part of a worldwide effort to offer seamless international astronomical data-discovery services, is exploring such techniques. It is developing an R-tree-based indexing scheme that supports fast, scalable access to massive databases of astronomical sources and imaging data sets.⁹ (R-trees are tree data structures used for indexing multidimensional information. They are commonly used to index database records and thereby speed up access times.)

In the current implementation, the indices are stored outside the database, in memory-mapped files that reside on a dedicated Linux cluster. It offers speedups of up to 1,000 times over database table scans and has been implemented on databases containing two billion records and TB-scale image sets. It is already in operation in the Spitzer Space Telescope Heritage Archive and the VAO Image and Catalog Discovery Service. Expanding techniques such as this to PB-scale data is an important next step.

Such custom solutions may prove more useful than adapting an expensive geographical information system (GIS) to astronomy. These systems are necessarily more complex than are needed in astronomy, where the celestial sphere is by definition a perfect sphere and the footprints on the sky of

instruments and data sets are generally simple geometric shapes.

Investigations of Emerging Technologies

A growing number of investigators are taking part in a concerted and rigorous effort to understand how archives and data centers can take advantage of new technologies to reduce computational and financial costs.

Benjamin Barsdell et al.¹ and Christopher Fluke et al.⁷ have investigated the applicability of GPUs to astronomy. Developed to accelerate the output of an image on a display device, GPUs consist of many floating-point processors. These authors point out that speed ups of more than 100 times promised by manufacturers strictly apply to graphics-like applications; GPUs support single-precision calculations rather than the double precisions often needed in astronomy; and their performance is often limited by data transfer to and from the GPUs. The two studies cited here indicate applications that submit to “brute-force parallelization” will give the best performance with minimum development effort; they show that code profiling will likely help optimization and provide a first list of the types of astronomical applications that may benefit from running on GPUs. These applications include fixed-resolution mesh simulations, as well as machine-learning and volume-rendering packages.

Others are investigating how to exploit cloud computing for astronomy. Applications best suited for commercial clouds are those that are processing and memory intensive, which take advantage of the relatively low cost of processing *under current fee structures*.² Applications that are I/O intensive, which in astronomy often involve processing large quantities of image data, are, however, uneconomical to run because of the high cost of data transfer and storage. They require high-throughput networks and parallel file systems to achieve best performance.

Under current fee structures, renting mass storage space on the Amazon cloud is more expensive than purchasing it. Neither option offers a solution to the fundamental business problem that storage costs scale with volume, while funding does not. Any use of com-

mercial clouds should be made after a thorough cost-benefit study. It may be that commercial clouds are best suited for short-term tasks, such as regression testing of applications and handling excessive server load, or to one-time bulk-processing tasks, as well as supporting end-user processing.

Implementing and managing new technologies always have a business cost, of course. Shane Canon⁴ and others have provided a realistic assessment of the business impact of cloud computing. Studies such as these are needed for all emerging technologies.

Despite the high costs often associated with clouds, the virtualization technologies used in commercial clouds may prove valuable when used within a data center. Indeed, the Canadian Astronomy Data Center (CADC) is moving its entire operation to an academic cloud called Canadian Advanced Network for Astronomical Research (CANFAR), “an operational system for the delivery, processing, storage, analysis, and distribution of very large astronomical datasets. The goal of CANFAR is to support large Canadian astronomy projects.”¹¹ To our knowledge, this is the first astronomy archive that has migrated to cloud technologies.⁸ It can be considered a first model of the archive of the future, and consequently the community

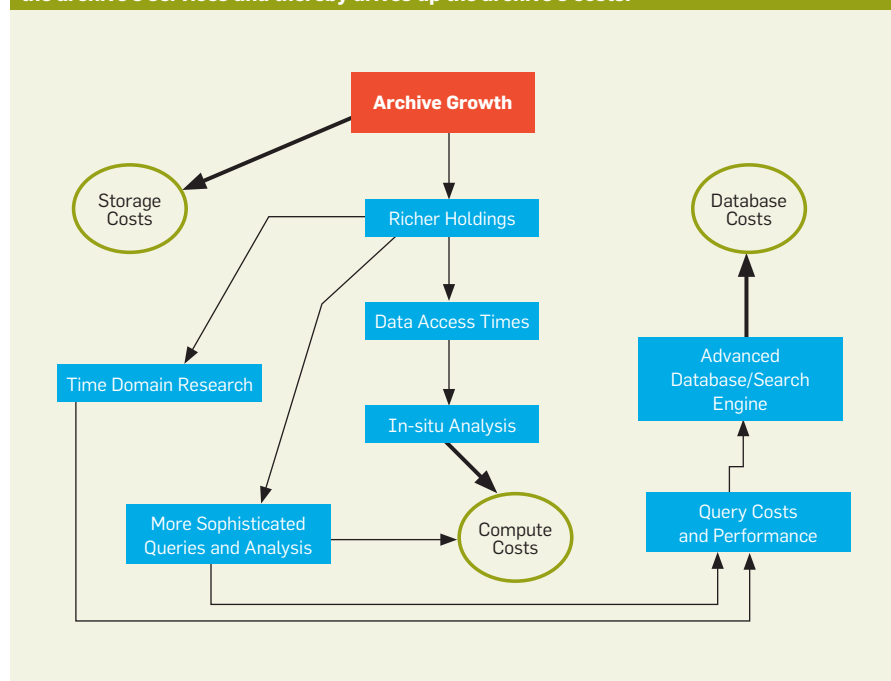
should monitor its performance.

The SKA has rejected the use of commercial cloud platforms. Instead, after a successful prototyping experiment, it proposes a design based on the open source Nereus V Cloud²⁰ computing technology, selected because of its Java codebase and security features. The prototype test bed used 200 clients at the University of Western Australia, Curtin University, and iVEC, with two servers deployed through management at a NereusCloud domain. The clients include Mac Minis and Linux-based desktop machines. When complete, “theskynet,” as it has been called, would provide open access to the SKA data sets for professionals and citizen scientists alike.¹² The design offers a cheaper and much greener alternative to earlier designs based exclusively on a centrally based GPU cluster.

Compute Infrastructure

Astronomy needs to engage and partner with national cyber infrastructure initiatives. Much of the infrastructure to optimize task scheduling and workflow performance and to support distributed processing of data is driven by the needs of science applications. Indeed, the IT community has adopted the Montage image mosaic engine³ to develop infrastructure (for example, task schedulers in distributed environ-

Figure 3. Schematic representation of how growth in data holdings drives up demands on the archive's services and thereby drives up the archive's costs.



ments and workflow optimization techniques). These efforts have not, however, been formally organized, and future efforts may well benefit from such.

Cultural changes. There is at present no effective means of disseminating the latest IT knowledge to the astronomical community. Information is scattered across numerous journals and conference proceedings. To rectify this, we propose an interactive online journal dedicated to information technology in astronomy or even physical sciences as a whole.

Even more important is the need to change the reward system in astronomy to offer recognition for computational work. This would help retain quality people in the field.

Finally, astronomers must engage the computer science community to develop science-driven infrastructure. The SciDB database,¹⁶ a PB-scale next-generation database optimized for science applications, is an excellent example of such collaboration.

Educational Changes. An archive model that includes processing of data on servers local to the data will have profound implications for end users, who generally lack the skills not only to manage and maintain software, but also to develop software that is environment agnostic and scalable to large data sets. Zeeya Merali¹⁴ and Igor Chilingarian and Ivan Zolotukhin⁵ have made compelling cases that self-teaching of software development is the root cause of this phenomenon. Chilingarian and Zolotukhin in particular present some telling examples of clumsy and inefficient design in astronomy.

One solution would be to make software engineering a mandatory part of graduate education, with a demonstration of competency as part of the formal requirements for graduation. Just as classes in instrumentation prepare students for a career in which they design experiments to obtain new data, so instruction in computer science prepares them for massive data-mining and processing tasks. Software has become, in effect, a scientific instrument.

The software engineering curriculum should include the principles of software requirements, design, and maintenance (version control, documentation, basics of design for ade-

quate testing); how a computer works and what limits its performance; at least one low-level language and one scripting language, development of portable code, parallel-processing techniques, principles of databases, and how to use high-performance platforms such as clouds, clusters, and grids. Teaching high-performance computing techniques is particularly important, as the load on servers needs to be kept under control. Such a curriculum would position astronomers to develop their own scalable code and to work with computer scientists in supporting next-generation applications.

Curricula designers can take advantage of existing teaching methods. Software Carpentry¹⁷ is an open source project that provides online classes in the basics of software engineering and encourages contributions from its user community. Frank Loffler et al.¹³ described a graduate class in high-performance computing at Louisiana State University in which they used the TeraGrid to instruct students in high-performance computing techniques that they could then use in day-to-day research. Students were given hands-on experience at running simulation codes on the TeraGrid, including codes to model black holes, predict the effects of hurricanes, and optimize oil and gas production from underground reservoirs.

Conclusion

The field of astronomy is starting to generate more data than can be managed, served, and processed by current techniques. This article has outlined practices for developing next-generation tools and techniques for surviving this data tsunami, including rigorous evaluation of new technologies, partnerships between astronomers and computer scientists, and training of scientists in high-end software engineering skills. ■

Related articles on queue.acm.org

Why Your Data Won't Mix

Alon Halevy

<http://queue.acm.org/detail.cfm?id=1103836>

If You Have Too Much Data, then "Good Enough" Is Good Enough

Pat Helland

<http://queue.acm.org/detail.cfm?id=1988603>

Information Extraction: Distilling Structured Data from Unstructured Text

Andrew McCallum

<http://queue.acm.org/detail.cfm?id=1105679>

References

1. Barsdell, B.R., Barnes, D.G. and Fluke, C. J. Analysing astronomy algorithms for graphics processing units and beyond. *Monthly Notices of the Royal Astronomical Society* 408, 3 (2010), 1936–1944.
2. Berriman, G.B., Deelman, E., Juve, G., Regelson, M. and Plavchan, P. The application of cloud computing to astronomy: A study of cost and performance. Accepted for publication in *Proceedings of the e-Science in Astronomy Conference* (Brisbane, AU, 2011).
3. Berriman, G.B., Good, J., Deelman, E. and Alexov, A. Ten years of software sustainability at the Infrared Processing and Analysis Center. *Phil. Trans. R. Soc.* 369 (2011), 3384–3397
4. Canon, S. Debunking some common misconceptions of science in the cloud. Presented at ScienceCloud2011: 2nd Workshop on Scientific Cloud Computing (San Jose, CA); <http://datasys.cs.it.edu/events/ScienceCloud2011/>.
5. Chilingarian, I. and Zolotukhin, I. The true bottleneck of modern scientific computing in astronomy. *Astronomical Data Analysis Software and Systems XX*. I. Evans et al., Eds. *ASP Conference Series* 442, 471 (2010).
6. Committee for a Decadal Survey of Astronomy and Astrophysics, National Research Council of the National Academy of Sciences. *New Worlds, New Horizons in Astronomy and Astrophysics*, 2010.
7. Fluke, C.J., Barnes, D.G., Barsdell, B.R., Hassan, A.H. Astrophysical supercomputing with GPUs: critical decisions for early adopters. *Publications of the Astronomical Society of Australia* 28, 15 (2011).
8. Gaudet, S. et al. CANFAR: The Canadian Advanced Network for Astronomical Research. *Software and Cyber Infrastructure for Astronomy*. N. Radziwill and A. Bridger, Eds. SPIE 7740, 1I (2010).
9. Good, J. Private communication, 2011.
10. Hanisch, R.J. Data discovery, access, and management with the virtual observatory. 2011, Paper presented at Innovations in Data-intensive Astronomy, Green Bank, WV, 2011; http://www.nrao.edu/meetings/bigdata/presentations/May5/1-Hanisch/Hanisch_VAO_GreenBank.ppt.
11. Hemsath, N. Canada explores new frontiers in astroinformatics. *HPC in the Cloud*; http://www.hpcinthecloud.com/hpccloud/2011-01-17/canada_explores_new_frontiers_in_astroinformatics.html.
12. Hutchinson, J. SKA bid looks to SkyNet for computing, 2011; http://www.cio.com.au/article/387097/exclusive_ska_bid_looks_sky_net_computing/.
13. Loffler, F., Allen, G., Bengler, W., Hutanu, A., Jha, S. and Schnetter, E. Using the TeraGrid to teach scientific computing. TeraGrid '11: Extreme Digital Discovery Conference (Salt Lake City, UT; July 18–21, 2011); <https://www.teragrid.org/web/tg11/home>.
14. Merali, Z. Why scientific programming does not compute. *Nature* 467 (2010) 775–777.
15. National Radio Astronomy Observatory. Innovations in Data-intensive Astronomy Workshop (Green Bank, WV, May 3–5, 2011); <http://www.nrao.edu/meetings/bigdata/index.shtml>.
16. SciDB Open Source Data Management and Analytics Software System. 2011; <http://scidb.org>.
17. Software Carpentry; <http://software-carpentry.org/>.
18. Space Telescope Science Institute. Hubble Space Telescope Publication Statistics 1991-2010 (2011); <http://archive.stsci.edu/hst/bibliography/pubstat.html>.
19. Virtual Astronomical Observatory; <http://us-vao.org>.
20. Nereus overview. http://www-nereus.physics.ox.ac.uk/about_overview.html

G. Bruce Berriman is a senior scientist at Infrared Processing and Analysis Center (IPAC). He is the program manager for the Virtual Astronomical Observatory and project manager for the W. M. Keck Observatory Archive and was formerly the manager of the NASA/IPAC Infrared Science Archive.

Steven L. Groom is a systems engineer at IPAC and is manager of the NASA/IPAC Infrared Science Archive. He has worked with mass storage, parallel processing, and data archiving in the space sciences, as well as commercial applications.

© 2011 ACM 0001-0782/11/12 \$10.00