

vate information is about the benefit of the innovation. These experiments eliminate that uncertainty by promising a known fixed payment for the best solution or for some part of it. Left open is how in practice the overall size of the reward is to be determined, or how it is to be allocated to different pieces. (In the experiment, each piece was allocated an equal share of the prize; in practice, some pieces are more difficult than others, so would ideally command a larger prize.) The patent system deals with this second source of uncertainty by tying the reward to the profit that can be obtained by the monopoly granted by the patent. There are various methods of tying the

reward to sales and price in the final goods market, as well as tapping other information about social value, but these need to be explicitly compared.

The standard economic approach to private information problems is mechanism design theory, which asks what is the best social outcome achievable for a given information structure. This theory has not yet been widely applied in the study of innovation. A good example of how to do so can be found in Hopenhayn and Mitchell (6). In addition to theoretical research, the experimental laboratory seems a good place to build on the research by Meloso *et al.* to investigate the

second half of the innovation puzzle: How should an innovation be valued?

References

1. M. Boldrin, D. Levine, *Against Intellectual Monopoly*, (Cambridge University Press, 2008).
2. D. Meloso, J. Copic, P. Bossaerts, *Science* **323**, 1335 (2009).
3. M. Kremer *Q. J. Econ.* **113**, 1137 (1998).
4. M. Kremer "Creating Markets for New Vaccines," NBER Working Paper No. W7716 (National Bureau of Economic Research, Cambridge, MA, 2000).
5. R. E. Lerner, J. Mokyr, "Intellectual Property Rights, the Industrial Revolution, and the Beginnings of Modern Economic Growth," presented at *American Economic Association Annual Meetings*, San Francisco, CA, January 2009.
6. H. A. Hopenhayn, M. F. Mitchell, *RAND J. Econ.* **32**, 152 (2001).

10.1126/science.1171406

COMPUTER SCIENCE

Beyond the Data Deluge

Gordon Bell,¹ Tony Hey,¹ Alex Szalay²

Since at least Newton's laws of motion in the 17th century, scientists have recognized experimental and theoretical science as the basic research paradigms for understanding nature. In recent decades, computer simulations have become an essential third paradigm: a standard tool for scientists to explore domains that are inaccessible to theory and experiment, such as the evolution of the universe, car passenger crash testing, and predicting climate change. As simulations and experiments yield ever more data, a fourth paradigm is emerging, consisting of the techniques and technologies needed to perform data-intensive science (1). For example, new types of computer clusters are emerging that are optimized for data movement and analysis rather than computing, while in astronomy and other sciences, integrated data systems allow data analysis and storage on site instead of requiring download of large amounts of data.

Today, some areas of science are facing hundred- to thousandfold increases in data volumes from satellites, telescopes, high-throughput instruments, sensor networks, accelerators, and supercomputers, compared to the volumes generated only a decade ago (2). In astronomy and particle physics, these new experiments generate petabytes (1 petabyte = 10^{15} bytes) of data per year. In bioinformatics, the increasing volume (3) and the extreme heterogeneity of the data are



Moon and Pleiades from the VO. Astronomy has been one of the first disciplines to embrace data-intensive science with the Virtual Observatory (VO), enabling highly efficient access to data and analysis tools at a centralized site. The image shows the Pleiades star cluster from the Digitized Sky Survey combined with an image of the moon, synthesized within the World Wide Telescope service.

challenging scientists (4). In contrast to the traditional hypothesis-led approach to biology, Venter and others have argued that a data-intensive inductive approach to genomics (such as shotgun sequencing) is necessary to address large-scale ecosystem questions (5, 6).

Other research fields also face major data management challenges. In almost every laboratory, "born digital" data proliferate in files, spreadsheets, or databases stored on hard drives, digital notebooks, Web sites, blogs, and wikis. The management, curation, and archiv-

ing of these digital data are becoming increasingly burdensome for research scientists.

Over the past 40 years or more, Moore's Law has enabled transistors on silicon chips to get smaller and processors to get faster. At the same time, technology improvements for disks for storage cannot keep up with the ever increasing flood of scientific data generated by the faster computers. In university research labs, Beowulf clusters—groups of usually identical, inexpensive PC computers that can be used for parallel computations—have

¹Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA. ²Department of Physics and Astronomy, Johns Hopkins University, 3701 San Martin Drive, Baltimore, MD 21218, USA. E-mail: szalay@jhu.edu

become ubiquitous. However, these cluster computing systems have limited connection to disks and lack database software. Scientists and computer scientists must now develop similarly cost-effective solutions for data-intensive research. Jim Gray was one of the first to anticipate this need. In 1995, he advocated building clusters of “storage bricks,” consisting of inexpensive, balanced systems of central processing units, memory, and storage for data-intensive research (7).

We have realized such an architecture in the GrayWulf system (8), which is built out of commodity components like Beowulf clusters. However, unlike Beowulf clusters, which are optimized for computation, the GrayWulf design emphasizes high-speed access to data residing on each node of the cluster, which supports a large database system; its performance scales well as the number of nodes is increased. GrayWulf won the Storage Challenge at the SC08 conference (9) by executing a query on the Sloan Digital Sky Survey (SDSS) database in 12 minutes; the same task took 13 days on a traditional (nonparallel) database system.

The bandwidth of inexpensive, commodity computer networks is also falling behind the data explosion. Copying large amounts of experimental data from a data center to personal workstations or distributing data to numerous independent centers is no longer tenable without recourse to extreme—and thus expensive—networking solutions. For research to be affordable, data analysis must increasingly be done where data sets reside, leaving academic research networks to handle low-bandwidth queries and analytic requests, including visualization.

The urgency for new tools and technologies to enable data-intensive research has been building for a decade or more (2, 7). In 2007, Jim Gray laid out his vision for a fourth research paradigm—data-intensive science—which he described as collaborative, networked, and data-driven (1, 10). He defined eScience as the synthesis of information technology and science that enables challenges on previously unimaginable scales to be tackled.

Despite the enormous potential of this approach, data-intensive science has been slow to develop due to the subtleties of databases, schemas, and ontologies, and a general lack of understanding of these topics by the scientific community. For example, virtually all large-scale models use databases to organize the vast array of files that hold data from computational modeling, but these databases rarely hold any data: They only hold pointers to the files that hold data, making direct analysis impractical. Indeed, many areas of

science lag commercial use and understanding of data analytics by at least a decade.

Astronomy has been among the first disciplines to undergo the paradigm shift to data-intensive science. The first step in this direction was made in 2001, when data from the SDSS were put into a publicly available database (11), with simple Web services offering the primary access to the multiterabyte (1 terabyte = 10^{12} bytes) data sets. Astronomers have embraced not only these services but have also frequently used the powerful Structured Query Language (SQL)—previously used almost exclusively by the financial and commercial sector—to gain direct access to data stored in a relational database. The site also offers an analysis workbench, where users can analyze data and store derived data sets next to the main database. About 15 to 20% of the world’s professional astronomers now have their own server-side database, and the SDSS servers are running close to saturation.

Now, astronomers have gone even further in embracing data-intensive science. An international grassroots effort has gone a long way toward integrating all astronomical data (hundreds of terabytes today) into the Virtual Observatory (VO) (see the figure), of which the SDSS is an integral part. In the VO, data are accessible through services, and ontological and semantic information is stored with each data set (12); this information provides crucial support for data searching, analysis, and reuse by using a standard vocabulary agreed on by the community and by recording semantic information about the structure and type of data, as well as the instrument that generated the data. Most major astronomical data providers have adopted a standardized interface for services, and there is a registry to find particular data sets.

A similar transformation is happening in many sciences: For high-energy physics, the CERN Large Hadron Collider (LHC) is set to create an integrated data system resembling the VO; in genomics, NCBI (National Center for Biotechnology Information) and GenBank play this part. Many day-to-day issues are the same, whether we deal with astronomy or oceanography data. The emerging solution to these challenges lies in more diverse computing system architectures—like the GrayWulf system—that are specialized for highly data-intensive computations. Such systems will offer specialized data-analysis facilities located next to the largest data sets, coexisting with and complementary to today’s supercomputers.

Data-intensive science will be integral to many future scientific endeavors, but demands specialized skills and analysis tools. In addition, the research community now has the

option of accessing storage and computing resources on demand. The IT industry is building huge data centers, far beyond the financial scope of universities and national laboratories (13). These “cloud services” provide high-bandwidth access to cost-effective storage and computing services. However, there are no clear examples of successful scientific applications of clouds yet; making optimum use of such services will require some radical rethinking in the research community. In the future, the rapidity with which any given discipline advances is likely to depend on how well the community acquires the necessary expertise in database, workflow management, visualization, and cloud computing technologies.

References and Notes

1. J. Gray, A. Szalay, *eScience—A Transformed Scientific Method*, Presentation to the Computer Science and Technology Board of the National Research Council, Mountain View, CA, 11 January 2007; see http://research.microsoft.com/en-us/um/people/gray/talks/NRC-CSTB_eScience.ppt.
2. A. Hey, A. Trefethen, in *Grid Computing: Making the Global Infrastructure a Reality*, F. Berman, G. Fox, T. Hey, Eds. (Wiley, Chichester, UK, 2003), pp. 809–824.
3. See www.ncbi.nlm.nih.gov.
4. J. C. Wooley, Herbert S. Lin, Eds., *Catalyzing Inquiry at the Interface of Computing and Biology* (Report of the National Research Council of the National Academies, 2005), chap. 3.
5. J. C. Venter *et al.*, *Science* **304**, 66 (2004).
6. D. B. Kell, S. G. Oliver, *BioEssays* **26**, 99 (2004).
7. Jim Gray talked about the need for such a “brick architecture” in his Mackay Lectures at Berkeley as long ago as 4 October 1994, under the title “SNAP—Scalable Networks and Platforms,” followed on 4 November 1995 by a lecture titled “Parallel Database Systems: A SNAP application.” These lectures are available in the Internet Archive www.archive.org under [www.archive.org](http://www.archive.org/details/~gray) under [www.archive.org](http://www.archive.org/details/~gray).
8. A. S. Szalay *et al.*, *GrayWulf: Scalable Clustered Architecture for Data-Intensive Computing*, Proceedings of the 42nd Hawaii International Conference on System Sciences, Hawaii, 5 to 8 January 2009, paper no. 720; available as Microsoft Tech Report MSR-TR-2008-187 at <http://research.microsoft.com/apps/pubs/default.aspx?id=79429>.
9. See <http://sc08.supercomputing.org/>.
10. Jim Gray was one of the founders of modern database technology and transaction processing and received the Turing Award in 1998. Recognizing that the greatest future data challenges would come not from commercial database applications, but from the next generation of scientific experiments, he spent 10 years systematically exploring his “Fourth Paradigm,” working with astronomers, biologists, oceanographers, and geographers. Jim was tragically lost at sea in January 2007 in the midst of his latest challenge.
11. J. Gray *et al.*, in *Distributed Data & Structures 4: Records of the 4th International Meeting*, W. Litwin, G. Levy, Eds. (Carleton Scientific, Waterloo, Ontario, 2003), pp. 189–210.
12. M. J. Graham, M. J. Fitzpatrick, T. A. McGlynn, Eds., *The National Virtual Observatory: Tools and Techniques for Astronomical Research*, ASP Conference Series **382** (Astronomical Society of the Pacific, Provo, UT, 2007).
13. R. Ozzie, Interview on Cloud Computing, *Wired Magazine*, December 2008.
14. We thank J. Fay (World Wide Telescope) for providing the figure.