Global Information Industry Center
at the School of International Relations and Pacific Studies
University of California, San Diego

The Department of Physics at MIT
The Scientific Data Flood: A Case Study of "How Much Information?"

Stuart Madnick, John Norris Maguire Professor of Information Technology, MIT Sloan School of Management & Professor of Engineering Systems, MIT School of Engineering

MacKenzie Smith, Associate Director of Technology, MIT Libraries

Kate Clopeck, Masters of Science, Technology and Policy Program, MIT

June 2009

Abstract:
MIT's Physics department has about 90 experimental physics faculty, who generate massive amounts of data. The nature and size of each project varies, but they tend to run continuously over months or years. For example, the Compact Muon Solenoid detector, housed at CERN, produces about 8000 terabytes per year of experimental data, plus a similar amount of simulation data, all of which is processed multiple times. The data is stored in an internationally distributed, tiered system that provides backup and sharing. Scientists at MIT pull off data in chunks of about 500 TB, which are then filtered and analyzed on campus. About 500 worldwide users store working data at MIT, with 1 TB in a RAID array allocated for each. Although the underlying physical sensors can remain fixed for years, the amount of raw experimental data still increases, based on upgrades in methods of collection and processing. A rough extrapolation is that the Physics department as a whole stores about 2 * 1018 bytes a year (2 exabytes) of new data. Other papers examine other labs at MIT.

☐ **Executive Briefing:** a summary of one or more research projects with preliminary findings for a non-academic audience.

☐ **Research Report:** a completed report drawing on one of more research projects that presents study data, findings and management implications.

☑ **Case Study:** an in-depth description of a firm's approach to an information management issue.

☐ **Research Article:** an academic research paper with sections on hypotheses tested, methods and data, analysis, findings and references.

Contents

The Department of Physics at MIT The Scientific Data Flood: A Case Study of "How Much Information?"

Stuart Madnick, MacKenzie Smith, and Kate Clopeck, MIT

June 2009

Background

The Massachusetts Institute of Technology's Department of Physics has been a national resource since the turn of the 20th century. This department is home to over 120 faculty members who conduct research on a wide variety of subject areas ranging from cosmology to string theory. These faculty members are divided into four major research divisions: Astrophysics; Atomic, Condensed Matter, and Plasma Physics; Experimental Nuclear and Particle Physics; and Theoretical Nuclear and Particle Physics. The largest of these divisions is Atomic, Condensed Matter and Plasma Physics, which spans a broad range of activities in physics, including atomic physics, optics, condensed matter experiment and theory, biophysics experiment and theory, and plasma physics. 41 faculty members and approximately 50% of the graduate students in the department conduct research in this Division. In addition to the laboratories located on campus, MIT is affiliated with over 20 other research centers and facilities, including the MIT–Harvard Center for Ultracold Atoms, the Plasma Science and Fusion Center, the Fermi National Accelerator Laboratory (Fermilab), the European Organization for Nuclear Research (CERN), the Brookhaven National Laboratory: High Flux Beam Reactor and the Laser Interferometer Gravitational-Wave Observatory (LIGO). These affiliated centers and facilities employee staff scientists who work together with MIT faculty and graduate students, as well as other universities, on joint research projects. For this case study, three scientists were interviewed from the Department of Physics, Divisions of Astrophysics and Experimental Nuclear and Particle Physics.

Data Generation

While there are approximately 30 theoretical physicists at MIT, the majority of the faculty members in the Physics Department are experimental, and therefore, most of the data generated in the department is experimental data. Many of these experiments are conducted off campus at one of the larger MIT-affiliated laboratories and run continuously for months, or even years, producing hundreds of megabytes of data per second. For example, one scientist is working with the heavy ion group of the Compact Muon Solenoid (CMS) detector experiment at CERN, located in Switzerland and one of the world's largest and most respected centers for scientific research. This scientist's research group is made up of 20 high energy or nuclear physicists who receive and analyze data produced at the CMS detector. In addition to this group, more than two thousand other scientists collaborate in CMS, coming from 155 institutes in 37 countries.

A CMS experiment at this facility runs for nine months every year, writing data continuously at the rate of about 300-400 Megabytes per second (approximately 8,165 terabytes per year). This raw data is then processed 2-3 times per year, which triples the amount of data produced by the experiment. In addition to this processing, the raw data is also combined with simulation data generated at computer centers around the world, including the lab at MIT. When the data from each of these centers is combined, the total amount of simulation data produced is about the same as the original CMS experiment (i.e. the simulations generate about 1-2 terabytes of data per week). As a result, one nine-month CMS experimental run will generate approximately 40,824 terabytes of data.

Another example of physics research being conducted at MIT is the work being done by scientists at a gravitational-wave observatory (LIGO). The purpose of this observatory is to detect cosmic gravitational waves and to develop gravitational-wave observations as an astronomical tool. The facility consists of two separate installations within the United States, operated in unison as a single observatory. This observatory is available for use by the world's scientific community, and is a vital member of a developing global network of gravitational wave observatories. The MIT scientist is part of an international scientific collaboration, a growing group of approximately 600 researchers at roughly 40 institutions working to analyze the data from the observatory and other detectors, and working toward more sensitive future detectors.

This second scientist's research group at MIT consists of ten researchers working to analyze the data produced at the observatory. Experiments at the observatory run continuously for years at a time and produce about 1 terabyte of data per day. Since his work began on this project, the observatory has generated approximately 1,095 terabytes (1 petabyte) of data. This data was produced by five years of "half-run" experiments followed by one two-year long experiment. Like the data produced by the CERN collaboration, the data analysis and processing steps increase the total amount of data produced by the gravitational-wave experiments. However, unlike the data processing and simulation steps at the European facility, that more than triple the total amount of data generated by the experiment, the analysis done at the observatory only increases the amount of data by about 10%. Therefore, the gravitational-wave observatory has generated approximately 1,204 terabytes of data over the past seven years.

The amount of data produced has been steadily increasing over time. For example, 5 years ago MIT was not a computer center for the European facility and was not producing any data simulations. As a result, the first scientist's research group was producing about 10 times less data. In the future, the group plans to improve the capabilities of their computing system, allowing them to produce about 1.5 times more data each year. According to this growth plan, in 2014 this group alone will be generating about 540 terabytes of simulation data and all of the collaboration's computer centers combined will be generating 61,230 terabytes (approximately 60 petabytes) of simulation data.

Despite the high growth rate for the simulation data, the amount of raw experimental data produced by the CMS detector will not increase as quickly. This is because the detector technology will not change for at least another ten years. The scientist predicts that the rate of raw experimental data generation will remain constant for the next three years and then increase steadily by a factor of two each year for the next ten years as scientists improve their methods for data collection and processing. At this rate, the CMS detector experiment will be produce about 98,000 terabytes of experimental (raw and processed) data in 2014.

The growth rates seen and predicted by this scientist and his colleagues seem to be typical for physics labs at and/or affiliated with MIT Even smaller projects, like the work done by one scientist at a linear accelerator lab, have experienced large increases in data generation capabilities over the past five years. Similar to the CMS detector, experiments at this lab can run continuously for up to three years. However, these experiments typically generate much less experimental data, averaging about 1 terabyte of data per week. Despite the smaller amount of data produced, this professor's group has experienced similar data generation growth rates to the professor working with the CMS detector data. He estimates that the experiments conducted at the linear accelerator currently generate anywhere from 5-10 times more data than they were generating five years ago. This professor does not attribute this increase to changes in experimental instruments (they typically use the same instrument for a number of years), but instead to improvements in computing power and data storage technologies, which allow researchers to run their experiments for longer periods of time and to perform more advanced analyses. If the amount of data generated by the other 93 experimental physicists at MIT is similar to the amounts produced by the scientists interviewed, the Physics department as a whole is generating over 1,900,000 terabytes of data each year.

Metadata

Like many other scientific fields, metadata is often a crucial component to physics research at MIT. However, standards for recording and saving metadata are either non-existent or only defined within the specific research group. For example, researchers at the gravitational-wave observatory use an electronic logbook to record their metadata. This metadata includes experimental conditions (like start time, end time, data collection channels) as well as records of external noise that could affect the output of the gravitational wave detector. For example, when a plane flies overhead while an experiment is running, the lab technician will note the date, time, and a description of this event in the electronic logbook. The details regarding the event can be subjective and often vary based on which technician is recording the information at the time.

Metadata is also important for scientists with CMS detector data. However, the amount of metadata produced is extremely small when compared to the amount of raw and processed experimental or simulation data (about 10-30 kilobytes per experiment or simulation). Examples of metadata for a CMS simulation run by this scientist include the configuration of the experiment, the beam energy used and the physics processes selected to simulate. The metadata is uploaded to the central CERN database in Europe, where it is linked up with raw experimental data.

Data Retention

Due to the large amounts of data generated by long-running physics experiments, data storage and retention becomes a challenge for both individual researchers and larger laboratory facilities. Usually, only a small amount of actual experimental data is kept at MIT. Instead, research groups use a tiered approach, where different amounts of raw data are stored at multiple facilities depending on their capacity and research goals.

The distributed data storage used by the European collaboration is a good example of the methods used by many MIT physicists and their affiliated laboratories for data analysis, storage, retention, and sharing. The entire CMS research group at this collaboration is broken into three tiers. One facility is considered "tier 0" where all of the raw experimental data is stored forever. No processed data is stored at tier 0. The raw data is then copied, divided and distributed to the "tier 1" sites to be processed. The tier 1 sites must permanently store both the portion of the raw data that they receive from tier 0, as well as the processed data that they produce. The tier 1 sites are located all around the world. Next, the processed data is copied, divided, and distributed to all of the "tier 2" sites. The MIT lab is considered a tier 2 site. The tier 2 sites will only receive data from one tier 1 site unless they contact a different site and request their data.

Since it is at the bottom of the tier structure, there is no permanent storage at MIT. Instead, the university provides space for users to analyze the portions of the CMS experimental data that they receive from their tier 1 site. There are approximately 500 users from MIT and many other local institutions that have space here at MIT. Each user is provided with 1 terabyte of storage. Their data is replicated in a RAID array but not backed up. Users will usually request about 100 terabytes of data from the tier 1 site at a time and then filter down to about 1 terabyte before beginning analysis.

The scientist working the gravitational-wave observatory's lab is also considered a "tier 2" site in the observatory's data storage and analysis structure, however more data is stored locally at MIT then was true of the CMS data. This scientist has 10 RAID arrays with about 2-3 terabytes of storage each (i.e. a total storage capacity of about 300 terabytes). Additionally, members in the lab have local drives with about 200-300 gigabytes of short-term storage to use during their data analysis. Since it is considered a tier 2 site, none of the raw data used by the lab is backed up. However, the RAID arrays do provide redundancy. The main raw data backup is located on the West Coast at a "tier 1" site. The gravitational-wave observatory's community is involved in all data retention decisions. The collaboration has established committees that oversee how data analysis is done at each site, and decides what data to delete if more storage is needed. Additionally, the data derived by this scientist and his colleagues at MIT (i.e. the output of the analysis conducted at his lab) is archived on tapes.

Data Sharing and Reuse

As the examples in the previous section shows, data sharing among physicists is very common, and often essential. However the extent of data sharing differs depending on the research group. For example, two of the scientists only share data within their research centers and rarely share raw data with scientists outside their collaboration. Another scientist's group has signed an agreement with a sister project in Europe and frequently shares raw experimental data with them.

Regardless of data sharing practices, most physicists agree that their data can be re-used and re-analyzed often. For example, one scientist explained that he could re-use the same raw data hundreds of times and often re-integrates old data into new analyses. Since the raw data can be used for such a long period of time, the tiered data storing structure is extremely useful, allowing researchers in smaller labs to have access to the raw data without having to permanently store it.

Key Trends and Indicators for Data Growth

While each physicist at MIT has distinct research goals and methods for dealing with data, we have identified several trends and indicators for data growth.

1. The majority of the faculty members in MIT's Physics Department are experimental physicists. These physicists are often affiliated with large international or inter-institutional research centers and perform their experiments off campus.

2. Based on interviews with faculty members in this department, MIT-affiliated research laboratories and centers are currently generating data at a rate of 1,900,000 terabytes of data each year.

3. While improvements in experimental instruments can cause a jump in data production every 5-10 years, these instruments take years to develop and are not replaced often. Instead, the steady increase in data generation over time can be attributed to faster computing power and cheaper data storage hardware, which allow researchers to run their instruments for longer periods of time and to perform better data analysis.

4. The rate of data generation has been steadily increasing over time. Experimental physicists at MIT are currently producing about 5-10 times more data than they were five years ago.

5. Based on historical data and predicted improvements in computing power and storage, the rate of data generation for MIT physicists in 2014 will be approximately 11,400,000 terabytes per year.

6. While metadata is a critical to understanding the experimental conditions, the amount of metadata produced and stored is insignificant when compared to the amount of raw experimental data generated.

7. Many of the MIT-affiliated labs and centers use tiered data storage and sharing structure where all of the raw data is stored permanently at the "tier 0" site, and then divided among other tiers for redundancy and analysis.

8. Data sharing among physicists working at the same laboratory is common; however sharing among scientists at different labs is rare.

MIT physicists commonly re-use raw data from the same experiment multiple times and often re-integrated into new analyses as they are developed.

About the HMI? Program

The How Much Information? (HMI?) research program is a multi-discipline, multi-university project, formed to investigate the nature of data and information generated and used by individuals and enterprises. The program is sponsored by seven companies, including AT&T, Cisco, IBM, Intel, LSI, Oracle, and Seagate, and involves multiple research universities. The Principal Investigator is Prof. Roger Bohn and the Research Director is Dr. James Short, at UC San Diego's Global Information Industry Center (http://giic.ucsd.edu). Founded in 1960, the University of California, San Diego is one of the nation's most accomplished research universities, widely acknowledged for its local impact, national influence and global reach.

Acknowledgements

This case study is the product of industry and university collaboration in applied research. We are grateful for the support of our industry partners, sponsor liaisons, university research partners, and administrative staff at the University of California, San Diego.

Financial support for HMI? research and the Global Information Industry Center is gratefully acknowledged. Our sponsors are:

AT&T

Cisco Systems

IBM

Intel

LSI

Oracle

Seagate Technology

Additional support was provided by the Alfred P. Sloan Foundation of New York.

Questions about this research may be addressed to the Global Information Industry Center at the School of International Relations and Pacific Studies, UC San Diego:

Roger Bohn, Principal Investigator, rbohn@ucsd.edu

Jim Short, Research Director, jshort@ucsd.edu

Pepper Lane, Program Coordinator, pelane@ucsd.edu