

The Data Deluge: An e-Science Perspective

Tony Hey^{*+} and Anne Trefethen^{*}

UK e-Science Core Programme

Tony.Hey@epsrc.ac.uk

Anne.Trefethen@epsrc.ac.uk

Abstract

This paper previews the imminent flood of scientific data expected from the next generation of experiments, simulations, sensors and satellites. In order to be exploited by search engines and data mining software tools, such experimental data needs to be annotated with relevant metadata giving information as to provenance, content, conditions and so on. The need to automate the process of going from raw data to information to knowledge is briefly discussed. The paper argues the case for creating new types of digital libraries for scientific data with the same sort of management services as conventional digital libraries in addition to other data-specific services. Some likely implications of both the Open Archives Initiative and e-Science data for the future role for university libraries are briefly mentioned. A substantial subset of this e-Science data needs to be archived and curated for long-term preservation. Some of the issues involved in the digital preservation of both scientific data and of the programs needed to interpret the data are reviewed. Finally, the implications of this wealth of e-Science data for the Grid middleware infrastructure are highlighted.

* Postal address: EPSRC, Polaris House, North Star Avenue, Swindon SN2 1 ET, UK

+ On secondment from the Department of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK

To be published in “Grid Computing – Making the Global Infrastructure a Reality”, Wiley, January 2003

1. Introduction

There are many issues that should be considered in examining the implications of the imminent flood of data that will be generated by both the present and by the next generation of global 'e-Science' experiments. The term e-Science is used to represent the increasingly global collaborations – of people and of shared resources – that will be needed to solve the new problems of science and engineering [1]. These e-Science problems range from the simulation of whole engineering or biological systems, to research in bioinformatics, proteomics and pharmacogenetics. In all these instances we will need to be able to pool resources and to access expertise distributed across the globe. The IT infrastructure that will make such collaboration possible in a secure and transparent manner is referred to as the 'Grid' [2]. Thus in this paper the term 'Grid' is used as a shorthand for the middleware infrastructure that is currently being developed to support global e-Science collaborations. When mature, this Grid middleware will enable the sharing of computing resources, data resources and experimental facilities in a much more routine and secure fashion than is possible at present. Needless to say, present Grid middleware falls far short of these ambitious goals. Both e-Science and the Grid have fascinating sociological aspects as well as technical: we shall consider only technological issues in this paper.

The two key technological drivers of the IT revolution are Moore's Law – the exponential increase in computing power and solid-state memory - and the dramatic increase in communication bandwidth made possible by optical fibre networks using optical amplifiers and wave division multiplexing. In a very real sense, the actual cost of any given amount of computation and/or sending a given amount of data is falling to zero. Needless to say, whilst this statement is true for any fixed amount of computation and for the transmission of any fixed amount of data, scientists are now attempting calculations requiring orders of magnitude more computing and communication than was possible only a few years ago. Moreover, in many currently planned and future experiments they are also planning to generate several orders of magnitude more data than has been collected in the whole of human history.

The highest performance supercomputing systems of today consist of several thousands of processors interconnected by a special-purpose, high-speed, low latency network. On appropriate problems it is now possible to achieve sustained performance of several Teraflop/s – a million million floating-point operations per second. In addition, there are experimental systems under construction aiming to reach Petaflop/s speeds within the next few years [3,4]. However, these very high-end systems are, and will remain, scarce resources located in relatively few sites. The vast majority of computational problems do not require such expensive massively parallel processing but can be satisfied by the widespread deployment of cheap clusters of computers at university, department and research group level.

The situation for data is somewhat similar. There are a relatively small number of centres around the world that act as major repositories of a variety of scientific data. Bioinformatics, with its development of gene and protein archives, is an obvious example. The Sanger Centre at Hinxton near Cambridge [5] currently hosts 20 Terabytes of key genomic data and has a cumulative installed processing power (in clusters - not a single supercomputer) of around $\frac{1}{2}$ Teraflop/s. Sanger estimate that genome sequence data is increasing at a rate of 4 times each year and that the associated computer power required to analyse this data will 'only' increase at a rate of 2 times per year – still significantly faster than Moore's Law. A different data/computing paradigm is apparent for the particle physics and astronomy communities. In the next decade we will see new experimental facilities coming online that will generate data sets ranging in size from 100's of Terabytes to 10's of Petabytes per year. Such enormous volumes of data exceed the largest commercial databases currently available by one or two orders of magnitude [6]. Particle physicists are energetically assisting in building Grid middleware that will not only allow them to distribute this data amongst the 100 or so sites and 1000 or so physicists collaborating in each experiment, but also will allow them to perform sophisticated distributed analysis, computation and visualisation on all or subsets of the data [7-11]. Particle physicists envisage a data/computing model with a hierarchy of data centers with associated computing resources distributed around the global collaboration.

The plan of this paper is as follows. The next section surveys the sources and magnitudes of the data deluge that will be imminently upon us. This survey is not intended to be exhaustive but rather to give numbers that will illustrate the likely volumes of scientific data that will be generated by scientists of all descriptions in the coming decade. Section 3 discusses issues connected with the annotation of this data with metadata as well as the process of moving from data to information and knowledge. The need for metadata that adequately annotates distributed collections of scientific data has been emphasized by the Data Intensive Computing Environment (DICE) Group at the San Diego Supercomputer Center [12]. Their Storage Resource Broker data management middleware addresses many of the issues raised here. The next section on Data Grids and Digital Libraries argues the case for scientific data digital libraries alongside conventional literature digital libraries and archives. We also include a brief description of some currently funded UK e-Science experiments that are addressing some of the related technology issues. In the next section we survey self-archiving initiatives for scholarly publications and look at a likely future role for university libraries in providing permanent repositories of the research output of their university. Finally in section 6 we discuss the need for 'curation' of this wealth of expensively obtained scientific data. Such digital preservation requires the preservation not only of the data but also of the programs that are required to manipulate and visualize it. Our concluding remarks stress the urgent need for Grid middleware focused more on data than on computation.

2. The Imminent Scientific Data Deluge

2.1 Introduction

There are many examples that illustrate the spectacular growth forecast for scientific data generation. As an exemplar in the field of engineering, consider the problem of health monitoring of industrial equipment. The UK e-Science programme has funded the DAME project [13] - a consortium analyzing sensor data generated by Rolls Royce aero-engines. It is estimated that there are around 100,000 Rolls Royce engines currently in service. Each trans-Atlantic flight made by each engine, for example, generates about a Gigabyte of data per engine – from pressure, temperature and vibration sensors. The goal of the project is to transmit a small subset of this primary data for analysis and comparison with engine data stored in three data centres around the world. By identifying the early onset of problems, Rolls Royce hope to be able to lengthen the period between scheduled maintenance periods thus increasing profitability. The engine sensors will generate many Petabytes of data per year and decisions need to be taken in real-time as to how much data to analyse, how much to transmit for further analysis and how much to archive. Similar (or larger) data volumes will be generated by other high-throughput sensor experiments in fields such as environmental and earth observation, and of course human health-care monitoring.

A second example from the field of bioinformatics will serve to underline the point [14]. It is estimated that human genome DNA contains around 3.2 Gbases which translates to only about a Gigabyte of information. However, when we add to this gene sequence data, data on the 100,000 or so translated proteins and the 32,000,000 amino acids, the relevant data volume expands to the order of 200 Gigabytes. If, in addition, we include X-ray structure measurements of these proteins, the data volume required expands dramatically to several Petabytes, assuming only one structure per protein. This volume expands yet again when we include data about the possible drug targets for each protein – to possibly as many as 1000 data sets per protein. And there is still another dimension of data required when genetic variations of the human genome are explored. To illustrate this bioinformatic data problem in another way, let us look at just one of the technologies involved in generating such data generation. Consider the production of X-ray data by the present generation of electron synchrotron accelerators. At 3 seconds per image and 1,200 images per hour, each experimental station generates about 1 Terabyte of X-ray data per day. At the next generation ‘DIAMOND’ synchrotron currently under construction [15], the planned ‘day 1’ beamlines will generate many Petabytes of data per year, most of which will need to be shipped, analysed and curated.

From these examples it is evident that e-Science data generated from sensors, satellites, high-performance computer simulations, high-throughput devices, scientific images and so on will soon dwarf all of the scientific data collected in the whole history of scientific exploration. Until very recently, commercial databases have been the largest data collections stored electronically for archiving and analysis. Such commercial data are usually stored in Relational Database Management Systems (RDBMS) such as Oracle,

DB2 or SQLServer. As of today, the largest commercial databases range from 10's of Tbytes up to 100 Tbytes. In the coming years, we expect that this situation will change dramatically in that the volume of data in scientific data archives will vastly exceed that of commercial systems. Inevitably this watershed will bring with it both challenges and opportunities. It is for this reason that we believe that the data access, integration and federation capabilities of the next generation of Grid middleware will play a key role for both e-science and e-business.

2.2 Normalization

To provide some sort of normalization for the large numbers of bytes of data we will be talking about, the following rough correspondences [16] provide a useful guide:

A Large Novel	1 Mbyte
The Bible	5 Mbytes
A Mozart Symphony (compressed)	10 Mbytes
OED on CD	500 Mbytes
Digital Movie (compressed)	10 Gbytes
Annual production of refereed journal literature (~20k journals; ~2M articles)	1 Tbyte
Library of Congress	20 Tbytes
The Internet Archive (10B pages) (From 1996 to 2002) [17]	100 Tbytes
Annual production of information (print, film, optical & magnetic media) [18]	1500 Pbytes

Note that it is estimated that printed information constitutes only 0.003% of the total stored information content [18].

2.3 Astronomy

The largest astronomy database at present is around 10 Terabytes. However, new telescopes soon to come online will radically change this picture. We list three types of new 'e-Astronomy' experiments now under way:

1. Virtual Observatories

E-Science experiments to create 'virtual observatories' containing astronomical data at many different wavelengths are now being funded in the USA (NVO [19]), in Europe (AVO [20]) and the UK (AstroGrid [21]). It is estimated that the NVO project alone will store 500 Terabytes per year from 2004.

2. Laser Interferometer Gravitational Observatory (LIGO)

LIGO is a gravitational wave observatory and it is estimated that it will generate 250 Terabytes per year beginning in 2002 [22].

3. VISTA

The VISTA visible and infrared survey telescope will be operational from 2004. This will generate 250 Gbytes of raw data per night and around 10 Terabytes of stored data per year [23]. By 2014 there will be several Petabytes of data in the VISTA archive.

2.4 Bioinformatics

There are many rapidly growing databases in the field of bioinformatics [5, 24]:

1. Protein Data Bank (PDB)

This is a database of 3D protein structures. At present there are around 20,000 entries and around 2000 new structures are being added very 12 months. The total database is quite small, of the order of Gigabytes.

2. SWISS-PROT

This is a protein sequence database currently containing around 100,000 different sequences with knowledge abstracted from around a 100,000 different scientific articles. The present size is of the order of 10's of Gigabytes with an 18% increase over the last 8 months.

3. TrEMBL

This is a computer-annotated supplement to SWISS-PROT. It was created to overcome the time lag between submission and appearance in the manually curated SWISS-PROT database. The entries in TrEMBL will eventually move to SWISS-PROT. The current release has over 600,000 entries and is updated weekly. The size is of the order of 100's of Gigabytes.

4. MEDLINE

This is a database of medical and life sciences literature (Author, Title, Abstract, Keywords, Classification). It is produced by the National Library of Medicine in the USA and has 11.3M entries. The size is of the order of 100's of Gigabytes.

5. EMBL Nucleotide Sequence Database

The EBI in the UK is one of three primary sites for the deposition of nucleotide sequence data. It contains around 14M entries of 15B bases. A new entry is received every 10 seconds and data at the 3 centres – in the USA, UK and Japan – is synchronized every 24 hours. The EMBL database has tripled in size in the last 11 months. About 50% of the data is for human DNA, 15% for mouse and the rest a mixture of organisms. The total size of the database is of the order of Terabytes.

6. GeneExpression Database

This is extremely data intensive as it involves image data produced from DNA chips and microarrays. In the next few years we are likely to see 100's of experiments in 1000's of laboratories world-wide. Data storage requirements are predicted to be in the range of Petabytes per year.

These figures give an indication of the volume and variety of data that is currently being created in the area of bioinformatics. The data in these cases, unlike in some other scientific disciplines, is a complex mix of numeric, textual and image data. Hence mechanisms for curation and access are necessarily complicated. In addition, new technologies are emerging that will dramatically accelerate this growth of data. Using such new technologies it is estimated that the human genome could be sequenced in days rather than the years it actually took using older technologies [25].

2.5 Environmental Science

The volume of data generated in environmental science is projected to increase dramatically over the next few years [26]. An example from the weather prediction community illustrates this point.

The European Centre for Medium Range Weather Forecasting (ECMWF) in Reading UK currently have 560 active users and handle 40,000 retrieval requests daily involving over 2,000,000 meteorological fields. About 4,000,000 new fields are added daily, amounting to about 0.5 Terabytes of new data. Their cumulative data store now contains 3×10^9 meteorological fields and occupies about 330 Terabytes. To 1998, the increase in the volume of meteorological data was about 57% per year: since 1998, the increase has been 82% per year. This increase in data volumes parallels the increase in computing capability of ECMWF supercomputers.

This pattern is mirrored in the USA and elsewhere. Taking only one agency, NASA, we see predicted rises of data volumes of more than 10 fold in the 5 year period 2000 to 2005. The Eros Data Center (EDC) predict that their data holdings will rise from 74 Terabytes in 2000 to over 3 Petabytes by 2005. Similarly, the Goddard Space Flight Center (GSFC) predicts that its holdings will increase by around a factor of 10, from 154 Terabytes in 2000 to about 1.5 Petabytes by 2005. Interestingly, this increase in data volumes at EDC and GSFC is matched by a doubling of their corresponding budgets during this period and steady-state staffing levels of around 100 at each site. It is estimated that NASA will be producing 15 Petabytes by 2007. The NASA EOSDIS data holdings already total 1.4 Petabytes.

In Europe, European Space Agency (ESA) satellites are currently generating around 100 Gigabytes of data per day. With the launch of Envisat and the forthcoming launches of the Meteosat Second Generation satellite and the new MetOp satellites, the daily data volume generated by ESA is likely to increase at an even faster rate than that of the NASA agencies.

2.6 Particle Physics

The BaBar experiment has created what is currently the world's largest database: this is 350 Terabytes of scientific data stored in an Objectivity database [27]. In the next few years these numbers will be greatly exceeded when the Large Hadron Collider (LHC) at CERN in Geneva begins to generate collision data in late 2006 or early 2007 [28]. The

ATLAS and CMS experiments at the LHC each involve some 2000 physicists from around 200 institutions in Europe, North America and Asia. These experiments will need to store, access and process around 10 Petabytes per year which will require the use of some 200 Teraflop/s of processing power. By 2015, particle physicists will be using Exabytes of storage and Petaflop/s of (non-Supercomputer) computation. At least initially, it is likely that most of this data will be stored in a distributed file system with the associated metadata stored in some sort of database.

2.7 Medicine and Health

With the introduction of electronic patient records and improvements in medical imaging techniques, the quantity of medical and health information that will be stored in digital form will increase dramatically. The development of sensor and monitoring techniques will also add significantly to the volume of digital patient information. Some examples will illustrate the scale of the problem.

The company InSiteOne [29] is a US company engaged in the storage of medical images. It states that the annual total of radiological images for the US exceeds 420 million and is increasing by 12% per year. Each image will typically constitute many Megabytes of digital data and is required to be archived for a minimum of 5 years.

In the UK, the e-Science programme is currently considering funding a project to create a digital mammographic archive [30]. Each mammogram has 100 Mbytes of data and must be stored along with appropriate metadata (see below for a discussion of metadata). There are currently about 3M mammograms generated per year in the UK. In the USA the comparable figure is 26M mammograms per year, corresponding to many Petabytes of data.

A critical issue for such medical images - and indeed digital health data as a whole - is that of data accuracy and integrity. This means that in many cases compression techniques that could significantly reduce the volume of the stored digital images may not be used. Another key issue for such medical data is security - since privacy and confidentiality of patient data is clearly pivotal to public confidence in such technologies.

2.8 Social Sciences

In the UK, the total storage requirement for the social sciences has grown from around 400 Gigabytes in 1995 to more than a Terabyte in 2001. Growth is predicted to grow in the next decade but the total volume is not thought likely to exceed 10 Terabytes by 2010 [31]. There is experience in archive management for social science at the ESRC Data Archive in Essex, the MIMAS service in Manchester [32] and the EDINA service in Edinburgh [33]. The MIMAS and EDINA services provide access to UK Census statistics, continuous government surveys, macro-economic time series databanks, digital map datasets, bibliographical databases and electronic journals. In addition, the Humanities Research Board and JISC organisations in the UK jointly fund the Arts and

Humanities Data Service [34]. Some large historical databases are now being created. A similar picture emerges in other countries.

3. Scientific Metadata, Information and Knowledge

Metadata is data about data. We are all familiar with metadata in the form of catalogues, indices and directories. Librarians work with books that have a metadata ‘schema’ containing information such as Title, Author, Publisher and Date of Publication at the minimum. On the World Wide Web, most web pages are coded in HTML. This ‘HyperText Mark-up Language’ contains instructions as to the appearance of the page – size of headings and so on – as well as hyperlinks to other web pages. Recently the ‘eXtensible Mark-up Language’ XML has been agreed by the W3C standards body. The mark-up language XML allows web pages and other documents to be tagged with computer-readable metadata. The XML tags give some information about the structure and type of data contained in the document rather than just instructions as to presentation. For example, XML tags could be used to give an electronic version of the book schema given above.

More generally, information consists of semantic tags applied to data. Metadata consists of semantically tagged data that are used to describe data. Metadata can be organized in a schema and implemented as attributes in a database. Information within a digital data set can be annotated using a mark-up language. The semantically tagged data can then be extracted and a collection of metadata attributes assembled, organized by a schema and stored in a database. This could be a relational database or a native XML database such as Xindice [35]. Such native XML databases offer a potentially attractive alternative for storing XML-encoded scientific metadata.

The quality of the metadata describing the data is important. We can construct search engines to extract meaningful information from the metadata that is annotated in documents stored in electronic form. Clearly, the quality of the search engine so-constructed will only be as good as the metadata that it references. There is now a movement to standardize other, ‘higher-level’ mark-up languages – such as DAML and OIL [36] – that would allow computers to extract more than semantic tags and to be able to reason about the ‘meaning or semantic relationships’ contained in a document. This is the ambitious goal of Tim Berners-Lee’s ‘semantic web’ [37].

Although we have given a simple example of metadata in relation to textual information, metadata will also be vital for storing and preserving scientific data. Such scientific data metadata will not only contain information about the annotation of data by semantic tags, but will also provide information about its provenance and its associated user access controls. These issues have been extensively explored by Reagan Moore, Arcot Rajasekar and Mike Wan in the DICE group at the San Diego Supercomputer Center [38]. Their Storage Resource Broker (SRB) middleware [39] organizes distributed digital objects as logical ‘collections’ distinct from the particular form of physical storage or the particular storage representation. A vital component of the SRB system is the metadata catalog

(MCAT) that manages the attributes of the digital objects in a collection. Moore and his colleagues distinguish four types of metadata for collection attributes:

- SRB metadata for storage and access operations
- Provenance metadata based on the Dublin Core [40]
- Resource metadata specifying user access arrangements
- Discipline metadata defined by the particular user community.

In order for an e-Science project such as the Virtual Observatory to be successful, there is a need for the astronomy community to work together to define agreed XML schemas and other standards. At a recent meeting, members of the NVO, AVO and AstroGrid projects agreed to work together to create common naming conventions for the physical quantities stored in astronomy catalogues. The semantic tags will be used to define equivalent catalogue entries across the multiple collections within the astronomy community. The existence of such standards for metadata will be vital for the interoperability and federation of astronomical data held in different formats in file systems, databases or other archival systems. In order to construct 'intelligent' search engines, each separate community and discipline needs to come together to define generally accepted metadata standards for their community data grids. Since some disciplines already support a variety of existing different metadata standards, we need to develop tools that can search and reason across these different standards. For reasons such as these, just as the web is attempting to move beyond information to knowledge, scientific communities will need to define relevant 'ontologies' – roughly speaking, relationships between the terms used in shared and well-defined vocabularies for their fields – that can allow the construction of genuine 'semantic grids' [41, 42].

With the imminent data deluge, the issue of how we handle this vast outpouring of scientific data becomes of paramount importance. Up to now, we have generally been able to manually manage the process of examining the experimental data to identify potentially interesting features and discover significant relationships between them. In the future, when we consider the massive amounts of data being created by simulations, experiments and sensors, it is clear that in many fields we will no longer have this luxury. We therefore need to automate the discovery process - from data to information to knowledge - as far as possible. At the lowest level, this requires automation of data management with the storage and organization of digital entities. At the next level we need to move towards automatic information management. This will require automatic annotation of scientific data with metadata that describes both interesting features of the data and of the storage and organization of the resulting information. Finally, we need to attempt to progress beyond structure information towards automated knowledge management of our scientific data. This will include the expression of relationships between information tags as well as information about the storage and organization of such relationships.

In a small first step towards these ambitious goals, the UK GEODISE project [43] is attempting to construct a knowledge repository for engineering design problems. Besides traditional engineering design tools such as CAD systems, CFD and FEM simulations on

high performance clusters, multi-dimensional optimization methods and interactive visualization techniques, the project is working with engineers at Rolls Royce and BAESystems to capture knowledge learnt in previous product design cycles. The combination of traditional engineering design methodologies together with advanced knowledge technologies makes for an exciting e-Science research project that has the potential to deliver significant industrial benefits. Several other UK e-Science projects - the myGrid project [44] and the Comb-*e*-Chem project [45] – are also concerned with automating some of the steps along the road from data to information to knowledge.

4. Data Grids and Digital Libraries

The DICE group propose the following hierarchical classification of scientific data management systems [46]:

1. Distributed Data Collection. In this case the data is physically distributed but described by a single name space.
2. Data Grid. This is the integration of multiple data collections each with a separate name space.
3. Federated Digital Library. This is a distributed data collection or data grid with services for the manipulation, presentation and discovery of digital objects.
4. Persistent archives. These are digital libraries that curate the data and manage the problem of the evolution of storage technologies.

In this paper we shall not need to be as precise in our terminology but this classification does illustrate some of the issues we wish to highlight. Certainly, in the future, we envisage that scientific data, whether generated by direct experimental observation or by *in silico* simulations on supercomputers or clusters, will be stored in a variety of ‘data grids’. Such data grids will involve data repositories together with the necessary computational resources required for analysis, distributed around the global e-Science community. The scientific data – held in filestores, databases or archival systems – together with a metadata catalogue, probably held in an industry standard relational database, will become a new type of distributed and federated digital library. Up to now the digital library community has been primarily concerned with the storage of text, audio and video data. The scientific digital libraries that are being created by global, collaborative e-Science experiments will need the same sort of facilities as conventional digital libraries - a set of services for manipulation, management, discovery and presentation. In addition, these scientific digital libraries will require new types of tools for data transformation, visualization and data mining. We return to the problem of the long term curation of such data and its ancillary data manipulation programs below.

The UK e-Science programme is funding a number of exciting e-Science pilot projects that will generate data for these new types of digital libraries. We have already described both the ‘AstroGrid’ Virtual Observatory project [21] and the GridPP project [10] that will be a part of a world-wide particle physics grid that will manage the flood of data to be generated by the CERN LHC accelerator under construction in Geneva. In other areas

of science and engineering besides the DAME [13] and e-Diamond [30] projects described above, there are three projects of particular interest for bioinformatics and drug discovery. These are the myGrid [44], Comb-*e*-Chem [47] and DiscoveryNet [48] projects. These projects emphasize data federation, integration and workflow and are concerned with the construction of middleware services that will automatically annotate the experimental data as it is produced. The new generation of hardware technology will generate data faster than humans can process it and it will be vital to develop software tools and middleware to support annotation and storage. A further project, RealityGrid [49], is concerned with supercomputer simulations of matter and emphasizes remote visualization and computational steering. Even in such a traditional HPC project, however, the issue of annotating and storing the vast quantities of simulation data will be an important aspect of the project.

5. Open Archives and Scholarly Publishing

In the UK, the Higher Education Funding Council, the organisation that provides core funding for UK universities, is looking at the implications of the flood of e-Science data for libraries on a ten year time scale. In such a ten year time frame e-Science data will routinely be automatically annotated and stored in a digital library offering the 'usual' digital library services for management, searching and so on, plus some more specialized 'scientific data' oriented services such as visualisation, transformation, other types of search engines and so on. In addition, scientific research in many fields will require the linking of data, images and text so that there will be a convergence of scientific data archives and text archives. Scientific papers will also routinely have active links to such things as the original data, other papers and electronic theses. At the moment such links tend to be transitory and prone to breaking - perhaps the research group web address '~tony' stops working when Tony leaves, and so on. The Open Archive Initiative [49] - which provides software and tools for self-archiving of their research papers by scientists - addresses this issue to some extent, but this is clearly a large issue with profound implications for the whole future of university libraries. On the matter of standards and interworking of scientific digital archives and conventional repositories of electronic textual resources, the recent move of Grid middleware towards Web Services [50, 51] is likely to greatly facilitate the interoperability of these architectures.

Scholarly publishing will presumably eventually make a transition from the present situation - in which the publishers own the copyright and are therefore able to restrict the group of people who can read your paper - to a model where publishers are funded not for the paper copy but for providing a refereeing service and a curated electronic journal archive with a permanent URL. The difference between this model (proposed by Stevan Harnad [52]) and Paul Ginsparg's 'Eprint' archive for physics papers [53] is that Ginsparg's model is central and discipline based, whereas Harnad's is distributed and institution based. Both models depend on publishers to implement the peer review for the papers. Peer review is essential in order to identify signal from noise in such public archives. In Harnad's model, researchers' institutions pay 'publishers' to organize the peer reviewing of their research output and to certify the outcome with their journal name

and its established quality standard. The institutions' research output, both pre-peer review 'preprints' and post-peer review 'postprints' are archived in distributed, interoperable institutional Eprint archives. The Open Archives Initiative is providing a metadata harvesting protocol that could enable this interoperability. Using open source archiving software partly sponsored by the Budapest Open Access Initiative of the Soros Foundation, a growing number of universities in the US and elsewhere are setting up Eprint Archives to provide permanent open access to their research. In addition to archiving their own research output, users also want to be able to search these archives for related work of others. Using the metadata associated with the archived paper, the OAI Metadata Harvesting Protocol [54] provides one solution to the problem of constructing suitable search engines. Any search engine produced in this manner will only be only as good as the metadata associated with the papers [55], so strengthening and extending the metadata tagging and standards is a task of very high priority.

It seems just a question of time before scholarly publishing makes the 'Harnad Switch' – the outcome that Harnad has for a decade been describing as both optimal and inevitable. Authors actually want to maximize the impact and uptake of their research findings by making them accessible to as many would-be users as possible, rather than having them restricted, as they were in the paper era, to the minority of wealthy research libraries that can afford the access tolls. The Web has changed publishing forever and such a transition is inevitable. A similar transformation is likely to affect university libraries. The logical role for a university library in ten years will surely be to be the responsible organisation that hosts and curates (digitally) all the research papers produced by the university. It will be the university library that is responsible for maintaining the digital archive so that the '~tony' link continues to work for posterity. The Caltech Library System Digital Collections project [56] and the MIT DSpace project with HP [57] are two interesting exemplars of such an approach. There is also the interesting issue of how much responsibility individual universities would undertake for hosting and curating the scientific data produced by their researchers. Presumably, some universities would act as repositories for the scientific data for a number of university e-Science 'collaboratories', as well as acting as mirror sites for other organisations in the collaboration. Of course, particular communities will support specialised data archives - such as those of the European Bioinformatics Institute [24] and some national research organizations – and no doubt there will be commercial archives as well. An important issue not considered is the question of ownership of data. Since much of the research in universities is funded by public bodies there is clearly room for debate as to the ownership – and the curation costs!

6. Digital Preservation and Data Curation

Generating the data is one thing, preserving it in a form so that it can be used by scientists other than the creators is entirely another issue. This is the process of 'curation'. For example, the SWISS-PROT database is generally regarded as the 'gold standard' for protein structure information [58]. Curation is done by a team of 25 full-time curators split between the Swiss Bioinformatics Institute and the EBI. This shows how expensive the curation process is and why it will be necessary to address this support issue –

involving extreme levels of automated, semi-automated and manual annotation and data cleansing . In addition, preservation of the data will be a crucial aspect of the work of a data repository. A recent EU/US study [59] recommended the establishment of a ‘Data Rescue Centre’ that would be concerned with research into the longevity of electronic data archives. The report envisaged that such a centre would examine the issues concerned with the refreshment, replication, repackaging and transformation of data and become a center of much-needed expertise in these technologies.

There are many technical challenges to be solved to ensure that the information generated today can survive long term changes in storage media, devices and digital formats. An introduction to the issues surrounding this problem has been given by Rothenberg [60]. To illustrate these issues we shall briefly summarize a novel approach to long term preservation recently suggested by Lorie [61]. Lorie distinguishes between the archiving of data files and the archiving of programs. The archiving of programs is necessary in order that their original behaviour with the original data set can be reproduced in the future. For example, it is likely that a significant percentage of the scientific digital data to be preserved will be generated directly via some program P. A simple example is a spreadsheet program. In order to make sense of the data in the future we need to save the original program P that was used to create and manipulate the data along with the data itself. Of course, in one sense the program P is just a bit stream like the data it produces - but the important difference is that the machine and operating system required to run P may no longer exist. Lorie discusses the pros and cons of two proposed solutions to this problem: ‘conversion’ – copying files and programs to each new system as new systems are introduced – and ‘emulation’ – saving the data and the program as a bit stream along with a detailed description of the original machine architecture and a textual description of what the original program P should do to the data. Lorie then proposes a third approach based on specifying the program P in terms of instructions for a ‘Universal Virtual Computer’ (UVC). When archiving data, the UVC would be used to archive the methods that are required interpret the stored data stream. For archiving a program, the UVC would be used to specify the functioning of the original computer. It is not clear which of these three approaches will turn out to be most feasible or reliable. Needless to say, a solution to these problems is much more than just a technical challenge: all parts of the community from digital librarians and scientists to computer scientists and IT companies need to be involved.

7. Concluding Remarks

From the above discussion, it can be seen that the coming digital data deluge will have profound effects on much of current scientific infrastructure. Data from a wide variety of new sources will need to be annotated with metadata, archived and curated so that both the data and the programs used to transform can be reproduced in the future. E-scientists will want to search distributed sources of diverse types of data and co-schedule computation time on the nearest appropriate resource to analyse or visualize their results. This vision of Grid middleware will require the present functionality of both SRB [39] and Globus [62] middleware systems and much more. The present move towards Grid

Services and Open Grid Services Architecture represents a unique opportunity to exploit synergies with commercial IT suppliers and make such a Grid vision a reality.

Acknowledgements

The vision of the Grid described in this paper - with its emphasis on the access and integration of distributed data resources combined with that of remote access to distributed compute resources - owes much to discussions with many people. We would particularly like to acknowledge the contributions to our understanding of these issues of Jim Gray, Jeff Nick, Bill Johnstone, Reagan Moore, Paul Messina and the Globus team in the USA and Malcolm Atkinson, Stevan Harnad, Jessie Hey, Liz Lyon, Norman Paton, and Paul Watson in the UK. We are also grateful to Malcolm Read of JISC for always innovative support, to Sir Brian Follet for his early insight into the implications of e-Science for libraries and for universities and to John Taylor for both his vision for e-Science and for obtaining funding for the UK e-Science programme.

The authors are also grateful to David Boyd, Reagan Moore and Stevan Harnad for some helpful detailed comments on an earlier version of this paper.

References and URLs

- [1] J.M.Taylor, see <http://www.e-science.clrc.ac.uk>
- [2] I. Foster and C. Kesselman (editors) *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, San Fransisco 1999.
- [3] F.Allen et al., *BlueGene: A vision for protein science using a petaflop computer*, IBM Systems Journal, Vol 40, No 2, 2001.
- [4] T.Sterling, The Gilgamesh MIND Processor-in-Memory Architecture for Petaflops-Scale Computing, ISHPC Conference, Kansai, Japan, May 2002
- [5] Sanger Institute, Hinxton, UK: <http://www.sanger.ac.uk>
- [6] Gray J. and Hey T., *In Search of Petabyte Databases*, talk at 2001 HPTS Workshop
Asilomar www.research.microsoft/~gray
- [7] EU DataGrid Project: <http://eu-datagrid.web.cern.ch>
- [8] NSF GriPyN Project: <http://www.griphyn.org>
- [9] DOE PPDataGrid Project: <http://www.ppdg.net>
- [10] UK GridPP Project: <http://www.gridpp.ac.uk>
- [11] NSF iVDGL Project: <http://www.ivdgl.org>
- [12] A. Rajasekar, M.Wan and R.Moore, MySRB & SRB – Components of a Data Grid, 11th International Symposium on High Performance Distributed Computing, Edinburgh, Scotland (2002).
- [13] DAME Project www.cs.york.ac.uk/DAME
- [14] D.Stuart, Presentation at NeSC Workshop, Edinburgh, June 2002
- [15] DIAMOND project: <http://www.diamond.ac.uk>
- [16] Lesk M., *Practical Digital Libraries*, Morgan Kaufmann, San Fransisco (1997).
- [17] Internet Archive: <http://www.archive.org>
- [18] Lyman P. and Varian H.R., *How Much Information?* UC Berkeley School of Information Management & Systems Report: <http://www.sims.berkeley.edu/how-much-info> (2000).
- [19] NVO: <http://www.nvo.org>
- [20] AVO: <http://www.eso.org/avo>
- [21] AstroGrid: <http://www.astrogrid.ac.uk>
- [22] LIGO: <http://www.ligo.caltech.edu>
- [23] VISTA: <http://www.vista.ac.uk>
- [24] European Bioinformatics Institute: <http://www.ebi.ac.uk>
- [25] J. Hassard, private communication (2002).
- [26] R.Gurney, private communication (2002)
- [27] BaBar Experiment: www.slac.stanford.edu/BFROOT/
- [28] LHC Computing Project: <http://lhcgird.web.cern.ch/LHCgrid>
- [29] InSiteOne digital image storing and archive service: <http://www.Insiteone.com>
- [30] Proposed e-Diamond project: See <http://e-science.ox.ac.uk/>
- [31] J. Neathey, private communication (2002).
- [32] MIMAS Service: <http://www.mimas.ac.uk>
- [33] EDINA Service: <http://edina.ac.uk>
- [34] Arts and Humanities Data Service: <http://www.ahds.ac.uk>
- [35] Xindice native XML database: <http://xml.apache.org/xindice>

- [36] DAML+OIL: <http://www.daml.org/2001/03/daml+oil-index.html>
- [37] Tim Berners-Lee with Mark Fischetti, *Weaving the Web*, Harper Collins, 1999
- [38] A. K. Rajasekar and R.W. Moore, *Data and Metadata Collections for Scientific Applications*, European High Performance Computing Conference, Amsterdam, Holland (2001).
- [39] The Storage Resource Broker: <http://www.npaci.edu/DICE/SRB>
- [40] Dublin Core Metadata Initiative: <http://dublincore.org>
- [41] D. DeRoure, N. Jennings and N.Shadbolt, *Towards a Semantic Grid*, Concurrency & Computation (to be published) and in this collection.
- [42] R.W. Moore, *Knowledge-Based Grids*, Proceeding of the 18th IEEE Symposium on Mass Storage Systems and Ninth Goddard Conference on Mass Storage Systems and Technologies, San Diego, April 2001.
- [43] The GEODISE Project: <http://www.geodise.org/>
- [44] The myGrid Project: <http://mygrid.man.ac.uk>
- [45] The Comb-e-Chem project: <http://www.combechem.org>
- [46] R.W.Moore, *Digital Libraries, Data Grids and Persistent Archives*, presentation at NARA December 2001.
- [47] The DiscoveryNet project: <http://www.discovery-on-the.net>
- [48] The RealityGrid project: <http://www.realitygrid.org>
- [49] C. Lagoze and H. Van De Sompel, *The Open Archives Initiative: Building a Low-Barrier Interoperability Framework*, JCDL'01, Roanoke, Virginia, ACM.54-62 (2001)
- [50] I.Foster, C. Kesselman and J. Nick, *Physiology of the Grid*, Concurrency and Computation (to be published) and in this collection.
- [51] T. Hey and L.Lyon, *Shaping the Future? Grids, Web Services and Digital Libraries*, International JISC/CNI Conference, Edinburgh, Scotland, June 2002.
- [52] S.Harnad and J.M.N. Hey, *Esoteric Knowledge: the Scholar and Scholarly Publishing on the Net*, in Proceedings of International Conference on Networking and the Future of Libraries: Managing the Intellectual Record, edited by L. Dempsey, D.Law and I. Mowlat, Bath, UK, April 1995.
- [53] Ginsparg e-print archive: <http://arxiv.org>
- [54] C. Lynch, *Metadata Harvesting and the Open Archives Initiative*, ARL Bimonthly Report 217:1-9 (2001).
- [55] Z. Lui, K. Maly, M. Zubair and M.Nelson, *Arc – An OAI Service Provider for Cross-Archive Searching*, JCDL'01, Roanoke, Virginia ACM 65-66 (2001)
- [56] The Caltech Library Systems Digital Collections project: <http://library.caltech.edu/digital/>
- [57] The MIT Dspace project: <http://www.mit.edu/dspace/>
- [58] SWISS-PROT: <http://www.ebi.ac.uk/swissprot>
- [59] EU/US Workshop on Large Scientific Databases: <http://www.cacr.caltech.edu/euus>
- [60] J.Rothenberg, *Ensuring the Longevity of Digital Documents*, Scientific American, 272(1), January 1995.
- [61] R.A. Lorie, *Long Term Preservation of Digital Information*, JCDL'01, Roanoke, Virginia, June 2001
- [62] The Globus project: <http://www.globus.org>