

# Manchester United in Supporting UK Academia: Recent developments at the UK's Largest Academic Datacentre

Ross MacIntyre,  
*Manchester Computing, University of Manchester, UK.*  
*Email: ross.macintyre@man.ac.uk*

## Abstract

*This is not a paper about strategic directions for Digital Library development. It deliberately focuses on what is being put in place right now in an effort to better support academic endeavour in particular using electronic materials. The examples are all drawn from the UK's largest academic datacentre, Manchester Computing, at the University of Manchester, which, in addition to supporting local needs, receives funding to host services for use by the UK academic community. These services include networking, visualisation, datasets and software provision. Many of the needs being met will be common to a wider community.*

## 1. Introduction

### 1.1. The University of Manchester

The University of Manchester[1] has its origins in Owens College, which was established in 1851. The College was granted a Royal Charter in April 1880 as the Victoria University, a federal institution which established colleges in Leeds and Liverpool. These Colleges were granted their independence in 1903 becoming the Universities of Leeds and Liverpool respectively, and creating the Victoria University of Manchester - the first UK civic university.

The University has more than 70 departments involved in teaching and research, with almost 3,000 academic staff. Manchester is the UK's most popular University with more than 18,000 full-time students, including 2,500 international students, from over 120 countries. All departments offer postgraduate qualifications via research or taught courses.

There are excellent laboratory and computer facilities with more than 6,000 PCs across the campus and every student has access to e-mail and the World Wide Web. With 3,600,000 books and more than

1,000,000 manuscript items, the John Rylands University Library of Manchester is the third largest in the UK.

Some of the key scientific developments of the century have taken place at the University; for example it was here that Rutherford split the atom. In particular, Manchester has a long tradition of involvement in the design and construction of computers and the development of computing techniques. In 1948, Williams and Kilburn invented the Williams tube, the first computer memory. They then built the world's first "stored program machine"; which was rebuilt for the 50th anniversary celebration and can be seen at the Science Museum. The world's first commercial computer was developed here with Ferranti. Virtual memory was invented at Manchester, then subsequently sold to IBM.

Manchester soon began to provide services and share its facilities with other universities. This role was formalised in 1969 when it became one of the two Regional Computing Centres and in 1987 it was designated a National Computing Centre. It eventually became known as just 'Manchester Computing'[2].

### 1.2. Manchester Computing

As would be expected, Manchester Computing (MC) offers a full range of local support services to the University, including faculty liaison, application support, printing services and administrative computing.

MC also provides a number of other services, where the scope extends beyond local boundaries, to a regional, national and international user base. These are Manchester Visualisation Centre (MVC), High Performance Computing (HPC) and National Data Services. Network systems, which comes under the Communications, Operations & Systems (COS) division, underpins all these services.

This paper focuses on those areas looking beyond the purely local and will start with the network, before highlighting some examples from the other services listed above.

## 2. Networking

### 2.1. The UK's Joint Academic Network - JANET

The UK education and research community benefits from its own network, referred to as JANET[3]. JANET is linked to other academic and commercial networks in the UK and abroad, forming part of the global internet. The current version is specifically called SuperJANET3 and links Universities and Higher Education Institutions in the UK via 4 main sites: Manchester, London, Leeds and Bristol, interconnected at 155Mbps, with the other sites linked in at speeds ranging from 64Kbps to 155Mbps. The situation is about to change dramatically with the introduction of SuperJANET4. This new network, to be implemented over the next six months, will be based around 8 sites: Manchester, Leeds, London, Bristol, Reading, Portsmouth, Glasgow and Edinburgh, all at speeds of 2.5Gbps, rising to 10Gbps in 2002. SuperJANET4 will also be linked to the Internet2 network in the US, offering a high-speed research network for collaborative research between geographically separated sites. In particular, this infrastructure enables what has become known as 'eScience'.

In all configurations, Manchester is a major node and supports gigabit connection to this national network from the campus. The COS division at MC are responsible for this substantial infrastructure, including connections to the Metropolitan Area Network called G-MING[4], described below.

### 2.2. Greater Manchester Information Network Group (G-MING)

G-MING originated in 1993 when a feasibility study was funded to examine the scope and potential for a high speed telecommunications infrastructure for the educational community within the Greater Manchester area. Behind the initial initiative were 6 Higher Education Institutions (HEIs) in Greater Manchester. Collectively, the consortium represents the largest educational precinct in Europe.

All the institutions faced a common need to provide distributed teaching services and to open up links for

research. In particular they all had teaching and accommodation buildings that were geographical dispersed, some of them sharing a site, and faced common difficulties in handling voice and data services between central and remote sites. Thus a high-speed backbone could be used to support data and cost effective telephone services within institutions, digitised audio-video services for remote teaching buildings together with security and other remote control systems (e.g. heating). Furthermore, access to central resources could be provided from halls of residence and even staff/student homes where locally based. Above all, the infrastructure had the potential to establish a pervasive multimedia service which could enable new distributed multimedia teaching and remote learning facilities to be set up and shared between institutions. A particular aspect would be the ability to link academic libraries. Furthermore, as well as serving the digital distribution needs of the six founding HEIs, G-MING was conceived as a network for the broader educational community of Greater Manchester.

Over the past seven years, the infrastructure, a high bandwidth ATM-based network, has been implemented that stretches across central and southern Manchester with links out to neighbouring Salford, interlinking the core, remote teaching and residential sites of the 6 HEIs together with key city locations, encompassing up to 48 sites. The development established a core network between the main sites of the 6 institutions and then expanded out to strategic sites. A key advantage provided through the G-MING infrastructure is that organisations that have remote buildings co-located within a particular site are able to benefit from the shared link back to the core site.

### 2.3. NetNorthWest

NetNorthWest[5] is a consortium of 18 Higher Education Institutions (HEIs) in the North West of England that is installing a pervasive and advanced high-speed telecommunications infrastructure. This goes beyond connecting the main campuses of the institutions and involves networking a whole range of subsidiary campuses and sites.

However, the HEIs see NetNorthWest as something much broader than an initiative to improve communications amongst its initial members. They regard it as providing an important opportunity for regional collaboration that will support not just

educational developments but also economic regeneration.

The core network (10 sites) went into service in July 1998. The links to each site are all at 155 Mb/s except that between Manchester and Liverpool which is at a speed of 622 Mb/s. This has created in the North West a basic high-speed network that can provide an unequalled opportunity to form the core for a regional network to support both economic and educational developments.

Recently 5 more sites were connected (including Jodrell Bank) each at 155 Mb/s. In addition to the main core sites 53 FE Colleges, 22 Hospitals, 6 Research Institutes, 5 Public Authorities, 6 Schools and 12 other sites (e.g. Science Parks, commercial companies, Museums etc) are connected to the Internet via NetNorthWest.

#### 2.4. The National Cache

For the past two years Manchester, Loughborough and London Universities have jointly provided the UK academic community's National Cache. The current service is provided by over 30 machines located at the three locations. An institution wishing to use the current service is allocated a machine at two of these nodes to peer with. For a given institution, the time taken to retrieve a Web object from a remote site is related to the loading on its two peer machines; as this loading increases during periods of high demand, retrieval times increase, sometimes to unacceptably high levels.

A new system makes the process of allocating individual machines to an institution obsolete. Instead it groups together the cache machines at each of the three National Cache nodes, enables load balancing within a group of machines and facilitates the seamless incorporation of extra caching capacity when required. This new system is based upon the Linux Virtual Server model (LVS)[6].

### 3. Service Provision

Having invested in such a comprehensive networking infrastructure, Manchester is well placed to exploit its position via collaborations and high-demand service provision. These are explored in the remaining sections.

#### 3.1. High Performance Computing (HPC)

At Manchester Computing the primary service for High Performance Computing is provided by CSAR[7], Computer Services for Academic Research, which consists of a consortium of the Cray Research division of Silicon Graphics Incorporated (SGI), the Supercomputing Operations division of Computer Sciences Corporation and Manchester Computing. CSAR offers access to an 816 processor Cray T3E-1200E, a 16 processor Origin2000 and an 8 processor Fujitsu VPP300. This is backed-up with a Storagetek Powderhorn Tape Management Library with a capacity up to 150 Terabytes. The service started in November 1998 as a £26 Million, 6 year investment which allows CSAR to provide "a world class service for world class science".

As part of the commitment to the continuous development of the CSAR service, SGI has agreed to locate the European ASCI Prototyping Centre at Manchester. The Accelerated Strategic Computing Initiative (ASCI) is an American Government Department of Energy programme, aimed at eliminating the need for nuclear testing through the simulation of the existing nuclear stockpile in the USA. ASCI is also an acronym for the type of programming being used.

Manchester has recently been successful in gaining funding to make available to the JANET community a global supercomputer spanning Europe, United States and Japan. The metacomputer will be based around two persistent hubs at Manchester and the Pittsburgh Supercomputer Centre in the US. The metacomputer will give access to 2.5 Teraflops of computing power, making it larger than any other machine available to the academic community in Europe. The project has the support of all of the sites that participated in SuperComputing'99, including the universities of Tsukuba and Tokyo, Japan; Pittsburgh and Portland, US; Stuttgart, Germany and Manchester, UK.

#### 3.2. National Datasets via MIMAS

MIMAS[8] is a National Data Centre funded primarily by government via its Joint Information Systems Committee (JISC)[9] of the Higher Education Funding Councils. It provides the UK higher education, further education and research community with networked access to key data and information resources to support teaching, learning and research across a wide range of disciplines. MIMAS is thus a

major resource centre within what is referred to as the DNER - Distributed National Electronic Resource[10].

MIMAS services are free of charge to users at eligible institutions - although site subscriptions are required for a number of the data and information resources available and some services require individual registration.

Some of the dataset services hosted are mentioned below.

### 3.3.1. Socio-Economic data

#### Census

The ESRC Census Dissemination Unit (CDU) is based at MIMAS. UK Census of Population Statistics for 1981 and 1991 and Samples of Anonymised Records are available, along with related information and data derived from Census data.

Casweb, the web-based interface to the 1991 Census statistics, was developed by the CDU at MIMAS to extend access to those engaged in short projects, and for teaching at undergraduate and post-graduate level. The interface is easy to learn and requires little prior knowledge of the content or structure of Census statistics. A number of visualization tools are also offered to facilitate visual exploration of Census data with reference to geographical location.

#### Surveys

MIMAS hosts several large and complex survey datasets, including repeated cross-sectional surveys, where different cases are sampled for each survey, e.g. Labour Force Survey and longitudinal (panel) surveys, where the same cases are repeatedly sampled, e.g. National Child Development Study.

#### Time Series

Information gathered over recent years (up to 20 years in some cases) on production, trade, sales, etc. can be extracted and downloaded, or explored on the MIMAS server using packages such as SAS and RATS. The time series come from the following sources: Office for National Statistics (ONS), Organisation for European Co-operation and Development (OECD), United Nations Industrial Development Organisation (UNIDO) and International Monetary Fund (IMF).

A web interface provides easy access to over 40,000 time series in the ONS Databank, and we are developing tools to access the other time series via the Web.

Numerous software packages are hosted at MIMAS to support socio-economic data analysis and spatial data manipulation.

### 3.3.2. Spatial data

MIMAS hosts a large spatial data resource including the Bartholomew digital map data, the '81 and '91 Census of Population digital boundary data, and extensive archives of satellite data from SPOT and Landsat.

### 3.3.3. Scientific data

#### Crossfire

The Beilstein CrossFire Information System is an integrated chemical information system offering a Reactions database for organic chemistry, the Gmelin database for inorganic and organometallic chemistry and an Abstracts database. MIMAS provides this to UK, Republic of Ireland and Scandinavian Universities.

Also available are the Cambridge Structural Database, containing crystal structure data for research chemists and the Mossbauer Effect Reference Database, containing information on spectroscopy for research chemists.

### 3.3.4. Electronic Journals

#### JSTOR

The UK JSTOR Mirror Service is run by MIMAS. JSTOR is a unique digital archive collection of core scholarly journals starting with the first issues (many of which date from the 1800s) and made available to participating UK higher education institutions via the World Wide Web. 117 journals in a variety of fields are accessible in the Arts & Sciences collection, which contains some 4.6 million pages of journal literature from key titles. A General Science collection is currently being developed covering: Philosophical Transactions, Proceedings of the Royal Society, Science and Proceedings of the National Academy of Science.

JSTOR is widely used by institutions in many countries including US, UK and Japan. The Japanese service is provided in conjunction with USACO and users benefit from documentation translated and published in Japanese.

NESLI - The National Electronic Site Licence Initiative

The NESLI service has been established by the JISC to promote the widespread delivery and use of electronic journals in the UK Higher Education and research community. NESLI aims to address the many issues that at present hinder the most effective use, access and purchase of electronic journals in the academic library community. These include technical issues of access control, both for on-site and 'remote' users, cost, site definitions, and archiving issues. JISC appointed a consortium of SwetsBlackwell and Manchester Computing as Managing Agent for NESLI. The role of the Managing Agent is to coordinate delivery of the electronic material, undertake negotiations with publishers and oversee the day-to-day operation of the programme.

The initiative has attracted world-wide interest and a recent visit from a Japanese delegation followed others from US, Canada, Australia and many European countries.

### 3.3.5. Bibliographic reference

ISI Web of Science Service for UK Education

MIMAS host ISI's citation databases for the UK education community. Users include Higher and Further Education Institutions and the Research Councils.

The service is the most widely used of all JISC-supported datasets. The service is supported in collaboration with colleagues in John Rylands University Library of Manchester and Manchester Metropolitan University library.

COPAC

COPAC is a union catalogue that gives free access to the merged online catalogues of members of the UK's Consortium of University Research Libraries (CURL). COPAC records represent documents dating from c.1100AD onwards, in many languages. A variety of materials are included and some records include links to the document full-text. In addition to the Web Interface you can access COPAC using the Text Interface (telnet) or via the Z39.50 Target. As the service is freely available, its user base is global.

ZETOC

MIMAS has recently started a joint development programme with the British Library. The newly launched ZETOC service provides Z39.50-compliant access to the British Library's Electronic Table of

Contents (ETOC) database. The database contains details of approximately 20,000 current journals and 16,000 conference proceedings published per year. With almost 15 million article and conference records, the database covers every imaginable subject in science, technology, medicine, engineering, business, law, finance and the humanities. The database covers the years from 1993 to date and is updated daily. Copies of all the articles and conference papers listed on the database are available from the British Library's Document Supply Centre. A Table of Contents (ToC) Alerting service has also been added, which sends users the ToC of their nominated journal issues as they are added via the daily update - this is typically within 72 hours of publication.

### 3.3.6. Research and Development

MIMAS is also active in a number of collaborative development projects and research initiatives, most notably in the area of electronic publishing. Consultancy is provided to many organisations, both academic and commercial, on metadata, SGML/XML, linking and other aspects of 'e-publishing'.

MIMAS is a member of the Dublin Core Metadata Working Groups[11], most recently attending DC8 in Ottawa (note that DC9 will be held in Tokyo next year). Dublin Core (DC) is a truly international initiative with translations of the basic metadata elements now being available in many languages, including Japanese.

Putting this into practice, a project is currently underway to develop a repository of DC metadata information relating to MIMAS data and information services. This includes offering a Z39.50 infrastructure, utilising the 'Cheshire II' software[12]. This project is in collaboration with the Universities of Liverpool and California - Berkeley.

MIMAS is also a member of the CEN/ISSS (European standards) Workshops on Metadata for Multimedia Information (MMI-DC).

### 3.3. The Manchester Visualization Centre (MVC)

The MVC[13] was originally established in 1974 as the Computer Graphics Unit, but 'outgrew' its title by providing services and undertaking R&D in high-performance and cluster computing, interactive computer graphics, multimedia, image processing and visualization. Some significant areas of activity are given below.

### 3.3.1. VIP Lab - Visualization Immersive Projection Laboratory

High quality visualisation and virtual reality images can be displayed in the newly opened state-of-the-art immersive projection theatre at MVC - known as the VIP lab. It offers a 6 metre wide, 2 metre high cylindrical screen, 3 projector systems, a 6-channel surround sound system, interactive control, stereographic glasses and emitters. The experience can be shared by audiences of up to 30 people, facilitating group teaching and dissemination activities.

A number of research areas in both medical and physical science will now take significant steps forward through the use of visual supercomputing. The centre combines high-performance graphics, computing and data management technology to solve complex research problems that were unapproachable only a few years ago.

Applications include:

- Interactive interpretation of 3D scanned images (CT, MRI, and Angiograms) in the operating theatre during actual surgery.
- Visualization of the results of finite element calculations for engineers.
- The understanding and prediction of complex 3D flows.
- Multi-dimensional visualization for astronomy.
- Numerical modelling in basin studies
- Visualization of complex flows in polymer processing machinery.
- Particle physics and the high performance visualisation facility.

The VIP lab is supported by Europe's most powerful graphics supercomputer, a 40-processor Origin2000 machine with 6 graphics pipelines.

### 3.3.2. National Video Service

Recognising the value of video as a visual aid in presentation, MVC offers the National Video Service, covering creation, conversion and serving.

Video provides a means to record sequential visualised data and in particular multi dimensional models of the data. The great advantage of video is that it is a relatively cheap method of presenting information whilst allowing tight control over the output.

The National Video Server is an SGI Origin 200 running MediaBase and RealServerG2 video server

software. The machine has 110Gb of disk space and dual 100BaseT network connections to the SuperJANET backbone, with a total available bandwidth of 180Mbits/sec. The MediaBase software allows the videosever to supply live and timed delivery, MPEG-1 and MPEG-2 video streams. While the Real G2 server allows the video server to broadcast live RealMedia G2 streams over the web.

### 3.3.4. Development Projects

As might be expected, MVC has numerous development projects underway, two are mentioned here.

#### IERAPSI - Integrated Environment for the Rehearsal and Planning of Surgical Interventions

In Europe, the integration of advanced technologies such as 3D visualisation, virtual reality, and computerised simulation for pre-operative planning, performance and training of surgical procedures is rapidly gaining in interest. Funded by the European Commission, a team of eight partners is setting up an IERAPSI at MVC. The IERAPSI consortium will particularly focus on the planning, simulation, and training of surgical procedures performed on the ear's petrous bone, a common site with a complex anatomy. The system will address frequently applied interventions, such as mastoidectomy, but equally very exacting ones, like cochlea implantation. In a later stage, IERAPSI might also be used for other surgery specialities.

#### WebSET - Web-based Standard Educational Tools

The WebSET project aims to produce a standardised suite of interactive three-dimensional educational tools, delivered across the WWW. The major focus is put on the use of open technology and standards, as well as the production of learning components which can be used as building blocks for further development in a wide range of application areas, such as surgical training and physiological education. WebSET is supported by the European Commission.

## 4. Conclusion

The key to successful operation is active co-operation and collaboration, acting as a team with a shared goal, something we have in common with another successful team in the area - Manchester United Football Club.

Manchester Computing began the millennium as it intends to go on, as Europe's premier academic computing facility, supporting world-class teaching and research. The services offered to UK academia continue to expand and MC is active in many high profile projects that will be of benefit to users of computing technologies, not just within the UK, but around the world.

#### References

- [1] <http://www.man.ac.uk>
- [2] <http://www.mcc.ac.uk>
- [3] <http://www.ja.net>
- [4] <http://www.g-ming.net.uk>
- [5] <http://www.netnw.net.uk>
- [6] <http://wwwcache.ja.net/LVS>
- [7] <http://www.csar.cfs.ac.uk>
- [8] <http://www.mimas.ac.uk>
- [9] <http://www.jisc.ac.uk>
- [10] [http://www.jisc.ac.uk/pub99/dner\\_desc.html](http://www.jisc.ac.uk/pub99/dner_desc.html)
- [11] <http://purl.org/dc>
- [12] <http://cheshire.lib.berkeley.edu>
- [13] <http://www.man.ac.uk/MVC>