Astro-grid

an e-science initiative

Paper for PPARC Astronomy Committee
October 31st 2000

Andy Lawrence, Richard McMahon, Mike Watson, Peter Allan, Fionn Murtagh, Mike Irwin, Clive Page, Bob Mann, Mike Lockwood, John Sherman, David Giaretta.

EXECUTIVE SUMMARY

A tidal wave of data is approaching Astronomy, requiring radical new approaches to database construction, management, and utilisation. The next two explosion points are in 2003 and 2005 when UKIRT WFCAM and VISTA (both UK facilities) come on-line and start producing a Terabyte of data *every night*. The UK also has a lead role in other key databases accumulating now in X-ray, solar, and space plasma astrophysics. There are rich scientific opportunities, but also present serious worries; current data manipulation methods will not scale to these volumes. The UK is in a position of strength internationally because of the importance of these datasets, but will remain in a strong position only if it solves the data-handling problems. Astro-grid proposes to deliver the tools needed to meet that challenge.

At the same time, a new style of science is emerging, which requires the searching, filtering, manipulation, or browsing of *entire* vast datasets. Sometimes this is because the scientific problem requires massive analysis (e.g. the spatial clustering power spectrum of a billion galaxies), sometimes because rare events or objects need to be found (e.g. locating the one in ten million faint stellar images that *might* be one of the first generation of pre-galactic stars) or sometimes because one wants to examine, explore, and get ideas. These systematic and "data-mining" styles of working are already a strength of a minority of astronomical specialists, but require a long difficult slog, and will become impossible as the data explosion hits unless we transform our technology. In fact our intention is to develop hardware, architecture and data management solutions based round a small number of expert data centres, along with query, analysis and exploration tools that will make it *easier* and *quicker* to take this approach from any desktop, thus democratising it. Everybody can be a power-user. We therefore expect an acceleration of real science.

The real fruits however will come from the ability to take this approach to multiple databases simultaneously. One wants to browse optical and X-ray images of the same piece of sky taken by different research teams, to click on an interesting object and to find to ones surprise that somebody just got an IR spectrum of it at UKIRT last month ... and of course to retrieve the calibrated spectrum immediately. This kind of science is already done, but slowly and painfully by hand, as the databases are spread all over the globe, with a variety of formats, access levels and policies, and it is hard to keep up to date with what's available. A key aim of astronomers world-wide is to stitch together essentially all such databases in a "Virtual Observatory" together with intelligent discovery agents so that the archive becomes like a second sky. This is a very ambitious aim, with many technical and organisational challenges.

The aim of the UK Astro-Grid project is to focus on short term deliverables, both relevant application tools and the federation, by the data centres that manage them, of key sky

survey datasets, namely: (a) SuperCOSMOS, Sloan, INT-WFC and UKIRT WFCAM; (b) XMM-Newton and Chandra; (c) and SOHO, Cluster and IMAGE. The differences between the data types involved in these federations means that each brings distinct challenges, whose solutions will shed light on generic problems to be faced by the developing global "Virtual Observatory". We set out below a three year programme, starting in April 2001 and estimated to cost £5.5M, which will add value to existing UK astronomy resources in the short term, as well as positioning the community strongly with respect to wider Grid initiatives, in astronomy and beyond. The new Government e-science initiative means that there is a good chance that PPARC will have the necessary extra funding to undertake such a programme.

(1) COLLECTIVISATION AND EMPOWERMENT OF THE INDIVIDUAL

Over three decades astronomy has moved inexorably towards a more communally organised approach, but paradoxically this has increased the power of individuals to undertake research. First, it became normal for observatories and satellites to be equipped with "common-user" or "facility class" instruments rather than private equipment, built to a high standard of robustness, simplicity of use, and documentation. Most astronomers did not need to build equipment, only to learn to drive it, and so could get straight to the science. More than anything else, this change led to the flowering of astronomy in the UK. The next step was the communal development of data reduction tools (Starlink, IRAF, Midas) which again shortened the route to scientific results. Then came the provision of calibrated electronic archives, with some simple tools, but mostly intended for file download and open-ended use. Over the last few years, another big change in our working lives has come from the development of on-line information services (ADS, NED, Simbad, astro-ph, LEDAS, HEASARC). The latest phase, still ongoing, is the collectivisation of data collection - i.e. large pre-planned projects producing coherent databases undertaken by organised consortia (MACHO, 2dF-z, SDSS, VISTA). The classic astronomer's individualist approach will re-assert itself as the databases are opened up and trawled.

This paper concerns itself with what we see as the next steps in this direction. The first step is for key databases, and eventually all databases, to become *interoperable*, i.e. to become seamlessly jointly browsable and queryable in a simple way. The second step is the communal development of database analysis and exploration tools, to enable easy *data mining*. The third step is the development of *information discovery tools*. The overarching idea is that data, computing power, and services are distributed but that the astronomer uses a single simple interface - she plugs into the information grid. We shall explain these ideas in a little more detail in later sections. These developments will make it easier for astronomers to do science, and will enable exciting new kinds of science. Furthermore, as we shall see, the explosion of data volume that we are facing means that developments of this kind are needed anyway if we are not to grind to a halt.

(2) GROWTH OF DATA AND ARCHIVES

Today a substantial fraction of astronomical research papers are based wholely or in part on archive data. The archive is a second sky. A recent ESA paper reports that 10% of the ISO archive is downloaded every month, and that the volume equivalent of the whole archive has been downloaded twice. For the HST archive the volume of archive downloads is growing faster than the accumulated data (Fig. 1), and is now at around 20 GB/day. The HST archive is growing at 1-2 TB/yr and currently holds ~ 7 TB. Likewise the SOHO archive is growing at 1TB/yr. This is also fairly typical of major ground-based observatories around the world. (Recall that Giga=10^9, Tera=10^{12}, Peta=10^{15} ; also note that B=bytes and b=bits). This data growth is largely a consequence of the growth of detector size, which has been exponential for three decades (Fig. 2). Quite soon now VISTA will be using a Gpixel array.

A benchmark data volume is that needed to store a 16 bit image of the whole sky with 0.1" pixels - 100 TB. The raw data may be substantially larger depending on the observing technique, but will rarely be needed by users. On the other hand, the normal "science database" may be an order of magnitude smaller. Astronomical object catalogues derived from current surveys are typically ten times smaller than the pixel data. The pixel image itself can also be compressed by about an order of magnitude without losing much information using a variable blocking technique such as the H-compress algorithm developed at StScI. So a simple survey may require PB-scale shelved-tape storage (raw data for occasional re-reduction), 10TB on-line storage (compressed pixels and objects), and 100TB near-line storage (full pixel data). For multi-passband surveys, large-scale monitoring programmes, and spectroscopic surveys, or sampled in-situ data (as for STP satellites) the volume increases correspondingly.

(3) THE NEW ERA OF BIG SURVEYS

Sky surveys have been of great historical importance, and a particular UK strength (3C, Ariel-V, IRAS, UK Schmidt/COSMOS/APM). Several US surveys are currently attracting attention as representing the new age of digital surveys with on-line archives - SDSS, 2MASS, FIRST, and NVSS, each of which are of multi-TB sizes. The real data explosion will come however with two UK projects which will be the premier survey facilities on the world scene during the coming decade - the UKIRT Wide Field Camera (WFCAM), and VISTA, each of which will accumulate PB size archives. Finally the most ambitious (as yet unfunded) project is the US Large Synoptic Survey Telescope (LSST) which aims to map most of the sky every night ! Here we provide a few details on some current UK projects, followed by plans for WFCAM and VISTA.

The SuperCOSMOS and APM Programmes

The SuperCOSMOS Sky Surveys programme (see www-wfau.roe.ac.uk/sss/) is providing on-line access to digitised scans of photographic plates covering the whole southern hemisphere, in three bands (BRI) and with two epochs at R. Currently 5000 sq. deg. of the Southern Galactic Cap are available, from which users can extract pixel data for regions up to 15x15 arcmin and/or user-designed object catalogues covering up to 100 sq. deg., while the eventual southern sky survey database will be ~2 TB in size. Multi-colour object catalogues from the APM high-latitude scans with 500 million objects are also available. The vast areal coverage and, in particular, the temporal baseline, of the survey data makes this a very valuable resource even in the era of CCD-based photometric surveys, but to make full use of this important legacy dataset within the context of the developing "Virtual Observatory" will require the database federation and data-mining tools that Astro-Grid would provide for the community.

The INT Wide Field Survey

The INT Wide Field Camera survey program uses 20% of UK and NL time on the 2.5m INT to carry out publicly accessible multi-colour and multi-epoch surveys. The survey was carefully designed by a consortium of astronomers to be appropriate for many different projects, and to be a long term resource. The raw and pipeline processed images have no proprietary period in the UK and NL. Broadly the goal is to cover 100 sq.deg of high galactic latitude sky in the u,g,r,i,z wavebands with multi-epoch data in two bands over 50 sq.deg. The survey is a factor of 5 deeper than the SDSS over the whole region .i.e. m(r)=24.5

(5sigma). Smaller regions of around 10 sq.deg. are being surveyed 3-4 magnitudes deeper. Infrared coverage of the same area of sky with CIRSI has been proposed. The survey is described, and initial data products are available at www.ast.cam.ac.uk/~wfcsur/.

The XMM Survey Science Centre

XMM-Newton, by virtue of its sensitivity and wide field of view, is producing the largest and deepest serendipitous X-ray source survey to date with a projected estimated content of 500,000 objects for an assumed 10 year mission lifetime. The XMM-Newton serendipitous source catalogue will clearly become a major resource for a wide range of astrophysical projects. The XMM-Newton Survey Science Centre (SSC), an international collaboration led by Leicester University, has the responsibility within the XMM-Newton project of producing and maintaining the XMM-Newton serendipitous source catalogue, as well as coordinating follow-up identification programmes aimed at delineating the nature of the faint X-ray source populations. See xmmssc-www.star.le.ac.uk/.

The SOHO data archive

The ESA/NASA Solar and Heliospheric Observatory (SOHO) spacecraft is the principal space-based research tool for the solar community world-wide. Official ESA data-archives have been set up at the NASA Goddard Space Flight Center near Washington, at the Institut d'Astrophysique et Spatial near Paris, and at the Rutherford Appleton Laboratory. The RAL archive contains all SOHO data, which can be accessed via a dedicated Web site. SOHO has been in operation for almost 5 years and operations are anticipated to continue to at least 2003. The mission produces about 1 terabyte per year. The RAL archive also includes all data from the associated NASA TRACE (Transition Region and Coronal Explorer). UK investment in SOHO has been very significant, with one UK-led instrument, another instrument with a significant UK hardware contribution, and scientific involvement in most of the 12 instrument payload. Many UK research groups regularly use SOHO and this has maintained a very strong position in world-wide solar physics for the UK. Data from the mission are regularly used by a dozen UK research groups spread around the UK. Along with SOHO, TRACE, and Yohkoh, the next few years will see more high profile solar missions with strong UK involvements (Solar B, Stereo, Solar Orbiter) with a wide variety of instruments and data types - there is an urgent need to be able to explore all these data coherently.

The Cluster data archive

The Cluster Joint Science Operations Centre (JSOC) and Co-ordinated Data Handling Facility (CDHF) for Solar-Terrestrial Physics data, including Cluster, both located at RAL, provide a natural location for a full Cluster archive. As yet, ESA has not selected official archive sites but, as with SOHO, we would anticipate an RAL facility. This would serve the strong UK Cluster community which includes a number of university groups, such as Imperial College, MSSL, QMW and Sheffield as well as RAL. This would not be a large archive, compared to SOHO, for example, producing about 1 terabyte for the entire mission. However, exciting prospects come from the prospect of being able to search and examine space-based in situ particle and wave data that are contemporaneous with ground based radar and lidar data, or shortly after coronal mass ejection events studied by solar observatories with imaging and spectroscopy, and in the future even heliospheric in-situ data. This involves many different datasets with a rich variety of data formats. This kind of work is attempted now but is extremely difficult in practice, and so presents a classic opportunity for Astro-grid.

UKIRT WFCAM

By early 2003, UKIRT will have a new wide-field camera (WFCAM: see www.roe.ac.uk/atc/projects/wfcam/), which will be the most capable JHK survey instrument in the world. Through a large allocation of UKIRT time (30-50% of the total, for several years), WFCAM will undertake a combination of large, public surveys and private programmes allocated via PATT, which will together provide the UK community with an archive of near-infrared photometric data unrivalled in the northern sky. Particularly exciting is the plan to image ~3000 sq. deg. of the Sloan Digital Sky Survey region to matching depths in JHK, yielding an 8-band UV/optical/near-infrared photometric database unique in the north, and only surpassed in the south by another UK project, VISTA: federation of the SDSS and WFCAM archives is a key science goal for Astro-Grid, and a concrete deliverable on the path to development of a global "Virtual Observatory".

VISTA

One of the most exciting prospects for ground-based astronomy in the UK in the coming decade is VISTA (see www.vista.ac.uk), a 4m optical/near-IR survey telescope funded by £24.8M from JIF and to be located in Chile. VISTA will dominate post-Sloan survey astronomy, thanks not only to its technical capabilities, but also to its sophisticated observational strategy: rather than produce a monolithic sky atlas to uniform depth, as traditional in wide-field astronomy, it will undertake a set of surveys of varying depth and areal coverage, as required to meet the needs of a series of key programmes, expressing the scientific goals of the 18 UK universities in the VISTA consortium. VISTA will produce ~1TB of raw data each night, yielding a vast (~300TB) science archive after ten years of operation: the VISTA database represents the sternest challenge for UK wide field astronomy, but, with the tools deliverable by Astro-Grid, it can provide the UK community with a world-beating resource, that will drive research for several decades.

(4) NEXT STEPS IN USE OF ARCHIVES

Current use of archives can be broken into three types. First, and most common, is *support of other observations*, i.e. examination of small images ("finding charts") and other datasets, for use as illustration, as exploration, and as assistance in other research programmes. Second is *directed research*, the download of specific observation datasets for off-line analysis by the user. Third is *discovery based programmes* - the "power-user", running statistical analyses of huge amounts of data, re-reducing the data in novel ways, searching for rare objects and so on. The aims of Astro-Grid are to enhance all three classes of use, but we have the power user particularly in mind, as we believe an increasing fraction of astronomers will want to become power users. To date, this sort of work has been the restricted domain of specialists, requiring privileged access to the datasets, detailed technical knowledge, and months of work. The advent of the new generation of databases will motivate many more astronomers to become power users, as new science becomes possible, through trawling the multidimensional data spaces made available by database federation. The aim is to make this kind of science much faster and easier and as standard as reducing your spectrum with FIGARO or IRAF.

The idea of *interoperability* is that one can browse and query multiple databases from a single point of contact. For example given a region of sky, one could log-on to Astro-grid and view the Sloan, WFCAM, and XMM-Newton images simultaneously. Then one could click on an interesting object and download the X-ray spectrum. Or one could issue joint queries, such as "give me a flux-ordered list of all the XMM-Newton sources in this RA and Dec range with optical colours redder than X that have no radio counterpart". Such things can be done now with the small current archives, and with considerable effort. The aim is to log on to

Astro-grid and undertake such explorations automatically without having to know where the databases are, or having to learn six different packages, and for all this to work fast enough to be interactive. More ambitiously, one would also want a more sophisticated level of *visualisation and exploration* - very large image plotting and scrolling (possibly on power-walls), interactive subset selection and plotting, projections of multi-dimensional datasets, 3D examination with CAVEs or VR headsets and so on. Next we want large data-set *manipulation tools* - Fourier transforms, kd-trees, Gaussian mixture models, fuzzy joins, finding of outliers and so on. We will never invent enough analysis tools for the most imaginative astronomers, and so we also need a method for *uploading user code* to run on the databases. Finally and most ambitiously we need *information discovery tools*. After finding that interesting object in the federated XMM/Sloan/WFCAM databases, one would want to send off a search agent which returns with a report that this is probably the same object observed at Kitt Peak last week. Or with a suggestion that a recent paper on astro-ph seems to be about objects with very similar properties. The route to such information discovery could either be through intelligent search agents of some kind - a challenging technical problem - or through a global "Virtual Observatory" club - a challenging sociological problem !

(5) LARGE DATABASE SCIENCE

At the very least, the kind of advances we have sketched above will greatly improve the general infrastructure of research, making science easier and more efficient (and indeed, possible at all, in the face of the huge expected data volumes). The most striking difference, however, will be that the ambitious data-mining projects become the norm, rather than the province of the occasional stubborn specialists. One can already feel the pressure as astronomers realise the value of this style of working. It is difficult to guess what may be thought of in the future - and, indeed, one of our key arguments is that these advances open up unexpected avenues of research - but some inferences may be made on the basis of exciting examples from current or planned work:

• *Rare object searches* involve trawling through large databases, looking for the (quite literally) one in a million object. Cool white dwarfs have been found in the SuperCOSMOS database, constraining the nature of Galactic dark matter (Fig. 3); z=4 quasars have been found in the APM data, and a z=5.8 quasar by a combined SDSS(optical)/2MASS(IR) search (see Fig 4). Adding UKIRT WFCAM data, or turning to VISTA, will produce quasars at z=6 or 7. Brown dwarfs have been found from joint analysis of SDSS/2MASS data (Fig.5). VISTA will probe further down the luminosity function. Astronomers will also look for the missing dwarf galaxies, Trans-Neptunian objects, and Near Earth asteroids.

• Huge *statistical manipulations* are required to deduce the structure and history of the Galaxy from stellar motions, or to determine cosmological parameters from the anisotropy of the microwave sky. The use of techniques from computational geometry to structure databases in sophisticated ways can make such computations much easier - improvements by orders of magnitude are reported in the time taken to compute the N-point correlation functions of galaxy catalogues when they are indexed using k-d trees, for example.

• *Population modelling* could yield galaxy cluster histories from ugrizJHK diagrams, with the same descriptive power for galaxy evolution as H-R diagrams of stellar clusters had for the evolution of stars, or detect ancient galaxy mergers from groupings in stellar phase space (tidal streamers produced by the Galaxy swallowing the Sagittarius dwarf have been found by Irwin et al using APM).

• *Massive photometry*, monitoring thousands or millions of objects simultaneously can yield exciting advances in: micro-lensing and dark matter; high-redshift supernovae and the cosmological constant; quasar variability and the growth of black holes; parallaxes and the solar neighbourhood problem; planet searches; and stellar seismology and the age problem.

• Finally, there is the lure of the *unknown* - what astrophysical phenomena lurk in hitherto unexplored corners of parameter space that can only be reached through the federation of large and disparate datasets?

(6) STORAGE and DATA MANAGEMENT REQUIREMENTS

As described earlier, data volumes are growing alarmingly, especially for the major surveys. Multi-TB databases are becoming normal now, and PB databases (WFCAM and VISTA) will be with us in a few years. Technology is keeping pace, but requires moderately specialised equipment. For example the SuperCOSMOS object catalogue and compressed pixels (2TB) are stored on-line on a highly reliable RAID array costing about £60K. (The raw data (15TB) are accumulating on tape but will probably migrate to DVD jukebox soon). We have made an estimate of the UK science data volumes (from SuperCOSMOS, AAO, UKIRT, ING, XMM, Gemini, SOHO, and Cluster) accumulating in the next few financial years and find +20TB in 01/02, +30TB in 02/03, and +35TB in 03/04. The ramp up comes from UKIRT WFCAM starting in early 2003. So by April 2004 we need to deal with 85TB. VISTA will start in 2005 and by the time it has been running for three years in 2008 it may have accumulated a full database of PB size and a science database of around 100TB. Over the last few years the world storage capacity (and by implication a characteristic system size) has doubled about every 11 months, so that by 2004 we can seriously anticipate on-line storage systems at the 50-100TB level, and hopefully some kind of near-line storage (tape robots, or FMDs) will make the full data available within minutes. (Our future technology forecasts come from a paper given by Jim Gray of Microsoft Research at the recent Caltech NVO conference : see http://research.Microsoft.com/~Gray)

There is then no gross technical problem, but we need to make sure we have the financial provision to keep pace with our data. As today, mass storage equipment will be moderately specialised, and so available only at a few specialist centres. More importantly, although the storage technology is cheap, the associated staff effort is not. The management of such mass storage systems is far from trivial, and maintenance of the databases will require specialised staff, including for example migrating the data to new storage technologies every few years. Because the storage management is intimately linked with the problems of understanding, calibrating and documenting the data, it is however not appropriate to store the data at a national "computer centre" which provides such services in abstract - the data should stay near the experts. This is one of several factors that leads us towards the idea of building a grid around *expert data centres*.

(7) DATA INTENSIVE COMPUTING

If one requests an image or catalogue of a random small piece of sky from the SuperCOSMOS database it will be available in seconds. If one searches through all the objects over 5000 square degrees looking for those that pass a particular parameter filter, this will take about 2 hours. Such a search is limited by the I/O bandwidth of a SCSI disk, currently about 10MB/s. In recent history, although CPU power and mass storage have increased at the rate of 100x/decade, device bandwidths have increased at only 10x/decade, and current technology forecasts expect similar trends. Some kinds of problems are limited by the "seek time" from a disk, and this has improved by only a factor of two over the last decade (from about 10msec to 5msec). Extrapolating to a 100TB science database, searching this even at 100MB/s would take 12 days - unacceptable as a general service. Considerable improvement for many tasks can be made by intelligent structuring and indexing of the database, and caching of recent search results. Such techniques are being actively pursued by the SDSS team now. However, there will always be some queries that require streaming all the data. The only way around this is *parallelism* - to partition the data and set many CPUs/discs off searching simultaneously. In other words, we need supercomputer database engines.

Fortunately, many simple queries are quite strictly parallelisable, so we don't need proper Cray-type supercomputers, but rather massive "PC farms" with commodity components and relatively simple message-passing software. Currently the most popular version is the "Beowulf" cluster, costing around £2K/node. To achieve reasonable turn-round on a 100TB database we will need such machines with >100 nodes however. As with mass storage, we expect the technology to be in place to solve our problems, but the equipment will be expensive and specialised. We do not expect that everybody has a Beowulf on their desk, but rather that everybody has access to one provided by a specialist centre. And again, the real issue is the staff effort involved in developing and maintaining such specialised services. Finally, we should note that although PC farms look very appealing for simple queries, other kinds of analysis (Gaussian mixture models, Fourier analysis, model fitting) require much more message passing and so may run optimally on a different (more expensive !) type of machine, such as a shared memory architecture, or a vector machine. This will probably require collaboration and timesharing with a supercomputer centre. One of the most urgent tasks for Astro-grid will be to predict and define the types of problems users will want to solve in order to design the best architecture for database engines.

(8) REMOTE ANALYSIS SERVICES

A key aim of Astro-grid is to design new data analysis and data mining techniques, and to implement them in applications packages that will be developed to the degree of completeness, robustness, ease of use, and documentation that is currently considered the norm in instrumentation, and in data reduction software. This is where we expect the greatest scientific benefit to flow, where a large part of the cost will fall, and where the greatest degree of community involvement is likely to be. However this will not just be a matter of constructing and delivering applications packages which the user installs on their desktop PC. Rather, it will require a commitment to services provided by the expert data centres. This is because of the problem of network bandwidth.

Prediction of future network capacity is considerably harder than for CPU or storage capacity. There is no technical problem. In principle a single fibre can carry up to 25Tb/s. Network speeds in practice are determined by economic factors, and Wide Area Network (WAN) prices have changed little over the last decade. Local Area Networks (LANs) have improved, and most departments have somewhere in the range 10 to 100 Mb/s. However the characteristic real end-to-end speed if one is for example downloading a file from a web page is nothing like this - one will be very lucky to get 1Mb/s. This is just good enough for downloading single observation files but copying large databases is impossible, even if one had the space to store them. The future price/performance ratio for networks is anybody's guess, but an optimistic view is that when the public start to demand downloadable movies on demand then the prices will at last come tumbling down. With luck then in a few years we might hope for end-to-end performance of 1 Gb/s. (This is the number being assumed by the CERN Grid plans). This would still make it impractical to download large databases.

The logic is clear - that large database exploration and data analysis calculations need to be performed on remote computers attached to the databases (in parallel). The motto is *shift* *the results not the data*. The necessity of offering a remote analysis service forces us into well designed and robust software tools, an aim which is also consistent with our belief that scientific gold lies in systematic large database analysis, and that we can liberate many more scientists to undertake this sort of work. As well as providing click-and-play packages, we will need to provide a toolkit for users to develop their own programs, as well as the facility to upload and run such code, including the monitoring and allocation of CPU time. Along with the constraints of storage, I/O speed, and database management, this is yet another factor leading us to the idea of expert data centres. Such centres will need to ingest, calibrate, and manage the databases, and provide services in data subset access, database queries, on-line analysis tools, and remote visualisation.

(9) THE GRID CONCEPT

We have argued that services need to be *remote*, but they also need to be *distributed* because the expertise is distributed around our community. In the last few years the idea of *Grids* has emerged in the computer science world, where many computers collaborate transparently on a problem. The original idea was one of a computational grid. Jobs or portions of a job may be allocated to machines over the network, choosing the optimum architecture for particular calculations, filling spare capacity in CPU cycles, or simply aggregating the CPU power of many machines at times of peak demand. The term "grid" is used by analogy with the electrical power grid - a user can simply plug into the grid and obtain computing power, without needing to know where the power station is or how it works. All then user needs is an appliance which uses the power from the grid. Prototype grids are in operation now, but the technique is developing and growing in importance. Networked users will have supercomputer power at their fingertips.

Such computational grids will be relevant to theoretical astronomy, but the data intensive problems we are concerned with here present different problems, being I/O bound rather than CPU bound, and being driven by average load rather than peak load. However the concept leads naturally on to the idea of a *service grid*, and ultimately to a *knowledge grid*. One can imagine an Astro-grid gateway of some kind (not necessarily a web page) where once a user is logged on, a variety of databases will be browsable, searchable, and manipulable, actually running on one of several database engines and analysis machines, but without the user needing to know where or how. Whether a typical job is simply remote or actually distributed in a computational-grid-like fashion will depend on the kind of job and on future technological and economical developments, but the surface concept remains the same.

(10) GRID TECHNOLOGY

Our themes are common ones in modern science (and commerce), requiring mass storage management, remote use of high-throughput database engines, distributed systems, and globalisation of data and meta-data standards. Some of the work we need to undertake is quite specific to astronomy. Some of it is fairly typical of several science areas, but is high level applications work that we would expect to undertake ourselves, with the possibility of our solutions migrating outwards. Much of the work required however is fundamental computing technology. We will not describe this work in detail here, but simply append a list of some of the key issues to give a flavour of some of the problems. Most of these technical issues will not be solved by Astro-grid, but neither will they soon be commercial off-the-shelf solutions. Rather, they will be addressed by computer science research, by the largest scientific-grid projects (especially the CERN/US LHC grid project), and by commercial concerns with similar data-mining problems. (The UPS database is already 17TB). Our task will be to interact with such larger grid-projects and adopt appropriate solutions, but also potentially to drive them as an interesting application area.

Technical issues :

- data format standards
- metadata and annotation standards
- information exchange protocols
- presentation service standards
- request translation middleware
- workload scheduling, resource allocation
- mass storage management
- computing fabric management
- differentiated service network technology
- distributed data management - caching, file replication, file migration
- visualisation technology and algorithms
- data discovery methods
- search agents and AI
- database structure and query methods
- data mining algorithms
- s/w libraries and tools for upload requests
- data quality assurance (levels of club membership ?)

(11) CONTEXT OF THIS PROPOSAL

POLITICAL CONTEXT

Ideas of the kind we have been describing here have been brewing for some while, and were for example put forward in the context of the PPARC Long Term Science Review (LTSR). Indeed large database science IT infrastructure was highlighted in the Astronomy Panel LTSR and placed in the priority-1 list. However, the opportunity for extra funding to PPARC to actually achieve this has come about as part of the Government's Comprehensive Spending Review. The White Paper has recommended considerable new investment in "e-science". A cross-research council committee (on which AL sat) produced a paper for the OST recommending a large e-science programme. PPARC has included elements of e-science for both particle physics and astronomy in its upward bid to the OST, and we await the result of PPARC's actual allocation. Of course the big-ticket items in all this are the LHC grid, and the genome, as well as generic computer science funding. However, the current consortium has actively promoted astronomical grid work both within PPARC and above, and moreover, grew out of the realisation that, to stand any chance of astronomy-tagged funding from Grid monies, there was a need for a focussed project, led by the key data centres, and producing concrete deliverables targetted to benefit the wider community, rather than a vague "us too!" expression of a desire for a loose bucket of money for unspecified use in astronomy.

GRID CONTEXT

Astro-grid needs to interact closely with other grid projects, as much of the technology will be generic as discussed above. The most important one, in that it is most developed, and is driving much of the technology to date, is the LHC grid project. Within both the US and Europe there are almost-funded projects which are loose umbrellas for physical sciences (GriPhyn and DataGrid respectively). However in structural terms, much of our problem-set may be closer to the needs of Earth Science and even Biology, so we need to pay close attention to what is happening in those areas. Finally of course we need active collaboration with computer scientists (both academic computer science, and the applied computer science of the major supercomputer centres) and with the commercial IT interests.

VIRTUAL OBSERVATORY CONTEXT

The concepts and aspirations we have been discussing here have of course been mirrored by developments elsewhere. Many of the most impressive developments to date have come from the US (eg NASA Sky View and the ADS) and there is a proto-project called the "National Virtual Observatory" (NVO), which is on the decadal survey shopping list. So far it has only study money, but it is hoping soon to achieve funding at the level of $60M over 10 years. In Europe there is an EU-funded project called "AstroVirtel" but it is really just a funded archival research project. Likewise OptiCon has funding only for working group travel and workshops, but it is expected to form the basis of a major application to EU funds in the near future. Recent conferences in Caltech and Munich had wide international attendance and there is a growing feeling that in truth we must be aiming at international solutions - an IVO. However there is no attempt so far to construct an organised global project, but rather to keep independent projects in close collaboration. Within the UK Astro-grid consortium, we feel it would not be sensible to attempt a kind of global takeover - rather, we should be looking at focussed work that will be a clear contribution to such global contributions. This goes along with a slight worry that the US and European plans might be a little too ambitious and top down. This is why our top priority is to achieve the federation of a handful of key datasets, and the development of some key applications packages.

UK COMMUNITY CONTEXT

Our picture is one of services provided by expert data centres, with the majority of users elsewhere using the system and doing science. The idea is analogous to the building of an instrument like the EPIC camera on XMM, or SPIRE on FIRST. A handful of organisations actually build the beast, but it is built for everybody to use. However the wider community is the expert on what it wants and needs, and also holds substantial expertise in data analysis. This is even more important than usual, as the likely pattern of usage will have a crucial influence on the solutions adopted and indeed the whole philosophy. There have been open international conferences on this topic, and so far one open workshop in the UK. The community was formally consulted during the LTSR process, resulting in a high priority for database developments. In recent months due to pressure of time in the CSR process, the discussion and consultation has inevitably been more erratic. This paper to Astronomy Committee is part of a more formal consultation of course, but we also intend a wider process, with a circulated position paper for comment, and possibly a well-publicised workshop in January. Our "Workpackage One" is an intense study of current and future required functionality.

(12) THE UK ASTRO-GRID PROPOSAL

We have been led to a picture of expert data centres, and indeed the UK niche over the next few years is in key datasets - SuperCOSMOS, the INT Wide Field Survey, UKIRT WFCAM, VISTA, the XMM SSC, SOHO, and Cluster. But we also have a tradition in communal development of astronomical software. For this reason, it is not surprising that the consortium that has arisen to propose Astro-grid is based on the groups with primary responsibility for those major datasets and/or experience in organised astronomical software provision. In addition we have a strong push on data mining algorithms through the consortium membership of the Computer Science Department of Queen's Belfast. We also expect to work very closely with both the Edinburgh Parallel Computing Center (EPCC) and the new RAL e-science centre, both of whom will be undertaking generic grid technology work and will link us to other grid projects. However those organisations are not requesting any funding through Astro-grid, and we feel it is important to keep science goals rather than computing infra-structure at the forefront of our project. For that reason we have described them as "partners" rather than "members".

Appendix-1 lists the expected workpackages. Descriptions of the contents of these workpackages are evolving rapidly, but details are available for those interested. The work we anticipate breaks down into four main areas. Firstly,an intensive study phase, studying current and required functionality, grid-enabling current packages, and running benchmark studies versus system architectures and database systems, software toolkits and so on. Obviously the detail of what we plan in later stages is provisional until we complete these studies. The second area is applications development, in browsers, visualisation, data mining algorithms, and search agents. The third area is the federation of specific datasets with the aim of making actual working grid systems. This work needs to start very quickly as we require negotiation with international partners, and solutions on data standards and DBMS compatibility and so on will start to freeze very soon. The fourth area is provision of the necessary hardware to undertake Astro-grid. This does not include the improved network bandwidth that we hope for - this will be bid for through the SRIF process. What it does mean is (a) multi-TB storage facilities and (b) parallel database engines.

Our aim here is an approximate costing for a three year programme, to match the bid against the current spending review. This also covers the start of UKIRT WFCAM operations, and is just before VISTA comes on-stream. Clearly some of the work packages are more ambitious than others and are unlikely to be completed in three years (notably in visualisation and intelligent search agents) but others are achievable on this time frame, and all the workpackages can make significant progress.

Mass storage costs are estimated using the data volume estimates of section 6, and assuming that the cost is 30K£/TB in 01/02, 20K£/TB in 02/03, and 15K£/TB in 03/04. This is for high-reliability SCSI hardware RAID including tape backup etc. (These are much the same figures assumed by the LHC Grid project). For Beowulf or similar PC farm systems we assume 2K£/node in 01/02 and 1K£/node in 03/04. In 01/02 we will buy experimental modest sized (16 node) PC farm systems to use on SuperCOSMOS and XMM SSC. In 03/04 we would purchase a 200 node system to become the WFCAM database engine. In total the special purpose hardware costs within three years are estimated at £2.0M.

Staff effort will require a mixture of PDRA scientists, programmers, system engineers, and out-sourced programming effort. We have used an average cost per staff year of £60K, similar to the RAL sy cost, and equivalent to a senior researcher on 30K plus employer cost and overhead plus a reasonable annual travel cost (which will be absolutely vital in this work). In practice most PDRA scientists will be much cheaper, but we will also want to pay more for programmers, especially when out-sourcing. WPs 15 and 16 have only travel costs. Project management and admin assistance requires three full time staff for such a fragmented and complex project. The other packages require 1-2 staff each averaged over the project lifetime. (Actually effort will be quite heavily profiled). We therefore estimate 59 staff years total, giving a staff + travel cost of 3.5M.

The total estimated cost in three years is therefore 5.5M. This should be profiled towards the later years - probably 0.8M, 1.5M, 3.2M.

It is not intended that Astro-grid be an ongoing infra-structural item. However, it is clear (a) that the hardware provided will cover only the first year of WFCAM data, so more hardware investment will certainly be needed for VISTA and WFCAM over the following years, and (b) the more ambitious applications development will almost certainly take more than five years. Probably the best approach to take is initially to fund only the three year programme, and at some review point to request a proposal for "Astro-grid II" for a second fixed term.

DETAILS OF PARTICIPATING ORGANISATIONS

Lead investigators

Institute for Astronomy, University of Edinburgh	Andy Lawrence
Institute of Astronomy, University of Cambridge	Richard McMahon
Department of Physics and Astronomy, University of Leicester	Mike Watson
Space Data Division, RAL	Peter Allan
School of Computer Science, Queens University of Belfast	Fionn Murtagh

Other key staff

Mike Irwin (Cambridge CASU), Clive Page (Leicester XRA), Bob Mann (Edinburgh WFAU), Mike Lockwood (RAL Astrophysics), John Sherman (RAL Starlink) , David Giaretta (RAL Starlink), Nigel Hambly (Edinburgh WFAU), Clive Davenhall (Edinburgh Starlink), Jim Lewis (Cambridge CASU), Julian Osborne (Leicester XMM-SSC).

Partners

Edinburgh Parallel Computing Centre	A.Trew, R.D.Kenway
e-science centre, RAL	P.Jeffreys

Key working links

Sloan Digital Sky Survey Database team
US National Virtual Observatory Project
Opticon Consortium
Chandra Science Centre
CDS Strasbourg

Astro-grid : WORK PACKAGE SUMMARY

WP1 PROJECT MANAGEMENT

WP2 REQUIREMENTS ANALYSIS : existing functionality and future requirements; community consultation

WP3 SYSTEM ARCHITECTURES: benchmark and implement

WP4 GRID-ENABLE CURRENT PACKAGES : implement and test performance

WP5 DATABASE SYSTEMS : requirements analysis and implementation; scalable federation tools.

WP6 DATA MINING ALGORITHMS : requirements analysis, development and implementation

WP7 BROWSER APPLICATIONS : requirements analysis and software development

WP8 VISUALISATION : concepts and requirements analysis, software development.

WP9 INFORMATION DISCOVERY : concepts and requirements analysis, software development

WP10 FEDERATION OF KEY CURRENT DATASETS : e.g.. SuperCOSMOS, INT-WFS, 2MASS, FIRST, 2dF

WP11 FEDERATION OF NEXT GENERATION OPTICAL-IR DATASETS : esp. Sloan, WFCAM

WP12 FEDERATION of HIGH ENERGY ASTROPHYSICS DATASETS : esp. Chandra, XMM

WP13 FEDERATION of SPACE PLASMA and SOLAR DATASETS : esp. SOHO, Cluster, IMAGE

WP14 COLLABORATIVE DEVELOPMENT OF VISTA, VST, and TERAPIX PIPELINES

WP15 COLLABORATION PROGRAMME WITH INTERNATIONAL PARTNERS

WP16 COLLABORATION PROGRAMME WITH OTHER DISCIPLINES