Astronomical Computing in the 1990s: The Bahcall Committee Report

Larry L. Smarr and William H. Press

A major decennial report highlights the importance of computation in astronomy and suggests innovative new directions

Once every ten years in astronomy, as in physics and some other fields, a committee chartered by the National Academy of Sciences and the National Research Council is convened to survey and prioritize the most important new programs of the coming decade. The work of the present Astronomy and Astrophysics Survey Committee (AASC), chaired by John N. Bahcall with the 1990s as its purview, was recently completed with the release of its report, "The Decade of Discovery." Compared with the reports of previous decades (most recently the Field Committee report, "Astronomy and Astrophysics for the 1980s"), computers and computer-based science play an unprecedentedly large role in the Bahcall Committee's report. Astronomy, both observational and theoretical, is on the verge of becoming an all-digital science.

In this article, we review some of the findings and recommendations of the Bahcall Committee in the area of computing. The Committee's deliberations on this subject were in large measure guided by the separate report of a subsidiary Panel on Astronomical Computing (see sidebar, *Membership, Panel on Astronomical Computing*), which is also now published (and on which much of this article is based). In what follows, we will distinguish between recommendations of the Committee, which thus represent agreement of the full astronomical community, and recommendations of the Panel, which represent consensus by a representative computer-oriented subset of astronomers and astrophysicists.

Larry L. Smarr is Professor of Astronomy and Physics at the University of Illinois at Urbana-Champaign and is the Director of NCSA. William H. Press is Professor of Astronomy and Physics at Harvard University.

Astronomy and Computing

Computational technologies have been central to advances in astronomy and astrophysics for at least the past four decades. Martin Schwarzchild's stellar evolution codes are said to have used roughly half the cycles of John von Neumann's MANIAC "supercomputer" during the early 1950s. The 1960s saw the first detailed supernova computations. The Einstein X-ray observatory and the VLA radio array, coming online in the 1970s, created images by the use of computers as intermediaries between the sensor and the observer. In the 1980s, theorists used supercomputers to model a wide variety of complex astrophysical phenomena.

The 1980s also saw three major improvements in researchers' access to computational resources. Firstly, the emergence of personal computers and workstations allowed many scientists to gain complete control of their personal computing environment. These complemented the departmental or observatory mainframes which arose in the 1970s. Secondly, supercomputers became available to the general research community with the creation of the NSF national supercomputer centers. Finally, the beginning of the national research network enabled researchers to connect their desktops with each other's or with supercomputer centers, national laboratories, or observatories.

In the 1990s, these three trends will be accelerated and more tightly integrated. A comprehensive national computing infrastructure will develop, which will support scientific and

engineering research in this country and, increasingly, internationally. This infrastructure will be inherently hierarchical. Desktop computers will become the standard working environment for both theoretical and observational research. Many departments and observatories will obtain relatively inexpensive minisupercomputers whose compute power by the mid-decade will exceed that of supercomputers in 1990. Telescopes will derive their power as much from specialized computers as from new sensors. Scientific visualization capabilities will be widely used. The national networks will contribute to increasing scientific collaboration.

In addition, distributed archives of observational and laboratory data will be accessed from the desktop over the network. The largest theoretical simulations and observational data sets will be handled at national centers with a wide variety of computer architectures, including supercomputers. The ability to bring together on the desktop the results of both complex simulations and detailed observations and to interact with each data set visually as well as quantitatively will profoundly influence the progress of the astronomical sciences.

Machines for the 1990s: Workstations to Supercomputers

Over the next decade, workstations will grow in performance to become comparable to present-day supercomputers. The present generation of RISC (Reduced Instruction Set Computers)-based machines is much more cost-effective in terms of dollars per million floating point operations per second (MFLOP) than the current generation of supercomputers. These systems allow for affordable computing locally. Such systems also make possible a close coupling of analytic, numerical and visual computing.

Essentially all active researchers need convenient access to such workstations, as well as "clear channel" coupling from their research office, or laboratory, into the national network. In addition, small colleges and universities that have a history of training the students who become the future generation of scientists should be encouraged in their efforts to offer undergraduates exposure to modern computing. Support for high-performance workstations and mini-supercomputers will be a cost-effective first step, by providing an accessible computational environment with modern software and graphics. It would also be a step toward geographical equality of resources and opportunity. Accepting this case, the Committee's strong recommendation to both NSF and NASA is that funds for individual workstations and departmental mini-supercomputers be made available, thus providing a distributed network for astronomical computing.

However, the advantages of supercomputers—large memory and disk capacity, vector or massively parallel processing, and extremely high input/output (I/O) rates—are crucial to an important subgroup of computer users with large or demanding codes. A small subset (perhaps 20%) of both the theoretical and observational computational community will attack problems whose computational requirements, in speed, memory or storage, exceed those that can be reasonably provided locally. Such users will need access to central national facilities. The national centers provide high-cost technologies which are available over the national network to the community. These high performance computing technologies are situated in an environment of ultra-high-speed networks, large disk caches, the latest visualization technologies, and teams of specialized experts for consulting and training. The Committee commended the use of NSF, NASA, and DOE supercomputer centers by astronomers, and explicitly stated the need for these centers to be upgraded, in order to remain at the forefront of technology.

Computers and Theoretical Astrophysics

Astrophysics depends on theory and modeling to a greater degree than other physical sciences, because astronomers can only observe remotely; active experimental intervention is not possible. Moreover, the observed phenomena are typically the result of complicated interactions among highly nonlinear processes occurring simultaneously. Thus, it is necessary to construct rather elaborate models to achieve a satisfactory interpretation of the observations.

There is a long tradition of using analytic simplified models to capture the essence of a complicated astrophysical phenomenon. Desktop computers are becoming increasingly important to support this work. Modern symbolic mathematics software allows the theorist to use more complex analytic formulations of the problem. Ordinary differential equations, which required a supercomputer to solve in the 1960s, are routinely evaluated and graphed by workstations and personal computers today. During the next decade, the power of desktop machines will become so great that many of the problems for which astronomers and astrophysicists are using supercomputers today, will also become soluble locally.

The panel estimated that nearly 10% of practicing astronomers are presently engaged in theoretical simulation of astrophysical phenomena. Much of this computational astrophysics is done using local workstations and mini-supercomputers. Of the total allocations of time for all areas of academic science and engineering on the NSF Supercomputer Centers facilities, roughly 10% of the resources are being used by researchers in the field of astronomy and astrophysics. This is equivalent to about three processors of a current generation supercomputer which would cost around $15 million to purchase.

As researchers build more and more complexity into their models, they outstrip the ability to compute locally with a reasonable turnaround time. This complexity arises for two fundamentally different reasons. Firstly, not only does the spatial dimensionality grow from one to two to three, but, as more and more realistic models are attempted, systems which are first studied as static become time-dependent. Typically, as the geometric complexity grows, so does the number of physical variables which must be solved for (e.g. from a radial velocity vector in spherical symmetry to all three components of the velocity vector in general).

Secondly, adding more physics to the problem increases both the number of equations and their coupling. For instance, one may add magnetic fields, nuclear or chemical

reactions, radiation transport, or viscosity, to an inviscid fluid flow code. In some cases, the introduction of new physics raises the effective dimensionality of the problem. For example, to describe the radiation flow in the most general case, one would add two angle variables and one frequency variable to the calculations. If there are scattering terms in the sources (Compton, Thomson, Rayleigh, etc.) the system to be solved is a seven-dimensional integro-partial differential equation. In addition, the calculation of realistic properties (opacities, equation of state) strain the resources of current machines almost to the breaking point. In short, both geometric complexity and additional physics can rapidly drive up total computational time and memory to values far exceeding the capacity of today's fastest machines.

The 1990s will be the decade in which a number of long-standing astrophysical problems will be solved, and computers will play an important role in these solutions. Areas which seem particularly ripe for rapid theoretical progress and comparison with observations, can loosely be categorized as follows: the large-scale structure of the universe and cosmology, active galaxies and jets, star formation and the interstellar medium, dynamics of stars and stellar atmospheres, supernovae, accretion onto compact objects, generation of gravitational radiation, and the microphysics and magnetohydrodynamics of astrophysical plasmas. Astronomy will see significant advances, particularly through modeling and numerical simulations approaching realistic complexity, which can be directly compared to observations.

As more powerful computers and community simulation codes become available, the incorporation of realistic physics and, where necessary, full three-dimensional time-dependent geometry, will greatly increase the ability of astrophysicists to directly compare their simulations with observational data. For example, during this decade we will see the addition of nonequilibrium physics to hydrodynamics codes. A few codes have already taken a step in this direction with the introduction of two or three temperature systems comprising, say, electrons, nuclei, and radiation. But none yet allow for nonequilibrium effects in the excitation and ionization distributions. When this is done, the radiation field and the state of the material become inextricably interwoven, making it impossible, even in principle, to calculate the thermodynamic properties of the material in terms of purely local variables. Rather, the system becomes fundamentally nonlocal, and we are forced to solve very large systems of globally interlocked equations, characterized by a wide range of typical spatial and time scales. These problems require the resources of massively parallel machines, and we should devote considerable efforts to algorithm development for such machines, forming an effective alliance with computer science experts working on such machines.

Computers and Observational Astronomy

Modern astronomy and new technology are closely coupled. New developments in technology have spawned breakthroughs in astronomy, and the promise of scientific discovery has often pushed technologies beyond their existing state-of-the-art. The charge-coupled device (CCD), new technology telescopes, active optics and computing technology are examples of areas currently rich in this synergism. Analog devices are being replaced in new instruments with digital devices based on digital signal processors with greater precision and stability.

The increases in sensitivity and resolution that computers will add to observations of the coming decade will be comparable in many cases to the increases due to new telescope and detector designs. (Of course, we cannot have one without the other!) Put differently, it will be the incorporation of high-performance computing that allows new observing facilities to achieve their full potential.

A good example is the use of image processing algorithms in radio and optical astronomy. Powerful deconvolution algorithms have been developed which can greatly enhance the power of both radio and optical imaging telescopes. The Very Large Array (VLA) can be seen as an evolving telescope, with today's version being an instrument orders of magnitude more powerful and flexible, due to improvements in computing algorithms and hardware speed, than the one which was originally designed—all without major telescope design modification.

"Astronomical" Data Rates

Future telescopes, whether based on the ground or in space, will have as an essential component a very high speed data link between the sensor and a computer, and sufficient power to handle the data processing requirements in near real time.

Space astrophysics, which is a part of the mandate of NASA, is entering a decade of rich observing opportunities. About 10 missions will be flown within the next five years, resulting in a large number of observations and large volumes of data. These data are digitally telemetered to the ground for processing and archiving.

The current technology in ground-based radio and optical detectors available on the world's largest telescopes is capable of producing truly enormous data rates—in excess of 10 GB per day—and this rate will grow higher by the end of the decade. We cannot push our current ground-based telescopes, which cost tens of millions of dollars each, to their full theoretical limits, because large-capacity mass storage devices are not currently available. In the near future, this situation will get rapidly worse, when the next generation of $100 million telescopes, using very large arrays of charge-coupled device (CCD) detectors, becomes operational.

In the early 1990s, most CCD detectors will produce two-dimensional images with sizes up to 2048×2048 pixels. A few of these instruments generate three-dimensional data sets, and can yield data transfer rates in excess of 16 Mb per second (16 Mbaud).

The data flow from a radio interferometer like the VLA, with 27 operational antennae or 351 independent complex interferometer pairs, in its full spectral line mode, can amount to about 8.6 GB per day. The current VLA computers cannot handle this data flow, so subsets of the data must always be selected for transmission and analysis. By comparison, the Hubble Space Telescope will generate

around 2 GB of archival data per day in initial operations, and this number may increase substantially with second generation instruments.

Intensive Data Processing

Increasingly, interposed between observation and actual understanding will be multiple stages of intensive data processing. Operation counts in the teraflop range (10^{12} floating point operations) per reduced data set will become more common. On the theoretical side, teraflop numerical simulations will be both necessary and practical in making the vital connections among astronomical observations, astrophysical theory and remote observations.

In the case of ground-based astronomy, the computer that is required to carry out astronomical data reduction depends on the type of instrumentation and its data-reduction protocol. Even when these instruments are used in their simplest mode, they will produce data flows of several GB per day, as discussed above. The calibration and subsequent processing needed for the majority of astronomical applications requires several GB of storage space, and can be done on workstations with typical speeds of a few Mflops, and about 16 MB of core memory.

Remote and Real-Time Observing

Today, there are a number of observatories which allow remote management of telescope projects. During the coming decade, we expect this trend to continue. High-speed networks will make transfer of large data sets feasible, and advanced storage technologies and computing capabilities will alleviate problems of data handling at the analysis sites. It is important that our national observatories, whose mission couples multiple remotely-based users to data-gathering telescope facilities (both ground-based and space-based), stay abreast of trends in this area of computing technology.

Radio astronomy will enter the mode of real-time observing during the coming decade. The present operational mode for radio telescopes is to acquire data and store them on magnetic tapes for processing at a later time. There is little opportunity for astronomers to see the results of their observations soon enough for the telescope to still be available for follow-up observations (which might be needed either because of poor data quality or due to an exciting, unexpected result). Particularly with a radio synthesis array telescope, the interval between data acquisition and working with the processed data ranges from weeks to years. Real-time radio astronomy would revolutionize the field by permitting an observer using a synthesis array to see an image of the radio sky as the observations were being made. Interactive observing could become a reality if the image processing could be done and the images transmitted fast enough.

In the context of space astronomy, where remote observing is the default, the existing communications and control infrastructure must be reexamined. It should be possible to simplify the operations concepts for astronomical observatories, allowing more distributed planning and scheduling capabilities, and direct data distribution to the observers. This should reduce the ultimate cost of missions. The Panel commended NASA for its current studies of the operations needs of space astronomy in the next decades, and hoped for the continuing involvement of the astrophysics community as it incorporates more advanced concepts.

A tightly-coupled system of telescope, network and computer will significantly raise the utilization of telescope systems and the productivity of astronomers, and if successful, will be a prototype for the future of all observational science.

Archiving

The Committee found compelling scientific reasons for archiving selected astronomical data from ground- and space-based telescopes. Firstly, astronomical processes occur on time scales longer than the lifetime of individual researchers. The history of astronomy contains many examples where archival material dating back decades, and sometimes centuries, has proven essential in solving modern problems. Astronomers have an obligation to preserve contemporary data in intelligible form for future generations of astronomers. Secondly, the large two-dimensional detectors of the 1990s will produce images of the sky taken at different wavelengths and at different times. It is widely recognized that the definitive interpretations of much space astrophysics data are not found in the first papers, but rather in more extensive archival studies.

Rapid advances in high-density data storage and network connectivity are reducing costs and facilitating access to large data archives. Access time and cost can now be traded off against each other, to create cost-effective archives well-suited to users. Both the Panel and the Committee recommended that a high priority for the coming decade should be the development of electronic archives for data obtained at ground-based observatories, and the implementation of an infrastructure to make these archives available to remote users over computer networks. (The NSF Subcommittee on Ground-Based Optical and IR Astronomy, of the Advisory Committee for Astronomical Sciences, has also highlighted digital archiving as one of their key recommendations.)

The Committee's recommendations specify that appropriate long-term goals are that all new major observatories, both public and private, should incorporate archiving in their design and in their standard operations; most public and private observatory surveys should be archived; observers should obtain their raw data, whenever possible, in the archive format. Going even farther, the Panel's recommendation was that *all* electronic data obtained from space-based and ground-based telescopes should be archived and catalogued using modern database management systems. Panel and Committee agreed that, in most cases, data should be available to the community after a sensible proprietary period, typically one to two years after completion of an observing program.

New archiving technologies are emerging which allow archiving of large image and catalog databases. It should be an assumption of the astrophysics data-archiving programs that the hardware of the archiving medium will continue to evolve as new technologies emerge. If the software architecture is chosen with this in mind, the archive hardware can be

replaced periodically without the necessity for modification of the entire archive structure.

The most widely-used and trusted astronomical data reside in catalogs which are critically evaluated. The collection and distribution tasks of catalog-making can best be done by national centers with appropriate expertise in the data and techniques. The critical analysis required for catalog generation should involve experts at the centers and elsewhere in the community. Critical analysis is most needed for the cataloging of ground-based data of widely varying format and quality.

The issue of international collaboration and coordination of archiving and networks should be addressed. Reciprocal agreements must be reached with other international agencies with regard to access to and support of data archives. Astronomy on an international scale will develop much faster than in the past due to network access. Also, a large percentage of the total useful astronomical database will be located outside the USA. NASA has already begun to address this question with missions such as GINGA, EXOSAT and ROSAT, but no standard solutions exist; rather, case-by-case solutions have been adopted. One goal should be to achieve a universal approach to archiving and data analysis. The Panel recommended that NSF should approach its foreign counterparts concerning ground-based archiving and data access.

Research using astronomical data archives can be very broadly based, both geographically and institutionally. A nationally-supported archiving program will allow researchers and students, and smaller colleges and universities to enter the mainstream of modern astronomical research in a way that is not now possible.

Community Code Development

In order to utilize effectively the enormous advances in computer hardware expected in the next decade, we must have an accompanying development of scientific software. This is actually more costly than the computers themselves, and should be of at least equal concern. Code development activities often require tens of man-years, followed by a sizable budget and group for their maintenance.

According to the Panel, the observational astronomy community has shown an admirable degree of leadership and coherence by developing software systems like AIPS (Astronomical Image Processing System) and IRAF (Image Reduction and Analysis Facility), and data standards such as FITS (Flexible Image Transport System), all of which have been adopted widely. These packages have saved an immense amount of time and duplicated effort. The Panel recommended that funding be specifically earmarked to maintain "critical mass" staff and computing support at the national observatories for this important community support role.

The FITS data format, now used internationally for the exchange and often the archiving of astronomical data, has been the IAU standard since 1982. FITS was developed in 1979 by NOAO (National Optical Astronomy Observatories), NRAO, and NFRA (Netherlands Foundation for Research in Astronomy), and is an openly published standard. Similarly, the development by the NRAO of the AIPS software system, by NOAO of the IRAF system, and by ESO of the MIDAS system, provides the community with a few very functional image-processing systems, which are portable to many computing platforms and distributed widely, with standards maintained by national centers. The advantage of such a coordinated approach is demonstrated by the willingness of the astronomy community to adopt these systems. STScI (Space Telescope Science Institute) adopted IRAF as the environment for the HST (Hubble Space Telescope) data analysis system (STSDAS), and SAO (Smithsonian Astrophysical Observatory) adopted a similar approach for the ROSAT system (PROS). Indeed, in the area of image-processing, it appears that astronomy has already taken a technical leadership position relative to other scientific disciplines.

Community simulation codes in other disciplines, such as aerodynamics, fusion research, and quantum chemistry, have proven to be extremely useful theoretical tools, and help to drive the development of the fields. There has been no support or encouragement for the development of analogous codes in theoretical astrophysics. A great deal of parallel development has been done, however, by astrophysicists working individually or in small research groups. Support should continue for fundamental algorithm development of general use to astronomers. In addition, the most useful of these algorithms should be incorporated into community codes which are made generally available. The Panel recommended that funding be provided for the long-term development, maintenance, documentation and user training of these codes.

Surveying and commending the role played by NSF in the development of community standards, and noting the importance of these standards to NASA missions, the Committee recommended that NASA increase its corresponding role in furthering the development of community software and software standards.

Recommendations Summarized

Here we summarize the major recommendations of the Astronomy and Astrophysics Survey Committee and its Panel on Astronomical Computing regarding computation.

The Committee recommended:

- Digital data archives, encompassing ground-based and space-based observations, should be established and made available over high-speed national networks.
- Individual workstations and departmental mini-supercomputers should be acquired, providing a distributed network for astronomical computing.
- NSF should encourage the development of high-speed national networks by funding links to widely-used observatories.
- NASA should increase its role in furthering the development of community software and standards.

The Committee commended the use of NSF, NASA, and DOE supercomputer centers by astronomers, and explicitly stated the need for these centers to be upgraded, in order to remain at the forefront of technology.

Regarding archiving, the Panel made additional recommendations, including:

- All major ground-based observator-

ies, both public and private, should incorporate the capability of archiving digital data.

- Archives should be designed to outlive any specific hardware, software, or medium.
- Archives should include all the raw data, calibration data and information necessary to remove instrumental signatures from the data.
- Policy should be established to require the archiving of data described in papers published in the refereed literature. This policy should be enforced and implemented through journals, time-allocation committees and the proposed reviewing process.

Regarding community code development, the Panel's recommendations were even more specific than the Committee's:

- NASA and the NSF should support the development of community codes by funding proposals selected by peer review in an AO/NRA process. NASA should expand the initial effort of the Astrophysics Software and Research AIDS NRA. There should also be support for the maintenance of codes found to be of general utility to the community.
- A community code program should fund both individuals who propose small programs, and larger groups, consisting of both computational astronomers and software professionals of high quality, who propose major efforts.
- A recommended level of effort could involve the funding of, say, three large groups (each with five to eight professionals) and a number of smaller groups.

Outside of astronomy, on the national scene, the strategic importance of high-performance computing to the future competitiveness of broad sectors of the U.S. economy is coming to be widely recognized. As a consequence, a major national focus (the High Performance Computing and Communications Program) is emerging. The proposed components of such a program include high-performance computing systems, advanced software technology and algorithms, a National Research and Education Network, and support of basic research and human resources. The program is to be driven by grand challenges in computational science and engineering.

Astronomy in particular stands poised, by virtue of its intrinsic data- and computation-intensive nature, its manageable size as a discipline, its past experience and future opportunity, to be the Grand Challenge application discipline in a number of major aspects of a national program. Astronomy's task is to build its own internal computer infrastructure, in such a way as to maximize its leverage *vis-à-vis* the national program—and simultaneously to bring to astronomers the computational technologies that will permit innovative astronomical discovery.

Membership, Panel on Astronomical Computing

Chair

Larry L. Smarr, University of Illinois, Urbana, IL.

Vice-Chair

William H. Press, Harvard University, Cambridge, MA

Members

David W. Arnett, University of Arizona, Tucson, AZ

Alastair G.W. Cameron, Harvard College Observatory, Cambridge, MA

Richard M. Crutcher, University of Illinois, Urbana, IL

David J. Helfand, Columbia University, New York, NY

Paul Horowitz, Harvard University, Cambridge, MA

Susan G. Kleinmann, University of Massachusetts, Amherst, MA

Jeffrey L. Linsky, University of Colorado, Boulder, CO

Barry F. Madore, California Institute of Technology, Pasadena, CA

Richard A. Matzner, University of Texas, Austin, TX

Claire E. Max, Lawrence Livermore National Laboratory, Livermore, CA

Dimitri Mihalas, University of Illinois, Urbana, IL

Richard A. Perley, National Radio Astronomy Observatory, Socorro, NM

Thomas A. Prince, California Institute of Technology, Pasadena, CA

Christopher T. Russell, University of California at Los Angeles, Los Angeles, CA

Ethan J. Schreier, Space Telescope Science Institute, Baltimore, MD

Stuart L. Shapiro, Cornell University, Ithaca, NY

John R. Stauffer, NASA Ames Research Center, Moffett Field, CA

Gerald Jay Sussman, MIT Cambridge, MA

Joseph H. Taylor, Princeton University, Princeton, NJ

Juri Toomre, University of Colorado, Boulder, CO

Anthony J. Tyson, AT&T Bell Laboratories, Murray Hill, NJ

Simon D.M. White, University of Arizona, Tucson, AZ

James R. Wilson, Lawrence Livermore National Laboratory, Livermore, CA

Karl-Heinz Winkler, Los Alamos National Laboratory, Los Alamos, NM

Stanford E. Woosley, University of California at Santa Cruz, Santa Cruz, CA

A 100 Million Pixel Imager and Automated Image Analysis Software

Image databases have been large (several Gb) in optical astronomy for years. The automated detection, classification and photometry package (FOCAS (Faint Object Classification and Analysis System), developed over the past ten years, has enabled statistical studies to take place involving large image databases. FOCAS is a collection of image analysis and automated pattern recognition programs designed for automated reduction and morphological classification of astronomical images. Recent FOCAS releases (now available as part of the IRAF package) include powerful image pre-filtering operators, along with interactive color-graphics display programs which allow the user to quickly identify objects with selected properties, such as color or two-dimensional shape.

Current silicon CCDs cover about 1% of the quality imaging area in the focal plane of large telescopes. It is now necessary and possible to construct a mosaic of CCDs which cover a larger focal plane area. Let us consider a 5×5 mosaic containing 25 CCD arrays, each 2048×2048 pixels. The peak raw data rate for the camera would be 100 Gb/night. It is almost possible, by current 1990 technology, to process, reduce, analyze and archive this imaging database in nearly real-time. A CCD mosaic imaging survey of a 100° patch of the sky would produce vast amounts of data which would have to be managed and processed. The resulting detected object rasters in several wavelength bands, and the FOCAS matched catalog, would comprise 600 Gb.

It is imperative that the design of large CCD imagers make adequate provision for easy and rapid data analysis. The magnitude of the data-processing tasks needed for this mosaic imager would require a special-purpose system. The characteristics of this system are dictated by the requirement for real-time image correction and automated analysis, but the same hardware would be capable of performing extensive image post-analysis. The design of the data system must emphasize computational power, fast data transfer paths, flexibility, and expandability. Almost as critical as processing the raw data, this would provide the astronomer with a powerful workstation for exploration of the reduced images and catalogs. Due to the need for instant access, image displays, and interactive image analysis, remote supercomputers are not a solution to the computational requirements in optical-IR image-processing and analysis. On the other hand, the data rate from these detectors will overwhelm the traditional mini- or micro-computer or workstation. Instead, we look to fast processors based on Digital Signal Processors with Gflop average speed and multiple wide buses, together with large RAM memory and fast multi-port disks, which are now becoming available.

Consider a multi-wavelength digital sky survey using a 10^8 pixel imager. The final CCD images for each band are passed to the FOCAS automated detector and classifier, creating a catalog of properties (isophotal, aperture and total magnitudes, centroid positions and several central moments) for each detected object. For this mosaic imaging survey, with a limiting magnitude < 22 magnitude, the sky is sparse and the resulting catalog would be very small (a few Gb) compared with the full processed image databank (8000 Gb).

In summary, we will soon have mosaic imager/computer systems capable of pushing the largest existing telescope to its performance limit. High-efficiency CCD imagers covering most of the usable focal plane, together with specialized on-line computers using automatic image classification software, will radically change our ability to observe the universe.

Supertelescopes

In radio astronomy, aperture synthesis arrays do not themselves form an image, but depend upon a digital computer as the image-forming element of the telescope. Powerful deconvolution algorithms have been developed which can greatly enhance the power of radio telescopes. The original design goal of the Very Large Array (VLA) of a dynamic range of 100:1 has now been increased routinely to 2,000:1, and a dynamic range of 100,000:1 is achieved for point sources. The lesson of the VLA is that a telescope is no longer to be considered a fixed capability instrument. Rather, it becomes a "supertelescope" which becomes more powerful with time, by virtue of its coupling to new generations of more capable digital computers and new computational algorithms.

Unfortunately, the NRAO has insufficient computer power available to allow the full potential of the VLA to be realized. The NRAO has proposed a distributed approach to meeting its computing challenges:

(1) By writing and supporting flexible, all-purpose, exportable code (e.g. AIPS), much or even most of the projects scheduled on the VLA can be properly reduced at the observatory on mini-supercomputers, or at the researcher's home institution, on desktop or departmental computers. This approach satisfies perhaps 85% of the individual observing runs, and has led to a revolution in our understanding of the radio sky.

(2) To provide the capacity for the remaining 15% of very large projects, the NRAO advocates a supercomputer access plan. The required software, perhaps specially coded to match various high-performance architechtures, would be available on a few very large capacity machines. Fast data links must be available to allow real-time interaction of the user with the results. This is required since so much of radio astronomical data reduction is iterative in nature, and an experienced eye

is required to judge the progress. Projects such as low-frequency imaging, snapshot programs, studies of the interstellar medium of individual galaxies, galactic absorption studies, OH and H_2O maser emission that fall into the supercomputing category, are also of great astronomical interest. Finally, the new Very Long Baseline Array computing needs are dominated by the same types of spectral line projects as listed above.

To grasp the enormity of the computing problem for VLA data, we describe one special problem—imaging low-frequency data—which is a continuum in nature, but nevertheless requires the spectral line observing mode.

The NRAO has recently completed installation of 327 MHz (90 cm) receivers on the VLA. Unique science addressable with this new capability includes steep spectrum objects and objects of large diameter and low surface brightness. Due to the two-dimensional geometry of the VLA, the samples of the visibility function are made throughout a three-dimensional volume, and the conditions under which a two-dimensional Fourier transform can be used to recover the source brightness fail, requiring much more expensive solutions. The simplest of these is a three-dimensional Fourier transform, producing a three-dimensional image "volume" whose axes are in direction cosines, and within which the desired image is found on a sphere of unit radius. Processing of the image follows the same procedures normally done in two-dimensional processing—for example, deconvolution proceeds in an entirely analogous fashion using the three-dimensional image with a three-dimensional beam.

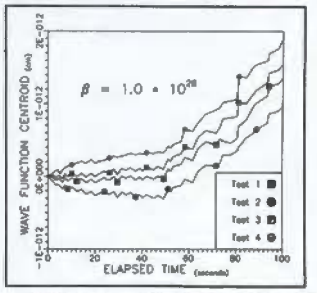
A typically large project at 327 MHz will use all four VLA configurations with perhaps 12 hours of observing in each. Because of chromatic aberration, the spectral line correlator must be employed to ensure that the bandwidth of each data channel remains small. The result of this is a 16-fold increase in data rate and volume over the continuum case. In this mode, radio frequency interference can be identified and purged without seriously corrupting adjacent channels. The integration time for each sample must be kept very short to prevent time-averaging smearing. The result is a very large database: typically 3.5 GB, containing over 500 million complex numbers.

Calibration of these basic data is straightforward, and can be accomplished with modest computer resources. But the imaging needs are exceptionally large. The simple 3D transform requires making a "dirty" map and beam, each $4096 \times 4096 \times 64$ pixels, which with 4 bytes per pixel requires 8.5 GB memory if they are made in the most efficient, straightforward way. Fortunately, another approach is more efficient in memory, although at the cost of I/O and CPU. If the image is built up through a series of large numbers of sub-images of limited depth, then the memory requirement is only about 135 MB per facet.

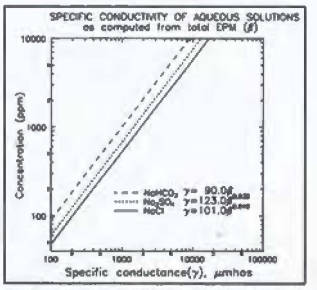
Deconvolution of the image is the next concern. In the simple, single cube approach, the procedure is straightforward, although highly consumptive of memory. No recall of the data is required. In the polyhedron imaging approach, deconvolution can be accomplished with much reduced memory, again at the cost of much increased I/O. A rough estimate is that perhaps 10,000 Fourier transform subtractions of clean components from the data are required, along with 10 re-imagings of the entire field. In other words, more than 2500 FFTs, each 1024×1024 pixels, will be required. But this does not finish the processing, since the ionospheric corruptions must be removed through self-calibration. Self-calibration and deconvolution are interlinked, the former using the results of the latter to generate a better estimate of the sky brightness. Typically, three loops of self-calibration and deconvolution are required before satisfactory convergence has been achieved. Thus, all the operations described in the previous paragraph must be multiplied by three.

A rough estimate of the time required can be made, as follows. Using the Cray Research Inc.'s Cray-2 supercomputer performance with a more modest case as a benchmark, and multiplying up by the ratio of database sizes and number of fields to process, gives a rough estimate of 250

SCIENTIFIC GRAPHICS



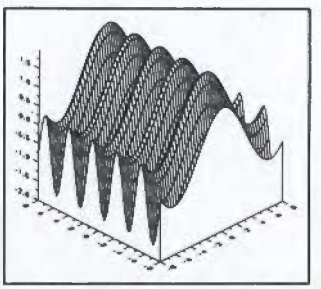
GRAPHER™ accepts your ASCII comma or space delimited file of up to 32000 XY pairs. You may combine an unlimited number of files on each graph. Choose from five types of error bars and six types of best-fit lines. Include automatic legends and unlimited text.



Use any combination of linear and logarithmic axes with automatic or user-specified tics and labels. Text may contain superscripts, subscripts, and mixed fonts from **GRAPHER™**'s complete symbol library, including Greek letters and special symbols.



SURFER® creates contour plots from your data quickly and easily. You may specify contour label frequency and format, irregular contour intervals, and data posting. Choose a rectangular border with tics and labels, or a user-defined shape.



SURFER® lets you display your data as a 3-D surface in perspective or orthographic projection, rotated and tilted to any degree or angle. Add axes, posting and titles to your plots. Stack surfaces for impressive results.

GRAPHER™ (PC Editor's Choice) . . **$199**
SURFER® (PC Editor's Choice) . . . **$499**
Demo Disk **$10**

FREE Brochure

Go ahead . . . . . give us a call.
1-800-333-1021
or (303-279-1021 - fax: 303-279-0909)

GOLDEN SOFTWARE, INC.
809 14th St., Golden, CO 80401

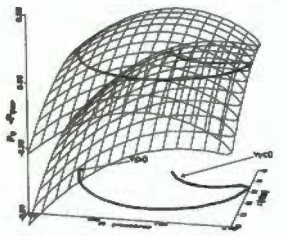
Purchase orders are welcome. for IBM-PC

Circle number 16 on Reader Service Card

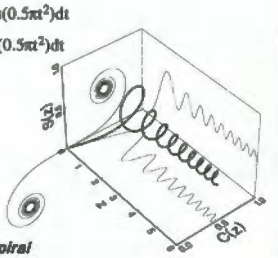
GraphiC™

"gives you all the C language routines you need to write an impressive scientific graphing program of your own. Highly recommended."
– PC Magazine
March 14, 1989

Topology of Tokamak Banana Orbits

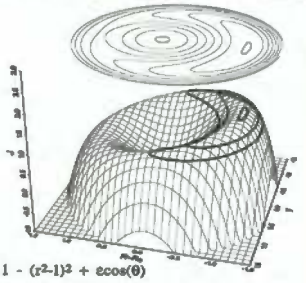


Fresnel Integrals
$C(z) = \int_0^z \cos(0.5\pi t^2)dt$
$S(z) = \int_0^z \sin(0.5\pi t^2)dt$



Cornu's Spiral

Orbits correspond to J = constant contours



$J = 1 - (r^2-1)^2 + \epsilon\cos(\theta)$

- 4096x3120 Tektronix format
- Convert files to PIC, GEM, HPGL and PostScript®
- Supports most hardware
- Publication quality output

IBM PC version
with C source code **$395.**
Macintosh version (ThinkC)
without source code **$295.**
Licensed for *personal* use only

Scientific Endeavors
508 N. Kentucky Street
Kingston, TN 37763
(615) 376-4146
FAX: (615) 376-1571

Circle number 14 on Reader Service Card

Supertelescopes (cont'd)

days for full processing. The time required is dominated by the gridding. We are confident that usable short-cuts will be found: for example, we can probably use a much less expensive gridding algorithm, which could cut processing by a factor of five.

It seems clear that the optimal approach will eventually involve massive parallelism. We could imagine 16 parallel machines, each comparable to, or faster than, the Cray-2 supercomputer. The data are then distributed to all of these machines, each of which is responsible for one sub-field. A central processor will be required to handle the component model subtractions—the model comes from all 16 parallel processors. Factoring in these expected savings, and imagining future, more powerful machines, predicts this particular problem to be soluble with a few hours of computing time.

The essential points of the above example are summarized here: (1) There is a need for very fast machines with very high I/O rates and extremely large memory to generate the data cubes. In many cases, parallel processing is clearly advantageous. (2) Assuming the NRAO will not be able to obtain such machines, the national supercomputer centers must provide access and support for astronomers requiring this computing. (3) Desktop computers must be supported, to allow proper interaction between astronomer and completed image. (4) Research and development of computing algorithms must be actively supported, both at national centers and at national observatories.

BIMA—A Real-Time Remote Observatory

Real-time radio observation should become a reality in the 1990s, and interactive observing is a possibility if image processing and transmission can be made sufficiently rapid.

One of the goals of the BLANCA networking experiment is to demonstrate such capabilities and to explore how they might improve, expand and extend the power of a telescope system. The BLANCA network testbed, which has a T3 (45Mbps) network stretching from AT&T Bell Laboratories in Murray Hill, N.J. to the University of California at Berkeley, connects with a 620 Mbps network (starting in early 1992) from NCSA at the University of Illinois to the University of Wisconsin. BLANCA is one of five testbeds being coordinated by the Corporation for National Research Initiatives, with funding from the National Science Foundation and the Defense Advanced Research Projects Agency. In addition to BLANCA, there are testbeds on the east coast (AURORA), in Pennsylvania (NECTAR), in North Carolina (VISTANET), and on the west coast (CASA).

The Berkeley-Illinois-Maryland Array (BIMA) is located at the Hat Creek Observatory in northern California, and is operated by the three universities. The telescope system is completely under computer control and is accessible via computer networks, so an astronomer can monitor in real-time the data acquisition process, edit the data, and set up command files for data processing.

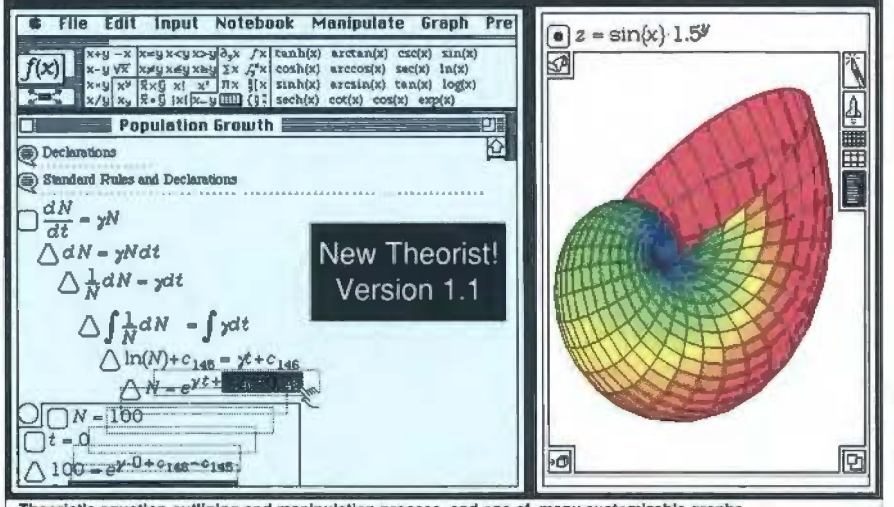
The Array is similar in concept to the VLA, but operates at millimeter rather than centimeter wavelengths. By early 1991, BIMA will consist of six antennas; there are plans for expansion to as many as 12. Although with six antennas BIMA has only about 5% of the number of simultaneous interferometers as the VLA, the BIMA spectrometer will have eight separate spectral windows available, of 256 channels each. In a typical eight-hour tracking, the visibility function will be sampled sufficiently to make useful data cubes (eight cubes of eight different spectral lines with right ascension, declination, and velocity axes) immediately. Since BIMA will be used in spectral-line mode virtually all the time, while

the VLA is often used in continuum mode, the BIMA data rate and computational requirements will be about one-third those of the VLA.

A typical BIMA data set will be in the 100 MB to 1 GB range; such data sets can be transferred from Berkeley to the supercomputer at Illinois at 45 Mbaud in considerably less than the period of a five-minute coffee break. The initial processing of the observed visibility data on the supercomputer will be automatic, under the control of an "expert system" with tunable parameters which may be set in advance by the astronomer. While the observations are in progress, calibration, map-making, and an initial deconvolution, self-calibration and mosaicing (if appropriate) of a partial data set may be carried out and the data cube returned to the astronomer at Berkeley for analysis on a workstation. The astronomer will be able to judge the quality of the data, to see if the signal is strong enough to proceed with the observations, to judge whether the area of sky being mosaiced is correct, and to begin to experiment with processing parameters. Instrumental or atmospheric problems can be detected quickly, and corrected, while reobservations can be carried out with the telescope still in the same configuration. Exciting or unexpected results can be pursued immediately.

As experience is gained with this new mode of radio observing, it is likely that special process computers executing the optimal computer code, developed on the general purpose supercomputers, will be installed at the site of the radio telescope. By that time, T3 network access to a wide base of national remote observers will be routine, and higher bandwidth testbed experiments in the gigabaud range will be underway to determine the next level of functionality that can be obtained for remote real-time observation.

Intuitive Tools For Mathematics



Theorist's equation outlining and manipulation process, and one of many customizable graphs.

Prescience (pronounced PRE-*shence*) brings you the complete mathematical solution for the **Macintosh**: Theorist and Expressionist. Theorist is the symbolic algebra and graphing program that is easy to use and powerful, but only requires one (1) megabyte of memory. You don't need to learn how to program, memorize syntax rules, or read a large manual since Theorist actually displays and interactively solves real equations on screen—step by step—the way you do on paper. 2-D and 3-D graphs, contour and density plots, solids, as well as animation files, are easily created and saved in PICT, EPS, or PICS formats for high quality output. Your equations can be exported to Expressionist, the leading equation editor, for typeset-quality results in your word processing and page layout documents. Both programs are simple enough for the student, yet powerful enough for the professional educator, scientist, and engineer. Our programs enable you to concentrate your time investigating and presenting your work, not learning how to! For more information or to place an order, call or write to the address below.

"[Expressionist] Equation manipulation has never been easier."
– Five mice review, MacUser *magazine*

Expressionist® $129.95
Typeset Quality Equation Editor

Theorist® $379.95
Symbolic Algebra & Graphing



"Theorist... surpasses the highly rated Mathematica... in interface and execution."
– MacUser *magazine, Editors' Choice Award, Best Math/Statistics Program of 1989*

Prescience Corporation
939 Howard Street #106, San Francisco, CA 94103
(415) 543-2252 Fax (415) 882-0530

Prescience
Intuitive Tools For Mathematics

Fourth Dimension Supercomputing

The "Fourth Dimension Supercomputer" is a system sufficiently powerful to calculate the evolution of complex, nonlinear systems in a fully three-dimensional space and in time. Presently, large memory supercomputers are just capable of providing this capacity, but with only crude resolution. Present trends in computer technology suggest that the next generations of machines will be serious "Fourth Dimension Supercomputers." We herein examine the scaling of time-explicit computer codes to solve such astrophysical problems.

Experience suggests that about 200 computational points per spatial dimension is the lowest resolution that begins to be interesting for a variety of problems involving simple physics, such as inviscid gas dynamics. In the following discussion, we will assume $200\times f$ to be the required number of computational points per spatial dimension. This may be thought of as a minimum acceptable resolution for a variety of problems of current interest; $f=2$ or $f=3$ may often be appropriate. At each point a 3D calculation would require 4000 or more floating point operations (e.g. multiplications, additions), giving a value of $3.2\times10^{10}f^3$ floating point operations per time step. Consideration of accuracy and stability of the numerical method often results in requiring a number of time steps for one dynamical time of the system, which is of the order of the number of space steps in one dimension—in this case $200f$. For ten natural times, this implies $6.4\times10^{13}f^4$ floating point operations.

How much memory is required? The physical state at a single point in the system may require 20 or more variables, each 8 bytes long. For two time slices (an old and a new state), this implies $3.2\times10^8f^3$ numbers or $2.56f^3$ GB. Assuming one had a supercomputer today capable of a sustained rate of a gigaflop and 4 gigabytes of memory, this code would run in less than 20 hours, which is acceptable. However, to move to 1000 cubed resolution (comparable in resolution to today's two-dimensional fluid flow codes) with the same turnaround, would require a sustained speed of 625 Gflops with a memory of over 300 GB. These values are likely to be achieved in massively parallel systems by mid-decade.

Using this value of $1.28f^3$ GB per state (time slice), and assuming that 100 of the $2000f$ time steps are saved for analysis, the storage needed per project is $128f^3$ GB. For 40 such projects, the storage requirement grows to 5 terabytes. This is a major limitation: the extensive storage at the NSF supercomputer centers for researchers in all areas of science and engineering is estimated to be of the order of a few terabytes. Data compaction and new storage technologies are required in order to alleviate this bottleneck.

How do we get these data to the scientist for analysis? A reasonable estimate may be obtained by requiring that the time for data transfer of the results must be less than the time for computation—otherwise the data flow becomes the bottleneck. From above, one time step takes $32f^3$ seconds to compute on a Gflop supercomputer and generates 1.28 GB of data. This implies a data transfer rate of 320 Mb/second if the user wishes are the variables at full precision. Data compaction and upgrade of the NSF backbone to Gb/second speeds are needed even for today's problems. It is particularly important to provide this national network resource to the wider community of scientists and students who do not reside at a supercomputer center site.

Assuming the data reach to the remote user (today by Federal Express on Exabyte tapes), there must be local facilities to deal with it when it arrives. For example, say one saves only one percent of the time slices from a $2000f$ time step evolution, generating only $20f$ saved states of $1.28f^3$ GB each. This would require $2.56f^4$ GB of local storage. Factors of two and four could be saved by using 32-bit and 16-bit wordlengths. Analysis of dynamical 3D systems requires a look at the time behavior of the system. This implies a requirement for local graphics capacity. For example, suppose we want 60 seconds of images at 10 screens/second. For 3 MB/screen (3 byte color on a 1000×1000 screen), this is 1.8 GB needed on a high-speed disk to feed the graphics engine. This graphics pipeline must be fed at 30 MB/second. This indicates once again the need for mini-supercomputers or graphics supercomputers at remote user sites.

Acknowledgements: We wish to express our thanks to the members of the Panel on Astronomical Computing for the large amount of time they spent preparing the Panel Report and the White Paper background document for the Panel Report. The sidebars are edited versions of essays written by individual Panel members. Larry Smarr would like to thank Radha Nandkumar, NCSA Research Staff Associate, who was instrumental in the creation of this report. Michael Norman, NCSA Research Scientist, was very helpful in many stages of the year-long process. ■

References

1. The Decade of Discovery: Report of the Astronomy and Astrophysics Committee, National Research Council, National Academy Press, Washington, 1991.
2. Astronomy and Astrophysics for the 1980s: Report of the Astronomy Survey Committee, National Research Council, National Academy Press, Washington, 1982.
3. "Array Telescope Computing Plan", a conceptual proposal submitted by the National Radio Astronomy Observatory (NRAO) to the National Science Foundation (NSF) in September 1987 and recently resubmitted.
4. "Grand Challenges: High Performance Computing and Communications", a report by the Committee on Physical, Mathematical, and Engineering Sciences, Federal Coordinating Council for Science, Engineering and Technology, Office of Science and Technology Policy, Feb. 1991.