

INTERIM REPORT

of The

1984 NSAC SUBCOMMITTEE ON COMPUTERS AND COMPUTING

October 1984

DOE/NSF NUCLEAR SCIENCE ADVISORY COMMITTEE

DOE/NSF NUCLEAR SCIENCE ADVISORY COMMITTEE

1984

Gordon A. Baym (University of Illinois)
D. Allan Bromley (Yale University)
Frank Calaprice (Princeton University)
Herbert H. Chen (University of California, Irvine)
Douglas Cline (University of Rochester)
Karl A. Erb (Oak Ridge National Laboratory)
Hermann A. Grunder (Lawrence Berkeley Laboratory)
Ole Hansen (Brookhaven National Laboratory)
Cyrus M. Hoffman (Los Alamos National Laboratory)
Arthur K. Kerman (Massachusetts Institute of Technology)
Malcolm H. Macfarlane (Indiana University)
Philip B. Roos (University of Maryland)
John P. Schiffer (Argonne National Laboratory), Chairman
Ingo Sick (University of Basel)
Erich W. Vogt (University of British Columbia)

1984 NSAC SUBCOMMITTEE ON COMPUTERS AND COMPUTING

Lawrence S. Cardman (University of Illinois)
Herbert H. Chen (University of California, Irvine), Chairman
John G. Cramer, Jr. (University of Washington)
Benjamin F. Gibson (Los Alamos National Laboratory)
Ed. V. Hungerford (University of Houston)
Steven E. Koonin (California Institute of Technology)
Stanley B. Kowalski (Massachusetts Institute of Technology)
William M. MacDonald (University of Maryland)

Ex-officio:

Enloe Ritter (Department of Energy)
John P. Schiffer (Argonne National Laboratory)
Harvey B. Willard (National Science Foundation)

SUMMARY AND RECOMMENDATIONS

THE 1984 NSAC Subcommittee on Computers and Computing was constituted in August, 1984 and charged with determining "the needs of the United States basic nuclear research program for computers and computing over the next decade. ..." In response to an urgent agency request for an immediate assessment of the supercomputer (CLASS VI) needs of the nuclear science community, we decided to concentrate upon three issues in this interim report: 1) the benefits of supercomputer use to the nuclear science community, 2) the approximate scale of this need in the next few years, and 3) the essential requirements for effective utilization of such systems.

Our initial action of assessing supercomputer related issues in the short term, i.e., two to three years, should not be interpreted as favoring supercomputers over the many other computer and computing needs in nuclear science. These will be addressed fully in our final report which should be completed by the end of this calendar year.

NSAC has previously identified the need for supercomputer access by the nuclear science community. This subcommittee has re-examined the current needs and requirements. Based upon information provided through oral presentations, written reports, telephone inquiries, a hearing of agency plans, and a review of requests to DOE for supercomputer time, we have reached the following conclusions:

1. There is an immediate need for supercomputer power of about one-half a CLASS VI machine.
2. Rapid growth of this need to the level of 1 1/2 CLASS VI machines is estimated over the next two to three years as experience and expertise develop.

To address these needs of the nuclear science community, we make the following recommendations:

1. Access to one-half of a CLASS VI computer with at least four million words of memory should be made available as soon as possible.
2. Such a computer center must provide a responsive, user-friendly interactive system to insure efficient use and effective algorithm development. Real time turnaround must be adequate for effective code testing. Batch processing for long, complex programs is also needed. Software support, sophisticated graphics capability, knowledgeable consultants (physics as well as computer), good documentation, and mass storage facilities are all required.

3. Reliable network connections to the center is essential. The present National Magnetic Fusion Energy Computer Center network provides a reasonable minimum standard. To adequately service remote users, network costs must be borne by the center.

4. Appropriate local computing connections to the network at remote user sites must be provided. Other local computing resources needed will range from personal computers which could function as smart terminals, to more sophisticated graphics terminals, to scientific work-stations, to super-minicomputers capable of significant scientific calculations. Hardcopy output requirements will range from simple screen printers to graphics quality printers.

I. INTRODUCTION

In response to the DOE/NSF request to the Nuclear Science Advisory Committee (NSAC) that it conduct a study of the needs of basic nuclear science in the United States for computers and computing, a subcommittee was formed in August, 1984. The charge to this subcommittee is:

"Please determine the needs of the United States basic nuclear research program for computers and computing over the next decade. This study should be developed within the scientific and budgetary framework of the 1983 Long Range Plan. Basic philosophy of the study should be identification of computer and computing levels required to implement the Long Range Plan. Evaluate which parts of the projected computational program would be best served by: on-site, locally controlled 'mini-computers' (with/without array processors); on-site, central main-frame computers (with/without array processors); and supercomputers. And finally, please develop a plan to meet nuclear physics computing needs for the next decade which involves any and all of the above classes of computers, which includes the utilization of existing computing capabilities, and which includes consideration of requirements for large scale networking. The plan should include an estimate of the costs of the proposed system."

The letter of request suggested that the study be completed and a report transmitted to DOE and NSF by October, 1984.

The subcommittee met August 31-September 1 in Washington, D.C. and heard presentations on DOE and NSF scientific computing initiatives. Agency plans include: installation of a CRAY XMP early in FY'85 at the National Magnetic Fusion Energy Computer Center (NMFEECC); DOE purchase of 70% of the time on the CYBER 205 to be operational early in 1985 at Florida State University; plans for about seven national supercomputer centers by NSF beginning with the funding of three such centers in FY'85. It also heard reports on the operation of the NMFEECC, parallel processing, software portability, scientific work-stations, and the needs of experimental and theoretical nuclear physics. In the discussions that followed, the subcommittee concluded that a timely response to the DOE/NSF request required an interim report which would concentrate upon the following supercomputer related issues:

1. the benefits of supercomputer use to the nuclear science community,
2. the approximate scale of this need over the next several years, and
3. the essential requirements for effective utilization of such systems.

Our first action of assessing supercomputer related issues in the short term, i.e. two to three years, should not be interpreted as favoring supercomputers over the many other computer and computing needs in nuclear science. These will be addressed fully in our final report, which should be completed by the end of this calendar year.

The subcommittee reconvened on September 26 in Washington, D.C. for a working session to begin drafting the interim report. Additional input was considered. For example, the DOE Division of Nuclear Physics provided requests from contractors for supercomputer time in FY'85 and FY'86. In addition, results of telephone inquiries by subcommittee members were reported and considered. The draft report was written and revised, using 'electronic mail', on a VAX facility at the University of California, Irvine in order to speed communication between committee members as well as to encourage access and input by the nuclear science community at large. This report was completed at the October 14-16 meeting in Houston, where the full spectrum of computer and computing needs were heard and discussed with interested members of the nuclear science community.

II. SUPERCOMPUTER APPLICATIONS

Limited computer resources have become a decisive factor in the selection of areas for investigation by many research scientists in this country. Lack of access to state-of-the-art computing facilities has directed the attention of scientists away from those topics where sophisticated supercomputers are essential to progress. When supercomputers are needed, they are indispensable. Thus, certain important fields have not been pursued with vigor, to the detriment of U.S. science and technology. Furthermore, many of the next generation of research scientists are being deprived of training in the use of the contemporary supercomputer. Efficient use of the sophisticated architecture of these modern machines is not automatic for those whose computing philosophy has been nurtured on standard scalar machines.

II.A. Theory

Several studies have documented the crucial role that computers have played in the development of modern physics. Numerical calculations provide the means of comparing model simulations of physical phenomena with experimental data. Nuclear physics has now matured beyond the stage in which pencil and paper results from essentially analytical models will suffice for incisive analysis of experiments. Advancement of our understanding of forefront nuclear physics problems requires precision calculations based upon complex, realistic models. Many degrees of freedom (dynamic variables) must be treated in order to describe the nucleus in more than just a rudimentary fashion. In many areas of nuclear science research, this can be accomplished

in a timely fashion only through use of contemporary super-computer capability.

A selection of frontier nuclear physics research areas where supercomputer power is essential for significant progress includes:

1. Large basis shell model studies of nuclear structure,
2. Interacting Boson Model studies of the structure of mid-shell nuclei,
3. Few-body structure and reaction studies,
4. Monte Carlo methods for calculating properties of Fermion systems,
5. Radiative correction calculations for electron scattering,
6. Coupled-channel reaction calculations,
7. Time Dependent Hartree Fock studies of heavy ion collisions,
8. Hydrodynamic studies of relativistic heavy ion collisions,
9. Properties of quark-gluon plasmas,
10. Nuclear astrophysics investigations of supernovae.

Large basis shell model calculations form the backbone of nuclear structure investigations. A prime example is nuclear double-beta decay, where one strives to determine whether the neutrino is Dirac or Majorana in character. Another is the calculation of core polarization effects due to valence nucleons as measured in inelastic electron scattering from nearly closed shell nuclei. These codes exceed the capacity of a VAX 11/780 even for simple, medium weight nuclei. They involve manipulation of large arrays and are therefore suited to modern super-computers. However, such a full shell model treatment of only valence nucleons in the A=154 isotope of samarium, for example, would exceed the capability of any known computer. Thus, Interacting Boson Model investigations (an approximation to the shell model) are essential to our understanding of the structure of heavier systems. This model, in its more sophisticated versions, holds promise of correlating large amounts of experimental data on excitation spectra in terms of a few parameters. The aim is to separate neutron and proton degrees of freedom such as disentangling the neutron and proton static and transition densities of the nucleus. Even these Interacting Boson Model calculations will require supercomputer resources.

Solving the bound-state and continuum-state few-body problem is essential to our understanding of the nature of the nuclear force in nuclear matter. Until the properties of few-nucleon systems can be calculated in terms of realistic force models which describe our experimental knowledge of the free nucleon-nucleon interaction, we cannot judge the validity of the standard model assumption that nuclear matter and reactions can be described in terms of measured two-body forces. Must

quark-gluon degrees of freedom be included in the model? Such calculations have large memory requirements because of the realistic force assumption. However, the matrix formulation of the problems and iterative nature of the solution algorithms make them well suited for supercomputers. Monte Carlo studies of the properties of few-body bound states are also a vital part of our modeling of nuclear structure in terms of realistic nuclear forces. Hundreds of hours of supercomputer time are needed to lower the variance of such Monte Carlo calculations to a level where significant physics can be extracted. Properties of nuclei heavier than about $A=6$ cannot be calculated via standard few-body techniques, whereas Monte Carlo methods may be applicable to large nuclei if sufficient computer power is available. The promising technique of Green's Function Monte Carlo, which has been successfully tested on large (256 particle) boson systems, has foundered on the "noise" inherent in any random number process when applied to fermion systems. New algorithms are essential if we are to apply this powerful technique to even light nuclei such as oxygen.

Further research into the theoretical description of the physical process of electron scattering is crucial to proper interpretation of these data. Exploring distorted wave calculations of radiative corrections and virtual excitations (coupled channel) effects consume hours of CRAY time at present. More complex phenomena such as electrofission and higher energy studies for new or upgraded electron machines will require even more supercomputer time.

Coupled-channel reaction calculations have been used to explore how the shapes of nuclei affect reaction processes. The importance of the coupling of inelastic channels in the enhancement of the fusion cross section below the Coulomb barrier is also investigated via coupled-channel reaction codes. The large systems of coupled equations required to describe heavy ion collisions have exceeded the capability of existing main frame computers. The strong coupling of the many channels makes the convergence of the calculation very slow. Time Dependent Hartree Fock studies of the evolution of heavy ion collisions teach us about the transparency and fission characteristics of nuclear matter. The time required to follow the collision process as a function of the many possible kinematic parameters exceeds the capacity of available main frame computers. Supercomputer calculations are essential to interpretation of the data from existing heavy ion accelerator facilities.

Hydrodynamic calculations of relativistic heavy ion reactions have been shown to consume hundreds of hours of CRAY time. It was the comparison of results from such calculations (and the alternative cascade model calculations) with data from the collision of two niobium nuclei that led to the conclusion that the observed "sidesplash" was evidence for compressed matter, matter that might exist in condensed stars. The compression and fluid flow of the hydrodynamical model was found

to yield such a sidesplash in simulations. Graphical display and real-time response output from these calculations are essential to the interpretation of such studies. Related investigations seek signatures for the formation of a quark-gluon plasma in ultra-relativistic nuclear collisions. Fluid dynamic or transport calculations require a very fine mesh (100 x 100 x 100 points). Searches for unique signals for the formation of such a novel form of nuclear matter will require a monumental effort in manpower and computer time but are vital to the success of any experimental program aimed in this direction.

Nuclear astrophysics calculations aimed at comprehending the dynamics of supernovae model the shock-explosion mechanism which follows the gravitational collapse of massive stars. Present studies are based upon spherically symmetric hydrodynamic calculations and involve direct integration of the Lagrangian equations describing the evolution of supernovae. The nuclear equation of state is a key element in any such calculation. The iterative nature of the calculation combined with its complexity makes main-frame calculations inadequate. Realistic modeling is possible only on a supercomputer.

II.B. Experiment

There is little supercomputer experience among experimental nuclear physicists because of the lack of access. However, there are at least three areas which will benefit substantially from supercomputer capabilities: accelerator design, Monte Carlo simulations of experiments, and data analysis requiring global searches. These activities progress most efficiently when results can be obtained quickly in real time. The next generation of detectors will place demands on analysis computers at least an order of magnitude greater than now experienced. However, the multiprocessor architectures presently under development may be more cost effective than supercomputers for processing and analyzing event-mode data.

The design of new accelerators as well as upgrading of existing accelerators are computer intensive processes. The computer is a basic tool of the designer, especially in the development of advanced accelerator concepts. Computer simulations can significantly reduce the cost and time involved in finding design imperfections and optimizing performance. Calculations in accelerator and beam transport design deal principally with properties of electromagnetic fields and the dynamics of particles in such fields. Present calculations are essentially two dimensional, utilizing various symmetries to describe accelerator components. Three dimensional codes are needed to realistically describe the problems that occur in such phenomena as nonlinear colliding beam instabilities as well as in microwave power generation and secondary beam transport systems. Machine design is an interactive process. Supercomputer capability is needed to achieve realistic modeling of both accelerators and beam dynamics.

Monte Carlo simulations of complex experiments is a standard procedure which will certainly benefit from supercomputer access. Multiparticle detectors have increased the memory and speed of computers required for their simulation and evaluation. The iterative nature of procedures required to determine the optimum experimental parameters necessitate use of the fastest machines available. Unfortunately, the lack of supercomputer access and the lack of a user friendly environment on such machines have forced experimentalists to accept the constraints of running limited calculations for extended periods on locally available VAX type computers. Simulations of future relativistic heavy ion collider experiments will require the use of supercomputers.

A fast machine with large memory is capable of performing global phenomenological fits to very large data arrays. Such calculations may be extended to systematize and evaluate data as well as provide global parameters for nuclear calculations.

III. SUPERCOMPUTER USE REQUIREMENTS

NSAC has previously identified the need for supercomputer access by both theoretical and experimental nuclear physicists. The 1979 NSAC Long Range Plan pointed out the continuing lack of adequate computing capability for Nuclear Theory. The December, 1981 report by the NSAC Subcommittee on Computational Capabilities for Nuclear Theory recommended a system designed around a CLASS VI central computer. The 1983 NSAC Instrumentation Subcommittee pointed out that the computing resources needed for rapid data analysis are significantly greater than was believed a few years ago, that analysis time can be 15 to 20 times longer than data acquisition time using the mid-sized computers presently available for data analysis.

We have re-examined the current nuclear science super-computer requirements and reviewed the recommendations of NSAC in prior years. We have studied the requests made to the DOE Division of Nuclear Physics for NMFEEC supercomputer time in FY'85 and FY'86, and we have conducted our own limited telephone inquiries of supercomputer requirements in theoretical and experimental nuclear research. The available evidence indicates:

1. There is a need for immediate access to supercomputer power on the order of one-half a CLASS VI machine.

2. The research needs of both theoretical and experimental nuclear science are expected to grow over the next two to three years to require computer resources equivalent to 1 1/2 CLASS VI machines. As the experience and expertise of the community with supercomputers advances, more realistic (and therefore more ambitious) nuclear physics problems will be attacked. Thus, the need for supercomputer resources will grow rapidly.

The supercomputer needs of nuclear science can be met by combining its needs with those of other areas of physics to provide the best available facility. One should aim to buy the best and then use it to the maximum feasible extent. Super-computer capability must be managed as a scarce resource, but it must also be made available in a timely manner to those researchers who have established a need.

To address the present needs of nuclear physics, access to a CLASS VI supercomputer with at least four million words of memory, which can be expanded as larger fast memories become available, should be provided as soon as possible. A responsive, user-friendly interactive system is necessary to insure efficient use and effective algorithm development. Real time turnaround must be adequate to insure thorough code testing. Batch processing for long, complex programs is also necessary. Software support, sophisticated graphics capability, knowledgeable consultants (physics as well as computer), good documentation, mass storage facilities, and a system for flexible allocation of resources are all required. Reliable network connections to the center are essential. The present NMFEC network provides a reasonable minimum standard: 24 hour availability, 1200 baud dial-in lines, 4800 and 9600 baud land-line links to large users, satellite links to major nodes. To adequately service users, network costs must be borne by the center. Appropriate local computing connections to the network at the remote user site must be established. Other local computing resource needs will range from personal computers which could function as smart terminals, to more sophisticated graphics terminals, to scientific workstations, to super-minicomputers capable of substantial scientific calculations. Hardcopy output requirements will range from simple screen printers to graphics quality printers.

IV. CONCLUSIONS

We have re-examined the need for supercomputer use by the nuclear science community. We strongly reaffirm the conclusions of previous NSAC bodies that there exists a need for supercomputers. To meet this need, a fast scalar machine with proven vector capability is necessary. Approximately 1/2 of a CLASS VI machine is required in FY'85. We estimate that this will grow to 1 1/2 CLASS VI machines in two to three years. A user friendly, interactive, nonsaturated environment is necessary at the supercomputer center along with a wide range of support services. Reliable network connections for remote computing are essential. Local user requirements at the remote sites must not be neglected.

The purpose of computing in nuclear science is not to generate numbers but to gain physical insight. Therefore, the closer coupling of the physicist to his computer models and algorithms, which would result from supercomputer access, is as important as the obvious gain in speed over computing on contemporary main frame computers. Furthermore, supercomputer use will reduce the degree of approximation required in the present modeling of nuclear phenomena. The investigation of new concepts will proceed at a faster pace and new opportunities for understanding will become available.