A COMPUTATIONAL PLAN FOR NUCLEAR SCIENCE

Final Report

of the

1984 NSAC Subcommittee on Computers and Computing

February 1985

DOE/NSF NUCLEAR SCIENCE ADVISORY COMMITTEE

850411

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue, Argonne, Illinois 60439

Telephone 312/972- 4066

April 17, 1985

Dr. Alvin Trivelpiece
Director of the Office of Energy Research
U.S. Department of Energy
Forrestal Building
1000 Independence Ave., S.W.
Washington, D.C. 20545

Dr. Marcel Bardon
Director of Physics
National Science Foundation
1800 G, N.W.
Washington, D.C. 20550

Gentlemen:

In your letter of June 18, 1984 you requested that our Committee determine the needs of the United States' basic nuclear research program for computers and computing. In response to your request, and on behalf of the Committee, I herewith transmit the report of the NSAC ad hoc Subcommittee, chaired by Professor Herbert H. Chen.

NSAC endorses the recommendations of the Subcommittee. However, the full implementation of these recommendations must be considered in the context of the present budgetary stringencies. The recommendation on the need for access to the equivalent of one-half of a class VI computer has very high absolute priority. Other recommendations are being implemented to a limited extent. For instance, nuclear scientists do have access to local computing stations (recommendation 2), though not to the full extent that would be desirable. Data analysis and acquisition systems are being upgraded (recommendation 3) and the development of multiprocessor systems is being supported (recommendation 4), though not at the highly desirable rate recommended in the subcommittee report.

In transmitting this report NSAC endorses the recommendations it contains, with the understanding that the priorities assigned to their implementation be considered in the overall context of the 1983 Long Range Plan. Recommendations of the Subcommittee are examples of the important class of activities for which NSAC is recommending a 10% increment to the base program in operating and equipment funds.

Sincerely yours,

John P. Schiffer
Chairman,
DOE/NSF Nuclear Science
Advisory Committee

JPS/bw

850411

Encl:

cc: H. B. Willard
E. T. Ritter

DOE/NSF NUCLEAR SCIENCE ADVISORY COMMITTEE

1984, 1985

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)
Robert J. Vandenbosch (University of Washington) - 1985
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 V (Los Alamos National Laboratory)
Ed V. Hungerford III (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:

John P. Schiffer (Argonne National Laboratory)

Liaison:

Anthony J. Baltz (Department of Energy)
Enloe T. Ritter (Department of Energy)
Harvey B. Willard (National Science Foundation)

Table of Contents

PREFACE

In response to the request from the Department of Energy (DOE) and the National Science Foundation (NSF) 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 over the next decade, the **1984 NSAC Subcommittee on Computers and Computing** was formed in August. 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 (see Appendix A) 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. The agenda and participants at this and subsequent meetings are given in Appendix B. The Subcommittee also heard reports on the operation of the National Magnetic Fusion Energy Computer Center (NMFECC), parallel processing, software portability, scientific workstations, and the computing 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.

A working group was set up to begin drafting the Interim Report on supercomputers. The membership of this working group is given in Appendix C. Our first action of assessing class VI computer related issues in the short term, i.e. two to three years, is not to be interpreted as a bias towards class VI computers over the many other computer and computing needs in nuclear science. The full spectrum of computing needs is considered here in detail.

The Subcommittee reconvened on September 26 in Washington, D.C. for a working session to begin drafting the Interim Report and to complete the organizational process for work on the Final Report. Additional working groups were set up to consider separately the computing needs of theoretical and experimental nuclear physics. The memberships of these working groups are also given in Appendix C.

For the Interim Report, additional input was considered. For example, the DOE Division of Nuclear Physics provided requests from contractors for class VI computer time in FY'85 and FY'86. Also, results of telephone inquiries by Subcommittee members were reported and considered. The Interim 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.

The Interim Report was essentially completed at the October 14-16 meeting in Houston, Texas, where the full range of computer and computing needs were heard and discussed with interested members of the nuclear science community. We also received written material from many others who were unable to attend due to the short time scales involved. A number of recommendations for the Final Report were considered and discussed. "Electronic mail", on a VAX facility at the Argonne National Laboratory, was used to forward the Interim Report to NSAC for its consideration.

The Subcommittee met December 5-7 in Irvine, California to discuss and write the Final Report and to consider suggestions for changes in the Interim Report. For the Final Report, a consensus was reached on the outlines of a computational plan for nuclear science, and draft recommendations were formulated and discussed. The Interim Report was completed and forwarded to NSAC for its endorsement and transmission to the agencies. The use of "electronic mail" continued through the drafting and editing phases between meetings. The need for a reliable network for improved communications between members of the Subcommittee, and the nuclear science community more generally, was made very apparent by this exercise. Also, personal computers played an essential role in the writing and communications process.

To obtain more complete information on the present nuclear science inventory of mini and larger computer systems, as well as requests and plans for acquisition of such systems, the Subcommittee initiated in December an informal telephone poll of all contractors and grantees.

The writing/editing of the Final Report was essentially completed at the January 24-25, 1985 Subcommittee meeting in Washington D.C., and the Final Report was submitted to NSAC for its consideration at its first meeting of 1985, on February 14-15, in Washington D.C.

SUMMARY

In response to the request by the Department of Energy (DOE) and the National Science Foundation (NSF) to "develop a plan to meet nuclear science computing needs for the next decade...", we have examined the current uses of computers and computing in nuclear science, and the directions toward which these are evolving in the solution of major scientific questions confronting nuclear science. The computer has become an essential and versatile tool for nuclear research. Computing facilities are at the heart of every nuclear physics laboratory. They are used in the recording, processing, and analysis of data, in the development of theoretical models, and in the confrontation between experimental results and theoretical prediction.

The rapid evolution of computer technology with the continuing rapid decrease in hardware costs has produced a number of significant opportunities for substantial enhancement of the computational tools available to the nuclear scientist. This includes increased computing power and improved communications. A continuing challenge which has become ever more serious due to the enormous growth in the use of computers and computing is to harness effectively this increased capability and to reduce duplication so that scarce resources will be efficiently utilized.

The computational plan presented here has as its primary objective the identification of the needs of the United States basic nuclear research program for computers and computing over the next decade within the scientific and budgetary framework of the 1983 **Long Range Plan for Nuclear Science.** We have examined the needs, opportunities and challenges in nuclear science in light of the significant advances in computer technology taking place now and anticipated in the near future. The computational plan, involving six inter-related recommendations, was developed through a series of meetings of this Subcommittee with input from the nuclear science community at large. This plan can be characterized by the following points: **Supercomputer Facilities; Local Facilities; Advanced Technology; Communications; and Coordination of Standards**. They are outlined in the following paragraphs.

Supercomputer Facilities

NSAC has previously identified the need for supercomputer access by the nuclear science community. This Subcommittee has re-examined the current and future needs and requirements. From information provided through oral presentations, written reports, telephone inquiries, a hearing of DOE and NSF plans, and a review of requests to DOE for class VI (Cray 1 equivalent) computer time, we find that these needs are substantial and pressing, and will increase. Therefore, we recommend that:

> RECOMMENDATION 1: **Access to the equivalent of one-half of a class VI computer should be made available as soon as possible. The agencies should prepare to meet the growth of this need, estimated to be about a factor of 5 within the next five years.**

The effective utilization of this computer time requires that certain standards be met by the centers supplying such class VI time and by the connecting networks. Any class VI computer center used for nuclear science should provide a responsive, user-friendly interactive environment to insure efficient use and effective algorithm development. Real-time turnaround must be adequate to promote thorough code testing. Batch processing for long complex programs, software support, knowledgeable consultants, good documentation, and sophisticated graphics are all essential elements.

The fact that DOE has already established two class VI computer centers (National Magnetic Fusion Energy Computer Center [NMFECCand Florida State University [FSU, and the expectation that NSF will establish about seven supercomputer centers for scientific computation is an impressive commitment to such computing capability. However, this proliferation of centers can create problems for users and can be wasteful of this scarce resource. The lack of commonality in many computer systems causes much duplication of effort. The difficulties associated with moving codes from one type of computer to another or in using a program developed on a different machine reduces productivity. These problems can be minimized by concentrating nuclear physics computing at a small number of supercomputer centers.

Reliable network connections to the centers should be provided. The present NMFECC network provides an excellent existing model which should be upgraded as communications technology improves. Appropriate local computing connections to the networks at remote user sites will be needed. Other local computing resource needs, discussed in the next section, will range from personal computers which could function as smart terminals, to more sophisticated graphics terminals, to scientific workstations with hard copy graphics output.

Local Facilities

The steadily declining size and cost of computers with greater capabilities have dramatically increased the computational power that can be made readily available to the individual scientist as well as to a small group of scientists. Increasing personnel costs require that new tools be developed and used to improve productivity, and the availability of powerful, low-cost personal computers allows greater personnel efficiency and improves utilization of overall resources. The dramatic ways in which computer power "on demand" change the way people work and think have become apparent in industry and national laboratories. However, many nuclear scientists do not yet have access to this technology. Therefore, we recommend that:

> RECOMMENDATION 2: **DOE/NSF supported research scientists with a demonstrated need should have access to a local computation/communication environment at least equal to that of a personal workstation.**

This computation/communication environment can be satisfied by a variety of options ranging from personal-scale machines to an adequate fraction of a VAX-scale facility. The local needs of many theorists can be satisfied by a variety of personal workstations with appropriate software now available commercially. Although large experimental groups and laboratories may have the flexibility in their research support to purchase such equipment as needed, theorists and smaller experimental groups generally do not. Additional funds are required to provide this needed equipment for these DOE/NSF supported nuclear scientists in order to take advantage of the emerging opportunities for enhanced productivity that personal-scale machines can generate.

Data acquisition and analysis systems play a central role in determining the complexity of the reactions that can be studied in a nuclear experiment. The majority of the needs of today's experiments for data acquisition, reduction, and analysis computers can be met by VAX-scale facilities supplemented by microcomputers and personal workstations. However, it is clear that the advanced technology, front-end and parallel processing systems, discussed in the next section, will be required for the generation of experiments currently on the drawing boards. It is equally clear that we must continually improve existing nuclear science computing facilities. The rapidly increasing capabilities and decreasing costs of new computing hardware allow significant upgrades to existing systems at a modest fraction of initial costs, and replacements of decade old systems near current operating and maintenance costs. Therefore, we recommend that:

> RECOMMENDATION 3: **Data acquisition and analysis systems should be upgraded or replaced in a timely manner. Upgrades leading to enhanced capability and improved productivity should be made regularly; based on past experience, replacements should occur about every 7 years in order that operation and maintenance of such systems remain cost effective.**

With upgraded data acquisition systems having enhanced computing capabilities, more efficient use of expensive accelerator beam time can result as a consequence of improved operations and on-line monitoring of experiments. This provides opportunities for increased productivity from limited accelerator schedules which are increasingly being constrained. A continued rejuvenation of such facilities will also ensure that opportunities for sophisticated and detailed studies of nuclei at new accelerator facilities coming on-line in the near future would not be limited by obsolete and unreliable computing systems.

Advanced Technology

Since data analysis time for present day experiments can exceed data collection time by more than a factor of ten, and since several advanced experiments having extremely complicated event patterns, high multiplicities, and event rates have already been approved and are being planned; it is imperative that faster and more cost effective solutions utilizing new higher density integrated-circuit technologies in distributed and parallel processing architectures be developed. The event-mode structure of almost all data from nuclear science is especially suited to multi processing but it is a problem of modest interest to computer scientists. The recent introduction of powerful microprocessors with current VAX capabilities provides an example where the opportunity and challenge now exist to develop high speed multi processor systems that would be more generally useful to the nuclear science community. The availability of such systems will avoid the projected need for more expensive mainframes and supercomputers in data acquisition and analysis systems. Therefore, we recommend that:

> RECOMMENDATION 4: **Development of multiprocessor systems and associated software designed for acquiring and analyzing event-mode data at high speed which would be of more general use should be supported.**

The development of such multiprocessor systems should proceed in close association with approved nuclear science experiments that require data collection and processing rates at or beyond the limits of existing systems in order to have strong input from the intended user community. The distributed and parallel aspects of multi processing should be considered as integral parts of the design of such systems to optimize their capabilities for both data acquisition and data analysis. The full potential of such systems needs to be explored, in both architecture and software, for efficient processing of event-mode data from nuclear science.

Communications

One of the clear benefits emerging from the on-going rapid evolution in computers and computing is the ability to improve telecommunications. Effective implementation of this capability has allowed the establishment and use of national centers with class VI or larger computers. Another emerging benefit is improved communications between people, and it is exemplified by the formation of "special interest groups" which exist in commercially available services like COMPUSERVE and THE SOURCE, and by the explosive growth of networks dedicated to particular fields, like CSNET for computer science. The advantages of improved communications within the nuclear science community are clear. Therefore, we recommend that:

> RECOMMENDATION 5: **A Nuclear Science Communications Center accessible by all members of the community should be established to improve scientific communication.**

The establishment of a Nuclear Science Communications Center which is accessible by all members of the nuclear science community will provide an effective mechanism for improving scientific communication. Initially this Center should provide electronic bulletin board, mail conferencing, and committee communications and report writing. The addition of a nuclear data and journal-article database to this Center may be a viable option and should be considered for implementation at a later time.

Coordination of Standards

The duplication of effort in many areas of computing caused by a lack of common standards has been wasteful of our scarce pool of talented personnel and other limited resources. This duplication could easily be avoided by the adoption of standards that are accepted and supported by the major experimental facilities of nuclear science.

A committee on Computer Applications in Nuclear and Plasma Science (CANPS) has recently been chartered as a technical committee of the Nuclear and Plasma Sciences Society of the IEEE. This committee has begun work in several important areas including graphics primitives and data tape interchange formats. The adoption of standard graphics primitives would greatly simplify both the transfer of software between facilities and the use of new and more powerful graphics output devices as they become available. The adoption of standards for data tapes would simplify the transfer of data between facilities, and reduce the software overhead for research groups that utilize many different facilities. Therefore, we recommend that:

> RECOMMENDATION 6: **NSF and DOE should support the activities of the CANPS committee in its efforts to define computing standards, and encourage the adoption of such standards for nuclear science.**

In some areas, such as graphics standards, it may well be that the nuclear physics community can simply agree to adopt standards such as the graphics kernel system (GKS) currently under development in the broader computing community. In other areas, such as data tape formats, our needs are probably unique. If we do not set standards, none will be established.

The development of a "standard workstation" is another area where coordinated effort could result in significant improvements in productivity. It could contribute to the efficient use of supercomputer centers, and provide a laboratory independent interface to data reduction and analysis computers. An advanced workstation could also provide a powerful tool for the development and testing of experimental equipment. The success of the Çaviar system at CERN, and recent efforts there to develop a second-generation system of this type, provide useful examples.

I. INTRODUCTION

The usage of computers and computing in nuclear science continues to evolve and expand with the rapid development of computers. With each new generation of computers, new applications are discovered and implemented. The first generation of general purpose electronic computers, using vacuum tubes, were developed for scientific computing about forty years ago. It was a significant advance over the mechanical calculator tools available at the time, but cost and reliability limited the range of applications. The second generation, the transistorized computers, followed about thirty years ago. Computers became more generally available via the formation of computer centers to share the cost of such machines. The primary application continued to be scientific computing. Development of the third generation of computers, based on small scale integrated-circuits (SSI), occurred about twenty years ago. This gave rise to a new class of "minicomputers". The relatively low cost of such machines opened the doors to new types of applications such as real-time monitoring and control of large scale equipment (e.g. accelerators and detectors), and management of data bases, as well as broadening the access to computers for scientific computing by decentralization. The fourth generation of computers, based on large-scale integrated-circuits (LSI), were built using microprocessors which first became available some ten to fifteen years ago. The new class of very low cost "microcomputers" has continued the explosive growth of computer applications, both new and old. New applications being developed include communications and graphics; older applications being expanded include intelligent instruments of all types and the introduction of scientific computing in desk-top sized computers. The early microcomputers had limited computational power. They were relatively slow and had limited information handling capability. Today, the most powerful of the microcomputers rivals the large computer of a few years ago, but costs only a small fraction of the earlier machine. This steady trend towards lower cost is anticipated to continue for some time.

Many frontiers are opened by the rapid evolution of computers, but one critical parameter, the maximum speed of computers, has recently remained relatively stable because physical limitations are being approached. As scientific computing increases in complexity, the time required to complete a computation increases unless the speed of the computer also increases. The present generation of fastest computers, known as supercomputers, are confronting limits set by the speed of light. The alternative to a single faster computer is a collection of fast computers working on a single problem. This approach, known as parallel or concurrent processing, is an area where much effort is directed today. Generally such computers cease to be general-purpose computers in that they can only process much faster those problems which are amenable to their parallel architecture, but they are multi-purpose computers in that many problems are amenable to this approach. Fortunately, a substantial fraction of nuclear science computing belongs to this category and can be adapted to take advantage of the increased speed. This capability does not come easily, and the price to be paid is increased programming effort.

With its broad range of applications, the computer has become an indispensable tool of nuclear science. However, from the smallest microcomputer to the largest supercomputer, the power of these tools must be continually developed. Making these tools useful requires that these tools as well as capable personnel be available. The trend toward lower hardware costs does not translate directly into lower computing costs since the effort to program computers has not decreased significantly. Thus, the fraction of computing resources applied towards personnel and software will continue to grow unless new approaches are developed. The historical emphasis toward tailored programs for each application is decreasing, and use of shared programs adaptable to a large variety of applications is increasing. This sharing can reduce overall computing cost where the application is not unique, but such sharing imposes constraints on the hardware and/or software.

The opportunities and challenges of using computers and computing as tools have to be combined with the opportunities and challenges in nuclear science, so that this community can produce maximal returns for the resources it manages. It is obvious that any strategy developed in this very dynamic field of research which uses tools that are evolving so rapidly must be examined and updated regularly.

The 1983 NSAC Long Range Plan for Nuclear Science clearly shows that nuclear physics is in an explosive stage of development, stimulated by new insights from quarks and quantum-chromodynamics and new evidence for the importance of other non-nucleonic degrees of freedom.

Massive theoretical efforts are required (1) to develop new descriptions of nuclear matter, the structure of finite nuclei, and nuclear reactions, (2) to produce reliable and quantitative calculations necessary for critical tests of these ideas, and (3) and to provide guidance essential to the planning of experiments which will generate truly significant data. At the same time, nuclear theory has the task, sometimes neglected but now urgent, of providing a critical evaluation and updating of the nuclear models developed over the past three decades to account for the vast amount of experimental data on nuclear structure and reactions.

As accelerator and experimental systems become more complex and more expensive, the overall design of these systems becomes critical, and detailed simulations are required to ascertain the optimal choice. In operation, the new accelerator facilities and modern detector systems also produce information at a prodigious rate, and improved methods of data acquisition, on-line analysis, and control must be developed and implemented in order to maximize the performance of these facilities, and to insure the high quality and usefulness of the resulting data. It is already not unusual to find that the computer time required for off-line analysis may exceed by an order of magnitude or more the time used for data acquisition. Therefore, it is evident that instrumentation, whether for accelerator control, data acquisition, or data analysis, must be optimized and must make full use of the latest applicable developments in computer technology to avoid gathering information of questionable value and to keep data analysis capability up with data acquisition.

II. COMPUTERS AND COMPUTING IN NUCLEAR THEORY

A traditional view of the nucleus has been one of individual nucleons interacting via two-body potentials. Indeed, this picture describes many of the general characteristics of nuclei. To achieve this goal, one has often resorted to effective potentials and phenomenological models. Still, the approach has been successful in correlating an impressive body of nuclear data. However, it has been known for some time that the long-range part of the nucleon-nucleon interaction can be well described by the exchange of mesons (primarily pions). In the last decade strong evidence has been amassed indicating that these mesons (as well as the resonances they form with the baryons) play a direct role in the description of nuclear structure and reactions. During this same era, a revolution was also taking place in particle physics. There now exists a candidate gauge theory of the strong interaction (based upon quarks and gluons) -- quantum chromodynamics or QCD. However, it is clear that exact calculations based directly upon QCD at energies relevant to nuclear physics are most difficult, if not impossible. Therefore, one must make use of secondary approaches involving experimentally observable, asymptotically free hadrons and effective interactions based upon the more fundamental theory.

A principal task of nuclear theorists is to delineate the relevant degrees of freedom needed to describe nuclear processes. At present these fall into three classes: i) baryons only, ii) baryons and mesons, and iii) quarks and gluons. The appropriateness of any one regime in describing nuclear phenomena varies according to the energy and observables. For example, nuclear structure described by the shell model involves primarily baryonic degrees of freedom, whereas deep inelastic scattering requires at least the inclusion of mesons or alternatively the consideration of quark-gluon degrees of freedom.

II.A. SUPERCOMPUTER FACILITIES

Regardless of the relevant regime for describing a particular nuclear phenomenon, one must construct a calculable model with which to confront the available experimental data. Numerical calculations provide the means of comparing theoretical model simulations with data. Advancement of our understanding of forefront nuclear physics problems requires precision calculations based upon complex, realistic mathematical models. Many dynamic variables must be treated in order to describe the nucleus in more than a rudimentary fashion. In many areas of nuclear research this can be accomplished in a timely manner only through use of contemporary supercomputer capability.

In the following paragraphs we describe some of the frontier nuclear research areas where supercomputer power is essential for significant progress.

FEW-NUCLEON SYSTEMS. Few-nucleon investigations are essential to our testing the limits of validity of the basic assumption of the traditional view of the nucleus as one composed of individual nucleons interacting via two-body forces determined from the asymptotic properties of free nucleon-nucleon scattering. Unless we can successfully calculate the properties of bound and continuum few-nucleon systems utilizing such realistic force models, we have little reason to pursue the many-body description of heavier nuclei in terms of a nucleon model assumption. Our understanding of the nature of the nuclear force in nuclear matter will have been shown to be incomplete. Questions of the size of three-body force effects have only been touched upon. Charge symmetry breaking in the helium-3/triton binding energy difference remains unexplained. The unequal photoproton and photoneutron cross sections resulting from the photodisintegration of the alpha particle remain an enigma. A first principles calculation of the breathing-mode excited state of the alpha particle remains just a goal of few-nucleon theorists. Lack of computer resources has been a major factor inhibiting rapid advancement in this area of fundamental research. Such realistic force calculations have large memory requirements, and hours of supercomputer time are needed to solve the continuum problem for the many partial waves required. However, the matrix formulation of the problems and the iterative nature of the solution algorithms make them well suited for supercomputers.

MONTE-CARLO STUDIES OF BOUND SYSTEMS. Monte Carlo calculations of the bound-state properties of physical nuclei are a vital element in our attempt to understand nuclear structure in terms of realistic two-body forces. Properties of nuclei heavier than about A=6 cannot be calculated using standard few-body techniques, whereas Monte Carlo techniques may be applicable to medium mass nuclei if sufficient computer power is available. Variational calculations, conveniently performed utilizing Monte Carlo algorithms to evaluate Hamiltonian expectation values for semi-realistic potential models, provided part of the evidence for the need for three-body forces to explain the binding energy per nucleon in systems ranging from helium to oxygen to calcium to infinite nuclear matter. Hundreds of hours of supercomputer time are needed to lower the variances of Monte Carlo estimates for realistic potential model calculations to levels where the techniques can be tested against benchmark few-nucleon calculations and where significant physics can be extracted from calculations for light nuclei. Nuclear physics is unique in that there exist competing calculational approaches valid in the same physics region. The promising technique of Green's function Monte Carlo has been successfully tested on large (256 particle) systems of bosons. However, it has foundered on the "noise" inherent in the solution of any fermion problem utilizing random number techniques. The development of new algorithms is essential if we are to apply this powerful technique to even light nuclei such as oxygen. Hundreds of hours of supercomputer time are needed for this crucial research.

STRUCTURE OF FINITE NUCLEI. Large basis shell model calculations remain essential both to fundamental theoretical studies and to the analysis of experimental data. A prime example of the former is the calculation of nuclear double-beta decay. Here one strives to determine whether the neutrino is Dirac or Majorana in character, and whether it has a mass. An example of the latter is the calculation of core polarization effects due to valence nucleons. Here one hopes to extract single particle densities by comparison with measurements of inelastic electron scattering from nearly closed shell nuclei. Such calculations can 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 supercomputers.

A full shell model treatment of just the valence nucleons in the A=154 isotope of samarium exceeds the capability of the present generation of supercomputers. The Interacting Boson Model (IBM) approximation to the shell model appears to be successful in correlating large amounts of experimental data in terms of a few collective parameters. The IBM offers hope for quantitative calculations of, for example, the neutron and proton static and transition densities in heavier nuclei. First, however, connections between the boson degrees of freedom in the IBM and the fermion degrees of freedom in the well-established shell model must be made. Even then, it is important to realize that IBM calculations for heavy nuclei will require supercomputer resources.

HEAVY-ION REACTIONS. Coupled-channel reaction calculations are used to explore the effects of the shapes of nuclei upon reaction processes. The importance of the coupling of inelastic channels in the enhancement of the fusion cross section below the Coulomb barrier is also being investigated by means of coupled-channel reaction codes. The large systems of coupled equations required to describe heavy-ion collisions have exceeded the capability of existing mainframe computers. The strong coupling of the many channels can make the convergence of the calculation very slow. Time-Dependent Hartree-Fock and three-dimensional relativistic fluid dynamic studies of the evolution of heavy-ion collisions have begun to 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 mainframe computers. Supercomputer calculations are essential to complete interpretation of the data from existing heavy-ion accelerator facilities because three- dimensional codes have yet to be used and effects such as viscosity and heat conduction are yet to be included.

Hydrodynamic calculations for 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 "side splash" was evidence for the production of compressed matter in the collision, matter which exists in condensed stars. The matter compression and fluid flow of the hydrodynamical model was found to yield such a side splash effect in the model simulations of the data. The graphical display of output from such calculations is essential to their interpretation. The real time response of a supercomputer is required.

ELECTRONUCLEAR REACTIONS. Continued research into the theoretical description of the physical process of electron scattering is crucial to proper interpretation of these data. Distorted wave calculations of radiative correction and virtual excitation (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.

The electromagnetic current densities extracted from electron scattering data are still not well understood because the single-particle currents calculated from, for example, Hartree-Fock wave functions do not conserve current. With the improvements in the calculation of nuclear wave functions, including core polarization effects, there is hope of obtaining information on meson exchange currents, delta-hole contributions, and possible relativistic effects. However, extensive improvements in the present self-consistent model wave function generation procedure is required, so that one is assured of having conserved one-body current operators; computer limitations have forced the use of current operators which are not conserved. Improvements in the treatment of exchange current contributions from short-range correlations to make the shell model wave function calculations consistent with the current density operators, necessary for proper interpretation of both elastic and inelastic electron scattering data, await supercomputer access.

NUCLEAR ASTROPHYSICS. Nuclear astrophysics investigations, aimed at comprehending the dynamics of supernovas, model the shock-explosion mechanism which follows the gravitational collapse of massive stars. The nuclear equation of state is a key element in any such calculation. Present studies are mostly limited to the spherically symmetric approximation in the direct integration of the Lagrangian equations describing the evolution of supernovas. The complexity and iterative nature of the calculation make supercomputer solution the proper approach. Realistic three-dimensional calculations are not possible on standard mainframe machines.

OTHER BARYONIC DEGREES OF FREEDOM. The pion-nucleon resonances, such as the delta (3,3) at 1232 MeV, are becoming recognized as indispensable in arriving at a quantitative description of the nucleon-nucleon interaction. The coupling of the NN channel to N-delta, delta-delta, and pi-delta-delta channels is being included in efforts to obtain a dynamical theory of the nucleon-nucleon interaction. The role of the delta in nuclear structure and reactions was discussed at length in the 1983 Long Range Plan, where it is pointed out that the delta may account for the unexpected suppression of the excitation cross sections for Gamow-Teller transitions in (p,n) reactions. The delta is known to play an important role in pion scattering and reactions. Without access to supercomputers, calculations for comparison with the available experimental data have been restricted to simplified model approximations.

HYPERNUCLEI. The study of hypernuclei adds a third dimension to our traditional proton/neutron picture of nuclear matter, that of strangeness. Utilizing a distinguishable probe, a hyperon or an s quark which truly differs from the components of normal nuclear matter, permits one to discriminate between probe-target interactions and target component correlations. One can explore fully the consequences of the developing pictures of nuclear matter, the complementarity of the descriptions in terms of baryons and mesons and of quarks and gluons. Scientists can investigate the consequences of the Pauli exclusion principle for hyperons, the nature of the hyperon spin-orbit interaction, and the unexplained narrow widths of sigma-hypernuclei. Sophisticated nuclear structure calculations are needed in each case if we are to gain real insight into each of these little understood phenomena. Even more detailed calculations will be required if we are to utilize hypernuclei to investigate the microstructure of multihadron systems. For example, how does the addition of a lambda to a deformed nucleus modify its rotational spectra? How does the addition of a lambda alter the moments of nuclei deduced from gamma transitions? Can changes in the Coulomb energy with the addition of a lambda shed light on the compressibility of the nuclear core? Can few-body calculations with realistic potentials reproduce the experimentally known underbinding of the supersymmetric A=5 hypernuclear ground state? Calculations for such systems will be even more demanding than those for conventional nuclear matter because of the additional species involved and because of the strong coupling between the lambda and sigma hyperons, which is absent in the case of normal nuclear matter.

RELATIVISTIC NUCLEAR THEORY. The apparent success of Dirac phenomenology in describing certain spin effects in nucleon-nucleus scattering has led to the search for additional relativistic effects overlooked in other areas of nuclear physics. The potentials which have been extracted from these studies contain large scalar and vector terms of opposite sign which result in significant negative energy components in the Dirac spinors, but which appear to be in reasonable agreement with those necessary to account for the properties of nuclear matter in current relativistic theory. A first order multiple scattering theory has been constructed to relate these nucleon-nucleus potentials to nucleon-nucleon potentials, but a fully relativistic theory of the latter remains to be developed. Generation of the required nucleon-nucleon interaction from the Bethe-Salpeter equation with the inclusion of the most significant meson exchanges has consumed hours of CDC-7600 time and some time on a CYBER-205 in Europe. Extensive calculations of nucleon-nucleus scattering using a Dirac momentum space code have required many hours of CRAY time.

QUARK-GLUON PHENOMENA. The current interest in searching for evidence of quark/gluon degrees of freedom in nuclei is comparable to the surge in interest generated by the discoveries of fission, fusion, the giant resonances, the pion (to mediate the nucleon-nucleon force), etc. First, are there significant consequences for nuclear structure and/or nuclear reactions due to the composite nature of the baryons and mesons? We know at low energy that assuming 6Li to be comprised of two elementary trinucleon objects leads to predictions which are inconsistent with calculations based upon the assumption that it is composed of six elementary nucleons. Experimental consequences of the composite nature of nucleons in nuclei are not so obvious. Second, can the study of nucleons in the nuclear medium lead us to understand better the nature of the quark confinement mechanism? The nuclear medium provides a unique opportunity to relax the confining force without violating the fundamental assumptions of QCD. Recent experiments involving the scattering of high-energy muons and then electrons from nuclear targets do suggest that quarks are bound differently within nuclei than in free nucleons. However, competing explanations in terms of meson degrees of freedom have now appeared.

To obtain convincing evidence for quark degrees of freedom in nuclei, it must be established that conventional theories in terms of the experimentally observed hadrons are inadequate to account for the effects seen. The experimental data must be confronted with results of complex calculations that follow from the sophisticated theoretical models of nuclear structure and reactions which incorporate our full knowledge of nucleons and mesons. These calculations, such as those described in the preceding paragraphs, will require hours of supercomputer time. Further improvements in the traditional nucleon/meson picture will clearly have to be pursued. The limits of validity of our understanding of nuclear matter in terms of the asymptotically free hadrons must be established to insure a believable signature for quark/gluon effects. It was such a careful procedure of comparison with state-of-the-art theory that led to the establishment of the need for meson exchange currents to explain neutron-proton radiative capture and the need for three-body forces to explain the ground-state binding energy systematics throughout the periodic table.

Lattice gauge calculations on small lattices have indicated a possible phase transition in pure gluon matter as one reaches energy densities that imply temperatures of 200 MeV/nucleon at matter densities near those of nuclei in their ground states. Lattice gauge calculations to date have required hundreds of hours of CRAY time. Larger lattice calculations will require thousands of hours more. Nonetheless, we must know what QCD effects might be observable, if we are to search with ultra-relativistic heavy-ion collisions for evidence of quark/gluon degrees of freedom. Fluid dynamical transport calculations of the collision process require very fine meshes (100x100x100 points) to adequately model the interaction volume. Hundreds of hours of supercomputer time will be needed to search for unique signals of the formation of a quark/gluon plasma. A monumental effort in manpower and machine time are needed, but are vital to the success of any experimental program directed toward finding such a novel form of matter.

II.B. MIDICOMPUTER FACILITIES

Clearly, not all numerical modelling requires supercomputer resources. The computational needs of nuclear theorists run the gamut from the scientific workstation to the supercomputer. In fact, midicomputers of the VAX-class have proved capable of solving problems which were not accessible to nuclear theorists a few years ago. The largest VAX-class computer (the 8600) offers a scalar processing speed greater than 10% of a CRAY-1S. Such a computer can provide the power necessary to support a major fraction of the research performed by a nuclear theory group. The sharing of such a machine offers a way to meet the computing needs of more than one research group. The DOE/NSF arrangements at LANL, MIT, ORNL, and StonyBrook (with the smaller VAX-11/780) have proved to be an effective way of meeting part of the computing needs of the associated nuclear theory groups. The local experimental groups maintain and operate the computers, because the theory group budgets are not sufficiently flexible to permit operation of a midicomputer. Where such opportunities for shared computational facilities exist, support should be provided for their acquisition, possibly also with attached processors (array processors).

Investigations suited to midicomputer solution include exploratory calculations on nuclear structure, heavy-ion collisions, pion-nucleus reactions, relativistic nucleon-nucleus scattering, and soliton models of nucleon structure. In addition, realistic models of pion-deuteron elastic scattering, simple nuclear and hypernuclear structure calculations, optical model studies of hadron-nucleus scattering and reactions, and two-body approximations to many-body theory are examples of forefront research which can be carried out on computers of this class.

II.C. PERSONAL WORKSTATIONS

The cost of personal workstations continues to drop, and the computing speed of these workstations is rapidly approaching that of midsize computers of a few years ago. Such workstations are ideally suited for: 1) attacking the smaller numerical problems of vital interest to nuclear physicists, 2) serving as sophisticated communication links to supercomputer centers when calculations become too complex and time consuming on the workstation, and 3) providing preprocessing and postprocessing for supercomputers such as screen editing and graphical display. (The NMFECC has adopted the IBM-PC as the standard personal workstation for the latter two purposes; an IBM-PC/AT would provide somewhat greater computing power in addition.) As the cost of workstations decreases and their computational capability increases, the appropriateness of a workstation for each active theoretical physicist, whose research involves computational physics, grows. The need for a personal-sized computer for each member of a research team is currently recognized by business, the programmatic areas of our national defense laboratories, and federal agencies such as the National Bureau of Standards. The personal workstation is becoming an indispensable tool for the working research scientist.

It must be understood, however, that the scientific workstation will not replace the supercomputer. There will not soon be a CRAY on a chip. Just as 20 VAX-11/780's do not provide the same quality of computing as a single CRAY-1S, the personal workstation will not replace the need for access to a supercomputer. These machines are complementary in the world of the nuclear theorist. Husbanding supercomputer access as a scarce resource necessitates making available to the computationally oriented theorist a reasonable scientific workstation environment.

III. COMPUTERS AND COMPUTING IN NUCLEAR EXPERIMENTS

Physics is based upon a foundation of experiments and experimental data. In nuclear physics our basic knowledge of the properties of particles, nuclei and their interactions is dependent, almost entirely, on complex measurements involving accelerators, and large detector systems as well as insight and ingenuity.

Computers have played a crucial role in measurements ever since they first became generally available. Most nuclear physics experiments today would not be possible without the use of computers. All phases of an experiment: design, data acquisition, data reduction and analysis, and data interpretation require them. The scale of the hardware ranges from microprocessors connected to front-end electronics on the one hand, to supercomputers which are needed to confront theoretical predictions with experimental data. In between, modern minicomputers and midicomputers are the workhorses for data acquisition, reduction and analysis.

However, we see serious shortfalls already developing in the computing capability needed to analyze experimental data in a timely and efficient manner. Some of the experiments now in the planning stages will require detector systems and computing capabilities an order of magnitude beyond the current state-of-the-art. If we are to take advantage of these new physics opportunities, a significant investment in the development of faster and more powerful computer systems will be required. This investment will lead to designs with new architectures involving parallel processors as well as software which can make effective use of this advanced technology.

In the following paragraphs, we review briefly the role of computers and computing in experimental nuclear science. The current capability and future needs in each phase of an experiment is examined. Conclusions are reached on the steps which should be taken to ensure that adequate computing capability is available to take advantage of the physics potential that exists with the new accelerator facilities now becoming available and those planned for the near future.

III.A. EXPERIMENTAL DESIGN

As detector systems in nuclear science become more complex, it is not only important to monitor and control them via computers or microprocessors, but also necessary to know what type of response would be considered normal. It is essential in the initial design of a detector system to assure a particular design is optimal. In addition this information is needed for proper calibration and monitoring of such a system after construction. Present diagnostic techniques compare the detector response to sophisticated Monte Carlo simulations. Such Monte Carlo simulations can require a considerable investment in software development and in computer time to trace all physically reasonable interactions of particles within the system. Most existing nuclear physics detectors can be well simulated on VAX-type computers, but with the advent of relativistic heavy-ion experiments having extremely complicated event patterns, faster and more powerful computer systems will be required. High energy physics codes such as EGS for electromagnetic shower simulation or GESHIA for hadronic showers require hours to produce a few events and are best used on supercomputers. Within the next several years some nuclear physics problems will require similar complexity in event simulation.

As requirements for simulation become more complex it is important to be able to draw on a library of subroutines so that features common to many systems do not have to be reprogrammed and verified for each new experiment. This will save considerable time, provide confidence that a program is correct, and allow common standards of comparison for detector performance between different detectors and laboratories. Such library routines need to be transportable, well documented, and widely distributed.

Computers are now capable of significantly aiding the electronic and mechanical design of experiments. Sophisticated front-end amplifier and discriminator hybrids for wire chambers have been designed commercially. Design aids of this type would be more generally useful in nuclear physics, particularly with the complex detectors now planned for next generation accelerator facilities. Certainly, improved graphics capability will allow the prototyping of detector construction designs on computers relieving time and expense of conventional methods.

Computers also play an important role in accelerator design. Large computer codes for machine design can be divided into three types:

1. Beam optics
2. Particle tracking
3. Magnet and cavity design

Programs of the first type are generally known and widely used. Computational times from a few seconds to a few minutes are typical. Programs of the second type are much more time intensive, but are needed for beam stability studies or more accurate transport calculations. In some instances as much as 1 week of CRAY-1 time is needed for simulations of large systems. Finally field calculations are needed for accelerator and spectrometer design; these are also quite computer intensive. Increasingly, greater demands will be placed on computational facilities of all sizes as experiments and accelerators become more complex and precise.

III.B. FRONT-END HARDWARE

In the immediate future CAMAC type systems will satisfy the majority of nuclear physics data acquisition needs. These systems, while not state-of-the-art, offer commercial availability and ease of use. CAMAC is currently being augmented by bit-slice and ECL devices that use the CAMAC backplane primarily for power and program transfer. These devices allow some decisions to be made at the front end. Nuclear data acquisition systems will range from simple multichannel analyzers and scalers to very large detector arrays. Forefront systems will demand more speed and larger event size, which will drive such systems toward distributed processing and large local memory storage. Although CAMAC systems can probably be designed to operate under most conditions, whether by adding more crates or by using a variety of bit-slice processors with very fast cycle times, problems arise in transferring this information to the host system at word rates faster than about 500 kHz. Above this rate, a FASTBUS-type system would be needed. As conceived FASTBUS is both faster (10 MHz vs 1 MHz) and wider (32 bits vs 24 bits) than CAMAC. It is also designed to facilitate distributed processing, and allows fast, complex decisions for event tags and triggers.

A substantial effort by the computer industry ia aimed toward the production of VAX-class VLSI products at the chip and board level. These single board computers (SBC's) can be used in a variety of interesting ways for data acquisition and reduction. Simple nuclear experiments may be satisfied by an inexpensive microcomputer using one of these boards. Alternatively, more complex experiments can be served by a distributed system with a large number of SBC's interconnected by a fast bus. Each SBC can have local memory and even auxiliary processors on a local bus system. Data flow and operation would be controlled by a master processor. In such a system, throughput rates can be increased dramatically by sending separate events to individual SBC's, or by dividing the event analysis into pieces for parallel processing. Prototype distributed systems of this type are under construction now. It is planned for these systems to recognize FORTRAN-77 as the programming language. Such systems will enhance the development of special purpose hardware and the use of a variety of particular algorithms. Hard-wired systems using fast logic remain available for ultimate speeds, but programmable triggers with large memory look-up tables are becoming increasingly competitive in many situations.

III.C. ON-LINE CONTROL AND EXPERIMENT MANAGEMENT

Data acquisition computers in nuclear physics are used to acquire and store data, to sample and reconstruct a fraction of the events on-line, and to control and monitor the experiment. In a modern complex experiment involving many detectors and detector elements, the computer is the most central element, insuring that the hardware is operating properly and that the measurements are accurate, reliable and meaningful.

Up to now, 16-bit minicomputers have been the standard for data acquisition. The main advantage of these computers is their relatively fast response to real-time interrupts. However, the 16-bit architecture with its restrictive address space precludes running large tasks easily even with large physical memories.

Modern 32-bit midicomputers have a very large address space which simplifies the development and operation of acquisition and analysis software. Their more elaborate prompting systems incur interrupt latencies approximately 10 times worse than the 16-bit machines. Inexpensive, large fast memories and block transfer capabilities can minimize these inadequacies in most applications. As our outdated and outmoded data acquisition hardware is replaced it is highly likely that 32-bit computers will become the new standard.

For more complex experiments which are even now in the planning stages, the expected event rates will be too high even for the fastest real-time minicomputers. Architectures allowing parallel processing offer one possible approach for more than an order of magnitude expansion of current capabilities. The availability of powerful, low-cost microprocessors and bit-slice processors as well as fast, standardized data buses provides the necessary elements for implementing such a system. Designs now being investigated connect several processors together via an asynchronous bus, with each processor handling separate events in parallel. An intelligent event handler is required to connect the parallel processor system to the CAMAC/FASTBUS front-end. The host computer manages the taping of events, graphics displays and equipment control. Time-critical jobs are handled by the parallel processors.

The operation of an experiment usually requires that a fraction of the events be analyzed on-line to ensure that the data are meaningful and that valuable beam time is being used effectively. This analysis typically involves trajectory reconstruction, histogramming, and displays, as well as the monitoring and calibration of the equipment on a frequent schedule. The host computer is also often used to set-up and control experimental equipment.

The availability of standard data bus interfaces has resulted in the use of inexpensive microprocessors to control major subsystems and thereby remove some of the load from the host computer. Dedicated modest speed applications such as experiment and beam line control and monitoring are a typical example of such use. The microprocessors are usually downloaded from the host and use the central data base for information storage. Software development using a high level language such as Pascal makes such systems rather straightforward to implement.

Optical storage devices (laser disks) hold promise of replacing magnetic tapes as a removable bulk storage medium. This new technology offers high data rates combined with compact storage at competitive costs. It is a relatively new development however, and industry-wide standards have not yet emerged. This will limit our ability to exploit these devices in the near future. It is important to exert pressure on industry to develop a set of minimal common standards as rapidly as possible.

The development of data acquisition software is costly and must be properly factored into the overall economics in evaluating options for choosing the best computer system. Real-time, multi-user, multi-tasking operating systems are available for most minicomputers, greatly reducing the effort required for the generation of complex data acquisition systems. The use of high level languages such as Fortran and Pascal should be encouraged where appropriate. Modern compilers have made their use efficient in carrying out data acquisition tasks reliably and effectively.

General purpose data acquisition software such as Q and MULTI are in general use. These codes, which originated at the national laboratories, are very well documented and supported. Other laboratories, and in particular small university groups, have found substantial economic benefits in standardizing on the same software. For physicists and graduate students who are involved in experiments at several sites the benefits of a common acquisition software are reflected in increased productivity and efficiency.

The demonstrated benefits of hardware (CAMAC, FASTBUS) and software (Q, MULTI) standards would indicate a need for developing a common standard for the event-mode recording of raw experimental data. As a minimum, basic header information could be used to define the recording mode and the character of the data stream. Modular software front-ends for the data analysis codes would then be able to easily deal with such data no matter from which facility it originated. Too much effort is expended in adapting analysis codes to the different tape formats of various accelerator facilities.

III.D. DATA REDUCTION AND ANALYSIS

Modern complex experiments require substantial computing resources for data reduction and analysis. For many experiments today analysis time is a factor of ten or more longer than the data acquisition time. The next generation of detectors will increase demands on analysis facilities by another order of magnitude. The construction of new accelerator facilities (for high energy, CW electrons and relativistic heavy-ions) as recommended by NSAC in the 1983 Long Range Plan for Nuclear Science sets a schedule of approximately 5 years to achieve such enhanced capability.

Over the past decade there has been a significant transition from accumulation of simple spectra to event-mode recording. The basic idea is to preserve as much of the raw data as possible and to leave the decision for making cuts to software control rather than to hardware. This conservative approach has been made possible by the availability of larger disks and higher density tapes at reasonable cost. Data analysis systems must be improved significantly to meet these increased needs.

Much of the computer intensive off-line analysis involves trajectory reconstruction using event-mode data. In some experiments the measured parameters exceed several hundred per event and this number is expected to grow substantially in the future. Since each event can in general be considered independent, parallel processing techniques may be used to increase throughput. It is very important to encourage and support the development of computer systems with multi-processor architectures which offer substantial increases in speed for data analysis.

Attached processors have not yet played a significant role in data analysis. To make effective use of their high speeds for specialized calculations often requires a substantial reprogramming effort which is not generally cost-effective. However, as mentioned elsewhere in this report, array processors can provide substantial enhancement to midicomputer systems for theoretical calculations. As compilers for these processors continue to improve, we believe that they will begin to be used for data reduction as well.

For the most part rather simple graphics capability (Tektronix 4010 equivalent) has served the needs of experimental nuclear physics for the past decade. This is far short of presently available technology and the capability usually found in a modern industrial research environment. In order to deal with the greater complexity of modern experiments, increased interactive graphics capabilities both for on-line and off-line applications are needed. High resolution and color are important features for dealing with large amounts of information effectively. Intelligent graphics display terminals significantly improve the man-machine interface.

This can often be more important to productivity than computer hardware alone.

A large data analysis system usually has several different types of graphics devices associated with it. In most cases the operation of each device depends on its own particular set of codes. This proliferation of similar software is unproductive. Manufacturers should be encouraged to agree on generic graphics standards which would aid the development of unified plotting packages.

Scientific workstations with VAX-11/780 capability, high quality graphics, large memories, and a substantial amount of hard disk storage will soon be available at low cost. When combined with a local area network link to common I/O facilities, such systems could prove to be a cost-effective approach for some data analysis needs. Benefits of such an approach include local autonomy and faster response times.

In data analysis, as the hardware costs have continued to decrease, expenses associated with software development have come to dominate overall costs. To continue to keep such costs manageable we must aim for more interaction of the nuclear physics community in a constructive program of software development and sharing. There is often needless duplication of effort due to lack of communication and/or information. Continued efforts must be made to develop general purpose software packages that are transportable and hardware independent. The reliance on high level languages such as Fortran and Pascal for this purpose should be encouraged.

III.E. SUPERCOMPUTER FACILITIES

The uses which experimental nuclear physicists find for supercomputer facilities parallels the usual activities of experimentalists: experiment design, data reduction, data analysis, and finally the confrontation of experimental data with theory through data interpretation. The first group of activities has been discussed in previous sections. The last activity "data interpretation", is the focus of this section. It may be carried out by either theorists or experimentalists and in either case requires about the same level of computing capability and support. In considering the pattern of usage of Class VI and similar facilities for this activity, it is important to distinguish between the primary use of a given program by the person or persons who developed it and who are familiar with its structure, peculiarities, and input requirements, and the secondary use of the same program by individuals who wish to treat the program as a "black box" for obtaining theoretical predictions.

The familiar pattern in our field is that theorists will develop computer programs for making model predictions to be compared with a small and specific set of experimental data, and that after program development has progressed to a certain point, the use of the program will be broadened as experimentalists begin to use the same program to produce predictions which can be directly compared with experimental data. We can expect to see this pattern repeated as more nuclear theorists gain access to more powerful Class VI computational facilities. In particular, in the next year or two, there will be intense use of the new facility by theorists and relatively little by experimentalists. But as appropriate programs are developed, their operation will be shifted to experimentalists who wish to make comparisons with data. When this pattern stabilizes, it can be anticipated that the experimentalist use of a given facility for data interpretation will be roughly equal to that of theorists on the same facility.

The problems of experimentalists in performing secondary calculations with a particular program will differ from those of the theorists who developed the program and performed the primary calculations with it. Experimentalists will have a greater need for documentation of the particular program, for well-commented example calculations performed with the program.

The use by experimentalists of elaborate programs operating on Class VI computers brings into sharp focus the need for a well organized program library, for careful management of the documentation associated with this library, and for good access to this documentation. The time investment of the occasional user of such programs is minimized, and the efficient and productive use of the facilities is greatly enhanced by good documentation and support. Since the majority of users will be accessing the facility from remote sites by data transmission links, the support strategy should focus on providing support, documentation, and information in this environment.

IV. COMPUTATIONAL PLANS, PAST AND PRESENT

Although computer and computing needs for nuclear science have been considered previously as a component of long range plans, as special needs by nuclear theorists, or as versatile and necessary parts of instrumentation for nuclear experiments, the computational needs for nuclear science have not been examined comprehensively as a whole. We review briefly here the computer related recommendations of previous nuclear science committees before presenting our plan.

The computational plan formulated here has as its primary objective the identification of the needs of the United States basic nuclear research program for computers and computing over the next decade within the scientific and budgetary framework of the **1983 Long Range Plan for Nuclear Science.** The computational plan is given in the following sections as six inter-related recommendations. The budgetary implications of this plan are considered following a detailed discussion of the plan.

IV.A. PREVIOUS COMPUTER RELATED RECOMMENDATIONS

The **1979 Long Range Plan for Nuclear Science** prepared by NSAC under the chairmanship of Herman Feshbach and submitted to NSF and DOE in December 1979 noted: "nuclear theory in the U.S. is hampered considerably by the lack of good computing capability required by modern theory". In referring to the National Academy of Sciences **Friedlander** Report of 1977, which stated: "Present computer facilities available to the nuclear science community are not adequate to handle all the required computation for theory", the **1979 Long Range Plan** also noted: "The problems alluded to in that report have not been resolved (indeed, if anything they have been exacerbated) and the recommendations have not been implemented."

The **1979 Long Range Plan** recommended that $ 10-15 M be budgeted for the establishment of computational facilities, possibly a national nuclear theory computer center, and that additional support be provided for the operating costs which amount to approximately $ 3 M/year. These recommendations were reinforced by the **Report of the 1981 NSAC Subcommittee on Computational Capabilities for Nuclear Theory**, chaired by Donald Robson. This report considered three levels of computational capability corresponding to: i) a "static" approach, ii) steady increases, and iii) a strong initiative. The report recommended a program between levels ii) and iii) which included the development of up to ten midicomputer centers with provision for networking to a center based on a CLASS VI computer.

Since then, the situation for computing by nuclear theorists, as noted in the **1983 Long Range Plan for Nuclear Science** written under the chairmanship of John P. Schiffer, has not improved: "Finally, important areas of our science in which U.S. nuclear theorists require capabilities for large-scale computations are being dangerously impeded by lack of adequate access to Class VI computers." The support for computational facilities provided nuclear theory during the period 1982-1984 has fallen short of the recommended program identified by the Robson committee.

With regard to computers in experimental applications, the 1979 NSAC Instrumentation Subcommittee, chaired by Gerald T. Garvey, identified a serious problem with aging instrumentation, particularly at universities. As a specific example, it found that the median age of real-time data acquisition computer systems at all university nuclear science laboratories was 9.0 years. The old machines, besides being unreliable and expensive to maintain, did not have the capabilities of newer machines of much lower cost. To alleviate this situation, the subcommittee recommended, as its highest priority, that equipment funds be diverted to address this problem.

In the Report **of the 1983 NSAC Instrumentation Subcommittee** written also under the chairmanship of Gerald T. Garvey, it was noted that the situation on computer systems has improved greatly since the problem was identified by the earlier instrumentation subcommittee, but that a potential new problem exists, i.e. that data analysis time has grown much faster than had been anticipated even very recently. The subcommittee indicated its concerns for the consequences of this trend and then recommended: "Support for development of computer systems with multiprocessor architectures especially suited to the task of processing and analyzing event-mode data from nuclear experiments." It also foresaw the advantages of standardization and the need to begin to plan for new complex detectors, and concluded that: "There is a definite need for small-scale systems which could be widely replicated as well as for a smaller number of large systems."

IV.B. A COMPUTATIONAL PLAN FOR NUCLEAR SCIENCE

Previous sections of this report have shown the varied uses of computers and computing in the many areas of nuclear research. Computers and computing at every scale, from personal computers to supercomputers, simple text terminals to high resolution 3D color graphics, 1200 baud modems to high speed satellite network links, local area networks, attached processors, symbolic processors, etc., are all being applied to extend the frontiers of our knowledge of the properties and interactions of nuclei and nuclear matter. The computer has become an indispensable tool of nuclear science. Computing tools are now at the heart of every nuclear physics laboratory.

The computational plan for nuclear science presented here acknowledges a central role for computers and computing in nuclear research. The opportunities, challenges and needs of the community identified above are recognized and incorporated in the considerations which follow. The plan is presented in the form of six recommendations which should not be viewed as separate recommendations. Instead, they are inter-related components of an overall plan to maintain and improve the current and future computational capabilities of nuclear science. Success in achieving the objectives of this plan will require an increase in the quantity and the quality of computational support. Additional financial resources are clearly important but organizational improvements also play a critical role. This plan can be characterized by the following points: **Supercomputer Facilities; Local Facilities; Advanced Technology; Communications; and Coordination of Standards**. They are discussed in detail in the following sections.

IV.B.1. Supercomputer Facilities

NSAC has previously identified the need for supercomputer access by both theoretical and experimental nuclear physicists. We have reviewed the results of previous surveys of nuclear theorists regarding their needs for computational support. We have also reviewed and analyzed the requests made to the DOE Division of Nuclear Physics for NMFECC supercomputer time in FY'85 and FY'86. Finally, in the time available for preparation of this report, we have conducted our own limited telephone inquiries concerning the supercomputer requirements in theoretical and experimental nuclear research. The available evidence leads to two conclusions.

1. Progress on some of the most significant problems in nuclear physics relating to the **1983 Long Range Plan for Nuclear Science** requires access to a supercomputer of CLASS VI or larger. The requests submitted to DOE for supercomputer time document an immediate need for approximately one-half of a CLASS VI machine (Cray 1 equivalent).
2. The research needs of theoretical and experimental nuclear science for supercomputer power are expected to grow by a factor of 5 over the next five years. As experience and expertise of the community with supercomputers develops, more realistic (and therefore more ambitious) problems will be attacked. Thus, the need for supercomputer resources will grow rapidly.

Therefore, we recommend that:

> RECOMMENDATION 1: **Access to the equivalent of one-half of a class VI computer should be made available as soon as possible. The agencies should prepare to meet the growth of this need, estimated to be about a factor of 5 within the next five years.**

The effective utilization of this computer time requires that certain standards be met by the centers supplying such class VI time and by the connecting networks. The supercomputer centers should provide a responsive, user-friendly interactive environment 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, good documentation, mass storage facilities, and a system for flexible allocation of resources are all key elements of any useful computer center.

The fact that DOE has already established two class VI computer centers (NMFECC and FSU), and the expectation that NSF will establish about seven supercomputer centers for scientific computation is an impressive commitment to such computing capability. However, this proliferation of centers can create problems for users and can be wasteful of this scarce resource. The lack of commonality of many computer systems causes much duplication of effort. The difficulties associated with moving codes from one type of computer to another or in using a program developed on a different machine reduces productivity. These problems can be minimized by concentrating nuclear physics computing at a small number of supercomputer centers.

Reliable network connections to the centers should be provided with a guaranteed low error rate. The present NMFECC network provides an excellent existing model (24 hour availability, 1200 baud dial-in lines, 4800 and 9600 baud land-line links to large users, satellite links to major nodes), and should be upgraded as communications technology improves. In addition, with the anticipated formation of a number of networks, it would be desirable if the number and variety of networks are minimized and interconnected.

Appropriate local computing connections to the networks at the remote user sites should be provided. Other local computing resource needs will range from personal computers which could function as smart terminals, to more sophisticated graphics terminals, to scientific workstations with hardcopy output. Allocation of additional funds for such equipment are part of the local computing facility needs discussed next.

IV.B.2. **Local Facilities**

Local facilities are essential components of the computational tools of nuclear science. The computer is at the heart of every nuclear physics laboratory. The bulk of nuclear science computing has been provided by these facilities which have ranged from minicomputers and VAX-scale machines to mainframes. Such facilities, particularly those located at universities, provide opportunities for "hands-on" experience which is especially useful for the training of students. It is widely acknowledged that the VAX-scale machines have had great impact on the education of recent generations of students. More recently, personal-scale machines with substantial computational power have begun to be adopted and used in applications that include but also extend beyond the traditional areas. Many of these new applications enhance productivity; they include accessing supercomputer centers, improved communications, etc. Key attractions of local facilities are the convenience and its responsiveness to local needs.

The steadily declining size and cost of computers with greater capabilities have increased dramatically the computational power that can be made readily available to the individual scientist as well as to a small group of scientists. Increasing personnel costs require that new tools be developed and used to improve productivity. The availability of powerful, low-cost computing tools increases personnel efficiency and improves utilization of overall resources. The dramatic ways in which computer power "on demand" changes the way people work and think have become apparent in business and engineering. However, many nuclear scientists do not yet have access to this technology. Therefore, we recommend that:

> **RECOMMENDATION 2: DOE/NSF supported research scientists with a demonstrated need should have access to a local computation/communication environment at least equal to that of a personal workstation.**

This computational environment can be satisfied by a variety of options ranging from personal-scale machines to an adequate fraction of a VAX-scale facility. The local needs of many theorists can be satisfied by a variety of personal workstations ranging from personal computers to scientific workstations. Personal computer systems are already available costing less than $ 3 k each. Scientific workstations with existing VAX-type capability are anticipated to be available soon. Although large experimental groups and laboratories may have the flexibility in their research support to purchase such equipment as needed, theorists and smaller experimental groups generally do not. Additional funds are required to provide this needed equipment for these DOE/NSF supported nuclear scientists in order to take advantage of the emerging opportunities for enhanced productivity that personal-scale machines can generate.

Data acquisition and analysis systems play a central role in determining the complexity of the reactions that can be studied in a nuclear experiment. The majority of the needs of today's experiments for data acquisition, reduction, and analysis computers can be met by VAX-scale facilities supplemented by microcomputers and personal workstations. However, it is clear that the advanced technology, front-end and parallel processing systems, discussed in the next section, will be required for the generation of experiments currently on the drawing boards. It is equally clear that we must continually improve existing nuclear science computing facilities. The rapidly increasing capabilities and decreasing costs of new computing hardware allow significant upgrades to existing systems at a modest fraction of initial costs, and replacements of decade old systems near current operating and maintenance costs. Therefore, we recommend that:

> RECOMMENDATION 3: **Data acquisition and analysis systems should be upgraded or replaced in a timely manner. Upgrades leading to enhanced capability and improved productivity should be made regularly; based on past experience, replacements should occur about every 7 years in order that operation and maintenance of such systems remain cost effective.**

With upgraded data acquisition systems having enhanced computing capabilities, more efficient use of expensive accelerator beam time can result as a consequence of improved operations and on-line monitoring of experiments. This provides opportunities for increased productivity from limited accelerator schedules which are increasingly being constrained. A continued rejuvenation of such facilities will also ensure that opportunities for sophisticated and detailed studies of nuclei, at upgraded accelerators coming on-line in the near future, would not be limited by obsolete and unreliable computing systems. New funding and support are also required for new major accelerator facilities such as CEBAF. The upgrade of fully saturated systems in theoretical programs (attached processors are cost effective for some problems) should not be overlooked.

IV.B.3. **Advanced Technology**

Since present data analysis time for experiments can exceed data collection time by more than a factor of ten, and since several advanced data acquisition systems for experiments having extremely complicated event patterns, high multiplicities and event rates are being planned, it is imperative that faster and more cost-effective solutions utilizing new higher-speed, higher-density integrated-circuit technologies in distributed and parallel processing architectures be developed for data acquisition and analysis.

The distributed and parallel processor architectures have the potential to be easily adapted to the demands of a wide range of experiments, from those with simple requirements by the use of a single processor, to those with many subsystems by the use of many processors, and/or to those with high multiplicities and event rates by the use of large farms of processors, provided that the appropriate software is also developed. The low cost, the enhanced capability, and the wide range of applications in nuclear data acquisition and analysis together offers the opportunity for a broad based acceptance of such multiprocessor systems. Such a possibility could encourage the formation of new standards. Over the next decade, a multiprocessor data acquisition and analysis system, applicable to both small and large scale experiments in nuclear science, has the greatest potential for major improvement in capability as well as for savings since these applications will continue as primary uses of computers and computing in nuclear science. Therefore, we recommend that:

> **RECOMMENDATION 4: Development of multiprocessor systems and associated software designed for acquiring and analyzing event-mode data at high speed which would be of more general use should be supported.**

The development of such multiprocessor systems should be in close association with the more complex nuclear science experiments now being planned in order to have strong input from the user community. The distributed and parallel aspects of multiprocessing should be considered as integral parts of the design of such systems to optimize their capabilities for both data acquisition and data analysis. The full potential of such systems needs to be explored, in both architecture and software for efficient processing of event-mode data from nuclear science.

Current efforts in this direction exist at several national laboratories but they have limited goals for application in a restricted environment and are minimally supported. To spur greater efforts and encourage competition in this direction with broader goals and a more open environment, additional resources should be made available for the development of several systems complementing corresponding efforts in high energy physics with the intention that these be used in the more complex experiments now being planned. A modest increase in support and in effort here can lead to the design and implementation of powerful systems that can be applied much more widely. The success of such efforts will be essential in order to satisfy, at a reasonable cost, projected data acquisition and analysis needs of the nuclear science community in the next decade.

IV.B.4. **Communications**

One of the clear benefits emerging from the on-going rapid evolution in computers and computing is improved telecommunications capability. Effective implementation of this capability has allowed the establishment and efficient use of national computer centers with class VI computers (See Recommendation 1). Another emerging benefit is better communications between people, and it is exemplified by the formation of "special interest groups" which exist in commercially available services like COMPUSERVE and THE SOURCE, and by the explosive growth of networks dedicated to particular fields, like CSNET for computer science. The advantages of improved communications, which would lead to a better exchange of information within the nuclear science community, are clear. Therefore, we recommend that:

RECOMMENDATION 5: **A Nuclear Science Communications Center accessible by all members of the community should be established to improve scientific communication.**

The Nuclear Science Communications Center accessible by all members of the community should be established as soon as possible, after the appropriate proposals are received and reviewed. A fraction of the community would already have access to the networks of supercomputer centers and so will be well equipped to communicate also with this Center. The remainder of the community will need modest additions to their local computing facilities for this purpose. Initially this Center should provide electronic bulletin board, mail conferencing, and committee communications and report writing. Such improved information exchange capability will increase productivity.

An example of the implementation of such a center might be a small, dedicated computer with modest computational power and memory, but with considerable disk storage space, with connections to several existing networks. In this case, an initial capital cost of about $ 200 k, an operating cost of about $ 50 k/yr, and a commercial communications cost of about $ 200 k/yr is estimated. A cost effective way to implement such a system would be to include it under the management of a group already responsible for several VAX-class machines. Initially, the use of an existing network with modest transmission rates, e.g. BITnet, can reduce communications costs considerably.

After the Nuclear Science Communications Center is well established, it may be desirable to expand the function of the Center to provide a readily accessible data base for nuclear science. It could furnish a central location providing access to recent data posted by experimentalists and to files of evaluated data provided by data compilation groups. Appendix E discusses this data base option in more detail.

IV.B.5. **Coordination of Standards**

As hardware costs have continued to decrease, expenses associated with software which is personnel intensive have come to dominate overall costs. This situation is not expected to change, and in order to keep computing costs manageable, greater efforts must be directed towards a constructive program of software development and sharing. There is often needless duplication of effort due to inadequate information and/or standards. The creation of a Nuclear Science Communications Center will improve the quality of information exchange (see Recommendation 5). Additional development of standards, both hardware and software, will facilitate the use of computers and computing via increased sharing of programs.

The duplication of effort in many areas of computing caused by a lack of common standards has been wasteful of our scarce pool of talented personnel and other limited resources. By reducing the proliferation of independent solutions to common computing problems, much progress can be made which would be of inestimable value to the entire community. This duplication could easily be avoided by the adoption of standards that are accepted and supported by the major experimental facilities of nuclear science.

A committee on Computer Applications in Nuclear and Plasma Science (CANPS) has recently been chartered as a technical committee of the Nuclear and Plasma Sciences Society of the IEEE. This committee has begun work in several important areas including graphics primitives and data tape interchange formats. The adoption of standard graphics primitives would greatly simplify both the transfer of software between facilities and the use of new and more powerful graphics output devices as they become available. The adoption of standards for data tapes would simplify the transfer of data between facilities, and reduce the software overhead for research groups that utilize many different facilities. Therefore, we recommend that:

> RECOMMENDATION 6: **NSF and DOE should support the activities of the CANPS committee in its efforts to define computing standards, and encourage the adoption of such standards for nuclear science.**

In some areas, such as graphics standards, it may well be that the nuclear physics community can simply agree to adopt standards such as the graphics kernel system (GKS) currently under development in the broader computing community. In other areas, such as data tape formats, our needs are probably unique. If we do not set standards, none will be established.

The development of a "standard workstation" is another area where coordinated effort could result in significant improvements in productivity. It could contribute to the efficient use of supercomputer centers, and provide a laboratory independent interface to data reduction and analysis computers. An advanced workstation could also provide a powerful tool for the development and testing of experimental equipment. The success of the Caviar system at CERN, and recent efforts there to develop a second-generation system of this type, provide useful examples. These are discussed in detail in Appendix F.

IV.C. FUNDING IMPLICATIONS OF PLAN

The computational plan presented here has as its primary objective the identification of the needs of the United States basic nuclear research program for computers and computing over the next decade within the scientific and budgetary framework of the **1983 Long Range Plan for Nuclear Science.** Although some identifiable long standing computing needs are not yet fulfilled, some recently fulfilled, and others anticipated to be fulfilled in the next several years our basic philosophy has not been to focus on these but to focus on specific opportunities and challenges which would contribute substantially to maintaining and improving the overall current and future computational capabilities for nuclear science. These opportunities and challenges have been discussed in the previous sections.

In attempting to estimate the costs associated with the implementation of this Plan, it is useful first to gauge the current and already planned levels of expenditures for computers and computing, then to determine which aspects are included as parts of these existing plans so as to identify specific additional costs, and the potential impact of any such new expenditures.

We recognize that it is impossible to determine precisely the overall costs associated with computers and computing due to the variety of sources of funds and the variety of possible expenditures. Considering only agency funds, our estimate is that for FY'84, 8-10% ($ 16-20 M) of all nuclear science expenditures are in computers and computing. Of this total, substantially more than half (60-65%) is the cost for personnel providing computing services at local facilities, i.e. computer maintenance, system managing, system programming, consultants, etc. This fraction varies greatly from one institution to another. It is larger at the national laboratories which service large user groups. Approximately 10% goes toward the acquisition of computer time, which is split about equally between mainframes mostly at universities and supercomputers at national centers such as NMFECC (the hours assigned to nuclear science have been converted to dollars for our present considerations). It should be noted that universities often subsidize local mainframe usage at substantial levels and the equivalent dollars are not included here. Only about 15-20% is spent towards the acquisition of new computer equipment often with matching grants at universities. The remaining 10% or so goes towards miscellaneous operating items such as maintenance contracts, materials and supplies, etc.

Full implementation of the computational plan presented here would require an increase in the overall level of funds for computers and computing as well as a change in the distribution of these funds in the next few years. A significant fraction of this increase is already part of existing agency plans. The remaining increases are relatively modest and consistent with the intent of the **1983 Long Range Plan for Nuclear Science**, but would have a substantial short term and a long term impact. Table I shows the estimated current expenditures and the suggested steady state expenditures associated with the recommendations here. The funding implications of each of the six inter-related recommendations are considered separately in the following sections.

IV.C.1. Supercomputer Facilities

The new initiatives for scientific computing developed recently by NSF and DOE show plans for substantial improvements in supercomputer availability and in communications to such centers. The present NSF OASC plans are to establish three supercomputer centers in FY'85 and four more in the following two years, and to provide access to these centers using existing networks in the short term and to establish SCIENCEnet in the long term. The present scientific computing plans at DOE are to: install a CRAY-XMP at NMFECC (November, 1984) as an interim computer; install a class VII machine there as a "permanent" machine in FY'86; initiate a supercomputer center at FSU in FY'85; and expand and upgrade communications.

If nuclear science is assigned the fraction of computer time at these centers proportional to the fraction of overall funds from DOE and NSF, then the long pressing need for supercomputer power would be satisfied provided that the support requirements discussed earlier are met by the

Table I

Present Recommendations	Current Level of Expenditures	Recommended Level of Expenditures
1. Supercomputer Facilities	[a]	[a]
2. Workstations	0	$ 1 M/yr [b]
3. Upgrades/Replacements of Existing Systems	$ 3-4 M/yr [c]	$ 5.25 M/yr [d]
4. Advanced Technology	$ 0.6 M/yr [e]	$ 2 M/yr
5. Communications	0	$ 0.25 M/yr [f]
6. Coordination of Standards	0	0 [g]

Notes:

[a] Supercomputer facilities are assigned to separate offices both at DOE, with the Office of Scientific Computing (OSC), and at NSF, with the Office of Advanced Scientific Computing (OASC).

[b] This is a recommended initial expenditure of 3 year duration after which annual expenditures of about $ 0.4 M/yr for upgrades and/or replacements would be needed for this as a component of local facilities.

[c] Estimates of current expenditures on computers by DOE at local facilities are relatively imprecise.

[d] These costs do not include contributions by universities or special discounts by vendors, both of which can be substantial.

[e] Present efforts at several laboratories are funded out of existing operating budgets.

[f] Communications costs of $ 200 k/yr can be reduced considerably by the use of existing, low bandwidth networks, e.g. BITnet. An initial capital cost of about $ 200 k is required. Other needs are included under recommendation 2.

[g] This involves primarily a redirection of some existing personnel and expenses associated with meetings.

supercomputer centers and that the distribution of this capability can be made equitably to those researchers with the greatest needs. Then the percentage of resources directed towards supercomputers in nuclear science would increase dramatically from the current 5% range, even though these funds are assigned out of separate offices.

IV.C.2. Local Facilities

The funding implication of each of the two recommendations on local facilities is considered separately. A reasonable estimate of the cost to provide the recommended computational environment to those members of the nuclear science research community with a demonstrated need would be an initial expenditure of \$ 3 M over some reasonable period, e.g. 3 years, or about \$ 1 M/yr. This would provide personal workstations, averaging \$ 10 k each, and access to telecommunications networks for those with inadequate local facilities.

The local experimental needs at new facilities are typically included with the cost of such facilities. What is often not adequately considered is the replacement and/or upgrades of existing machines. Many data acquisition computers at the major national laboratories were acquired in excess of a decade ago and are not only obsolete, but are also unreliable and expensive to maintain, with maintenance costs approaching replacement costs for an equivalent machine. Such systems should be retired and replaced by modern, more reliable, and more cost effective machines now available. Upgrades of existing systems, particularly for data analysis but including acquisition of attached processors (array processors) which are cost effective in some theoretical applications, should be carried out on a regular basis. An informal survey carried out recently by this Subcommittee showed that the present inventory of computers in nuclear science for data acquisition and analysis, and for local theoretical computation has a replacement value in excess of \$ 35 M. If annual upgrades averaging 5% of acquisition costs and replacements over a ten year cycle are assumed, then annual expenditures of \$ 5.25 M/yr for these purposes are required. Upgrades extend the useful life and provide a cost effective approach to increasing the computational power of existing computers. Its inclusion in the above calculation of annual expenditures allows the ten year replacement cycle used. Otherwise, a seven year replacement cycle would be more appropriate.

The expenditure of \$ 1.4 M in FY'84 by NSF for computers, supplemented by matching funds from universities and special discounts by vendors, have been helpful in improving the local computing environment for grantees. However, to satisfy all the equipment needs for upgrades and replacements would require an increase of total DOE/NSF expenditures in this direction, estimated to be about \$ 3-4 M/yr in FY'84, to the level of \$ 5.25 M/yr. Total incremental computing costs, however, would not grow in proportion since operation and maintenance costs for the newer machines are lower.

IV.C.3. Advanced Technology

The primary resource needed for development of multiprocessor systems is talented and committed people who should be provided with adequate financial resources for the effort. Each development effort would require about a half dozen full-time personnel dedicated to this effort over at least two years to begin to produce a useful system. Two to three such efforts will cost about $ 2 M/yr, or about 10% of current expenditures in computers and computing in nuclear science. The total existing efforts are estimated to be about a third of the recommended scale and they are too small to produce the needed systems. The suggested scale of investment might be expected to produce tangible results in both hardware and software in a few years, and if such a system is adopted for more general implementation by the community, it should begin to have significant impact in three to five years by providing several orders of magnitude increase in local computational capability without a corresponding increase in cost.

IV.C.4. **Communications**

The costs for accessing supercomputer centers by remote users via networks have been included with the overall costs required for the operation of such centers. Using the NMFECC network as a reasonable existing standard, one finds that network costs average about 6% of the total. Both agencies have accepted this approach to service remote users and should be applauded for its implementation.

An example of the implementation of the Nuclear Science Communications Center might be a small, dedicated computer with modest computational power and memory, but with considerable disk storage space, with connections to several networks. In this case, an initial capital cost of about $ 200 k, an operating cost of about $ 50 k/yr, and a commercial communications cost of about $ 200 k/yr is estimated. A cost effective way to implement such a system would be to include this system under the management of a group already responsible for several VAX-class machines. Initially, the use of an existing network with modest transmission rates, e.g. BITnet, can reduce communications costs considerably.

IV.C.5. **Coordination of Standards**

The required support for the on-going CANPS group to facilitate the definition and implementation of standards for computing is primarily qualitative in that such efforts are considered worthwhile and would bear fruit. Beyond that the required resources are primarily human. The willingness of individuals to contribute their time and effort, plus travel and associated expenditures, would be enhanced if there is the perception that substantial benefits will come from the use of these new standards, and encouragement that the new standards will be accepted. The impact of any new standards would likely not be felt quickly, but over a period of years as such standards are adopted and implemented in the community. Improvements in efficiency and productivity, though always difficult to gauge, would be readily apparent as enhanced sharing of information and reduced duplication of effort occur.

V. CONCLUSIONS

As pointed out in previous long range plans, instrumentation reports, and emphasized again here, progress in nuclear science has become increasingly dependent on the use of computers and computing. Adequate computational capability has become essential to all aspects of theoretical and experimental nuclear research. Most major contributions to the field now require extensive data processing, computation, or both, and the expectation is that future advances will require even greater computational capability. Computers and computing have become necessary tools of our field, and rapid advances in nuclear science in the future will only develop with significant additional computer resources.

This report has identified the need for and estimated the scale of the uses of class VI machines for the remainder of this decade. It projects that most theorists and some experimentalists will require access to supercomputers in the next few years, and this issue is addressed in recommendation 1. In addition, it assumes that research scientists must use the opportunities being introduced by current developments in computer related technologies in order to maximize their individual productivity. The availability of personal workstations, addressed in recommendation 2, and continual upgrades and replacements of existing local facilities, addressed in recommendation 3, are components of this goal. In addition, developments in higher-density, higher-speed integrated circuits are projected to provide the only viable solution to the many problems connected with data acquisition and data analysis at proposed new accelerator facilities. The low per unit cost of powerful new microprocessors and bit-slice processors will allow the development and wide spread application of multiprocessor architectures which can provide substantial increases in computing speed available to most nuclear scientists without a corresponding increase in cost. We have suggested that this opportunity be fully exploited in recommendation 4.

Of comparable but often neglected importance, however, is improved communications within the field. Another rapid technological evolution is in progress which significantly enhances telecommunications. Since the capability to electronically transfer data, video, and mail at kilobits to megabits per second now exists, the nuclear science community should take advantage of this opportunity. Increased information flow will result in greater productivity, not only in providing more information quicker, but also in reducing the duplication of effort for many common tasks. These opportunities are addressed in recommendations 5 and 6.

New insights and deeper understanding will come as more detailed and more sophisticated calculations and experiments, requiring greater computational capability, are carried out. Fortunately, present progress in computer and computing technology can provide the means to satisfy these new requirements. The nuclear science community must take the challenges offered by new developments in the field, and use the opportunities provided by the rapid, on-going evolution in computers and computing toward the study of the properties and interactions of particles and nuclei.

REFERENCES

1. A Long Range Plan for Nuclear Science,
   A report by the DOE/NSF Nuclear Science Advisory Committee,
   John P. Schiffer, Chairman,
   December, 1983.

2. A National Computing Environment for Academic Research,
   Prepared under the direction of Marcel Bardon,
   by the NSF working group on computers for research,
   Kent K. Curtis, Chairman,
   July, 1983.

3. Report of the 1983 NSAC Instrumentation Subcommittee,
   Gerald T. Garvey, Chairman,
   December, 1983.

4. Report of the Panel on Large Scale Computing in Science and Engineering,
   Sponsored by DOD and NSF in cooperation with DOE and NASA,
   Peter D. Lax, Chairman,
   December, 1982.

5. Report by the Subcommittee on Computational Capabilities for Nuclear Theory to the DOE/NSF Nuclear Science Advisory Committee,
   Donald Robson, Chairman,
   December, 1981.

6. A Long Range Plan for Nuclear Science,
   The DOE/NSF Nuclear Science Advisory Committee,
   Herman Feshbach, Chairman,
   December, 1979.

7. A Report on Computer Needs for Nuclear Theory,
   William M. MacDonald, Chairman,
   July, 1979.

8. Report of the 1979 Instrumentation Subcommittee of the DOE/NSF Nuclear Science Advisory Committee,
   Gerald T. Garvey, Chairman,
   July, 1979.

Appendix A

Letter from Alvin W. Trivelpiece and Marcel Bardon

Department of Energy
Washington, D.C. 20545

June 18, 1984

Dr. John P. Schiffer
Chairman, DOE/NSF Nuclear
Science Advisory Committee
Division of Physics
Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60437

Dear John:

I would like to request that the DOE/NSF Nuclear Science Advisory Committee conduct a study of the needs of basic nuclear research in the United States for computers and access to computing capability. Please complete the study and submit a final report on this study to the Department of Energy and the National Science Foundation by October 1984.

The recently submitted "Long Range Plan for Nuclear Science" (December 1983) points to exciting new research opportunities for nuclear physics in the next decade. It is imperative that a logical plan be developed to provide the nuclear physics community with convenient access to the modern computers and computing which will adequately support the research in these exciting new areas. The previous nuclear physics computer study (Report by the Subcommittee on Computational Capabilities for Nuclear Theory to the DOE/NSF Nuclear Science Advisory Committee, December 1981) concentrated on the computer needs of nuclear theory. The requested study should give full consideration to all nuclear physics computer needs--theoretical and experimental. A proposed charge for this study is as follows:

CHARGE

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 or 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.

We recognize that studies such as the one requested here demand a large effort by a large number of scientists. We appreciate this effort, and we hope that the discussions and advice that result from this study will significantly enhance the ability of the nuclear physics community to capitalize on the exciting opportunities that appear at the nuclear frontier.

We would like to be put on the agenda for a meeting of NSAC in the near future to discuss this task with you.

Sincerely,

/s/

Alvin W. Trivelpiece
Director, Office of
Energy Research
Department of Energy

ORIGINAL SIGNED BY:
DR. MARCEL BARDON

Marcel Bardon
Acting Assistant Director for
Mathematical and Physical Sciences
National Science Foundation

Appendix B

Participants and Agenda at Subcommittee Meetings

First Meeting: Washington D.C.

(August 31-September 1, 1984)

Agenda

Aug. 31	9:00	DOE Plans	J. Decker
	9:45	NSF Plans	D. Vanbellegham
	10:30	Parallel Processing	C. Maples
	11:15	Software Portability	C. Maples
	11:30	Scientific Workstations/ Local Area Networks	R. Judd
	12:00	Lunch	
	13:00	NMFECC	D. Fuss
	13:45	Experimental Needs	D. Hensley
	14:30	Theoretical Needs	S. Koonin
	15:15	Discussions	
	17:00	Adjourn	
Sept. 1	9:00	Discussions	
	13:00	Adjourn	

Participants

Committee:
L.S. Cardman (Illinois)
B.F. Gibson (LANL)
S.E. Koonin (Caltech)
W.M. MacDonald (Maryland)
H.H. Chen (Irvine)
E.V. Hungerford III (Houston)
S.B. Kowalski (MIT)

Liaison:
E.T. Ritter (DOE)
H.B. Willard (NSF)

Others:
J. Decker (DOE)
D.C. Hensley (ORNL)
R. Judd (LLL)
P.J. Riley (NSF)
L.C. Welch (ANL)
members of the public
D. Fuss (LLL)
E.W. Hoffman (LANL)
C. Maples, Jr. (LBL)
D. Vanbellegham (NSF)

Second Meeting: Washington D.C.

(September 26, 1984)

Agenda

Sep. 25	19:00	Dinner Discussions
	22:00	Adjourn
Sep. 26	9:00	Discuss/organize Final Report
		Discuss/organize Working Groups
	12:00	Lunch
	13:00	Discuss/organize Interim Report
		Discuss/organize October Meeting
	17:00	Adjourn

Participants

Committee:

L.S. Cardman (Illinois)
J.G. Cramer, Jr. (Washington)
E.V. Hungerford III (Houston)
W.M. MacDonald (Maryland)
H.H. Chen (Irvine)
B.F. Gibson (LANL)
S.B. Kowalski (MIT)

Liaison:

A.J. Baltz (DOE)
P.J. Riley (NSF)
H.B. Willard (NSF)

Third Meeting: Houston

(October 14-16, 1984)

Agenda

Oct. 14	19:00	Dinner	
	20:30	Work on Interim Report	
	23:00	Adjourn	
Oct. 15	9:00	Organizational Remarks	H. Chen
	9:15	Theoretical Considerations:	
		Computer use at UNH	J. Dawson
		FPS/MAX at ORNL	M. Strayer
		Class VI applications	D. Strottman
		Computing at ANL	B. Wiringa
	12:00	Lunch	
	13:15	Experimental Considerations:	
		Next generation data acquisition	N. Di Giacomo
		Data acquisition at national labs	E. Hoffman
		High energy physics needs	I. Gaines
		Compatibility	L. Welch
	17:00	Dinner	
	19:00	Discussion of Plan	
	21:00	Work on Interim Report	
	23:00	Adjourn	

Oct. 16	9:00	"Other" considerations:	
		Attached processors at TRL	W. Thompson
		BNL nuclear data center	S. Pearlstein
	10:30	Discussion of Recommendations	
	12:00	Lunch	
	13:15	Discussion of Recommendations	
	14:00	Working Groups (Theory/Experiment)	
	15:30	Working/Planning Session	
	17:00	Dinner	
	19:00	Work on Interim Report	
	22:00	Adjourn	

Participants

Committee:
L.S. Cardman (Illinois)
J.G. Cramer, Jr. (Washington)
E.V. Hungerford III (Houston)
W.M. MacDonald (Maryland)
H.H. Chen (Irvine)
B.F. Gibson (LANL)
S.B. Kowalski (MIT)

Liaison:
A.J. Baltz (DOE)

Others:
J. Buchanan (Rice)
J.F. Dawson (New Hampshire)
I. Gaines (FNAL)
W.W. Kinnison (LANL)
S. Pearlstein (BNL)
D. Strottman (LANL)
L.C. Welch (ANL)
J. Clement (Rice)
N.J. Di Giacomo (LANL)
E.W. Hoffman (LANL)
C. Maples, Jr. (LBL)
M.R. Strayer (ORNL)
W.J. Thompson (TRL)
R.B. Wiringa (ANL)

Written contributions (not at Houston):
D.H. Fitzgerald (LANL)
D. Grisham (LANL)
T. Hunter (LANL)
R.E. Mischke (LANL)
D. Robson (Florida)
K.E. Schmidt (NYU)
J.W. Sunier (LANL)
S. Greene (LANL)
W.C. Haxton (Washington)
R.A. Jameson (LANL)
J.M. Moss (LANL)
M.E. Schillaci (LANL)
E.B. Shera (LANL)

Fourth Meeting: Irvine
(December 5-7, 1984)

Agenda

Dec. 5	19:00	Dinner
	20:30	Discussion of Plan/Recommendations
Dec. 6	9:00	Present Expenditures in Computing
	10:00	Discussion of Interim Report
	11:00	Overview of Plan
	12:00	Lunch

	13:00	Working Groups (Theory/Experiment)
	15:00	Discussion of Plan/Recommendations
	17:30	Adjourn
	18:00	Dinner
	20:00	Discussion of Plan/Recommendations
Dec. 7	9:00	Discussion of Plan
	10:00	Drafting of Recommendations
	12:00	Lunch
	13:00	Discussion of Recommendations
	14:30	Other matters (letter to Schiffer on Interim Report, next meeting)
	15:00	Writing/Editing of Report
	17:00	Adjourn

Participants

Committee members:
H.H. Chen (Irvine)
B.F. Gibson (LANL)
S.E. Koonin (Caltech)
W.M. MacDonald (Maryland)
J.G. Cramer, Jr. (Washington)
E.V. Hungerford III (Houston)
S.B. Kowalski (MIT)

Liaison:
A.J. Baltz (DOE)
H.B. Willard (NSF)

Fifth Meeting: Washington D.C.
(January 24-25, 1985)

Agenda

Jan. 24	9:00	Discussion of Plan/Recommendations
	13:00	Lunch
	14:00	Writing/Editing of Report
	19:30	Adjourn
Jan. 25	9:00	Discussion of Plan/Recommendations
	10:00	Editing of Recommendations
	12:00	Lunch
	13:00	Discussion of Recommendations
	14:30	Writing/Editing of Report
	16:00	Adjourn

Participants

Committee:
L.S. Cardman (Illinois)
J.G. Cramer, Jr. (Washington)
E.V. Hungerford III (Houston)
W.M. MacDonald (Maryland)
H.H. Chen (Irvine)
B.F. Gibson (LANL)
S.B. Kowalski (MIT)

Liaison:
A.J. Baltz (DOE)
H.B. Willard (NSF)
E.T. Ritter (DOE)

Appendix C

Working Groups

I. Interim Report	B.F. Gibson (chairman)
	J.G. Cramer, Jr.
	W.M. MacDonald
II. Theoretical Aspects of Final Report	W.M. MacDonald (chairman)
	B.F. Gibson
	S.E. Koonin
III. Experimental Aspects of Final Report	L.S. Cardman (chairman)
	J.G. Cramer, Jr.
	E.V. Hungerford III
	S.B. Kowalski

Appendix D

Interim Report

INTERIM REPORT

of the

1984 NSAC SUBCOMMITTEE ON COMPUTERS AND COMPUTING

December 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 V (Los Alamos National Laboratory)
Ed V. Hungerford III (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:

John P. Schiffer (Argonne National Laboratory)

Liaison:

Anthony J. Baltz (Department of Energy)
Enloe T. Ritter (Department of Energy)
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 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 class VI computer related issues in the short term, i.e., two to three years, should not be interpreted as favoring class VI computers 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 class VI computer time, we have reached the following conclusions:

1. There is an immediate need for supercomputer power at about one-half a class VI machine.
2. Rapid growth of this need to the level of about 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 recommendation:

Access to the equivalent of one-half of a class VI computer should be made available as soon as possible.

Effective utilization of this computer time will require that certain minimum standards be met by any computer center supplying such class VI time and by the connecting network. In addition, remote user requirements should not be ignored. Therefore, we make the following related recommendations:

1. Any class VI computer center used for nuclear science computing should provide a responsive, user-friendly interactive environment to insure efficient use and effective algorithm development. Real-time turnaround must be adequate for effective code testing. Batch processing for long complex programs, software support, sophisticated graphics capability, knowledgeable consultants, and good documentation are all essential elements.

2. Reliable network connections to such centers should be provided. To service remote users adequately, network costs should be borne by the

centers. The present National Magnetic Fusion Energy Computer Center network provides an excellent existing standard.

3. Appropriate local computing connections to the networks at remote user sites will be needed. Other local computing resource needs will range from personal computers which could function as smart terminals, to more sophisticated graphics terminals, to scientific work-stations with hardcopy graphics output. Allocation of funds for such equipment should be based on the evaluation by the funding agencies of specific proposals.

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 (NMFECC); 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 advanced scientific computer centers by NSF beginning with the funding of three such centers in FY'85. It also heard reports on the operation of the NMFECC, 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 class VI computer related issues in the short term, i.e. two to three years, should not be interpreted as favoring class VI computers 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 class VI computer 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 essentially 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 computers are essential to progress. When such computers 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 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 supercomputer capability.

A selection of frontier nuclear physics research areas where supercomputer power would have a significant impact 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 supercomputers. 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 class VI computer 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. The nuclear equation of state is a key element in any such calculation. Present studies are based primarily upon spherically symmetric hydrodynamic calculations and involve direct integration of the Lagrangian equations describing the evolution of supernovae. Future studies especially involving non-sphericity are needed. 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 midsized computers presently available for data analysis.

We have re-examined the current nuclear science class VI 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 NMFECC class VI computer time in FY'85 and FY'86, and we have conducted our own limited telephone inquiries of class VI computer requirements in theoretical and experimental nuclear research. The available evidence indicates:

1. There is a need for immediate access to computer 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 the level of about 1 1/2 class VI machines.

As the experience and expertise of the community with these computers advances, more realistic (and therefore more ambitious) nuclear physics problems will be attacked. Thus, the need for class VI computer resources will grow rapidly.

The supercomputer needs of nuclear science can be met by combining its needs with those of other areas of physics so as to provide the best possible facilities. One should aim to buy the best and then use it to the maximum feasible extent. Supercomputer 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 class VI computers 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. Responsive, user-friendly interactive systems are 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, good documentation, mass storage facilities, and a system for flexible allocation of resources are all key elements of any useful computer center.

Reliable network connections to the centers should be provided. To service remote users adequately, network costs should be borne by the centers. The present NMFECC network provides an excellent existing standard: 24 hour availability, 1200 baud dial-in lines, 4800 and 9600 baud land-line links to large users, satellite links to major nodes.

Appropriate local computing connections to the networks at the remote user sites should be provided. Other local computing resource needs will range from personal computers which could function as smart terminals, to more sophisticated graphics terminals, to scientific workstations with hardcopy output. Allocation of additional funds for such equipment should be based on the evaluation by the funding agencies of specific proposals.

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. The equivalent of approximately 1/2 of a class VI machine is required in FY'85. We estimate that this will grow to about 1 1/2 class VI machines in two to three years. A useful supercomputer center will have to provide a responsive, interactive, non-saturated environment along with a wide range of support services. Reliable network connections for remote computing are essential. Local user requirements at the remote sites should 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.

Appendix E

A Common Database for Nuclear Science

There has been considerable effort in certain areas of nuclear physics at developing databases. However, these have been of limited use because of the way in which this activity has been handled. It is not centralized; it is focused in only a few areas (e.g. neutron cross sections) which only minimally reflect the interests of the community; it is normally accessible only by mail over a two week time scale; and much of the information is not available in a computer compatible format.

The creation of the Nuclear Science Communications Center provides an important opportunity to develop the growing body of nuclear information into a truly up-to-date database which is immediately accessible to the nuclear community. A large fraction of the nuclear science community will already have access to and will be well equipped to communicate with the Center. Therefore, the same communications capabilities can, in principle, provide ready access to tabular or graphical information, for example, energy levels, cross sections, angular distributions, or extracted phase shift potentials.

We suggest that, concurrent with the creation of the Nuclear Science Communications Center, a committee be constituted to investigate the possibility of using the Center as the basis for a comprehensive nuclear science database and to make recommendations to NSAC concerning the costs and benefits of establishing such a capability.

Appendix F

Useful New Standards for Nuclear Science

I. Standard Graphics Interchange Format

The evolution of graphics devices has created an unwieldy situation in this area. Many types of devices are available (raster displays, vector displays, dot-matrix printers, pen plotters, etc.), all of which are used to present the same graphical information. Each device requires a different set of commands and often a different data structure for its operation. Each plotting program is required to generate appropriate commands and data for each graphical device. This situation is inflexible, inefficient, and programmer-intensive.

There is a well known strategy to solve this problem which has been successfully applied to similar situations. This is to develop and use an intermediate interchange format for graphics information. It can be done as follows: (1) design a graphics interchange "language" (or command/data structure) which permits the user to command all graphical operations of interest; (2) revise all plotting programs so that they produce only an intermediate file using this graphics language; and (3) develop software drivers which will convert the intermediate file to appropriate graphical output for each device of interest.

There presently exists no such standard graphics interchange language. In fact, a debate is raging over whether to include enhancements such as color, three-dimensional representations, coordinate rotations, etc. into the language. It would, of course, be desirable to adopt a graphic interchange standard for nuclear physics which is compatible with an universal standard (if such is ever agreed upon). However, even an interim standard for our field would be of great value. Moreover, if an universal standard is adopted by the entire industry, a single program could be written to convert an intermediate file to the new standard.

We suggest that the CANPS group begin efforts to define and maintain (if necessary) a graphics interchange standard which would be employed in all graphics software developed for use with data acquisition and analysis systems at the national laboratories and which would be recommended for general use by the nuclear science community.

II. Standard Data Tape Interchange Format

A similar situation has arisen within nuclear science in data collection and analysis. Each laboratory has developed data collection and analysis capabilities essentially independently of other facilities so that each laboratory has its own format for data tapes. This becomes a problem when data collected at one laboratory is analyzed at another, as is often the case for national facilities. With each occurrence of this situation, it is necessary to write a program for converting data tapes from one format to the other.

Again the interchange language strategy can solve the problem. In this case one must design and use an intermediate data interchange format. This could be done by: (1) designing a data interchange "language" which accommodates all (or most) of the structures commonly used on data tapes; and (2) requesting each laboratory to develop (as need demands) programs for converting data tapes to and from this intermediate form. If the language/structure is well designed, existing and new data collection programs will gradually evolve to use the interchange format directly rather than requiring conversions.

We suggest that the CANPS group begin efforts to define and maintain an interchange standard for data tapes which would be used with any data acquisition software developed at the national laboratories and which would be recommended for general use by the nuclear science community.

III. Standard Workstation

Electronic devices designed for data communication and display are in a period of explosive growth and development. In nuclear physics such devices are widely used in data collection/analysis and scientific computing to communicate with a variety of computer systems. To specify a "standard" workstation device of this kind now might appear to be counterproductive. However, experience at the **European Center for Nuclear Research** (CERN) and at the **National Magnetic Fusion Energy Computing Center** (NMFECC) has shown that there are significant advantages to defining a "standard workstation" at major computer centers, and supporting software (co-editors, graphics programs, etc.) can be targeted on the specified workstation. This does not preclude the use of other devices by the sophisticated user, but it provides the entry-level user with an easy (and usually inexpensive) route to the best-developed software.

Because of the rapid evolution of technology, any standard defined for a workstation will be obsolete within a few years. Therefore any standard which is created must be maintained and periodically modified or enhanced. We suggest that the CANPS group begin efforts to define and maintain a workstation standard for the nuclear science community for data collection and analysis, and graphics.