AMD-Vol. 75

COMPUTATIONAL MECHANICS

ADVANCES AND TRENDS

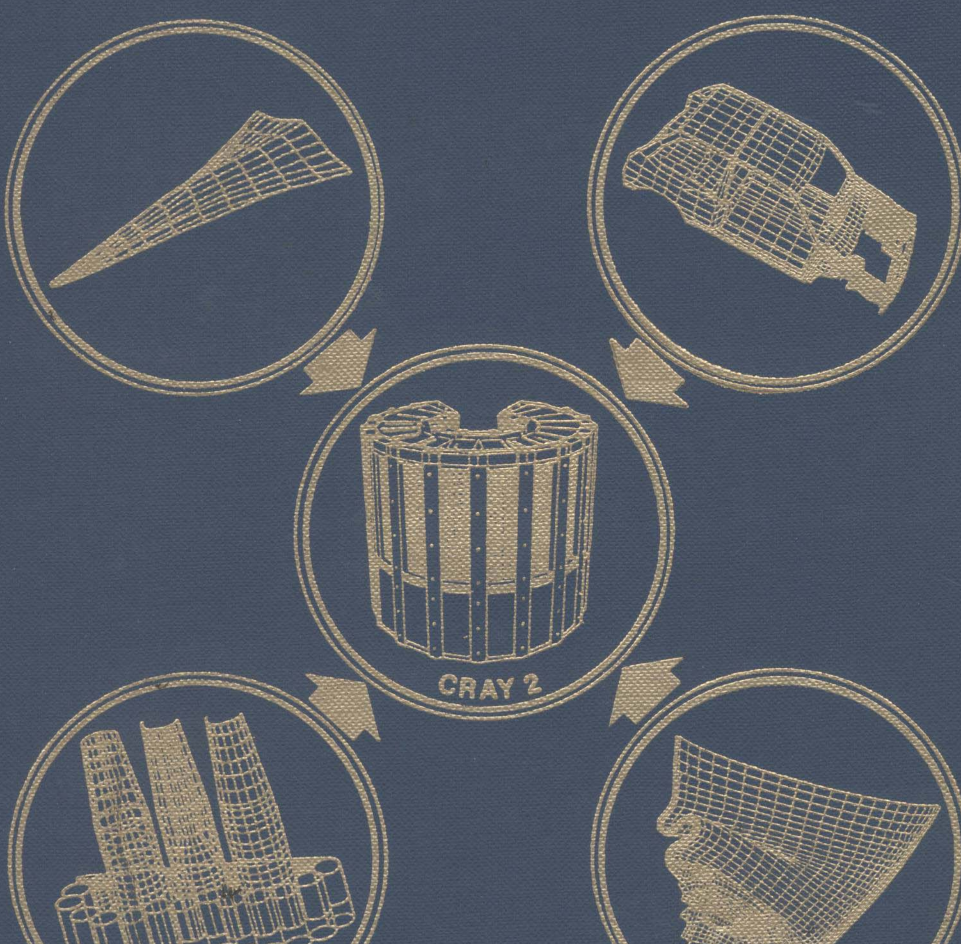

AMD-Vol. 75

COMPUTATIONAL MECHANICS — ADVANCES AND TRENDS

presented at

THE WINTER ANNUAL MEETING OF
THE AMERICAN SOCIETY OF MECHANICAL ENGINEERS
ANAHEIM, CALIFORNIA
DECEMBER 7–12, 1986

sponsored by

THE APPLIED MECHANICS DIVISION,
THE PRESSURE VESSELS AND PIPING DIVISION,
THE COMPUTER ENGINEERING DIVISION, ASME

edited by

AHMED K. NOOR
GEORGE WASHINGTON UNIVERSITY
NASA LANGLEY RESEARCH CENTER

THE AMERICAN SOCIETY OF MECHANICAL ENGINEERS
United Engineering Center • 345 East 47th Street • New York, N.Y. 10017

Library of Congress Catalog Card Number 86-072277

PREFACE

Computational mechanics has recently emerged as a new discipline involving a coordinated blend of insightful modeling of mechanical phenomena with the development of computational methods. It is a rapidly evolving discipline that is having a major impact on the development of mechanics as well as on its application to modern technology.

Computational mechanics is being used today in the numerical simulation of a wide variety of physical phenomena, and in the utilization of these numerical simulations in the analysis and design of engineering systems. The fields of application of computational mechanics are very broad and include aerospace, nuclear, automotive, shipbuilding, offshore oil, petroleum, building and electronic industries. In each of these application areas computational mechanics has transformed much of the theoretical mechanics and abstract science into practical and essential tools for a multitude of technological developments which affect many facets of our life. Computational mechanics is also used in a *heuristic mode* to explore complex physical phenomena which are difficult, if not impossible, to study by alternate approaches. Examples of these situations are provided by the study of vortex dynamics and evolution of material response across many levels of material scales (from microscopic to structural level).

The importance of computational mechanics to our national economy is reflected in the increased funding of different government agencies to both the computing equipment and the research on various aspects of computational mechanics. Of particular significance to the research community is the establishment of: a) the Supercomputing Research Center in Lanham, Maryland (a division of the Institute of Defense Analyses); b) the Numerical Aerodynamic Simulation (NAS) project at NASA Ames; and c) the National Science Foundation (NSF) five supercomputer centers. The Supercomputing Research Center and the NAS Project are described in these proceedings. The primary objective of the NSF supercomputer centers is to increase the access of advanced computing resources for the scientific and engineering research community. The centers, which became operational in 1986, are jointly sponsored by NSF, state governments, participating universities, and industrial participants. The five centers and their hardware configurations are:

1. National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign - CRAY X-MP/24, 32 MW SSD to be upgraded to a CRAY X-MP/48 with 128 MW SSD, CTSS operating system.
2. John von Neumann Center for Scientific Computing, Princeton, New Jersey - CYBER 205, VSOS 2.2 operating system to be upgraded to ETA-10, with a UNIX operating system.
3. Center for Theory and Simulation in Science and Engineering, Cornell University, Ithaca, New York - IBM 3084 QX to be upgraded to an IBM 3090-400, FPS 264 (4), FPS 164 (2), FPS 164/MAX (1), VM/CMS operating system.
4. San Diego Supercomputing Center, University of California, San Diego - CRAY X-MP/48, CTSS operating system.
5. Pittsburgh Supercomputer Center, Advanced Computing for Engineering and the Sciences, near University of Pittsburgh and Carnegie-Mellon University - CRAY X-MP/48, 128 MW SSD, COS operating system.

The widespread availability of these advanced computing facilities will result in increasing the scope of the engineering problems considered, as well as the level of sophistication of the models used. This, in turn, can lead to enhancing our understanding of complex physical phenomena.

A number of recent studies have attempted to identify the trends, needs and opportunities in computational mechanics and related fields. Of special interest to the research community are the following six reports:

1. "Large-Scale Computing in Science and Engineering," sponsored by DOD and NSF in cooperation with DOE and NASA, P. D. Lax, et al., NSF.
2. "A National Computing Environment for Academic Research," NSF Working Group on Computers for Research, K. K. Curtis, et al., NSF, 1983.
3. "Computational Modeling and Mathematics Applied to the Physical Sciences," National Academy of Sciences Committee on the Applications of Mathematics, W. C. Rheinboldt, et al., National Academy Press, 1984.

4. "Computational Mechanics - A Perspective on Problems and Opportunities for Increasing Productivity and Quality of Engineering," Computational Mechanics Committee, National Academy Press, 1984.

5. "Solid Mechanics Research Trends and Opportunities," ed. by J. R. Rice, AMD Vol. 70, American Society of Mechanical Engineers, New York, 1985.

6. "Future Directions in Computational Mathematics, Algorithms and Scientific Software," W. C. Rheinboldt, et al., Society of Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 1985.

There is a growing interest in taking advantage of the opportunities provided by the new and emerging computing systems to make computational mechanics more widely applicable and economical for use by industry. This can lead to increasing the productivity of engineers and the quality of future engineering systems and devices. Future progress of computational mechanics requires a synergestic interaction between the various disciplines constituting computational mechanics (viz., theoretical and applied mechanics, approximation theory, numerical analysis, and computer science), as well as effective mechanisms for setting research priorities and goals, and for rapid dissemination of research results. Also, efficient technology-transfer mechanisms are needed to make algorithms, software, and specialized hardware developed in one field of computational mechanics available to other computational mechanics fields.

As a means of communicating recent and projected advances, and as a step towards developing a coherent plan for research in computational mechanics, a six-session, two-day symposium entitled "Future Directions of Computational Mechanics," was organized as part of the ASME Winter Annual Meeting held in Anaheim, California on December 10-11, 1986. The symposium was sponsored by the Applied Mechanics Division, the Pressure Vessels and Piping Division, and the Computer Engineering Division.

The twenty-four papers contained in this volume document clearly the strides made in a number of aspects of computational mechanics, identify some of the anticipated industry needs in this area, discuss the opportunities provided by new hardware and parallel algorithms, and outline some of the current government programs in computational mechanics. The topic headings in the symposium are largely represented by the section headings of this volume, namely:

1. Opportunities provided by new hardware and parallel algorithms
2. Industry needs and applications of computational mechanics
3. Recent advances in computational structural and solid mechanics
4. Recent advances in computational fluid dynamics
5. Future directions of commercial software systems, and
6. Government programs in computational mechanics.

The fields covered by this symposium are rapidly changing and if the results and anticipated future directions are to have maximum impact and use, it is imperative that they reach the workers in the field as soon as possible. This consideration led to the decision to publish these proceedings prior to the symposium.

The editor is indebted to the many individuals who contributed to the symposium. In particular, special thanks go to the authors of the papers for their contributions; to the members of the Committee on Computing in Applied Mechanics, the Pressure Vessels and Piping Division, and the Computer Engineering Division for their helpful suggestions and encouragement; and to Mrs. Mary Torian for her secretarial support throughout this project.

Ahmed K. Noor
George Washington University Center
NASA Langley Research Center
Hampton, Virginia

CONTENTS

1. OPPORTUNITIES PROVIDED BY NEW HARDWARE AND PARALLEL ALGORITHMS

THE SUPERCOMPUTING RESEARCH CENTER

P. B. Schneck
The Supercomputing Research Center
Lanham, Maryland

ABSTRACT

The Supercomputing Research Center will pursue research to advance the state-of-the-art in supercomputing. As a secure center of excellence in parallel processing, the SRC will carry out research and development on national security programs, as well as unclassified research in parallel processing algorithms and systems. Participation of the academic and industrial communities is recognized as essential, and the Center plans a program to work with visiting scholars. The SRC will host unclassified conferences and share research results through the scientific literature and through participation in technical conferences and workshops. This summer, with participants from the academic community, the Center will hold a second workshop on research directions. The activities envisioned for the Center include basic research in parallel algorithms, operating systems, languages and compilers, and computer architectures.

THE SUPERCOMPUTING RESEARCH CENTER

The Supercomputing Research Center was established to advance the state-of-the-art in supercomputing. A group of 100 researchers in computer science, mathematics, and electrical engineering will be seeking new knowledge in algorithms, systems, and their applications. Additionally, prototype machines with new architectures will be developed at the Center. The research mission is intended to increase computational capability for national security applications.

INTRODUCTION

The Supercomputing Research Center, presently in Lanham, Maryland, is the newest division of the Institute for Defense Analyses, a federal contract research center. In several years the Center will move to the Maryland Science and Technology Center and a new permanent facility of about 100,000 square feet. We are in the growth phase and are currently recruiting permanent staff.

As the word "supercomputing" in our name indicates, we are concerned with not only hardware systems and machines, but with all aspects of high speed computation. This is reflected in our parallel thrusts in algorithms, software systems, applications software, as well as hardware and architecture.

About half of the Center's work will be classified, leading to improvements in computational capability for national security applications. The other half will be unclassified work and will be shared with the computer science community. Our plans are to publish actively in peer-reviewed journals and to encourage interchange and work with the computing community.

A BRIEF RESEARCH PROSPECTUS

The Center will address a broad set of activities in supercomputing. It will be organized along "matrix management" lines. That is, teams will be formed to work on projects (prototypes) and at each project's conclusion the team members will return either to their original or another group. This approach allows us the flexibility to bring different people together depending on the skills and requirements needed for a specific project and offers the flexibility of changes in assignment at the conclusion of a project. The organization is described in further detail below.

Algorithms Research

This activity will have the "flavor" of a university department in theoretical computer science. Its members will primarily be computer scientists and mathematicians. They will be investigating fundamental issues in parallelism including: computational complexity, algorithm transformations from sequential to parallel, algorithm representation, information theoretic computation, and data structures. The research in this group will not necessarily be directed toward a specific problem or application.

Systems Research

This group will concern itself with areas that are primarily "software" in nature. The first area is that of operating systems. Fundamental issues in operating systems for parallel processors are the focus. Problem partitioning, memory concurrency, resource allocation, synchronization, and load leveling are examples. The second area is compilers and programming languages. Here the question of representation again occurs. A key research topic is the design of a new language or schema for representing algorithms on parallel processor systems. Techniques for compiling parallel programs and for obtaining efficient execution are also of interest in this area. The third area is that of performance measurement. Two facets in performance measurement are: first, the ability to predict performance of an application on a given machine. This will allow us to size the resources necessary to accomplish a task. Second, the ability to predict the performance of a machine in given situations. For existing machines, this knowledge enables selection of the appropriate computing system. Schwartz (1) has suggested numerous possibilities for parallel processor architectures - so many that no one can try all of them. A good performance prediction mechanism may shed insight into the value of architectural alternatives and yield guidance toward selection of which machines should be prototyped.

Applications

The applications mission is somewhat more developmental. Although not generic research, the presence of applications at SRC is essential. The nature of applications serves to provide a goal against which accomplishments can be measured; to provide a context in which to develop machines and systems; and to provide a backdrop and a common culture for the in-house research community.

The presence of in-house applications will allow both a push and a pull of SRC research in the following sense: applications "pull" research in that their characteristics create needs for computation and cause research to flow in certain directions; they "push" research when the algorithms and approaches used in applications suggest further methods or systems. As an example, the massively parallel processor, built by Goodyear Aerospace for NASA, was originally conceived to do image processing problems for NASA satellites and has seen use in numerical weather prediction.

Projects

The project group is permanently staffed as a support organization which will receive influxes of team members associated with specific project activities. Several teams may exist simultaneously. Each team will be concerned with a specific implementation and will negotiate within the project group to obtain the supporting resources necessary.

Our plans are to build small scale prototypes of new architectures for computing systems in-house. These prototypes will be built with off-the-shelf components and will be of small size (as measured by number of processors or number of components), will be operated at conservative clock speeds, be conservatively engineered (e.g., air cooled), and be easily modifiable.

After the small scale prototype has been completed, tested with a skeletal operating system and compiler, it will be placed in use to verify the design concept, correct any design problems, and seek omissions. If the prototype machine remains interesting at this stage, then consideration will be given to contracting for an (industrial strength) prototype. This machine would be planned to be larger in number of processors, have a smaller cycle time, be more aggressive in packaging and components, and be equipped with a more sophisticated operating system and compiler. After delivery, this larger prototype would be used as a research tool for further development as well as be the basis for future algorithms and systems research.

EXTERNAL RELATIONSHIPS

The Supercomputing Research Center is planning a visiting scientist program for university and industrial researchers. We expect these visitors will work with us on projects of mutual interest. After clearances are obtained, we will make available our computational resources, including a current supercomputer and mainframe and access to our planned station network as well as to any prototypes which have been developed. Unclassified results during this period will be publishable by our visitors. We also expect to sponsor programs for students, some of whom will continue as employees at SRC. We will be working with faculty members in the creation of such a program.

CONCLUSION

The Supercomputing Research Center is a new research institution aimed at solving some of the fundamental problems in parallel computation. With its long-term goals, its computing tools, and its program sparked by applications, we look forward to a productive enterprise. We plan to bring together the top computer science researchers in the country from the industrial, government, and academic communities and expect to fulfill our mission of advancing the state-of-the-art of supercomputing.

REFERENCE

1. Schwartz, J.T., "A Taxonomic Table of Parallel Computers, Based on 55 Designs," Untralcomputer Note #69, New York University, 1983.

STATUS AND PROJECTIONS OF THE NAS PROGRAM

F. R. Bailey
NASA Ames Research Center
Moffett Field, California

ABSTRACT

NASA's Numerical Aerodynamic Simulation (NAS) Program has completed development of the initial operating configuration of the NAS Processing System Network (NPSN). This is the first milestone in the continuing and pathfinding effort to provide state-of-the-art supercomputing for aeronautics research and development. The NPSN, available to a nation-wide community of remote users, provides a uniform UNIX[1] environment over a network of host computers ranging from the new Cray-2 supercomputer to advanced scientific workstations. This system, coupled with a vendor-independent base of common user interface and network software, presents a new paradigm for supercomputing environments. Presented here is the background leading to the NAS Program, its programmatic goals and strategies, technical goals and objectives, and the development activities leading to the current NPSN configuration. Program status, near-term plans and plans for the next major milestone, the extended operating configuration, are also discussed.

INTRODUCTION

The Numerical Aerodynamic Simulation (NAS) Program is a major NASA effort to establish a national supercomputing center to support aeronautics research and development. The NAS Program had its genesis in 1975 at the Ames Research Center (ARC) when a small group of researchers associated with the computational fluid dynamics program proposed the development of a special purpose Navier-Stokes processing facility. The group had witnessed the rapid growth in the capability of computational methods to treat problems of practical interest in the early 1970s when computers became large enough to treat inviscid, three-dimensional transonic flows requiring the use of nonlinear forms of the governing fluid-dynamic equations. The first transonic solutions for a

---

[1]UNIX is a trademark for AT&T.

two-dimensional lifting airfoil with embedded shock waves appeared in the literature in 1970, and three-dimensional wing solutions were first obtained in 1972. By the mid 1970s computers were large enough to solve, in a practical amount of time, inviscid forms of the governing equations that provide satisfactory results for relatively complex (but aerodynamically clean) configurations operating at or near cruise conditions. However, many of the more critical design situations are difficult, if not impossible, to treat without including the effects of viscosity in a fully coupled fashion. Examples of these are inlet, engine, and exhaust flows; airframe propulsion system integration; stall and buffet; vortex enhanced lift; maneuvering loads; and performance near performance boundaries. The Reynolds-averaged form of the Navier-Stokes equations is suitable for treating most of these design problems that are dominated by viscous effects. Revolutionary advances in this technology were made in the 1970s, beginning with the investigation of shock-wave interaction with a laminar boundary layer in 1971 (1). Subsequent work has been largely limited to component simulations however, because of the large amount of computer time required. For example, 1 to 20 hours of computation were needed for all but the most basic aircraft components, even with the use of Cray-1-class computers.

Having recognized the need for more computing capability, the group at ARC began a series of Navier-Stokes processing facility studies. The first (1976 through 1979) produced preliminary configurations, a functional design, and rough estimates of cost and schedule. In September 1980, two parallel design-definition contracts were awarded. The contractors' proposals for detailed design, development, and construction were submitted for evaluation in April 1982. After an initial evaluation of the proposals, it became clear that the approach needed to be changed. First, the application and essential importance of computational aerodynamics had advanced rapidly during the previous few years. It was now deemed necessary to establish a state-of-the-art computational aerodynamics facility and also to initiate an aggressive on-going development effort to keep the facility at the leading edge in supercomputer and related technologies. Thus, the end-item NAS Project was redefined as an ongoing, continually upgraded, NAS Program. In addition, the development of supercomputers had advanced to the point that it was no longer considered necessary to directly subsidize their future development. As a result of these changes in approach, a new NAS Program Plan was developed at ARC in February 1983. (See Ref. 2 for a more detailed history of the NAS Program.)

A basic task in formulating the NAS Program Plan was estimating the computer-performance and memory-capacity requirements for computational aerodynamics. During the period of NAS Program formulation and advocacy, the state-of-the-art in computational aerodynamics progressed to the point where problems involving complex geometries could be treated with simple physics and those involving simple geometry could be treated with complex physics. The program was structured to address the fact that more powerful computers with more memory capacity are required to solve problems involving both complex geometries and complex physics. An estimate of the computer requirements (3) is shown in Fig. 1. Illustrated is the actual computer performance required, measured in millions of floating-point operations/sec (MFLOPS) and memory capacity in 64-bit words, to solve problems of varying complexity in fluid physics, and configuration model complexity. The degree of physics complexity being considered is represented by three approximations to the governing equations: nonlinear inviscid, Reynolds-averaged Navier-Stokes, and large-eddy simulation. The degree of configuration complexity is represented by airfoils, wings, and complete aircraft configurations. The performance requirements, based on a solution in 15 min of central-processor time, are compared to several existing and projected supercomputers. The comparisons show that current supercomputers can

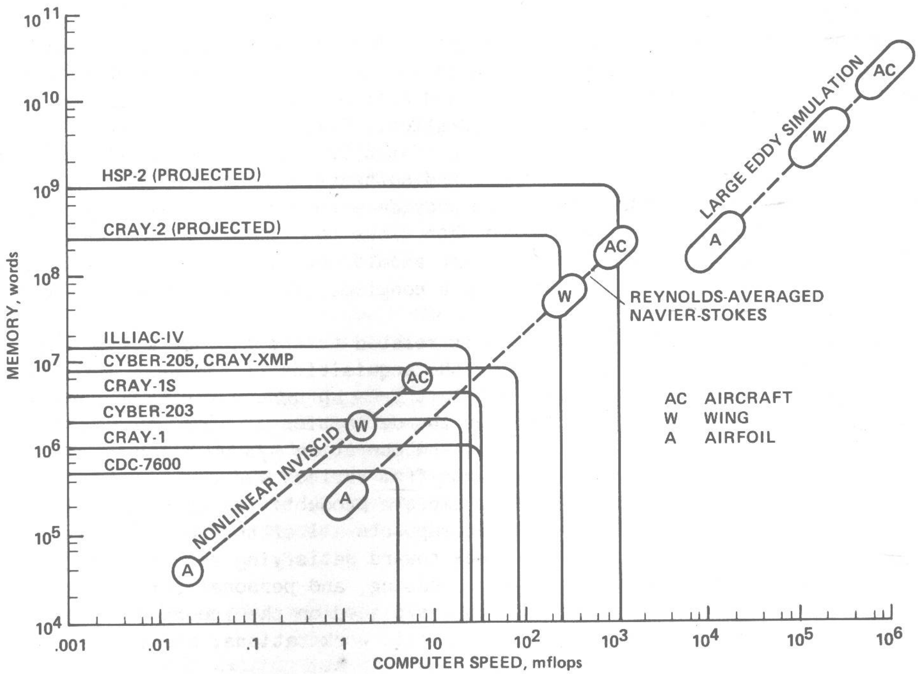

Fig. 1 Computer speed and memory requirements for aerodynamic calculations compared with the capability of various supercomputers; 15 min runs with 1985 algorithms

adequately address inviscid flows, but the computers needed to adequately address more complex Reynolds-averaged approximations to turbulent flows about full aircraft configurations will not be available until the end of this decade. Adequate treatment of large-eddy simulations approximations must wait future, more powerful computers.

NAS PROGRAM GOALS AND OBJECTIVES

The NAS Program Plan is defined with three principal goals: (1) provide a national computational capability, available to NASA, Department of Defence (DOD) and other Government agencies, industry, and universities, as a necessary element in insuring continuing leadership in computational fluid dynamics and related disciplines; (2) act as a pathfinder in advanced, large-scale computer system capability through systematic incorporation of state-of-the-art improvements in computer hardware and software technologies; and (3) provide a strong research tool for NASA's Office of Aeronautics and Space Technology.

The plan establishes an overall approach for meeting these goals that includes the design, implementation, test, and integration of the initial configuration of a large-scale computer system, called the NAS processing system network (NPSN). Future improvements to the NPSN are to be implemented by the systematic incorporation of advanced computer technologies into the system. The NAS Projects Office at ARC was given the responsibility to plan and implement the NPSN and manage its operation. The plan also called for the design and construction of a new building, the NAS Facility, to house the NPSN, the NAS

Projects Office, resident computational aerodynamics users, and associated contractors.

A thorough review of user requirements led to four key design objectives which have influenced the NPSN system architecture: The first objective was to permit the introduction of new hardware and software systems with minimum disruption to the current operating configuration. The approach to meeting this objective is reflected in the structuring of the NPSN into functional subsystems that communicate over a common hardware and software network.

The second design objective was to provide a common user level interface for the entire NPSN. By doing so, the user views the NPSN as a unified system. It has become clear that the user should not, and need not, utilize a functionally rich distributed system via a conglomerate of different operating system environments.

The third design objective, which is related to the second, was to support a multivendor environment which enables the acquisition of the best possible hardware without being constrained to select from an unnecessarily restricted set of vendors. This objective requires the definition of communication hardware interfaces, communication software, and operating system software that can be implemented on a wide range of products from various vendors.

The fourth design objective was to improve productivity by providing a comprehensive computing environment that supports all of the many computer-intensive user activities. Modern trends toward satisfying requirements for graphical display, highly interactive processing, and personal computing are taken into account. Particular emphasis was placed on the incorporation of the expected rapid emergence of powerful scientific workstations, and the importance of rapid on-line graphic display of results.

Implementation of the NPSN is based on the evolutionary strategy illustrated in Fig. 2. The strategy is to incorporate into the NPSN a sequence of successively more powerful prototype or early production model supercomputers as high-speed processors. The sequence begins with the installation of the first generation system, followed by a period of test and integration into the NPSN, and finally, production operations. The first generation is the Cray-2 which will reach production operation in July 1986. Installation of second generation will follow with operational status targeted for 1988. The sequence continues with successive generations replacing the older of the two installed systems. For example, the third generation is anticipated to replace the first in the 1989-90 time period.

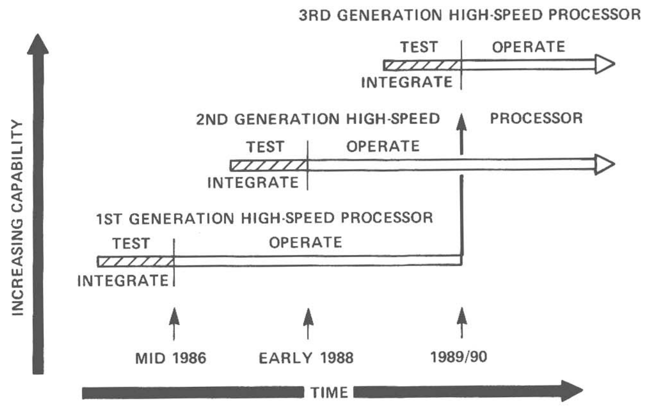

Fig. 2 Evolution of the NAS Processing System Network (NPSN)