Status and Projections of the NAS Program

Frank R. Bailey

July 1986

Status and Projections of the NAS Program

Frank R. Bailey, Ames Research Center, Moffett Field, California

July 1986

NASA

STATUS AND PROJECTIONS OF THE NAS PROGRAM

by Dr. F. Ron Bailey

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
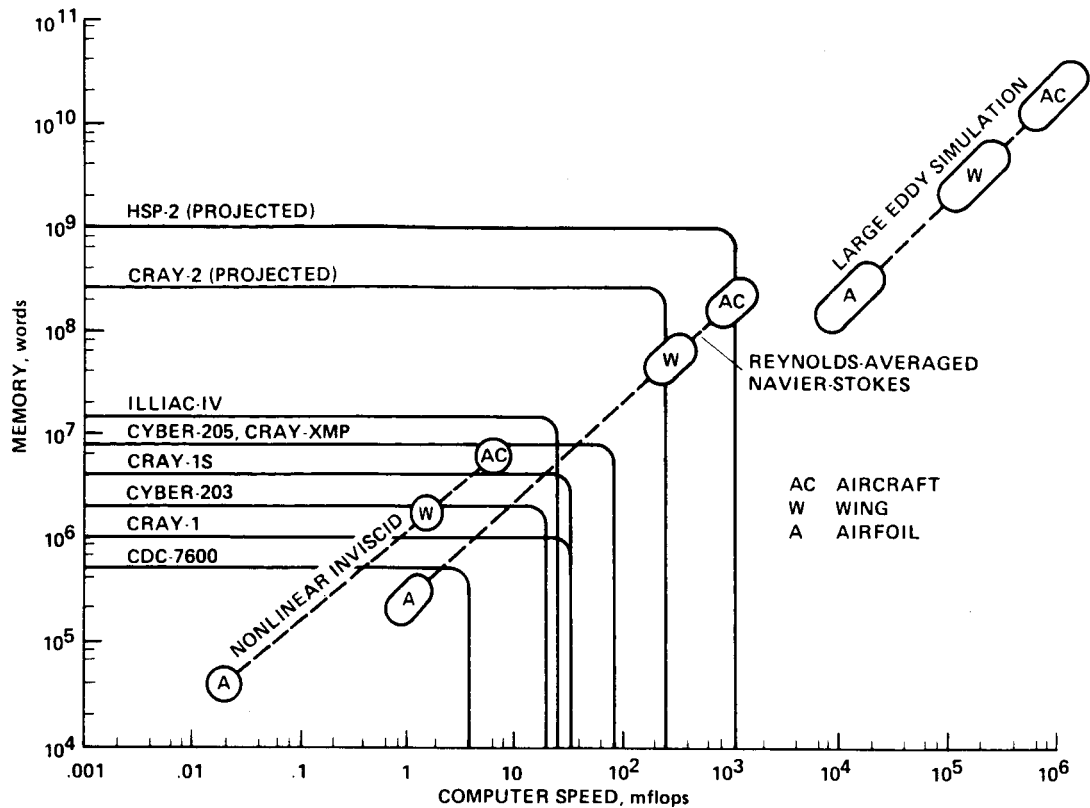
2

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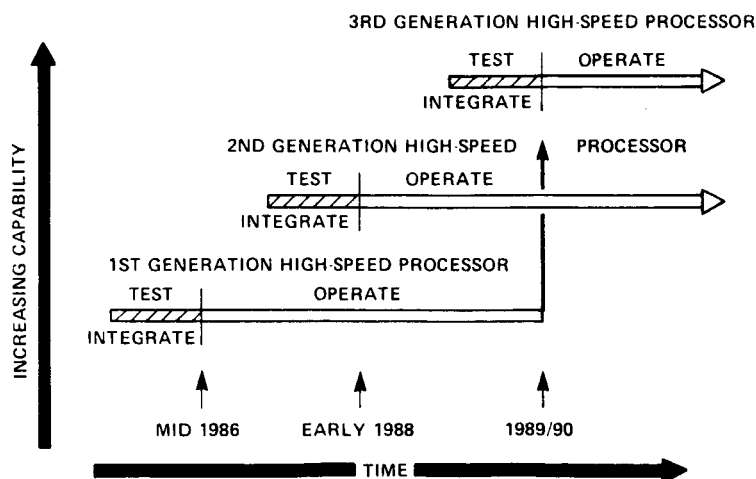
Fig. 2 Evolution of the NAS Processing System Network (NPSN)

4

Coupled with the introduction of each generation of high-speed processor is an NPSN system-wide implementation phase. The first two phases of it are initial operating configuration (IOC) and extended operating configuration (EOC). The key technical goals for each configuration follow:

Initial operating configuration (IOC)
    High-speed processor-1 (HSP-1)
        250 million floating-point operations/sec computing rate for optimized large-scale computational aerodynamic applications
        265 million 64-bit words of central memory
    Integrated and expandable NPSN configuration
    Common user interface for all subsystems
    Mediumband (56 kilobits/sec) communications for remote sites
    Wideband (1.544 megabits/sec) communications to NASA centers

Extended operating configuration (EOC)
    Additional high-speed processor-2 (HSP-2)
        1000 million floating-point operations/sec computing rate for optimized large-scale computational aerodynamic applications
    Upgraded subsystems to accommodate HSP-2 and graphics subsystems
    6.2 megabits/sec communications to NASA centers

Further extensions
    Replace HSP-1 with advanced supercomputer at 4-5 times HSP-2 performance

EOC has been established as the target baseline configuration. This configuration is shown in Fig. 3 and consists of the following seven functional subsystems:
1. High-speed processor subsystem (HSP)
2. Support processing subsystem (SPS)
3. Workstation subsystem (WKS)
4. Graphics subsystem (GRS)
5. Mass storage subsystem (MSS)
6. Long-haul communication subsystem (LHCS)
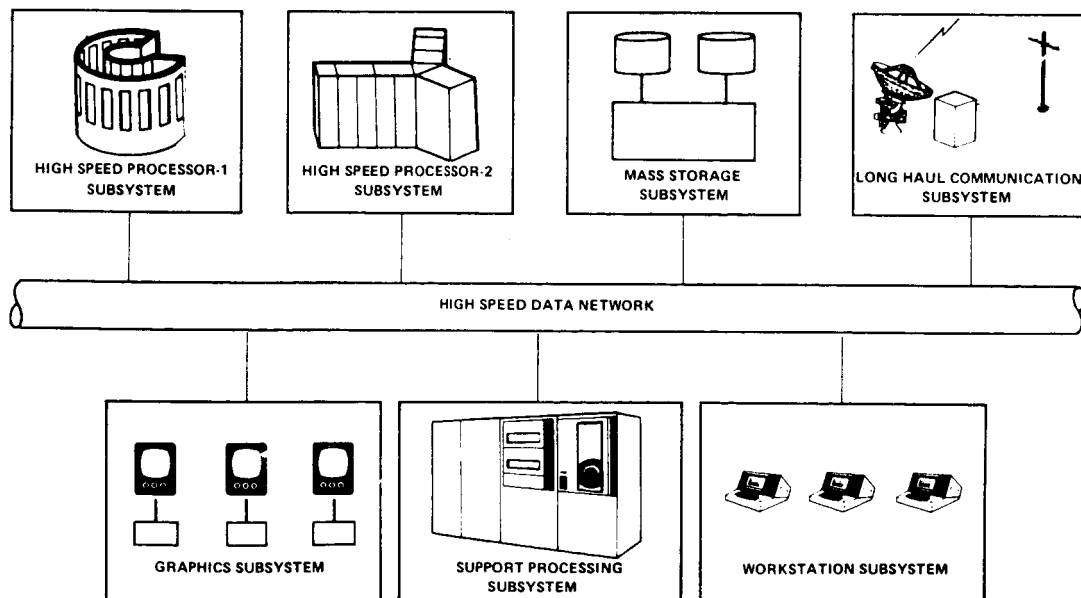7. High-speed data network (HSDN)

Fig. 3 NAS processing system network

5

The IOC is treated as an intermediate step toward the EOC target. IOC, described in the next section, includes all the EOC functional subsystems with two exceptions: GRS and the second generation HSP.

NPSN INITIAL OPERATING CONFIGURATION

Implementation of the initial operating configuration phase of the NPSN began in the summer of 1983 with the acquisition of the network test bed (NTB). The NTB was composed of two DEC VAX 11/780 computers configured with Berkeley 4.1 C UNIX operating systems, and with HYPERchannel[2] and Ethernet[3] communications. The NTB was used to gain experience with UNIX and to test candidate communication protocols including Livermore Interactive Network Communication System (LINCS), NETEX,[4] and the DOD Internet Protocol (usually referred to as TCP/IP). As a result of the experience gained, UNIX and TCP/IP were chosen as the software base upon which to implement the NPSN. Since the *summer of 1983*, the NPSN has grown to the configuration shown in Fig. 4.
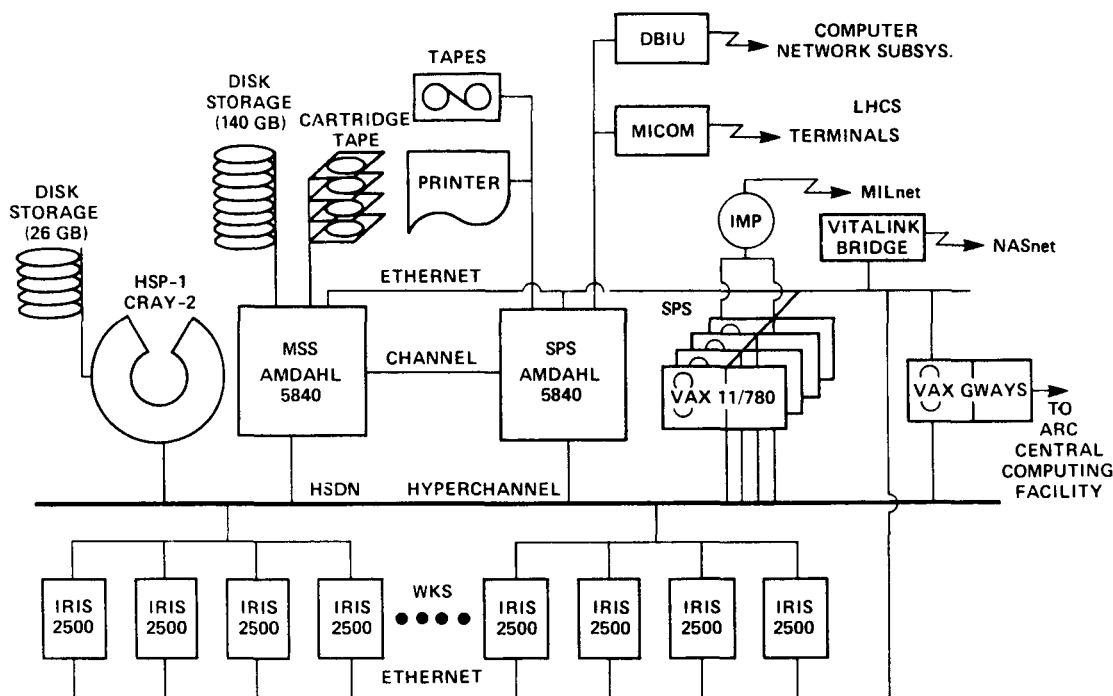
Fig. 4 Initial operating configuration of the NPSN (July 1986)

HSP-1

In April 1984, NASA signed a contract with Cray Research, Inc. to deliver a Cray-2 to the NAS Program for the HSP-1 subsystem. The Cray-2, which was delivered to ARC in the fall of 1985, has a clock-cycle time of 4.1 nanosec and

---

[2]HYPERchannel is a trademark of Network Systems Corporation.
[3]Ethernet is a trademark of Xerox Corp.
[4]NETEX is a trademark of Network Systems Corporation.

consists of an integral foreground processor, four independent background processors, and a large common memory.

The four processors can operate on separate programs or, concurrently, on a single program. The common memory consists of 268,435,456 words (64-bit), randomly accessible from any of the four background processors and from any of the high-speed and common-data channels. The common memory of the Cray-2 is significantly larger than that offered on any other commercially available computer system and allows users to run programs impossible to run on any other system. The foreground processor supervises overall system activity among itself, the background processors, common memory and peripheral controllers. When first delivered, the Cray-2 was initially configured with 12 gigabytes of local disk storage attached to the foreground processor. Disk storage has been incrementally increased to 26 gigabytes with installation of new DD-49 disk units. Early in 1987, disk storage will be further increased to 40 gigabytes. Cray-2 software includes an operating system based on AT&T UNIX System V, an automatic vectorizing FORTRAN compiler, a C language compiler, a large set of batch and interactive utilities, a large set of libraries including a multitasking library, TCP/IP networking, and Berkeley R commands.

At the time of contract award, the Cray-2 design was complete and the cooling technology well tested, but work was continuing in module design, implementation, and testing. Because the Cray-2 was not available for immediate shipment, a CRAY-XMP/12 was delivered to the NAS Program in August 1984 to be used as an early supercomputer for initial system integration and UNIX testing. This system has served its purpose and will be removed in August 1986. Meanwhile, Cray-2 hardware development was proceeding at Chippewa Falls, Wisconsin, and software developed at Mendota Heights, Minnesota, during which time a series of hardware and software demonstration tests were successfully completed. In April 1985, a single processor Cray-2 was shipped to National Magnetic Fusion Energy Computer Center (NMFECC), followed by serial number one of the Cray-2 in June. This 64-million word system and the uniprocessor system were shared by Cray Research Inc., NMFECC, and the NAS Program, and ran under both CTSS and UNIX operating systems. Cray-2 serial number 2 arrived at ARC on September 30, 1985. This was the first system with the full 268-million word memory configuration. After installation and a period of test and hardware debugging, the system began pilot operations in early December 1985. Because TCP/IP was not available, early access was gained through the VAX 11/780 and Amdahl 5840 systems using the Cray Research Inc.-provided SEP protocol software. Currently, the Cray-2 runs UNICOS 1.1 which supports TCP/IP protocol over HYPERchannel for communications to the other NPSN subsystems as a fully participating network host.

Workstation Subsystem (WKS)

The workstation subsystem (WKS) consists of 25 Silicon Graphics IRIS[5] (Integrated Raster Imaging System) Model 2500 and two model 2500 Turbo graphics workstations. These UNIX based workstations provide a "scientists workbench" for users to perform text and data editing, to run small application programs, and most importantly, to process and view graphic files.

Powerful graphics processing is provided by a special purpose microprocessor based Geometry Engine[6] that is interfaced to a high-resolution color raster display monitor. The Geometry Engine displays and performs hidden

---

[5]IRIS is a trademark of Silicon Graphics, Inc.
[6]Geometry Engine is a trademark of Silicon Graphics, Inc.

surface removal, rotation, zooming, and clipping of raster images at rates in excess of 50,000 points/sec. The typical model 2500 workstation includes a Motorola 68010 microprocessor (16 bit), floating-point accelerator, 2.5 megabytes of central memory, and a 474 megabyte disk. The turbo model incorporates the Motorola 68020 microprocessor (32 bit) and a Weitek floating-point chip set accelerator to provide minicomputer (VAX 11/780 equivalent) performance.

The objective of the WKS is to provide a powerful local workstation environment, shared among a few users, that is closely connected to the globally shared supercomputer environment. The WKS started as the user interface prototype in the summer of 1984. The purpose of this prototype effort was to introduce high-performance graphics workstations into the users' supercomputing environment, to explore the potential benefits of workstation to supercomputer interaction, and to introduce users to the UNIX environment. The initial configuration consisted of IRIS Model 1500 workstations co-located with Computational Fluid Dynamics (CFD) researchers, and interfaced to the NAS Program Cray XMP/12 via VAX 11/780 gateways. Later the ARC Central Computing Center Cray XMP/22 and Cyber 205 were also connected to these VAX gateways. After a month or so of familiarization, the CFD users began to routinely utilize the powerful graphics capabilities offered and to demonstrate the benefit of graphical display techniques for analyzing the output of large three-dimensional models.

The successful user interface prototype led to the current workstation subsystem with the Model 1500 upgraded to Model 2500 and 2500 turbos. The VAX gateway was eliminated so that each workstation is a participating host on a TCP/IP HYPERchannel network with direct access to all the other NPSN subsystems, including the Cray-2 and Cray XMP/12. The workstation also communicates to the other subsystems via a TCP/IP over Ethernet. Plans are to provide Ethernet direct access to the Cray-2 in 1987.

Data files can be transferred from the Cray-2 to the workstation for image generation and display. In addition, user programs executing on the Cray-2 can output directly and immediately to a workstation using the IRIS Remote Graphics Library implemented on the Cray-2. This software consists of the IRIS library of graphic calls and permits very computing-intensive image rendering to be performed on the Cray-2 with user interface and control identical to the workstation stand-alone environment.

Quality film output is provided by two Dunn camera systems. Users can make high quality still pictures with the Model 655 Dunn camera system. The Dunn system works directly off the RGB output of the workstations. For 16mm movies, the Model 632 Dunn workstation functions with the aid of a patch panel. The Dunn cameras and an extra monitor are kept in an area common to all users. When a user needs film hardcopy, he patches his workstation to the monitor connected to the Dunn systems.

Video animation capabilities are provided by a Lyon-Lamb video animation system. This system takes an RS170A signal from the IRIS, converts it to a television signal, and records it on video tape. To assist the user in making animation sequences, a program called graphics animation system (GAS) has been developed. GAS lets the user create scripts on the IRIS which can be played back at a later time on the animation system. GAS is a fairly "user friendly" program. It records a user's actions as graphic images are manipulated on the IRIS screen, automatically generating a playback script. If needed, the user can edit this script to either correct the errors or to add steps to smooth out the animation.

Support Processing Subsystem (SPS)

The support processing subsystem (SPS) provides an intermediate performance, general purpose interactive processing which complements the capabilities provided by the HSP and WKS. It is the principal location for electronic mail service, text editing, and other various utilities, system configuration management files, and various system accounting and administration data bases. In addition, it provides the interface to local and remote terminal users, and to printers, plotters, and tape drives. The SPS functions are shared among an Amdahl 5840 operating UTS, and four VAX 11/780 computers, two operating Berkeley 4.2 UNIX and two operating AT&T system V UNIX. Along with WKS, these support processors replace the "dedicated front-end environment." Each of these processors serves as a user gateway for interactive, file transfer, and batch job access to the Cray-2, and provides independent processing capabilities within an interconnected, distributed processing environment.

Mass Storage Subsystem (MSS)

The mass storage subsystem (MSS) functions as the file server for all client hosts on the NPSN and provides both global on-line and archival file storage. This subsystem validates and coordinates file requests; maintains a directory of all contained files; and performs file duplication, file migration across the various storage media, storage allocation, accounting, and file-management functions. The MSS is an Amdahl 5840 which was installed in the summer of 1985. The configuration includes 16 megabytes of central memory, 140 gigabytes of on-line 6380 disk storage, and a dual IBM 3480 tape cartridge system. To provide redundancy in case of failure, the MSS mainframe is connected via direct channels to a second Amdahl 5840 that provides support processing and remote communications functions.

The MSS software configuration consists of two interconnected, cooperating operating systems functioning within the IBM VM environment. The first is Amdahl's UTS which is based on AT&T UNIX System V. UTS provides the UNIX file director which manages all logical files, and is the only operating system visible to the user. The second operating system is IBM's MVS which in conjunction with data facility hierarchical storage system (DFHMS) provides the storage device management functions such as allocation of disk space to datasets, file migration, and file back-up. Future plans include an increase to 300 gigabytes of disk storage and an on-line automated tape cartridge device in 1987.

Long-Haul Communication Subsystem (LHCS)

The long-haul communication subsystem (LHCS) provides communication interfaces to the nationwide remote user community. The NAS Program remote communication design objectives are to:

(a) provide a broad range of data communication bandwidths and services in order to allow remote users to access the NPSN via terminals, workstations or host computer systems;

(b) provide access to the geographically dispersed user sites representing a diverse user community (NASA Centers, DOD and the Government research installations, aerospace industry sites, and universities);

(c) reduce costs by making use of existing and planned Government-supported remote communication networks; and

(d) establish a single data communications protocol for all NAS Program implemented communications.

To meet these objectives, the NAS Program has implemented a dual approach. First, extend the NPSN local-area network environment to remote user

sites in order to exploit the scientific workstation-to-supercomputer interconnection environment. Implementation of this approach began with the establishment of the long-haul communication prototype which consisted of connecting the NPSN Ethernet to remote site Ethernets by terrestrial and satellite communication links via Vitalink TransLAN communication bridges. These intelligent bridges permit NPSN Ethernet packets with remote addresses to be sent to destination hosts on the remote site Ethernet. Conversely, Ethernet packets on the remote sites own Ethernet can be routed to NPSN Ethernet hosts. In this way, the remote Ethernet host appears to the user to be just another node on the NPSN local Ethernet. The DOD internet protocol (TCP/IP) was selected as the protocol suite since it is the NPSN local-area network protocol, it offers a rich set of applications (mail, file transfer, remote login, etc.) and implementations exist for a large number of various computers ranging from IBM PCs to Cray supercomputers. The first prototype implementation was a 56 kbit/sec terrestrial link to a Sun workstation at Colorado State University, Fort Collins, Colorado, July 1985. This was followed by 224 kbit/sec satellite links to Langley Research Center, Hampton, Virginia, and Lewis Research Center, Cleveland, Ohio, in October 1985. The prototype has transitioned into the operational NASnet which is being expanded to include 56 kbit/sec links to 13 additional NASA and aerospace industry sites. These communication links are being provided by NASA's newly implemented Program Support Communication Network (PSCN). PSCN provides both packet and circuit switch service over 1.544 megabit/sec backbone satellite and terrestrial circuits connecting all NASA centers and installations. In most cases, traffic will be carried over T1 (1.544 mbit/sec) circuits between NASA Centers, and from there, fanout over 56 kbit/sec tail circuits to industrial sites.

The second approach to providing NPSN access from remote user sites is to establish the NPSN as a participating host on existing and planned Government-sponsored networks. The NPSN is currently a host on MILnet, the NASA Packet Switch System (NPSS), the NASA Computer Network Subsystem (CNS) and in late 1986, will be a host on the National Science Foundation's newly implemented NSFnet.

As a member of MILnet, access to the NPSN is provided to DOD research installations, and selected DOD contractors. Access to MILnet is provided by attachment of NPSN VAX 11/780 hosts to a MILnet packet switching node located at ARC. Access is also provided to ARPAnet sites (principally universities) via MILnet/ARPAnet gateways located across the country.

NPSS provides 9600 bit/sec terminal access to the NPSN from NASA Centers. Since NPSS is also connected to GTE's Telenet packet switched network, remote users outside NASA Centers can gain access to the NPSN at rates of 1200-2400 bits/sec.

NASA computer network subsystem (CNS) provides batch oriented bulk file transfer service among NASA OAST Research Centers (Ames-Moffett, Ames-Dryden, Lewis, and Langley). This capability as well as NPSS is provided by PSCN. CNS is based on a store-and-forward model of operation. Files are moved from the local NASA host computers to the CNS Cyber 810 based data buffer interface unit (DBIU) located at each OAST Center. The DBIU at the sending Center then negotiates with the PSCN's network management processor at the PSCN network operations center at Marshall Space Flight Center for access to a T1 satellite channel to the target receiving DBIU at an OAST Center. Files destined for a host at the target Center are moved to the receiving DBIU using Network Systems Corporation's proprietary HYPERchannel Network Exchange protocol (NETEX). The DBIU's interface to the OAST host computer systems using either NETEX on HYPERchannel or CDC's Remote Host Facility (RHF) on CDC's loosely coupled network (LCN). In the NPSN implementation, the DBIU interfaces to the SPS Amdahl

and is accessed by the user through CNS provided commands in the UTS environment.

The NAS Program is a participating member of the NSF-funded Bay Area Regional Research Network (BARRN). This NSF supported regional network includes the University of California campuses (Berkeley, Davis, San Francisco, and Santa Cruz), Stanford University, and Ames Research Center. The network, scheduled to be operational in the fourth quarter of 1986, will feature communications links between participating member sites operating at T1 and will be interfaced to DOD-internet IP routers which, in turn, are connected to one or more member sites local area networks. BARRN will be connected to ARPAnet at Stanford and will offer university users access to NPSN. Additionally, BARRN will be connected to the NSF funded San Diego Supercomputer Center (SDSC) through the services of a remote user access computer (RUAC) at Berkeley.

High-Speed Data Network (HSDN)

The HSDN provides the medium over which data and control messages are exchanged among the NPSN subsystems. The HSDN configuration consists of a Network System Corporation HYPERchannel network and an Ethernet network. The HYPERchannel network interconnects all NPSN subsystem via four trunks, each capable of a maximum transfer rate of 50 megabits/sec. Ethernet, with a maximum transfer rate of 10 megabit/sec, connects the WKS, SPS, LHCS, and MSS. It provides a means of connecting workstations to other ARC computing systems, and an alternate to HYPERchannel for short message traffic. Outside of device-driver level, all the communications software is the same on both networks, i.e., TCP/IP and Berkeley "r" commands.

NPSN Software

Early in the planning phase, it was recognized that the NAS Program was given the unique opportunity to design a completely new supercomputer-facility environment responsive to the needs of the modern user and unencombered by constraints dictated by an existing production environment. It was also recognized that the facility needed to be designed to accommodate an aggressive program of new technology upgrades after production operation was achieved. Fortunately, ARC had a powerful supercomputer installation in the Central Computing Facility (Cray-1 and Cyber 205), experience in a new untested supercomputer (Illiac IV), a large local area network, significant graphics capabilities, and a very sophisticated supercomputer-user community. Given these opportunities, constraints, and experience, the NAS Program established several key software design objectives:

(a) Implementation of a ventor-independent software environment to facilitate the incorporation of new technology into the local area network environment.

(b) Implementation of a common and consistent user environment across the NPSN.

(c) Provision for both interactive and batch operation of high-speed processors and other user accessible computers.

(d) Provision for a large, functionally rich library of common application tools for the user.

To meet these objectives, two critical software-design decisions were made. The first was to base the NPSN operating system software on UNIX, and the second was to use the DOD internet protocol for local-area network and remote-network communications. In retrospect, the decisions made in 1983 appear even more valid today. UNIX is the only operating system and TCP/IP is the only protocol that are even close to being vendor independent. Specifications for

both are readily available and they have been widely implemented on a large number of computers.

There were concerns in 1983. These were largely with regard to TCP/IP performance, and UNIX suitably for supercomputers, its batch processing environment, and its security. The TCP/IP performance issue was settled through experience in the network test bed project referred to earlier. The conclusion was that the key to network performance was the actual protocol implementation and the ability to fine tune performance parameters. In these respects, the TCP/IP open architecture provided the best choice. Confidence that UNIX could be efficiently implemented on a supercomputer (a UNIX shell or guest-operating system implemented on top of existing operating systems was a strong consideration at the time) was greatly enhanced when, just prior to contract award, Cray Research, Inc. decided to implement a version of UNIX (now called UNICOS) on the Cray-2 and Cray X-MP. An effective batch-processing environment was recognized as a development area that has since resulted in the network queueing system (NQS). Security in UNIX is still a concern. UNIX does not have particularly bad security features as compared to many vendor-unique operating systems, but its vulnerabilities are well known. It is an on-going effort to keep improving UNIX security. With these concerns in mind, an implementation strategy has been developed that has resulted in what the NAS Program calls vendor-independent network operating system (VINOS). The current and first release of VINOS is based on the integration of vendor-based implementations of UNIX System V, C. Shell program, U.C. Berkeley developed socket calls, TCP/IP, U.C. Berkeley "r" commands and NQS. (In addition to VINOS, the NPSN also incorporates two Berkeley 4.2 UNIX configured VAX 11/780 computers.)

The TCP/IP implementation on the VAX 11/780s, the Amdahls and the Cray-2 have been accomplished by the Wollongong Group in cooperation with the computer vendors. The TCP/IP implementation on the workstations was done by Sterling Software, Inc. under contract to the NAS Program. With the exception of the Cray-2, TCP/IP is implemented over both Ethernet and HYPERchannel. The Cray-2 implementation is for HYPERchannel only and is included in the Cray-2 UNICOS 1.1 release.

NQS is a specially developed addition to UNIX, providing consistent batch-management functions on each subsystem as well as on a network-wide basis. The salient features of NQS are batch-job-queue management, and output file staging to user-selected computers and devices. NQS is derived from the Army Ballistics Research Center's multidevice queueing system and was written by Sterling Software under contract to the NAS Program.

Future releases of VINOS (planned for later in 1986 and early in 1987) include significant new features. Network-wide administration, and diagnostic and performance-monitoring facilities will be added to aid in achieving high availability and throughput over the HYPERchannel/Ethernet LAN of 35 hosts.

A common set of graphical display utilities called common graphics services (CGS) will be implemented across the NPSN. CGS is made up of three libraries. The first is the common graphics library composed of the general purpose DI-3000 graphics package provided by Precision Visuals. The second is the IRIS graphics library which is expanded to provide preparation of IRIS specific images on other subsystems. The third is an application library consisting of programs written specifically for graphical presentation of computational fluid dynamic flow fields.

A network-distributed file system is planned. Such a file system, based on the Newcastle connection developed at the University of Newcastle upon Tyne, has been developed for the NPSN. This facility allows for file access across any NPSN subsystem and machine boundaries without the need to explicitly provide file or machine addresses. A single file name is recognized by all subsystems and a separate directory at each subsystem is not required. Also, through a

network-wide shell program, users who have logged onto one subsystem can access data within files on another subsystem without having to copy the entire file. An example of such use would be viewing frames from a graphical output file which is stored on the mass-storage subsystem, and then viewing portions of it on a workstation (even if the entire file is too large to copy to the workstation). Another advantage of the Newcastle connection is that it allows any program to access files at another subsystem simply by reloading the program into the Newcastle library; no other changes are required to make use of remote access. Its incorporation in VINOS is currently being deferred until more experience is gained with the current version of VINOS. Since other distributed file systems (notable Sun Microsystems NFS), have recently appeared, an assessment will be made as to wisdom of a unique NPSN implementation versus a more widely available one.

OPERATIONAL EXPERIENCE AND PLANS

The NAS Program initiated pilot operations for a select group of users in early December 1985, shortly after the Cray-2 had completed acceptance testing. When the Cray-2 arrived, initial access was obtained from a VAX 11/780 operating under System V UNIX using the simple effective protocol (SEP) provided by Cray Research, Inc. SEP was subsequently installed on additional VAX 11/780s, the Amdahl 5840s and the IRIS workstations. The ease with which this was accomplished illustrates one of the benefits of using the UNIX operating system throughout the NPSN. SEP has since been replaced by TCP/IP.

Access from the remote sites was provided by the NAS long haul communication prototype, MILnet/ARPAnet, and GTE Telenet. During the pilot-operation phase, December 1985 to July 1986, the NPSN operated on a 24 hour/day, 7 day/week basis with scheduled operating periods divided between scientific user production, system development, and test. The scientific user community rapidly grew from 5 to 60 users (over a 3-month period) as system stability improved. The Cray-2 FORTRAN compiler (CFT) was a major item in achieving both stability and performance improvement over this period. A set of 13 selected codes was used to continually monitor the improvements in CFT. Only 3 of these codes were running at the time the Cray-2 was shipped in September 1985; by October, the number had risen to 5, in November to 8, in December to 9, in January 1986 to 11, and by April 1986 all 13 codes were running.

As CFT improved and the user base increased, it was apparent that the Cray-2 performed particularly well for large-memory applications that maximize vector references. Many users have scaled-up their applications and 100-million-word jobs are now routine. Codes that do poorly on the Cray-2 include those that are scalar-computation intensive with frequent fetches from and stores to main memory. However, the performance of many scalar-intensive codes has improved significantly with the introduction of the most recent CFT compiler, Version 2.71. This version places all scalars from the currently executing subroutine into local CPU memory, unless they have been passed by that subroutine as parameters. In addition, the memory-access performance is greatly improved if power of two memory strides are avoided. For instance, an array previously dimensioned (128, 128) should be reset to (129, 128). Although the Cray XMP exceeds Cray-2 performance for many small memory, scalar-intensive applications, the gap is expected to narrow as the Cray-2 FORTRAN compiler improves. The large applications on the Cray-2 put a significant strain on the local disk storage of 12 gigabytes. The recent upgrade to 26 gigabytes and the introduction of MSS capability are expected to relieve this problem.

The IOC will begin interim operations in July 1986. The scientific user community will increase from 60 to 200 major users, with the number of remote user sites increasing from 10 to 27. In December, the entire NPSN will be moved

13

to the NAS Facility now under construction. The new NAS Facility is a 90,000 ft^2 building having more than 30,000 ft^2 of main computer floor and laboratories as well as office space for the NAS Program and Computational Fluid Dynamics Research staff. The relocation of computers, communications, and staff is anticipated to occur over a four-week period spanning the Christmas and New Year holidays. The NPSN is planning to be back on-line in January and, after a period of shake-down, full operation of the IOC will start in March 1987.

In preparation for full IOC operations, the NAS Program is implementing a usage-allocation procedure which includes an invitation for NPSN resource requests from the scientific user community. NAS is intended to be utilized for pioneering work in the uses of computation for basic physics research, for developing new advanced application programs, and for processing large-scale applications in support of advanced research and development. It is expected that 90% of the available usage will be allocated to research and advanced applications in the areas of fluid dynamics, aerodynamics, structures, material science, aerothermal dynamics, and other aeronautics applications. The remaining 10% will support advances in such disciplines as chemistry, atmospheric modeling, astrophysics, and other subjects of interest to NASA. About 55% of the usage will support NASA programs and be used by in-house staff and organizations involved in NASA grants, contracts, and joint programs. Roughly 20% of the usage will support DOD activities, 15% will be allocated to industry, 5% to other Government agencies, and 5% to universities. The cost for industry proprietary work would be reimbursed to the Government in accordance with a formula that includes operating and capital improvement costs.

The NASA Office of Aeronautics and Space Technology will receive approximately 45% of the NASA allocation to be shared among Ames, Langley, and Lewis Research Centers. Each Center will suballocate its share to individual in-house and supported research projects. The university allocation will be administered by the National Science Foundation, and all others by a NASA peer-review process. Overall resource allocation will be administered by the NAS Projects Office. To obtain an NPSN allocation potential users must demonstrate a requirement for the unique capabilities offered by the NPSN. In addition, project selection will be guided by the following criteria: technical quality, national need, timeliness, and suitability to NAS resources. Selection criteria and allocation guidelines will be periodically reviewed to ensure the best possible benefit to the user community. The guidelines for usage allocation described here apply only to NAS computers in an operational status. During periods when new computers are being integrated into the NPSN, they will be devoted primarily to systems integration, testing, and software development.

FUTURE PLANS

In addition to near-term improvements to the IOC discussed earlier, the NAS Program is moving forward in longer range planning for implementation of the EOC and its follow-on configurations. EOC is planned to be operational in 1988 and will feature the addition of an advanced supercomputer, HSP-2, with performance and memory capacity targeted at four times that of the Cray-2. The GRS will be included in the EOC and GRS prototype development will be initiated in late 1986. GRS is targeted to exceed current workstation capabilities by an order of magnitude increase in rates for image generation, three dimensional coordinate transformation, and image display. On-line mass storage capacity, the number of remote sites, local network bandwidth, and remote communication bandwidth will also be significantly increased.

Major prototype efforts are being planned to support NPSN implementation for EOC and subsequent future configurations. A high-speed bus prototype has been initiated with the goal to demonstrate an 800 mbit/sec transfer rate from a

high-speed processor to advanced on-line graphics devices and eventually to mass storage devices. The NAS Program maintains an interest in the development of high-capacity and high-performance storage technologies, such as magneto-optics, although no prototype plans have been formulated yet. An advanced remote-communications gateway prototype is also planned. The objective of this effort is to demonstrate more efficient utilization of high-speed switched circuit capacity in accordance with actual throughput demand. A key part of the effort will be to develop a type-of-service routing gateway that permits various application services (interactive character stream, on-demand data transfer, bulk file transfer, etc.) to request specific bandwidth and circuit delay characteristics. The bandwidth request is met by adding or deleting subchannel allocations of the backbone circuit and the delay request is met by choosing between low-delay terrestrial or high-delay satellite circuits. Prototypes for high bandwidth backbone circuits up to 6.2 mbits/sec are also planned.

A major software direction will be to transition from NPSN fostered implementations (made necessary to meet technical goals) to future industry-defined standard implementations where required functionality and performance can be achieved. It is expected that this approach will make software meeting NAS Program requirements more readily available on new systems. Candidate software examples include transition from the UNIX System V/Berkeley Socket implementation to AT&T System V release three with streams networking, and the transition from TCP/IP to the International Standards Organization defined open system architecture. Supplementing these will be enhancements in several areas including system security, back-up and recovery, job control, performance monitoring and tuning, and network management, diagnostics, and fault tolerance. Software development will also continue to provide better tools to the end user in areas such as graphics, programming and debugging for concurrent processors, data management and multimedia communication.

SUMMARY

Founded on 10 years of study and advocacy, the NAS Program has reached its first major milestone--the NPSN initial operating configuration. The program has met its goal of providing a national supercomputer facility capable of greatly enhancing the Nation's research and development efforts. Furthermore, the program is fulfilling its pathfinder role by defining and implementing a new paradigm for supercomputing system environments. The IOC is only the beginning and the NAS Program will continue aggressively to implement new emerging supercomputer, communications, storage and software technologies to strengthen computations vital role in supporting the Nation's aeronautics research and development goals.

REFERENCES

1. MacCormack, R. W., "Numerical Solutions of the Interaction of a Shock Wave with a Laminar Boundary Layer," in Lecture Notes in Physics, Vol. 8, Springer Verlag, 1971, pp. 151-163.

2. Peterson, V. L., Ballhaus, W. F., Jr., and Bailey, F. R., "Numerical Aerodynamic Simulation (NAS)," in Large Scale Scientific Computation, edited by Seymour V. Parter, Academic Press, Inc., 1984.

3. Peterson, V. L., and Arnold, J. O., "Impact of Supercomputers on Experimentation: A View from a National Laboratory," Proceedings of the American Society for Engineering Symposium, Atlanta, GA, June 17-19, 1985.

1. Report No. NASA TM 88339	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle STATUS AND PROJECTIONS OF THE NAS PROGRAM		5. Report Date July 1986
		6. Performing Organization Code
7. Author(s) Frank R. Bailey		8. Performing Organization Report No. A-86325
		10. Work Unit No.
9. Performing Organization Name and Address Ames Research Center Moffett Field, CA 94035		
		11. Contract or Grant No.
		13. Type of Report and Period Covered Technical Memorandum
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546		
		14. Sponsoring Agency Code 536-10-11

15. Supplementary Notes

Point of Contact: F. R. Bailey, Ames Research Center, MS 233-1, Moffett Field, CA 94035
(415)694-6419 or FTS 464-6419

16. Abstract NASA has initiated at Ames Research Center the Numerical Aerodynamic Simulation (NAS) Program. The objective of the Program is to develop a leading-edge, large-scale computer facility, and make it available to NASA, DOD, other Government agencies, industry, and universities as a necessary element in ensuring continuing leadership in computational aeronautics and related disciplines. The Program established an initial operational capability in 1986 and will systematically enhance that capability by incorporating evolving improvements in state-of-the-art computer system technologies as required to maintain a leadership role.

The NAS Program had its genesis in the efforts of theoretical aerodynamicists at Ames Research Center during the mid-1970s. Its purpose was to exploit computational aerodynamic research to improve the basic technology of aircraft and aerospace vehicle design. The combination of improvements in aerodynamic flow simulation modeling and supercomputer performance have permitted the rapid advancement of computational aerodynamics to the stage where it is practical to address nonlinear inviscid flow simulations about complex aerodynamic configurations, and detailed turbulent flow simulations about simple configurations. The necessary next step in the advancement of computational aerodynamics is to combine detailed turbulent flow models and complex configurations, in order to provide the capability to accurately simulate realistic designs over the entire flight envelope. The goal of the NAS Program is to provide the necessary computer throughput, memory capacity, and system architecture to reach this next step. In particular the NAS Program will achieve a supercomputer capability of 250 million floating-point operations per second sustained throughput and a 256 million word memory capacity on 1986, and a four-fold increase in these capabilities by 1987. This new supercomputer capability includes the Cray-2 which is integrated into the NAS Processing System Network (NPSN). The NPSN is designed to support supercomputer generated large-scale aerodynamic flow simulations with powerful scientific workstations, mass storage facility, high-resolution graphics facility, and high-bandwidth remote communication links to remote sites nationwide. To aid in its productive use the system will provide unique software features including a UNIX^{TM} operating system environment across all systems, common graphics services, a library of aerodynamic simulation models, and a rich set of local area network data communication services.

17. Key Words (Suggested by Author(s)) Computers Computer systems Numerical simulation Aerodynamics	18. Distribution Statement Unlimited Subject Category - 62	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 17	22. Price* A02