NASA TM-84386

NASA-TM-84386 19830024126

Numerical Aerodynamic Simulation (NAS)

V. L. Peterson, W. F. Ballhaus, Jr., and F. R. Bailey

July 1983

NASA

National Aeronautics and
Space Administration

N83-32397#

Numerical Aerodynamic Simulation (NAS)

V. L. Peterson,
W. F. Ballhaus, Jr.,
F. R. Bailey, Ames Research Center, Moffett Field, California

NUMERICAL AERODYNAMIC SIMULATION (NAS)

V. L. Peterson, W. F. Ballhaus, Jr., and F. R. Bailey

Ames Research Center

1. INTRODUCTION

Computational aerodynamics has emerged in the last decade as a powerful tool for the design of all types of aerospace vehicles, but it is one of several technical disciplines whose development is being restrained by the lack of suitably powerful computational systems. The equations of fluid flow are well-known, but the methods for solving them without approximation require unrealistically large amounts of computation for all but the most elementary flow situations. Thus, the aerodynamicist must resort to the use of approximate forms of the equations to treat complete aircraft geometries. In fact, for three-dimensional configurations, only the equations for inviscid flow can be solved in a practical amount of time with currently available computers. This limitation on the treatment of fluid-flow physics severely restricts a designer's ability to analyze aerodynamic phenomena and the resulting effects on aircraft performance.

The Numerical Aerodynamic Simulation (NAS) Program is designed to provide a leading-edge computational capability to the aerospace community. It was recognized early in the program that, in addition to more advanced computers, the entire computational process ranging from problem formulation to publication of results needed to be improved to realize the full impact of computational aerodynamics. Therefore, the NAS Program has been structured to focus on the development of a

1

complete system that can be upgraded periodically with minimum impact on the user and on the inventory of applications software. The implementation phase of the program is now under way. It is based upon nearly 8 yr of study and should culminate in an initial operational capability before 1986.

The objective of this paper is fivefold: (1) to discuss the factors motivating the NAS program, (2) to provide a history of the activity, (3) to describe each of the elements of the processing-system network, (4) to outline the proposed allocation of time to users of the facility, and (5) to describe some of the candidate problems being considered for the first benchmark codes.

2. MOTIVATING FACTORS

The discipline of aerodynamics plays an important role in the design of aerospace vehicles. Beginning nearly 200 yr before the Wright Brothers' historic powered flight in 1903 and continuing until the 1950s when electronic digital computers first became widely available, aerodynamic research and development was based largely on experimentation. This was augmented with some guidance from closed-form analytical solutions of highly approximate forms of the now well-known governing equations of fluid flow, after they were first published by Claude Luis M. H. Navier and Sir George G. Stokes in 1823 and 1845, respectively. Since the 1950s, advancements in computers have offered the possibility of obtaining solutions for successively refined approximations of the Navier-Stokes equations. These solutions have provided an increasingly important complement to experiments being conducted in wind tunnels and in flight. Today, the two principal tools of the aerodynamicist are wind tunnels and computers, and the balance between the use of these tools is beginning to shift in the direction of the computer.

Both technical and economic factors are driving the aerodynamicists toward computers, even though achievements in aeronautics continue to be made largely on the basis of experiment. As the speed and complexity of aircraft have increased, the inherent limitations of wind tunnels as the primary tool for developing and verifying the design of new aircraft and aerospace vehicles have become increasingly

2

apparent. For example, the walls in wind tunnels and the supports used to hold models severely restrict the accuracy of measurements obtained at near-sonic speeds. Aeroelastic distortions, which are known to always be present in flight, cannot be accurately simulated with subscale models tested in wind tunnels. Furthermore, limitations on temperature and on the working fluid restrict the ability of the wind tunnel to simulate flight entry into Earth's atmosphere and entry of probes into other planetary atmospheres. These limitations did not seriously affect the usefulness of wind tunnels during the era of simpler, low-speed flight. The increased complexity and broadened performance envelopes of aircraft have increased the number of wind-tunnel hours expended in the development of a new aircraft, and the cost per hour of testing has been increasing substantially at the same time. On the other hand, the cost of numerically simulating a given flow is decreasing at a rapid rate because of improvements in algorithms and computer cost effectiveness.

Even with improvements in wind tunnels and testing procedures over the years and with the use of the more sophisticated theoretical approaches made possible by computers, the available simulation capability has been inadequate to uncover design problems with a number of aircraft before full-scale prototypes were tested in flight. Some examples of these problems and their consequences are summarized for several aircraft in Table 1. All of these aircraft were developed a

Table 1. Examples of design problems resulting from inadequate simulation capability

AIRCRAFT	PROBLEMS DISCOVERED IN FLIGHT TEST	CONSEQUENCE
MILITARY TRANSPORT-1	INCORRECTLY PREDICTED WING FLOW	COMPROMISED PERFORMANCE, COSTLY MODIFICATIONS
MILITARY TRANSPORT-2	INCORRECTLY PREDICTED DRAG-RISE MACH NUMBER	REDUCED WING FATIGUE LIFE
FIGHTER/ BOMBER	INCORRECTLY PREDICTED TRANSONIC AIRFRAME DRAG	COSTLY MODIFICATIONS
BOMBER-A BOMBER-B	INCORRECTLY PREDICTED TRANSONIC PERFORMANCE	REDUCED AIRCRAFT EFFECTIVENESS
FIGHTER-1 FIGHTER-2	INCORRECTLY PREDICTED TRANSONIC DRAG	REDUCED PERFORMANCE
2 CIVIL TRANSPORTS	INCORRECTLY PREDICTED NACELLE-WING INTERFERENCE	REDESIGN REQUIRED

number of years ago, before computers were sufficiently power-
ful to have a significant impact on aerodynamic design. Some
of the problems could have been avoided using today's level of
computational capability, whereas others could not have been
uncovered without computers considerably larger than those
currently available. Successful design of next-generation
aircraft will be even more difficult without improvements in
simulation capabilities. Therefore, there are compelling rea-
sons to advance the state of the art in computational aerody-
namics. This would require more powerful computers and a con-
tinuing vigorous program to improve computational methodology.

The rapid growth in the capability of computational
methods to treat problems of practical importance began in the
early 1970s when computers became large enough to treat three-
dimensional transonic flows requiring the use of nonlinear
forms of the governing equations. The first transonic solu-
tions for a lifting airfoil with an embedded shock wave, which
could only be solved with a computer, appeared in the litera-
ture in 1970 [1]. One of the first solutions for a wing in
transonic flow was obtained in 1972; it took 18 hr of computa-
tion on an IBM 360 machine. Just 8 yr later, the results pre-
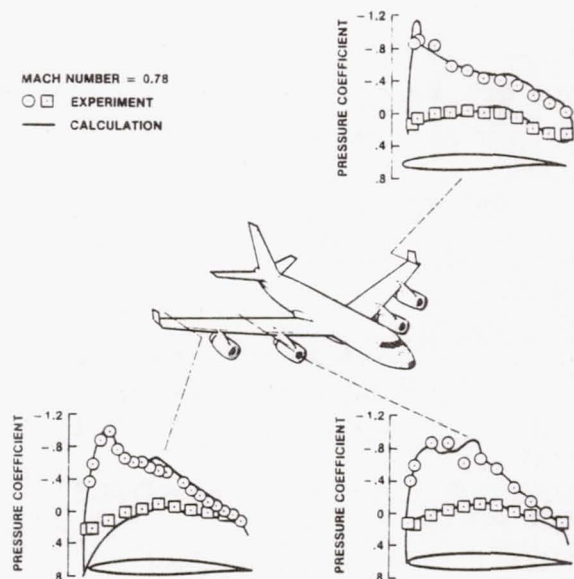sented in Fig. 1 for the flow about a complete KC 135 aircraft
with winglets at Mach 0.78
were published [2]. Chord-
wise pressure coefficients at
selected wing stations are
shown to compare favorably
with measurements for this
situation where neglected
viscous effects are not domi-
nant. Remarkably, the calcu-
lations required only about
15 min of time on a CDC 7600
computer.

Computers of the 1970s
were large enough to solve,
in a practical amount of time,
inviscid forms of the govern-
ing equations that provide

MACH NUMBER = 0.78
○□ EXPERIMENT
— CALCULATION

Fig. 1. Example Application of
Nonlinear Inviscid Equations to
Simulate Flow About a Modified
Boeing KC-135 Airplane.

4

satisfactory results for relatively complex but aerodynamically clean configurations operating at or near cruise conditions. However, many of the more crucial 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. These are also very difficult to treat effectively in wind-tunnel tests.

The Reynolds-averaged form of the Navier-Stokes equations is suitable for treating most of the 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 [3] and continuing into the 1980s with the treatment of high-Reynolds-number airfoils, bodies of revolution, and corner flows. This work has been limited to so-called building-block flows, however, because of the large amounts of computer time required for all but the most basic aircraft components. One to 20 hr of computation are needed for many applications, even with the use of the Class VI machines. An example of the application of viscous-flow technology to a turbulent supersonic flow over an axisymmetric conical afterbody containing a centered propulsive exhaust flow [4] is presented in Fig. 2. Computed density contours in the vicinity of the boattail region and the exhaust plume show excellent qualitative agreement with features of the flow made visible with schlieren photography in an experimental investigation reported for the same flow conditions [5].

A conclusion that may be drawn from the aforementioned examples is that problems involving complex geometries can be treated with simple physics and those involving simple geometries can be treated with complex physics, but more powerful computers with larger memories are required to solve problems involving both complex geometries and complex physics. The amount of computer power required to advance the capability of computational aerodynamics to treat increasingly complex aerodynamic flows is indicated by the data in Table 2. Solutions
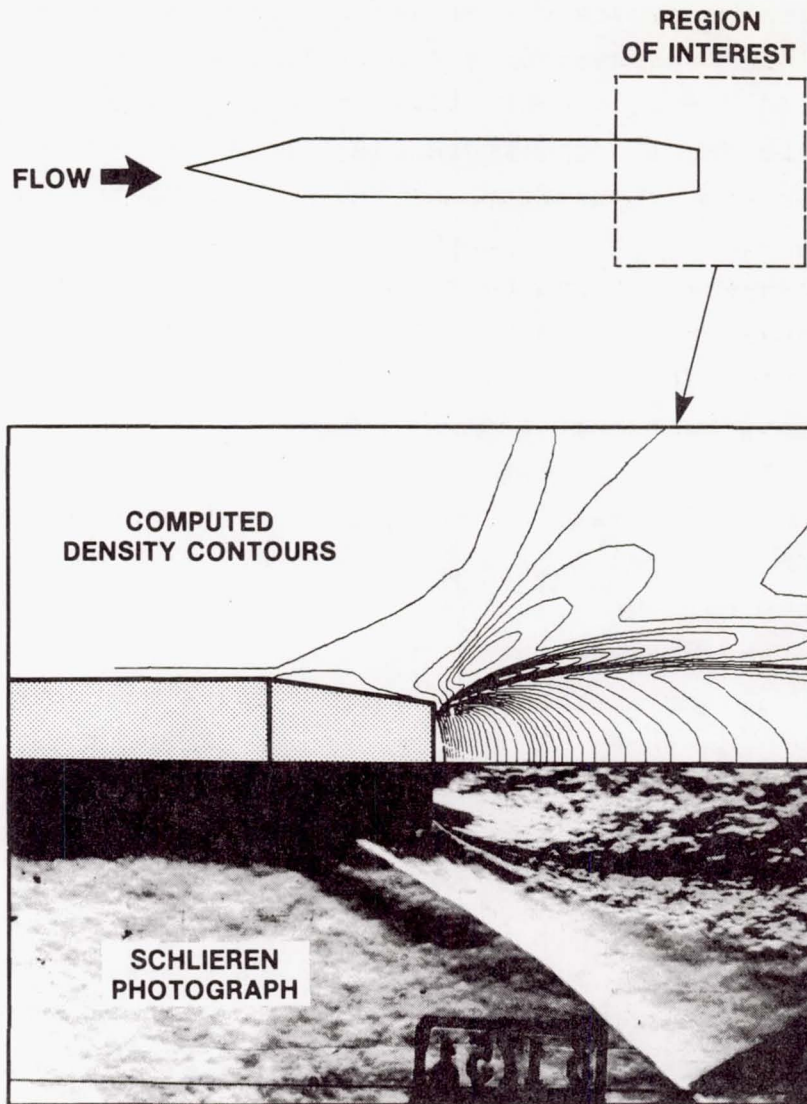
5

Fig. 2. Comparison of Numerically Simulated and Observed Turbulent Flow over an Axisymmetric Conical Afterbody Containing a Centered Propulsive Jet: Stream Mach Number = 2, Jet Mach Number = 2.5, Jet Static Pressure — Nine Times Free-stream Static Pressure.

of increasingly refined approximations to the full Navier-Stokes equations provide additional information of value to the designer, but they also require increasing refinements to the computational mesh to capture the additional detail in the flow physics. This places an increasing burden on the computer used to solve the equations. The computer requirements given in the table are based on obtaining solutions to problems involving the number of grid points shown in 10 to 15 min with algorithms expected to be available in 1985. Class VI computers are adequate for the inviscid-flow technology, but

Table 2. Computer requirements for treating increasingly
complex forms of the governing equations

APPROXIMATION	CAPABILITY	GRID POINTS REQUIRED	COMPUTER REQUIREMENT
LINEARIZED INVISCID	SUBSONIC/SUPERSONIC PRESSURE LOADS VORTEX DRAG	3×10^3 PANELS	1/10 CLASS VI
NONLINEAR INVISCID	ABOVE PLUS: TRANSONIC PRESSURE LOADS WAVE DRAG	10^5	CLASS VI
REYNOLDS AVERAGED NAVIER-STOKES	ABOVE PLUS: SEPARATION/REATTACHMENT STALL/BUFFET/FLUTTER TOTAL DRAG	10^7	30 x CLASS VI
LARGE EDDY SIMULATION	ABOVE PLUS: TURBULENCE STRUCTURE AERODYNAMIC NOISE	10^9	3000 x CLASS VI
FULL NAVIER-STOKES	ABOVE PLUS: LAMINAR/TURBULENT TRANSITION TURBULENCE DISSIPATION	10^{12} TO 10^{15}	3 MILLION TO 3 BILLION CLASS VI

they fall far short of meeting the requirements for problems
dominated by strongly coupled viscous effects. The mext major
step forward for the discipline of computational aerodynamics
can be taken when computers about 30 times more powerful than
the Class VI machines become available.

3. PROGRAM HISTORY

The Numerical Aerodynamic Simulation Program had its
genesis in 1975 at the NASA Ames Research Center. A small
group of people associated with the computational fluid dynam-
ics program proposed the development of a special-purpose
Navier-Stokes Processing Facility which would have several
orders of magnitude more speed than the ILLIAC IV computer for
solving the equations of fluid physics. The proposal called
for the central processor to have a minimum effective speed of
one billion floating point operations per second when operat-
ing on the three-dimensional Reynolds-averaged Navier-Stokes
equations and have performance comparable to the best general-
purpose computers when used for processing the equations of
other scientific disciplines. Its working memory had to
accommodate a problem data base of at least 31×10^6 64-bit
words. To keep development risks low, the goal of the project

7

was to assemble existing technologies into a specialized architecture rather than to develop new electronic components. Finally, the machine had to be user-oriented, easy to program, and capable of detecting systematic errors when they occurred. The proposal was endorsed in principle by NASA management in Nov. 1975 and the name of the project was changed to the Computational Aerodynamic Design Facility.

3.1 Computational Aerodynamic Design Facility Project

The first formal exposure of NASA's objectives occurred in October 1976 when proposals were requested from industry to "perform analysis and definition of candidate configurations for a computational facility in order to arrive at the best match between computational aerodynamic solution methods and processor system design." These analyses were to be directed toward the selection, preliminary design, and evaluation of candidate system configurations that would be best suited to the solution of given aerodynamic flow models. Design requirements that were established for this study included (1) the capability to complete selected numerical solutions of the Navier-Stokes equations for grid sizes ranging from 5×10^5 to 1×10^6 points and wall clock times (exclusive of input-data preparation and output-data analysis) ranging from 5 to 15 min; (2) a working memory of 40×10^6 words; (3) an archival store of at least 10×10^9 words; and (4) 120 hr/wk of availability to the users.

Two parallel contracts were awarded in Feb. 1977 to develop preliminary designs for the most promising configurations and to develop performance estimates, risk analyses, and preliminary implementation cost and schedule estimates for each of these designs. During these initial studies, which lasted about 12 mo, it became apparent that the overall approach to developing the facility was sound and that performance goals could be reached with new architectural concepts and proven electronic components.

3.2 Numerical Aerodynamic Simulation Facility Project

When it was recognized that the facility would be used primarily for computational research rather than for actual aircraft design, the name was changed to the Numerical

8

Aerodynamic Simulation Facility shortly after the first study contracts were awarded. It also became apparent that a computational resource of this magnitude would be a valuable tool for the solution of complex problems in other technical areas of interest such as meteorology, chemistry, astrophysics, and structures. Plans were made at that time to survey the aerodynamics community and technical universities to ascertain the level of interest in the projected facility and to ensure that design goals were consistent with the needs of the projected users of the facility.

Before the conclusion of the first round of contracted efforts, the need for further studies with greater emphasis on a computer suitable for a broader range of disciplines was recognized. Accordingly, 12-mo follow-on feasibility study contracts were awarded in March 1978. The results of these efforts were expected to provide data of sufficient accuracy to permit formulation of a definitive plan for development of the facility. Several events occurred during the period of these studies which resulted in some revisions to the basic performance specifications and a deeper involvement of the user community in the project activities.

The discipline of computational aerodynamics had matured significantly in the 3 yr since the project was first conceived. New numerical methods were developed and existing methods were refined. This led to the realization that if the size of the on-line or working memory was increased to 240×10^6 words, the facility could be used not only to estimate the performance of relatively complete aircraft configurations, but also to serve as an effective tool to study the physics of turbulent flows, a subject eluding researchers for more than 80 yr. A corresponding increase in the off-line file storage from 10×10^9 to approximately 100×10^9 words was required to accommodate the larger data sets.

A User Steering Group was formed in 1978 to provide a channel for the dissemination of information regarding project status, a forum for user-oriented issues needing discussion, and a sounding board by which the project office could obtain feedback from future user organizations. Examples of user-oriented issues of interest are (1) selection of user

9

languages, (2) management policy, (3) equipment required for remote access, and (4) data protection. The User Steering Group was composed of representatives of the aerospace industry, universities, and other government agencies. It is still active although its name was changed recently to the User Interface Group to more accurately reflect its current role.

The feasibility studies were completed in the spring of 1979. Each study produced a refined baseline configuration, a functional design, and rough estimates of cost and schedule. Both studies concluded that ~5 yr would be required to complete the detailed design and to develop, integrate, and test the facility. While preparations were being made to continue the contracted development process, the name of the project was changed once again to the Numerical Aerodynamic Simulator (NAS) Project.

3.3 Numerical Aerodynamic Simulator Project

A detailed plan for the design-definition phase of activity was prepared during the winter of 1979 by the NAS Project Office, which was established at Ames earlier in the year. This included refining (1) the specifications for the computing engine, (2) the support processing system, and (3) the collection of other peripherals, including intelligent terminals, graphical display devices, and data communication interfaces to local and remote users. Two 40-wk parallel design-definition contracts were awarded in Sept. 1980. Upon their completion in July 1981, the contractors were awarded follow-on contracts related to further design definition. These were concluded in April 1982 when the proposals for the detailed design, development, and construction were submitted by the contractors for evaluation.

After an initial evaluation of the proposals, the decision was made to abandon the previously defined approach to meeting the requirements of computational aerodynamics and to chart a new course. Several factors contributed to this decision. First, the application and essential importance of computational aerodynamics had grown significantly since the mid-1970s. Thus, it was deemed important to establish and sustain a leading-edge computational capability as an essential step toward maintaining the nation's leadership in

10

aeronautics. To achieve this goal the NAS Project was to be restructured as an on-going NAS Program in which significant advances in high-speed computer technology would be continuously incorporated as they become available.

Second, the supercomputer environment had changed since the inception of the NAS activity in the 1970s. Increased interest in supercomputing, along with increased foreign competition, changed the environment to the extent that it no longer appeared necessary to directly subsidize the development of the next generation of scientific computers. Rather, the current environment permits a more systematic, evolutionary approach toward developing and maintaining an advanced NAS computational capability.

Finally, it was recognized that advancement in the state of the art of supercomputers must be coupled with advanced system networks and software architectures. This is necessary to accommodate successive generations of supercomputers from different vendors and to provide the capabilities needed to enhance productivity for the user. This leads to a strategy that minimizes the dependence of the entire system on single vendors and leads to the establishment of a strong in-house technical capability to direct the initial and ongoing development efforts.

3.4 Numerical Aerodynamic Simulation Program

A plan for the redefined program was approved in Feb. 1983. It includes (1) the design, implementation, testing and integration of an initial operating configuration of the NAS Processing System Network; (2) the systematic and evolutionary incorporation of advanced computer-system technologies to maintain a leading-edge performance capability; and (3) the management and operation of the complex. The principal goals of the program are summarized as follows.

General goals

Optimize the computational process from problem formulation to publication of results.

11

Specific goals

    1985

        Speed — 250 MFLOPS, sustained

        Memory — at least 64×10^6 64-bit words

        Users — local and remote

    1987

        Speed — 1 GFLOPS, sustained

        Memory — 256×10^6 64-bit words

        Users — support at least 100 simultaneously on
           a time-sharing interactive basis

        Operating system and network — capable of
           accommodating a multivendor hardware
           environment

    Beyond 1987

        Continue to expand capability

4. NAS PROCESSING SYSTEM NETWORK

The planned NAS Processing System Network (NPSN) will be
a large-scale, distributed resource computer network at Ames
Research Center. This network will provide the full end-to-
end capabilities needed to support computational aerodynamics,
will span the performance range from supercomputers to
microprocessor-based work stations, and will offer functional
capabilities ranging from "number-crunching" interactive
aerodynamic-flow-model solutions to real-time graphical-
output-display manipulation. The NPSN resources will be made
available to a nationwide community of users via interfaces to
landline and satellite data communications links.

The NAS program is structured to accommodate the continu-
ing development of the NPSN as a leading-edge computer-system
resource for computational aerodynamics. This development
process is dependent upon the acquisition and integration of
the most advanced supercomputers industry can provide that are
consistent with computational aerodynamics requirements.
Figure 3 illustrates the continuing development of the NPSN
functional and performance capabilities while successively
introducing advanced high-speed processors into the network.
The introduction of each new high-speed processor involves an
integration phase in which new software and interfaces are
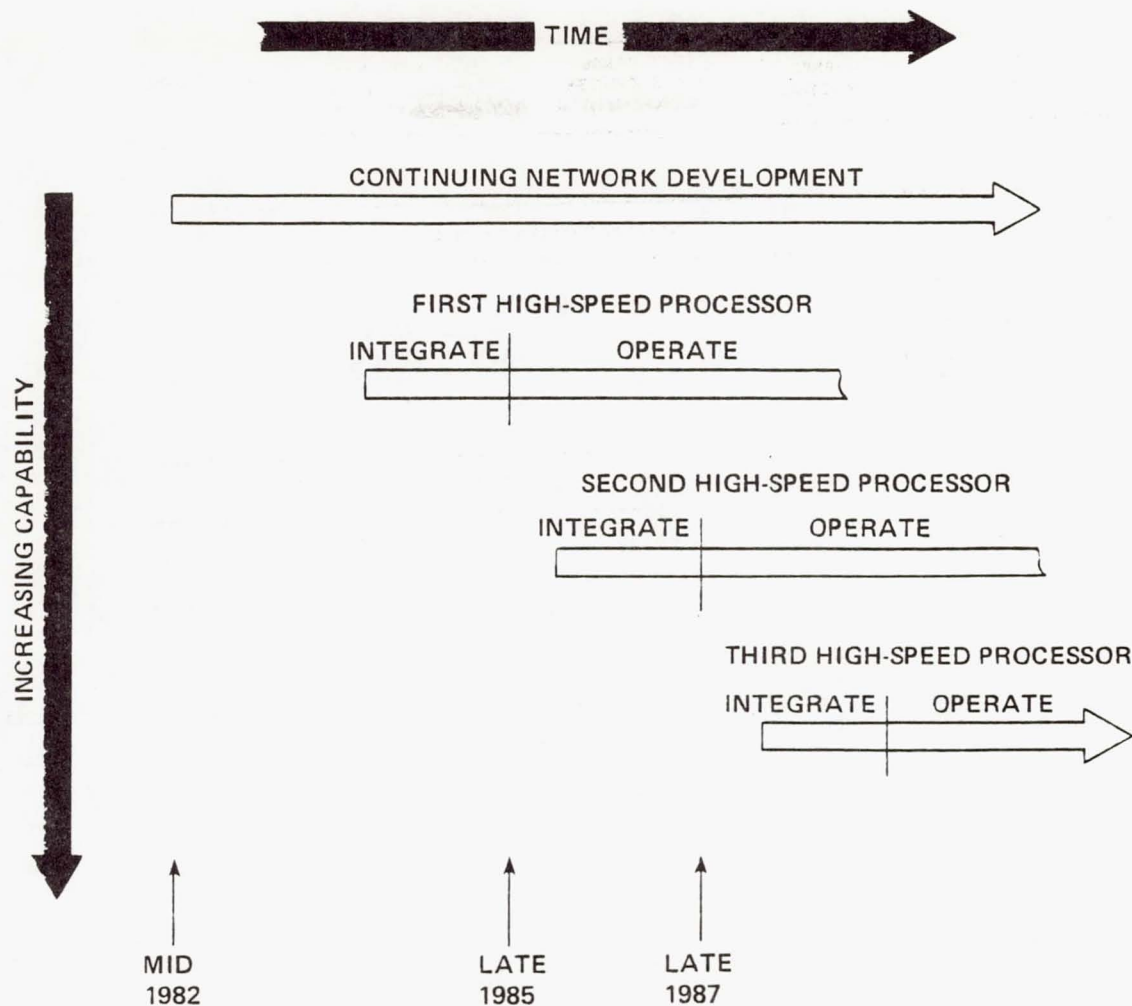
12

Fig. 3. Evolution of the NAS System.

implemented and tested, followed by an operational phase. An
important element in this evolutionary strategy is the early
implementation and testing of a fully functional NPSN designed
to accommodate new supercomputers from different vendors with
a minimum impact on the existing network architecture and on
operational use.

As shown in Figure 4, the NPSN will consist of the fol-
lowing eight 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 Communications Subsystem (LHCS)

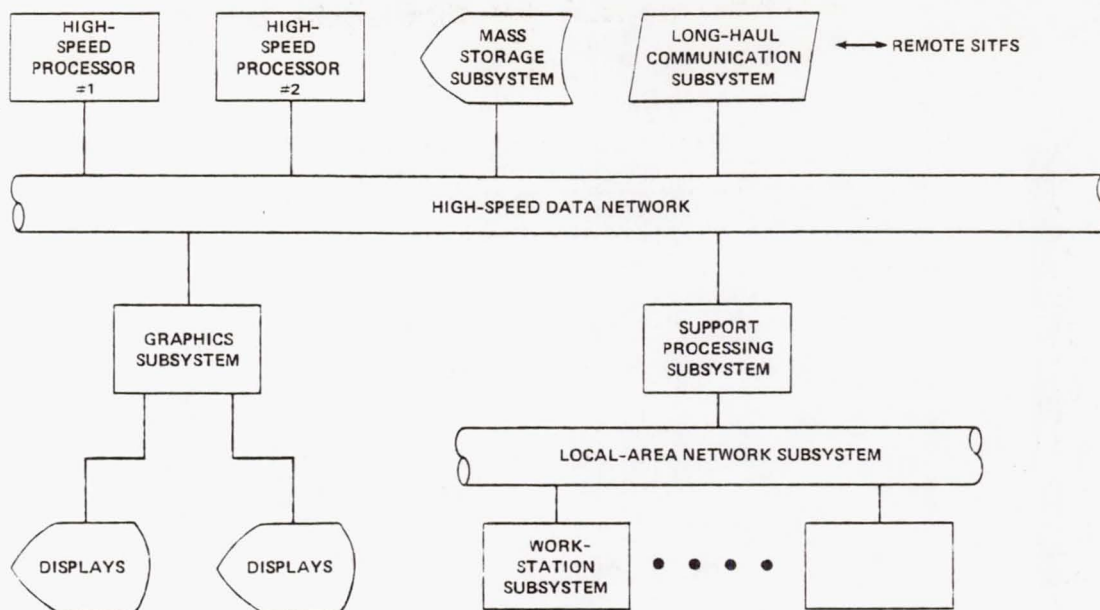7. High-Speed Data-Network Subsystem (HSDN)

13

Fig. 4.   Schematic of Fully Developed NAS Processing System
Network.

8.   Local-Area Network Subsystem (LANS)

Only the HSP, SPS, WKS, and GRS will be programmed by users.
The MSS and LHCS will provide a mass storage facility and a
remote data-communications interface, respectively.   The HSDN
and LANS will provide a data-transport function for the other
subsystems.

4.1   High-Speed Processor Subsystem (HSP)

The HSP is the advanced scientific computing resource
within the NPSN.   The purpose of this subsystem is to provide
the computational throughput and memory capacity to compute
computational aerodynamics simulation models.   In addition to
batch processing, interactive time-sharing processing will be
provided to aid in application debugging, result editing, and
other activities that depend on close user-processing coupling
to achieve optimum overall productivity.

Present plans call for two generations of HSP computers
to be in the system at any one time.   The first (HSP-1),
planned for integration in 1985, will provide a capability to
process optimally structured computational aerodynamics appli-
cations at a sustained rate of 250 MFLOPS with a minimum
64×10^6 word memory capacity.   The second (HSP-2), planned for
integration in 1987, will increase these values to 1000 MFLOPS
and a 256×10^6 word memory capacity.

14

4.2 Workstation Subsystem (WKS)

Whereas the HSP is the ultra-high-speed, large-scale computer resource serving the global user community, the WKS is the microprocessor-based resource used by the individual researcher. The WKS will provide a "scientist's work-bench" for local users to perform text and data editing, to process and view graphics files, and to perform small-scale applications processing. Each individual workstation will have the appropriate memory, disk storage, and hard-copy resources to fit the local user's needs. Individual clusters of workstations will be networked together via the LANS for use within local user groups. In addition to local processing, the WKS will provide terminal access to other user-programmed systems and a file-transfer capability via the LANS and HSDN.

4.3 Support Processing Subsystem (SPS)

The SPS is a multi-super-minicomputer-based system providing a number of important functions. The SPS will provide general-purpose interactive processing for local and remote terminal-based users (i.e., those without workstations), and provide an intermediate performance resource between the HSP and WKS performance as a WKS backup. The SPS will be a gateway between the HSDN and the LANS, the location for unit record input/output devices such as high-speed printers and microfilm, and the focal point for network monitoring and system operation.

4.4 Graphics Subsystem (GRS)

The GRS is a super-minicomputer-based system that will provide a sophisticated state-of-the-art graphical-display capability for those applications requiring highly interactive, high-density graphics for input preparation and result analysis. The GRS will provide a level of performance and storage capability beyond that provided by workstations and will be shared by first-level user organizations.

4.5 Mass Storage Subsystem (MSS)

The MSS will provide the global on-line and archival file storage capability for the NPSN. This subsystem will validate and coordinate requests for files to be stored or retrieved within the NPSN and maintain a directory of all contained data.

The MSS will act as a file server for other NPSN subsystems; control its own internal devices; and perform file duplication, media migration, storage allocation, accounting, and file management functions.

Users of the NPSN will create and use files on various subsystems, e.g., HSP, SPS, GRS, or WKS. However, after the user has exited the NPSN, the main repository of these files will be the MSS. This subsystem will hold those very large files that will be used as input to, or generated as output from, the largest tasks that will be processed on the HSP and GRS. It will contain user source and object codes, and parameter and data files that are kept for any significant length of time. The MSS will also contain the backup files that are created to improve the probability that long-lasting or high-value files are accessible when needed.

4.6 Long-Haul Communication Subsystem (LHCS)

The LHCS will provide the data-communication interface between the NPSN and data-communication links to sites geographically remote from Ames Research Center. This subsystem will provide the necessary hardware/software interfaces; modulation and demodulation devices; and recording, processing, data buffering, and management functions to support data transfers and job control by remote users.

In the sense that the MSS is a back-end resource for the entire NPSN, the LHCS is a front-end resource. It provides for access by remote users to the HSP, SPS, GRS, and WKS but it is not specifically addressed or programmed by the user. The LHCS processor functions as a data-communications front-end providing store and forward, protocol conversion, and data-concentrator service.

Current plans call for the LHCS to interface with data links capable of providing 9600 bits/sec to over 1.5×10^6 bits/sec transmission rates in order to interface with government-sponsored networks (e.g., ARPANET, and the proposed NASA Program Support Communication Network) and commercial tariffed services. Candidate data-communication protocols to be supported include ARPANET, IBM System Network Architecture (SNA), and Digital Equipment Corporation's network (DECnet).

16

4.7 High-Speed Data-Network Subsystem (HSDN)

The HSDN provides the medium over which data and control messages are exchanged among the HSP, SPS, GRS, MSS, and LHCS. Major design emphasis will be placed on the ability to support large file transfers among these systems. The HSDN will include high-speed (minimum 50 megabits/sec) interface devices and driver-level network software to support NPSN internal data communications.

4.8 Local-Area Network Subsystem (LANS)

The LANS will provide the physical data transfer path between the SPS and WKS, and between various workstations within a WKS cluster. The LANS will be designed to support up to 40 workstations and to provide a hardware data-communications rate of at least 10 megabits/sec (e.g., Ethernet).

The LANS differs from the HSDN in data-communication bandwidth because of the smaller size of files transferred on the LANS and the lower cost per LANS network interface device. The HSDN and LANS will use the same network protocol.

5. NPSN SOFTWARE

The NPSN will include a rich set of systems and utility software aimed at providing the most efficient and user-productive environment practical. Major software objectives are as follows:

1. A vendor-independent environment that allows the incorporation of new technologies without sacrificing existing software commitments

2. Common and consistent user environments across the NPSN

3. The maximum transparency to the heterogeneous subsystem nature of the NPSN

4. The optimal performance for critical resources

5. A rich set of user tools

The strategy taken to satisfy these objectives is to use a UNIX^{TM*} or UNIX^{TM} look-alike operating system on the user

---

*The use of trade names of manufacturers in this report does not constitute an official endorsement of such products or manufacturers, either expressed or implied, by the National Aeronautics and Space Administration.

17

programmable subsystems (HSP, SPS, GRS, and WKS). This
approach is the first attempt to achieve vendor independence
and a common user environment, and will be implemented by a
combination of native UNIX^{TM}, vendor UNIX^{TM} look-alikes, and
virtual operating system approaches using Software Tools. The
UNIX^{TM} approach provides for the implementation of a rich set
of user tools from text editors and compilers to graphics
packages that are transportable among systems. This approach
also provides a degree of transparency among subsystems. Ven-
dor independence and high performance are aided by the imple-
mentation of a highly functional and efficient network proto-
col, such as the Livermore Interactive Network Communication
System protocol.

As the NPSN evolution continues, further gains in meeting
these objectives in the areas of automatic-file and data-
format conversion, common network directory, architecture-
independent programming languages, and architecture-specific
optimization from ANSI languages will be forthcoming.

6. PROJECTED USAGE ALLOCATION

A NAS system usage policy has been proposed and will be
implemented when the initial configuration becomes operational
in 1985. It is expected that 90% of the available usage would
be allocated to the disciplines of fluid dynamics and aerody-
namics. The remaining 10% would support advances in such dis-
ciplines as chemistry, atmospheric modeling, aircraft struc-
tural analysis, and astrophysics. About 55% of the high-speed
processor time would support NASA programs and would be used
by the in-house staff and organizations involved in NASA
grants, contracts, and joint programs with other organizations.
Roughly 20% of the high-speed processor time would support
Department of Defense activities, 15% would be allocated to
industry, and 5% would be allocated to universities. Costs
for industry proprietary work would be reimbursed to the gov-
ernment. The reimbursement formula would include operating
and capital improvement costs. This policy is consistent with
similar NASA guidelines governing wind-tunnel usage. Usage
requests would be screened by Ames line management in a manner
similar to the review process governing the Unitary Plan Wind
Tunnel operations. Potential users must demonstrate a

requirement for the unique capabilities offered by the NAS system. In addition, task selection will be guided by the following criteria: technical quality, national need, timeliness, and funding support.

The guidelines for usage allocation described here apply only to NAS computers in an operational status. During the time that new computers are being integrated into the system, they will be devoted primarily to systems integration, testing, and software development.

7. BENCHMARK TEST SET

Sets of benchmark codes will be constructed to support the system integration and test activity. They will be designed to stress each of the subsystems, which is their primary purpose. However, these codes will also be designed to support pioneering scientific and engineering investigations. This opportunity provides a strong impetus for computational physics researchers to take an active interest in the development of the Benchmark Test Set and in the systems development in general.

Examples of candidate problems under consideration for the Benchmark Test Set include the following: (1) vortex formation, roll-up, and interaction with lifting surfaces; (2) effects of wall motion and geometry on turbulent skin friction; (3) simulation of boundary-layer transition; (4) simulation of afterbody aerodynamics, including effects of real-gas exhaust plumes; (5) simulation of observed nonlinear transonic aeroelastic instabilities of the B-1 wing; (6) numerical optimization of the wing for a transonic transport aircraft; (7) Aero-assisted Orbital Transfer Vehicle (AOTV) flow-field simulation, including real-gas and nonequilibrium radiation physics; (8) complete shuttle orbiter flow-field characteristics; and (9) simulation of crack initiation and propagation. These problems represent a balance of fundamental fluid physics research and applied computational aerodynamics and aerothermodynamics, except the last example listed, which is a computational chemistry investigation.

One of the fluid physics investigations is highlighted in Fig. 5. The basic concept involves the manipulation of the turbulent structure in the vicinity of a wall (or aerodynamic
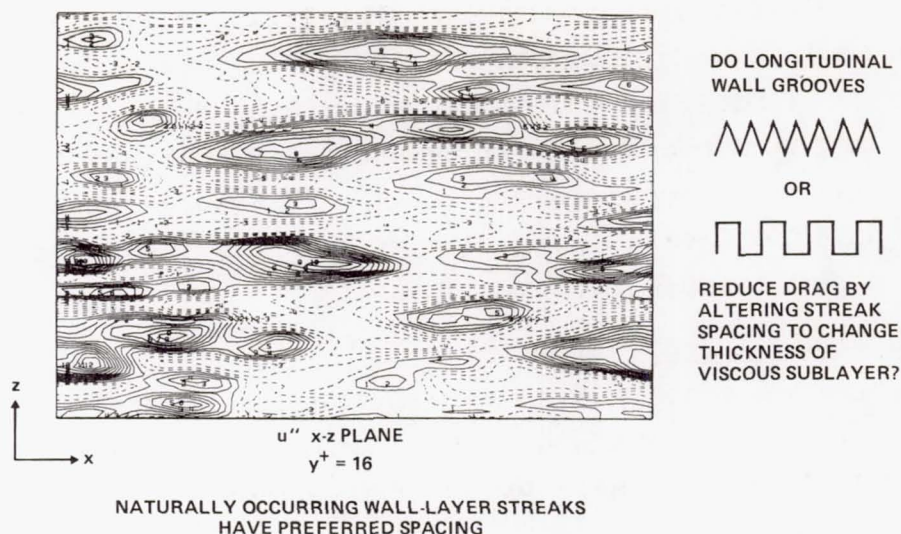
19

DO LONGITUDINAL
WALL GROOVES

OR

REDUCE DRAG BY
ALTERING STREAK
SPACING TO CHANGE
THICKNESS OF
VISCOUS SUBLAYER?

u" x-z PLANE
$y^+ = 16$

NATURALLY OCCURRING WALL-LAYER STREAKS
HAVE PREFERRED SPACING

Fig. 5. Candidate Fluid Physics Problem Involving the Manip-
ulation of Turbulence Structure Near a Boundary.

surface) to achieve net skin-friction reduction. Of particu-
lar interest is the riblet or wall-groove approach, in which
finely spaced longitudinal grooves are inserted on the wall
surface to alter the formation and growth of turbulent sub-
layer streaks. Experimental investigations indicate up to a
20% drag reduction, which would produce enormous benefits in
fuel savings and an increased range for transport aircraft.
To alter the formation of sublayer streaks, the spacing of the
grooves should be of the same order as the individual streak
dimensions. Even at moderate Reynolds numbers, a large number
of grid points are needed to resolve the riblets and the flow
about them, requiring an order-of-magnitude increase in com-
puter power over the Class VI machines. The code would be
virtually entirely vectorizable. Storage requirements are
23 words/grid point for a total of 120×10^6 words of data
that must be transferred to central memory at each integration
time step. A three-dimensional, large-memory color-graphics
capability is also needed.

One of the candidate problem examples is illustrated in
Fig. 2. The object in this work is to compute the flow field
about boattailed afterbodies containing propulsive jets and to
study the interactive influence of highly underexpanded

20

exhaust plumes on afterbody drag. The aerodynamic drag associated with jet-engine afterbodies represents a significant fraction of the recoverable drag inherent in the flight of modern aircraft. Because of the very strong interactive and complex nature of the flow, in-depth analysis has not been possible with current computer power. Expensive experimental investigations have been conducted to provide an understanding of the complex flow patterns involved. Even these studies, however, are not sufficiently detailed to adequately describe the strong interaction between the supersonic propulsive jet and the external flow at the base of the afterbody. This interaction produces flow regions of very rapid expansion accompanied by complex recompression patterns and contact discontinuities inside the exhaust-plume outer boundary. The flow over the afterbody separates in a three-dimensional pattern, contributing to the drag. Powerful computational resources would be required to simulate this flow field at flight Reynolds numbers. Typically, discretization of the computational control volume would require about 1×10^6 grid points and 17×10^6 words of memory. The coding is virtually entirely vectorizable, with vector lengths of about 55,000.

These two examples are representative of the candidate list for the NAS Benchmark Test Set. The final test set will be identified at a later time, after extensive analysis has been made to determine the suitability of each of the candidate codes for testing the various NAS subsystems.

8. CONCLUDING REMARKS

Computers are beginning to have a profound effect on aerodynamics, a technical discipline important to the design of all aerospace vehicles. Computational aerodynamics, although severely limited by available computer power, is already emerging as an important aircraft design tool. Successively more refined aerodynamic simulations are resulting from increases in computer power. Strong motivating factors to accelerate the development of computational aerodynamics exist, but a computational system having performance far in excess of that available today is required. The goal of NASA's NAS program is to provide an advanced capability by the mid to late 1980s that will be available for use by

21

government laboratories, industry, and academia. This capability will be continually upgraded beyond that date as computer technology advances.

The currently planned NAS processing system network is based on about 8 yr of study. Each of eight elements of the system has been defined by a set of initial performance goals. In addition, network software strategies for achieving vendor independence and common user environments have been developed. A NAS system usage policy has been proposed for implementation when the initial configuration becomes operational before 1986. Finally, candidate problems in several disciplines besides aerodynamics have been identified for the benchmark test codes.

REFERENCES

1. Magnus, R. and Yoshihara, H., Inviscid transonic flow over airfoils, AIAA J. 8 (1970), 2157-2162.

2. Boppe, C. W. and Stern, M. A., Simulated transonic flows for aircraft with nacelles, pylons and winglets, AIAA Paper 80-0130, Jan. 1980.

3. 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, 151-163.

4. Deiwert, G. S., A computational investigation of supersonic axisymmetric flow over boattails containing a centered propulsive jet, AIAA Paper 83-0462, Jan. 1983.

5. Agrell, J. and White, R. A., An experimental investigation of supersonic axisymmetric flow over boattails containing a centered propulsive jet, FFA Tech. Note AU-913, 1974.

NASA Ames Research Center
Moffett Field, California 94035

1. Report No. NASA TM-84386	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle NUMERICAL AERODYNAMIC SIMULATION (NAS)		5. Report Date July 1983
		6. Performing Organization Code
7. Author(s) V. L. Peterson, W. F. Ballhaus, Jr., and F. R. Bailey		8. Performing Organization Report No. A-9404
9. Performing Organization Name and Address NASA Ames Research Center Moffett Field, Calif. 94035		10. Work Unit No. T-4243
		11. Contract or Grant No.
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, D.C. 20546		13. Type of Report and Period Covered Technical Memorandum
		14. Sponsoring Agency Code 506-51-11

16. Abstract

The history of the Numerical Aerodynamic Simulation Program, which is designed to provide a leading-edge capability to computational aerodynamicists, is traced back to its origin in 1975. Factors motivating its development and examples of solutions to successively refined forms of the governing equations are presented. The NAS Processing System Network and each of its eight subsystems are described in terms of function and initial performance goals. A proposed usage allocation policy is discussed and some initial problems being readied for solution on the NAS system are identified.

17. Key Words (Suggested by Author(s)) Computational fluid dynamics Computational aerodynamics Computer systems	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 25	22. Price* A02

*For sale by the National Technical Information Service, Springfield, Virginia 22161