NASA Technical Memorandum **85965**

NASA-TM-85965

Application of Supercomputers to Computational Aerodynamics

Victor L. Peterson

June 1984

NASA

National Aeronautics and
Space Administration

NF00824

Application of Supercomputers to Computational Aerodynamics

Victor L. Peterson, Ames Research Center, Moffett Field, California

N84-27680

SUMMARY

Computers are playing an increasingly important role in the field of aerody-
namics such that they now serve as a major complement to wind tunnels in aerospace
research and development. Factors pacing advances in computational aerodynamics are
identified, including the amount of computational power required to take the next
major step in the discipline. Example results obtained from the successively
refined forms of the governing equations are discussed, both in the context of
levels of computer power required and the degree to which they either further the
frontiers of research or apply to problems of practical importance. Finally, the
Numerical Aerodynamic Simulation (NAS) Program — with its 1988 target of achieving
a sustained computational rate of 1 billion floating point operations per second and
operating with a memory of 240 million words — is discussed in terms of its goals
and its projected effect on the future of computational aerodynamics.

INTRODUCTION

Computational aerodynamics is emerging as an important new discipline even
though it is still limited by currently available computer power. The amount of
computer power (speed and memory) required for computational fluid dynamics depends
on three factors: (1) the complexity of the problem, which includes the fluid
physics governing the problem to be solved and the complexity of the geometry about
which the fluid moves; (2) the efficiency of the algorithm available for solving
the governing equations numerically; and (3) the amount of computational time that
can be invested to solve the problem at hand. The power of the largest currently
available supercomputers (Class VI machines), such as the CYBER 205 and the
CRAY X-MP, is far less than that required to solve the most complex problems involv-
ing the full Navier-Stokes equations and realistic aircraft shapes in a practical
amount of computational time, even though these machines can sustain computational
rates greater than 100 million floating point operations per second (MFLOPS) and
can have electronic memories in excess of 10 million words. Thus, the computational
fluid dynamicist is limited by available computer power and must resort to making
approximations either to the physics, to the geometry, or to both in order to obtain
numerical solutions to this problem.

The relationships between problem complexity, algorithm efficiency, and compu-
tational time are reviewed in terms of the power of computers of the past and
present. This review will serve to highlight the factors influencing computer
requirements and to provide some insight into the prospects for the future. In addi-
tion, example results of current capabilities are presented to provide a reference
upon which to gauge the current capabilities of computers to solve aero- and fluid-
dynamic problems. Finally, the Numerical Aerodynamic Simulation Program undertaken
by NASA to accelerate the development and application of computational aerodynamics
is discussed.

PROBLEM COMPLEXITY

The Navier-Stokes equations, although known to contain approximations when they are compared to the more basic nonlinear Boltzman equation which governs molecular transport, are definitely suitable for solving fluid dynamic and aerodynamic problems in the continuum flow regime. They contain the physics governing all scales of turbulence of which between four and five decades of scales are of practical importance, depending upon the Reynolds number. The equations are highly nonlinear and are strongly coupled, and their solution is further complicated by the boundary conditions associated with complex aircraft shapes.

It is not possible to obtain closed-form solutions to these full equations for situations of interest in aircraft design. Even with the level of computational power available today, the complete equations cannot, within practical constraints, be solved numerically in their full form, except in highly simplified situations. Therefore, various degrees of approximation have been evoked over the years to obtain useful results.

Long before the advent of digital computers, mathematicians and aerodynamicists devised methods for solving the linearized equations governing inviscid flows that are universally either subsonic or moderately supersonic. This was possible because only three partial derivative terms are retained in the equations. The first solutions for two-dimensional airfoils were obtained in about 1930 and for three-dimensional wings in about 1940. With the advent during the 1960s of computers of the IBM-360 and CDC-6600 class, it became practical to compute inviscid flows about somewhat idealized complete aircraft configurations. This level of approximation is still heavily used today, but because the equations neglect all viscous terms as well as inviscid nonlinear terms, such calculations provide only somewhat limited help in the design process. For many years, it has been customary to increase the value of this level of approximation by making corrections for viscous effects based on the so-called boundary-layer theory. This approach works reasonably well when flows do not encounter severe adverse pressure gradients and remain consistently attached to surfaces over which they pass.

The advent of the computer brought the next level of approximation into the realm of practical usefulness. By including the nonlinear terms in the equations, it is possible to treat flows at all Mach numbers, including transonic and hypersonic, although viscous terms are still neglected. Hand-calculated solutions for the transonic flow over a nonlifting airfoil were obtained in the late 1940s and over a lifting airfoil with detached bow wave at transonic speeds in the mid-1950s. However, the first transonic solutions for a practical lifting airfoil with embedded shock wave required the use of a computer and were not obtained until about 1970. This latter achievement can be interpreted as the turning point in computational aerodynamics, for it marked the beginning of a series of advances that would not have been possible without computers.

The next level of approximation, now in the stages of vigorous development, does not neglect any terms in the Navier-Stokes equations. The basic equations are averaged over a time interval that is long relative to turbulent eddy fluctuations, yet small relative to macroscopic flow changes. Such a process introduces new terms representing the time-averaged transport of momentum and energy which must be modeled using empirical information. Computers are required to work with this level of approximation, but the potential advantages are enormous. Realistic simulations

of separated flows; of unsteady flows, such as buffeting; and of total drag should be possible as the ability to model turbulence matures. Combined with computer optimization methods, these simulations should make it possible to develop aerodynamically optimum designs. Landmark advances in the decade of the 1970s include the investigation of shock-wave interaction with a laminar boundary layer, the treatment of high Reynolds-number transonic airfoil flows, and the first three-dimensional laminar flow over an inclined body of revolution. Relatively large amounts of computer time are required using this level of approximation. Although two-dimensional flows can be computed in a matter of minutes on a Class VI computer with current numerical methods, the routine computation of three-dimensional flows is not yet practicable; they require about 100 times more calculations than two-dimensional flows.

The final level of approximation involves the direct numerical simulation of turbulent eddies over a range of scales sufficiently broad to capture the transport of nearly all momentum and energy. Only the small eddies that dissipate energy, transport little energy or momentum, tend toward isotropy, and are nearly "universal" in character are modeled. Under such conditions, the computed results involve essentially no empiricism. This level of approximation is being used today on a limited basis to research the physics of turbulence at a level of detail not possible through experiment. Computer times for simulations of very basic flows, such as those in channels, range up to 80 hr on current supercomputers. Ultimately, this work should provide information that will lead to improved methods for modeling turbulence in the level III approximation, as well as to unlock the secrets of the generation of aerodynamic noise and the manipulation of friction drag.

Presently available machines are adequate for treating relatively complex configurations with the inviscid flow equations. Of course, the type of information derived from the computations is limited (e.g., no effects of flow separation). The viscous flow equations, being more complex and requiring fine computational meshes to resolve large flow gradients near surfaces and small scales of motion in regions of turbulence, demand substantially greater computational power to solve. Thus, the types of problems that can be solved with a given computer must be less complex. In effect, a designer has to make the choice between treating simple configurations with complex physics and treating complex configurations with simple physics. It is obvious, though, that in both the inviscid- and viscous-flow situations, each new generation of computers has resulted in corresponding advances in the value of computational aerodynamics as a design tool. The discipline will begin to mature when both complex configurations and complex physics can be treated together in a practical amount of time.

ALGORITHM EFFICIENCY

Considerable effort is being expended to develop improved numerical methods for solving the governing equations, particularly for levels of approximation including effects of nonlinear and viscous phenomena. Approaches based on finite-difference algorithms are the most popular, although work with the linearized inviscid equations is largely based on paneling methods, and work for the direct simulation of turbulent eddies frequently involves consideration of spectral methods.

Finite-difference methods require the use of a computational mesh engulfing the aerodynamic configuration and extending far enough in all directions to a point

3

at which outer boundary conditions can be expressed in terms of known quantities. Most meshes being used today are highly complicated nets that conform to geometric contours and involve stretching and clustering designed to alter grid spacings commensurate with the physical detail that must be captured to provide an accurate solution. The concept is to solve the governing equations for all of the flow variables for the parcels of fluid within each little volume formed by the grid. Obviously, if computers had unlimited speeds and memories, the grid could be refined to the point of resolving all scales of motion without approximation.

The thrust of the work in algorithm development in the past has been aimed at increasing (1) the degree of vectorization, (2) the economy of calculation at each time step, (3) the interval length between time steps, (4) numerical stability, (5) accuracy, and (6) ease of computer memory management. A considerable amount of attention is now being directed toward developing mathematical representations of complex aircraft geometries and automated computational grid generators. Advancements in algorithm development are keeping pace with advances being made in computer speed. Improvements in computers in the past 15 yr have reduced the time and cost of doing a computation with a given algorithm by a factor of almost 100. Over the same period of time, improvements in algorithms have reduced the time and cost of doing a computation with a given computer by almost a thousandfold. These advances compound to result in an overall increase in cost effectiveness and reduction in time required to obtain solutions of nearly 10^5.

Improvements in algorithm efficiency cannot be expected to continue indefinitely since there are absolute physical limits. At least one iteration will be required to resolve steady flows, and time-accurate solutions of unsteady flows will require the selection of time steps commensurate with the frequencies involved. Considering these limits, only one to two orders of magnitude of further improvement can be expected. Other factors are beginning to emerge, however, to create optimism about further improvements in software efficiency beyond these limits.

Heretofore, improvements in the efficiency of algorithms for solving the various forms of the equations have resulted in almost corresponding reductions in the time required to solve problems of interest. This correspondence is beginning to become less direct as problems being tackled become more complex. Numerical representations of complex three-dimensional geometries, automated optimal grid generation, and post-processing of the enormous result files for display and analysis are beginning to place significant burdens on the same computer that is used to solve the flow equations. These pre- and post-processing tasks are expected to account for a growing fraction of the computer power required to solve the important problems of the future. Because of these factors, we can expect that improvements in the more global problem-solving methodology will continue to keep pace with improvements in computer power for some time to come, even though limits to improvements in algorithm efficiency are in sight.

PRACTICAL SOLUTION TIMES

A perspective on the amount of computer time that is reasonable to invest in the solution of a problem can be obtained by reviewing information on the historical use of computers for computational fluid dynamics. Experience has shown that computational methods are not routinely used in the aerodynamic design environment unless the time required to obtain a simulation is about 10 to 15 min. Short computation

4

times are required for it to be practical to sort through many configurations early in the design cycle when aerodynamic factors can have the largest effect on the shape of a new aircraft. Simpler forms of the aerodynamic equations, such as the two-dimensional inviscid nonlinear or the three-dimensional inviscid linear, can be solved in 10 min or less on machines of the IBM-360 or CDC-6600 class. When machines of this class were first made available to the research community, where it is not uncommon to invest several hours of computer time to obtain a single solution, these machines were used to pioneer solution methods for the more complicated three-dimensional inviscid nonlinear equations. Then, as more powerful machines of the CDC-7600 class became available, it became possible for industrial designers to routinely use the two-dimensional inviscid nonlinear methods while the researchers moved on to develop methods for solving the next higher level of approximation of the governing equations. Class VI computers, such as the CRAY-1S and the CYBER-205, are adequate to routinely solve the two-dimensional, Reynolds-averaged, Navier-Stokes equations and to research extensions to three dimensions, but they fall short of making the three-dimensional viscous-flow simulations practical for design work.

FUTURE COMPUTER REQUIREMENTS

More powerful computers are required to realize the full potential of computational aerodynamics. An estimate of the computational speed that is necessary to compute a flow with the three-dimensional Reynolds-averaged Navier-Stokes equations using 10^7 grid points and based on the use of improved algorithms projected to be available in 1985, show that a computer must perform at least 1 billion floating-point operations per second on a sustained basis in order to complete the simulation in 10 min. It is interesting to note that the solution of the same problem would take about a month on an IBM-360, a day on a CDC-7600, and several hours on the Class-VI machines. Implicit in these results is the assumption that data-transfer rates to and from memory are fast enough to keep pace with the processing units.

The other aspect of computational power that must be considered is memory. About 27 words of memory at each iteration are required for each grid point used in Reynolds-averaged Navier-Stokes simulations. Thus, a calculation involving 10 million grid points would require high-speed access to about 270 million words.

A computer having a high-speed memory ranging in size from 250 to 300 million words and a sustained computational rate of at least 1 billion floating-point operations per second will bring the use of the Reynolds-averaged Navier-Stokes equations into the design environment. In addition, it would provide the tool required to research applications of direct numerical simulation of turbulent eddies, the highest level of approximation to the complete Navier-Stokes equations.

Additional background information on the value of computations to the discipline of aerodynamics and the importance of developing more computational power is presented in reference 1.

THE NUMERICAL AERODYNAMIC SIMULATION PROGRAM

Definition of the Numerical Aerodynamic Simulation (NAS) Program was undertaken by NASA in 1975 to provide not only the computer power required to take the

next major step in the development of computational aerodynamics, but also to provide a strategy for meeting continually increasing requirements for the indefinite future. After intervening years of study, this program was approved in 1983 and implementation is now under way. The principal goals of the program are summarized as follows:

General Goal: optimize the computational process from problem formulation to publication of results.

Specific Goals:

    1986:   Speed – 250 MFLOPS, sustained
             Memory – at least 64×10^6 64-bit words
             Users – local and remote

    1988:   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 1988:   continue to expand capability

The approach to meeting these goals is to develop a NAS Processing System Network (NPSN) which will be a large-scale, distributed resource computer network at the 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 land-line and satellite data communications links.

The NAS program is structured to accommodate the continuing 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. Provisions are being made for 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 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.

The NPSN will consist of the following 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)

8. Local-Area Network Subsystem (LANS)

Each of these subsystems is briefly described in the following paragraphs.

High-Speed Processor Subsystem

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 1986, will provide a capability to process optimally structured computational aerodynamics applications at a sustained rate of 250 MFLOPS with a minimum 64×10^6 word memory capacity. The second (HSP-2), planned for integration in 1988, will increase these values to 1000 MFLOPS and a 256×10^6 word memory capacity.

Support Processing Subsystem

The SPS is a multiple computer 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 the 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.

Workstation Subsystem

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 workbench" 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.

Graphics Subsystem

The GRS is a superminicomputer-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.

Mass Storage Subsystem

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 will be 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.

Long-Haul Communications Subsystem

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 transfer 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 commerical tariffed services. Candidate data-communication protocols to be supported include ARPANET, IBM System Network Architecture (SNA), and Digital Equipment Corporation's network (DECnet).

High-Speed Data Network Subsystem

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

8

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.

Local-Area Network Subsystem

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.

SUMMARY

The advent of computers marked the beginning of a systematic thrust to solve
increasingly complex forms of the Navier-Stokes equations which govern the motion
of fluids. By 1970 computers became powerful enough to begin to provide solutions
to problems not practical to treat either analytically or with human-powered calcula-
tors. Since that time, advances in solution methodology and computer power have per-
mitted computational aerodynamics to emerge as an important aircraft design tool.
The largest available computers and the most advanced methods now permit the flow
about a complete aircraft to be calculated in 10 to 15 min with the inviscid non-
linear form of the governing equations. However, much more powerful computers are
required to routinely solve three-dimensional problems for which viscous effects
are important. The goal of NASA's Numerical Aerodynamic Simulation Program is to
develop the most advanced computational capability possible by the mid to late
1980s and to make this capability available for use by government laboratories,
industry, and academia. This capability will be continually upgraded as computer
technology advances.

REFERENCE

1. Peterson, Victor L.: Impact of Computers on Aerodynamics Research and
       Development. Proceedings of the IEEE, vol. 72, no. 1, Jan. 1984, pp. 68-79.

1. Report No. NASA TM-85965	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle APPLICATION OF SUPERCOMPUTERS TO COMPUTATIONAL AERODYNAMICS		5. Report Date June 1984
		6. Performing Organization Code ATP
7. Author(s) Victor L. Peterson		8. Performing Organization Report No. A-9762
9. Performing Organization Name and Address Ames Research Center Moffett Field, CA 94035		10. Work Unit No. T-4243
		11. Contract or Grant No.
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546		13. Type of Report and Period Covered Technical Memorandum
		14. Sponsoring Agency Code 505-51-11-06-00-21

16. Abstract

Computers are playing an increasingly important role in the field of aerodynamics such that they now serve as a major complement to wind tunnels in aerospace research and development. Factors pacing advances in computational aerodynamics are identified, including the amount of computational power required to take the next major step in the discipline. Example results obtained from the successively refined forms of the governing equations are discussed, both in the context of levels of computer power required and the degree to which they either further the frontiers of research or apply to problems of practical importance. Finally, the Numerical Aerodynamic Simulation (NAS) Program — with its 1988 target of achieving a sustained computational rate of 1 billion floating point operations per second and operating with a memory of 240 million words — is discussed in terms of its goals and its projected effect on the future of computational aerodynamics.

17. Key Words (Suggested by Author(s)) Computational fluid dynamics Computational aerodynamics Computer systems Numerical Aerodynamic Simulation Program	18. Distribution Statement Unlimited Subject Category - 02
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price*
Unclassified	Unclassified	13	A02

End of Document