NASA TM-85887

NASA Technical Memorandum 85887

Computational Aerodynamics and Supercomputers

W.F. Ballhaus, Jr.

January 1984

National Aeronautics and
Space Administration

NASA Technical Memorandum 85887

Computational Aerodynamics and Supercomputers

W. F. Ballhaus, Jr., Ames Research Center, Moffett Field, California

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

N84-16139#

COMPUTATIONAL AERODYNAMICS AND SUPERCOMPUTERS

W. F. Ballhaus, Jr.
NASA Ames Research Center
Moffett Field, California 94035

Abstract

This paper reviews some of the progress in computational aerodynamics over the last decade and describes what remains to be done in the years ahead. The discussion includes a description of the Numerical Aerodynamic Simulation Program objectives, computational goals, and implementation plans.

1. Introduction

Computational aerodynamics has emerged in the last decade as an essential element in the design process for all types of aerospace vehicles. This revolution has been driven by advances in computer power which, in turn, have led to advances in numerical solution techniques. The result is better computational simulations achieved more cost effectively. Further advances in computer power are required for this revolution to continue.

NASA has initiated the Numerical Aerodynamic Simulation Program, an FY '84 new start, to provide this needed capability and to ensure continued U.S. leadership in computational aerodynamics. The program has been structured to focus on the development of a complete, balanced computational system that can be upgraded periodically to take advantage of advances in technology.

This paper reviews some of the progress in computational aerodynamics over the last decade and describes what remains to be done in the years ahead. The discussion includes a description of the Numerical Aerodynamic Simulation Program objectives, computational goals, and implementation plans.

2. Aerodynamic Simulation

The two basic tools of the aircraft designer, the wind tunnel and the computer, are depicted schematically in Figure 1. Both are simulations, one analog and the other digital. In spite of the continuing rapid advancement in the performance of supercomputers and numerical methods, it is not expected that computational simulations will completely replace wind-tunnel testing in the foreseeable future. Their roles, instead, are complementary. The wind tunnel is superior in providing detailed performance data once a final configuration is selected, especially for cases with complex geometry and complex aerodynamic phenomena. Computational simulations are especially useful for the following applications: (1) making detailed fluid physics studies, such as simulations designed to shed light on the basic structure of turbulent flows; (2) developing new design concepts, such as swept forward wings or jet flaps for lift augmentation; (3) sorting through many candidate configurations and eliminating all but the most promising before wind-tunnel testing; (4) assisting the aerodynamicist in instrumenting test models to improve resolution of the physical phenomena of interest; and (5) correcting wind-tunnel data for scaling and interference errors.

Aircraft configuration design is an interactive process of trial and error. Computational simulations substantially reduce the time and cost involved in detecting and correcting design errors in much the same fashion that word processing reduces the time and cost involved in correcting errors in office correspondence. Computational

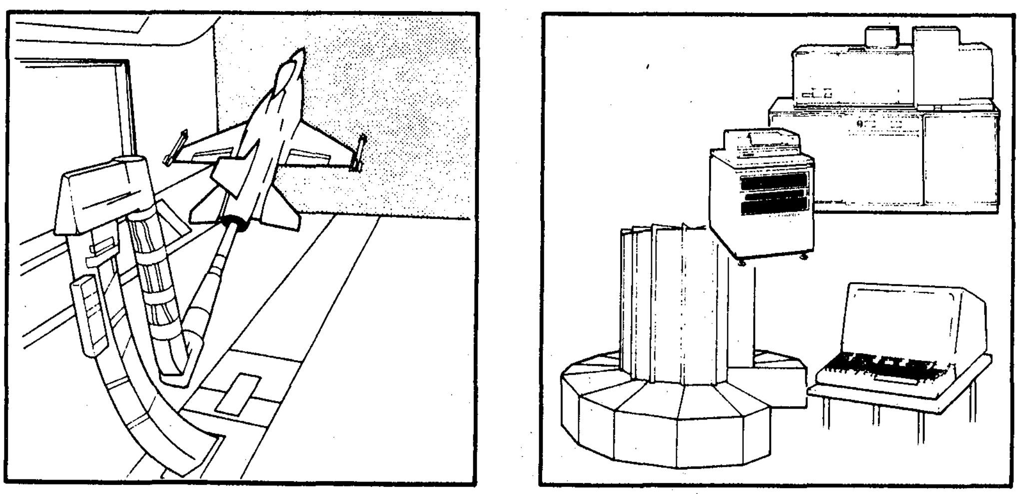

Wind Tunnel (Analog Simulation) Computer (Digital Simulation)

Wind Tunnels and Computers Complement Each Other

Figure 1. Two Basic Tools for Configuration Design

simulations can also provide data for conditions that are outside the operating range of existing experimental facilities. An example is a high-speed planetary probe entry condition, as in the case of the Galileo Probe scheduled to enter the atmosphere of Jupiter in the late 1980s.

Inadequacies in testing are associated with limitations in operating range (such as Mach number, Reynolds number, gas composition, and enthalpy level) and the control of boundary conditions (such as flow nonuniformity, wall- and support-interference effects, and model fidelity). These factors must be properly controlled for the simulation to accurately represent the desired free-flight conditions.

The inadequacies in computational simulations are primarily associated with poor resolution of physical phenomena, and this is a direct result of insufficient computer power. Because of limited available computer power, aerodynamicists are forced to solve approximate forms of the Navier-Stokes equations. These approximations introduce phenomenological errors, with the consequence that certain aspects of flow-field physics are not properly represented. A further discussion of these errors is contained in the next section. For any given mathematical formulation, the approximating procedures used to solve the governing equations and boundary conditions introduce numerical errors. These errors are compounded by inadequate grid refinement or incomplete treatment of complex aerodynamic configurations, both of which result from inadequate computer power. The consequences are, again, that physical phenomena are not represented properly.

The penalties associated with inadequate simulation capability are illustrated in Table 1. This is a list of major aircraft development projects which, in spite of the use of the best simulation tools available at the time, encountered major aerodynamic design problems that went undetected until flight testing. The consequences were severe penalties in cost, schedule, and/or performance. Several of these problems would have been detected early in the design cycle, and the consequent severe penalties avoided, with the computational capability available today. The first listed, design of a military transport, is an example. Other examples listed would require a computational capability beyond that currently available. Specifically, these cases would require the capability to treat complex, three-dimensional viscous flow fields. This is beyond current capabilities, but within reach in this decade. Successful design of next-generation aircraft will be even more difficult without further improvements in simulation capability.

Improvements in computational efficiency are illustrated in Figure 2, which is an updated version of a figure in Chapman.[1] The cost of a given computation has decreased at the rate of about two orders of magnitude over a 15-year period because of increases in computer power. There has been a similar two-order-of-magnitude decrease in the same time frame caused by improvements in solution methods. The net effect is a reduction in cost for a given computation of a factor of about 10,000. For example, the two-dimensional viscous flow field about an airfoil can be computed today in about 15 minutes of computation time at a cost of about $1,000. Twenty-five years ago, such a simulation would have cost $10 million on an IBM 704 computer and would have required 30 years of computer run time.

Until recently, the United States enjoyed a significant advantage in simulation capability over its military and commercial competitors. This was primarily due to an advantage in the availability of large-scale computers. Aerodynamicists in Europe and the USSR were forced to shoehorn large problems into small computers. Today, the United States still maintains its advantage over the USSR. However, the U.S. advantage over Europe has eroded. For example, of the 16 orders for Cray machines in 1982, 10 were from Europe.

3. Mathematical Stages of Approximation

The coupled set of partial differential equations governing aerodynamic flow fields, the Navier-Stokes equations, have been known for over a century. Because of their complex form, closed-form solutions have been obtained for only a few simple cases. Even a numerical solution using advanced

Table 1 Penalties Due to Inadequate Simulation

AIRCRAFT	PROBLEMS DISCOVERED IN FLIGHT TEST	CONSEQUENCE
C-141	INCORRECTLY PREDICTED WING FLOW	COMPROMISED PERFORMANCE, COSTLY MODIFICATIONS
C-5A	INCORRECTLY PREDICTED DRAG-RISE MACH NUMBER	REDUCED WING FATIGUE LIFE
F-111	INCORRECTLY PREDICTED TRANSONIC AIRFRAME DRAG	COSTLY MODIFICATIONS
B-58 B-70 YF-12	INCORRECTLY PREDICTED TRANSONIC PERFORMANCE	REDUCED AIRCRAFT EFFECTIVENESS
F-102 F-106	INCORRECTLY PREDICTED TRANSONIC DRAG	REDUCED PERFORMANCE
2 CIVIL TRANSPORTS	INCORRECTLY PREDICTED NACELLE-WING INTERFERENCE	REDESIGN REQUIRED

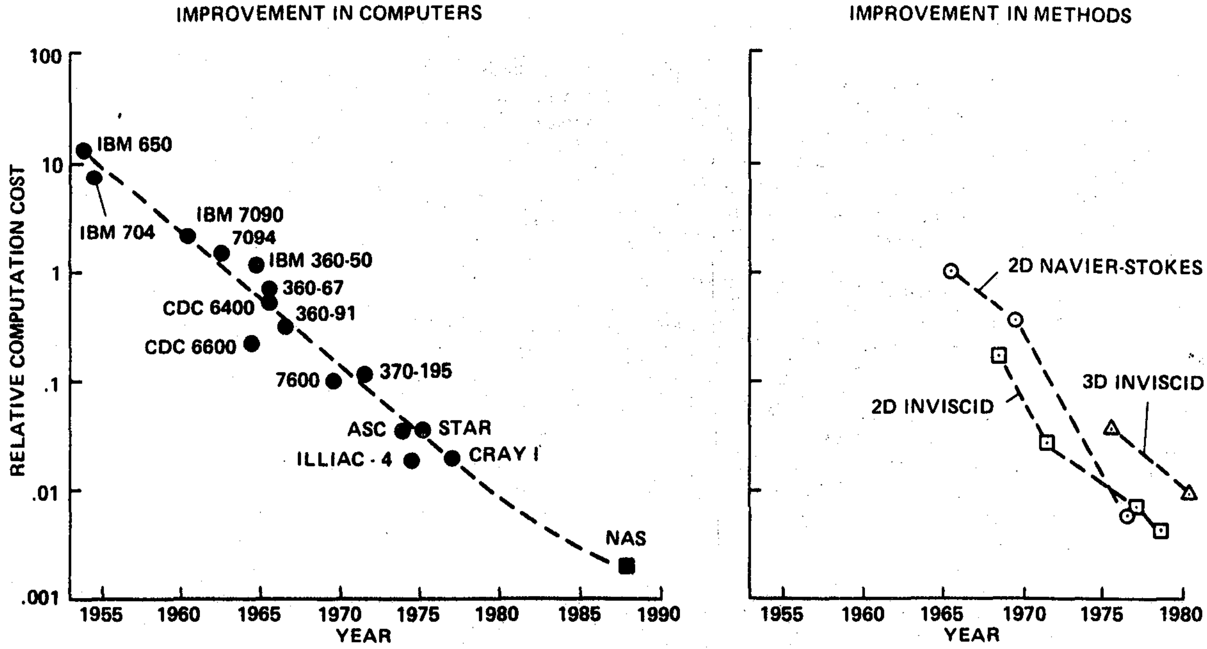

Figure 2. Computation Cost Trend for Computer Flow Simulations

computers is generally infeasible because current computer power is inadequate to resolve the wide range of length scales active in high-Reynolds-number turbulent flows of practical interest. Consequently, nearly all the simulations to date have involved solutions of mathematical formulations that are approximations to the Navier-Stokes equations.

Table 2 lists the four major stages of approximation in order of their evolution and complexity (updated from Chapman[1]). Advances in computational aerodynamics have produced a variety of computer codes ranging from Stage I codes, for complex aircraft configurations, to Stage IV codes for very simple geometry. Each succeeding stage requires an increase in computer power for a given geometry, but allows a new class of physical phenomena to be simulated; for example, subsonic lift distribution in Stage I, transonic wave drag in Stage II, airfoil drag and buffet in Stage III, and boundary layer transition and aerodynamic noise in Stage IV. Computer codes based on Stages I and II are used now throughout the aircraft industry, while Stages III and IV are the bases for a number of pioneering research efforts.

Numerical computation methods for Stage I are called "panel methods" because complex aircraft geometries are modeled by a large number of contiguous surface panels. Whereas the full Navier-Stokes equations representing conservation of mass, momentum, and energy contain altogether 60 partial-derivative terms when expressed in three Cartesian coordinates, the linearized inviscid (Stage I) approximation truncates these 60 to only three

Table 2 Governing Equations and Computer Requirements for Computational Aerodynamics

STAGE	APPROXIMATION	CAPABILITY	GRID POINTS REQUIRED	COMPUTER REQUIREMENT
I	LINEARIZED INVISCID	SUBSONIC/SUPERSONIC PRESSURE LOADS VORTEX DRAG	3×10^3 PANELS	1/10 CLASS VI
II	NONLINEAR INVISCID	ABOVE PLUS: TRANSONIC PRESSURE LOADS WAVE DRAG	10^5	CLASS VI
III	REYNOLDS AVERAGED NAVIER-STOKES	ABOVE PLUS: SEPARATION/REATTACHMENT STALL/BUFFET/FLUTTER TOTAL DRAG	10^7	30 x CLASS VI
IV	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

terms: the Laplace equation for subsonic flows and the wave equation for supersonic flows. The principal advantage of this stage is that the governing three-term equation is linear, and, hence, the forces on any panel are obtained by summing the influences from all the other panels. Thus, very complex geometries can be treated relatively simply.

A major complication associated with approximations other than Stage I is that the entire flow field, instead of just the aircraft surface, must be divided into a very large number of small cells or grid points for which finite-difference approximations to the partial-derivative terms can be expressed. This is a formidable task for treating general aerodynamic configurations because no straightforward, automated technique for grid generation has yet been devised. Considerable human effort on a computer graphics terminal is required.

The Stage I approximation, although crude, provides surprisingly useful simulations for purely subsonic or purely supersonic flows in which viscous effects are not dominant. A situation in which the Stage I approximation is not applicable is for flight near the speed of sound, i.e., transonic flows, wherein there are mixed subsonic-supersonic flow and embedded shock waves. Another example is a fighter aircraft in a gross maneuver, for which viscous effects would be important. The effects of viscosity, which include boundary layers on surfaces, turbulence, and a flow separation and reattachment, strongly affect aerodynamic performance, especially near the limits of performance.

Stage II, in its complete form, neglects only viscous terms and contains 27 of the 60 partial-derivative terms in the complete Navier-Stokes equations. There are subsets of the Stage II approximation that involve fewer than the 27 terms. This stage of approximation is applicable for subsonic, transonic, supersonic, and hypersonic flows as long as viscous effects do not dominate.

Stage III neglects none of the terms in the Navier-Stokes equations. However, the equations are time-averaged over an interval that is long compared with turbulent eddy fluctuations, yet short compared with macroscopic aerodynamic flow changes. Such an averaging process introduces various new terms, representing the time-averaged transport of turbulent momentum and energy, that must be modeled. No entirely suitable model for all flow types of engineering interest has yet been discovered. The current predominant thinking among researchers in the field is that no such universal turbulence model exists. Hence, attention is now focused on developing menus of turbulence models. The process consists of the synergistic use of computation and experiment to develop and test models of various types of flows that are considered building blocks to more complete aerodynamic configurations. The primary merit of the Stage III approximation is that it provides realistic simulations of separated flows, of unsteady flows such as buffeting, and of total drag and the other aerodynamic forces. Predictions of these types of flows are fundamentally important in a wide range of aerodynamic problems for aircraft, missiles, helicopter rotors, and turbine and compressor blades. Generally, considerably more computer speed and storage are required for Stage III because of the large number of grid points needed to resolve viscous phenomena, even with the time averaging of turbulent fluctuations.

Stage IV, in its full complexity, involves the direct numerical simulation of large-scale turbulent eddies from the complete time-dependent Navier-Stokes equations. The main physical concepts are that large eddies absorb energy from the mean flow, are highly anisotropic, vary from flow to flow, and transport the principal turbulent momentum and energy; on the other hand, small eddies dissipate energy, tend toward isotropy, are nearly universal in character, and transport relatively little turbulent energy or momentum. Thus, the large eddies are computed, and the small subgrid-scale eddies are modeled. Such simulations can be extremely demanding on computer memory and speed. Given sufficient computer power, however, numerical simulations from essentially first principles could be made of phenomena such as laminar/turbulent transition, aerodynamic noise, surface pressure fluctuations, and all relevant quantities characterizing turbulence. This stage, though it is in a relatively primitive research phase, has already provided some information about the nature of turbulent flows that has long been intractable to experimental measurement.

4. Example Inviscid Computations

Stage I and II approximations, in their purest forms, are frequently referred to as "inviscid formulations." These approximations are applicable to cases for which viscous effects can be either neglected or accounted for by adding boundary-layer corrections. An example application is the cruise design point for a commercial transport.

Examples of current inviscid formulation capability are provided in Figures 3 and 4. The first illustrates a panel method (Stage I) computation for the Space Shuttle Orbiter mounted on top of the B-747 carrier aircraft.[2] Accurate determinations were made of the lift characteristics for the combined configurations and for each vehicle during separation of the orbiter from the carrier. Configuration orientations selected from the computational design phase were tested in the wind tunnel for verification. Figure 4 illustrates the progress that has been made during the last decade in Stage II computations. In 1972, when the first three-dimensional nonlinear transonic computations were performed, about 18 hours of computer time per case were required on the IBM 360-67. Now, more detailed simulations can be performed in substantially less time because of advances in numerical solution techniques and the availability of more powerful computers. Since design is normally a trial and error (iterative) process, such reductions in required computer run time can substantially reduce the time required for configuration optimization.

Inviscid computations have become an integral part of the design process in the aircraft industry today, and there are numerous examples of successful applications. One is the design of the European Airbus A-310. According to Jupp,[3] the use of these methods was primarily responsible for the claimed 20% improvement in fuel efficiency of the A-310 over the A-300, which was designed before these methods were widely available. Less wind-tunnel testing of wing designs was required, leaving more test time to address complex aerodynamic component interference problems. A simple back-of-the-envelope calculation, assuming fuel costs of $1.30 per gallon for an assumed 15-year life of a 400-airplane fleet, indicates a potential saving in fuel costs of

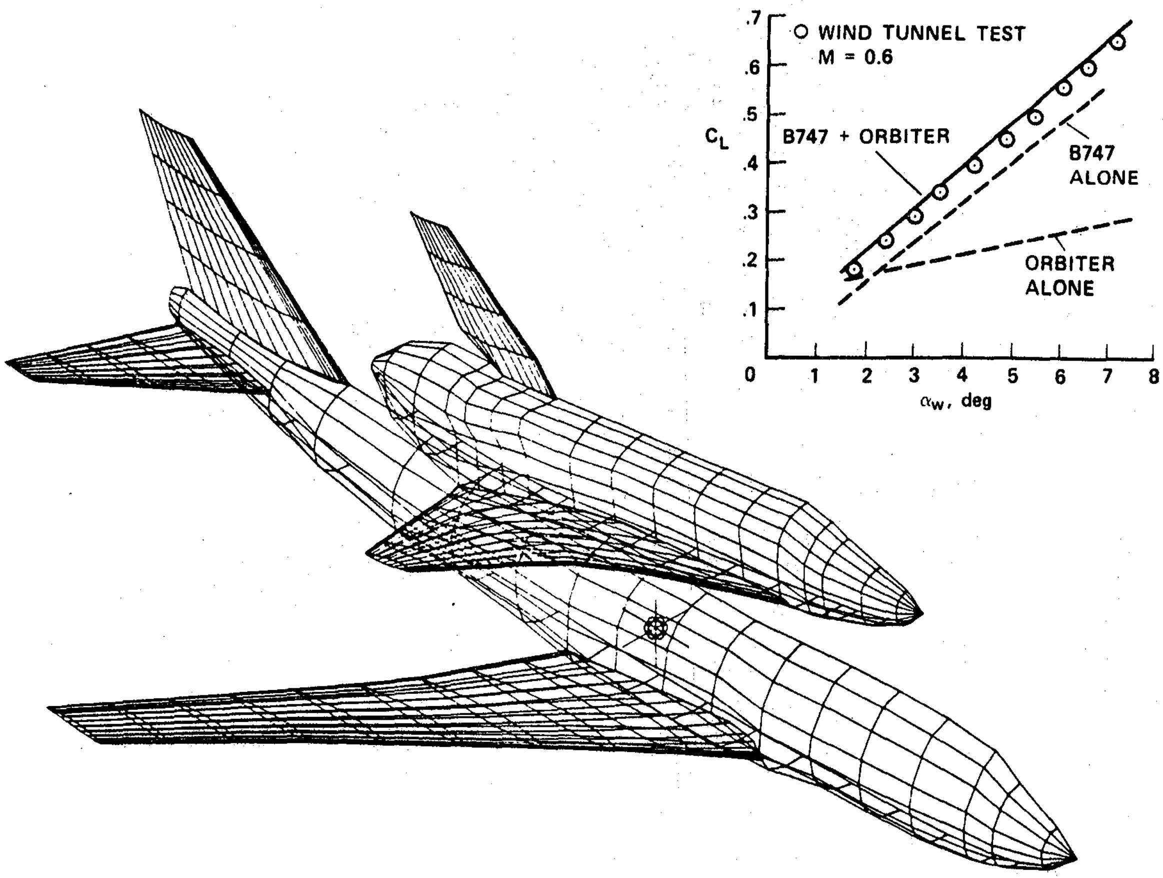

Figure 3. Linearized Inviscid Panel Method Applied to Space Shuttle Orbiter/B747 Combination

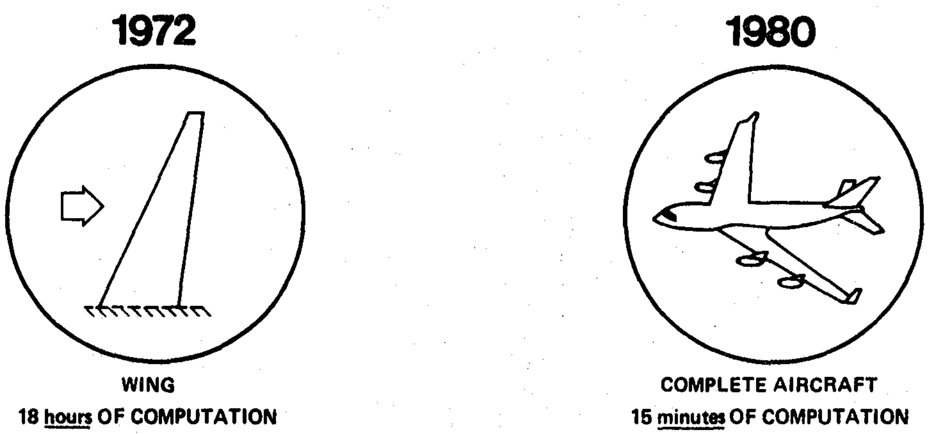

Figure 4. Revolutionary Advances in Inviscid-Flow Technology During the 1970s

$10 billion. Incidentally, Airbus has captured a substantial percentage of the wide-body jet market from its American competitors in a very short period of operation.)

5. Example Viscous Computations

Examples of viscous-dominated aerodynamics problems, for which inviscid formulations generally are not applicable, are provided in Table 3. These problem areas can be addressed using the Stage III approximations, the Reynolds-averaged Navier-Stokes formulation. Currently, this stage is the basis for vigorous research focusing on the computation of

Table 3 Examples of Important Viscous-Dominated Problems Amenable to Solution on NAS

- PREDICTION OF AERODYNAMIC FORCES
- INLET FLOWS
- COMPRESSOR STALL
- AIRFRAME/PROPULSION SYSTEM INTEGRATION
- STRAKE DESIGN
- STALL/BUFFET
- PERFORMANCE NEAR PERFORMANCE BOUNDARIES

complex flow-field phenomena. Thus far, investigations have been limited to simple geometries. Figure 5 illustrates some of the progress that was made during the 1970s. These building block computations, verified by comparison with experiment, indicate that the basic physics of these flow fields can be computed, at least qualitatively. Good quantitative agreement will come in time as turbulence models (to reduce phenomenological errors) and solution-adaptive grid generation techniques (to reduce numerical errors) mature.

Currently, for three-dimensional cases such as the after-body drag computation for a fuselage with propulsive jet, shown in Figure 5, as much as 20 hours of Class VI computer time are required. Clearly, for computations about complete aircraft, additional computer power is required. The specific requirements are quantified in Figure 6, which is modified from Chapman[1] to include recent performance predictions. These requirements are also compared with the capabilities of existing machines, and with the Cray 2, expected to be available in 1985, and the Cyber 2XX and Cray 3, expected to be available in the late 1980s. The requirements are stated in terms of computing speed and memory for a 15-minute computation based on solution algorithm efficiency anticipated for 1985. Experience indicates that 10-15 minutes of run time are typically required for adequate turnaround in an applications-oriented environment. Algorithm efficiency for 1985 is set by extrapolation, assuming the same rate of progress in algorithm development experienced during the 1970s. The computer speed requirements are stated in terms of millions of floating point operations per second sustained, i.e., the average speed expected for a typical flow-field computation. Note that two-dimensional airfoil computation requirements are well within the capabilities of the Cray 1S and Cyber 205 Class VI machines. Requirements for three-dimensional flow past a wing are beyond the capabilities of these machines, at the limits of the capability of the Cray 2. Machine capability expected in 1988 is sufficient to tackle the problem of a complete aircraft or a helicopter rotor with a 1-million-point computational grid system. This requires a computing speed of about 1 billion floating point operations per second and a memory of 250 million words.

There are few examples of the Stage IV approximation, Large-Eddy Simulation (LES), computations because of the enormous computing power required even for the simplest geometries. Figure 7 shows vorticity contours from an LES computation of a three-dimensional compressible jet.[4] The contours illustrate transition from laminar to turbulent flow.

6. The Numerical Aerodynamic Simulation (NAS) Program

The NAS program is an FY '84 NASA aeronautics new start that will be an ongoing evolutionary and pathfinding effort to provide a national computational capability available to NASA, DOD, industry, other government agencies, and universities. This program is a necessary element in insuring continuing leadership in computational fluid dynamics and related disciplines. For additional information on the NAS Program, including a summary of the 8-year evolution of the concept, see Peterson et al.[5]

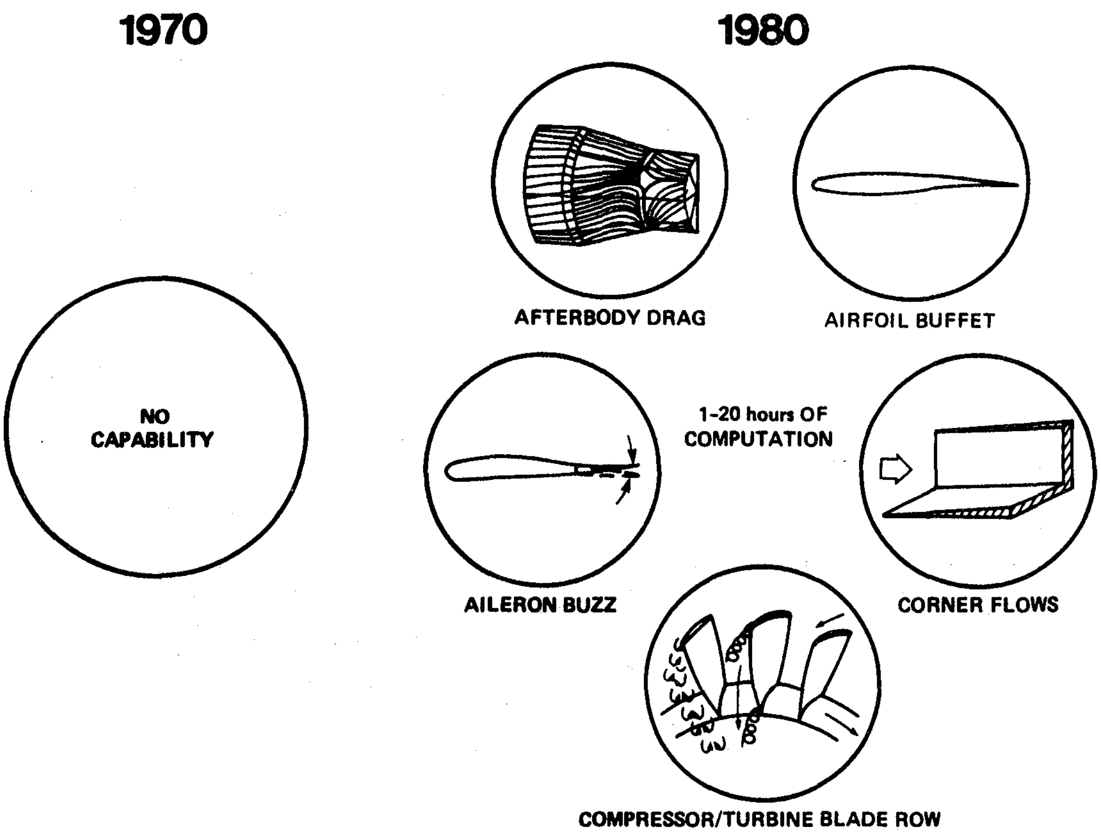

Figure 5. Revolutionary Advances in Viscous-Flow Building Block Technology During the 1970s

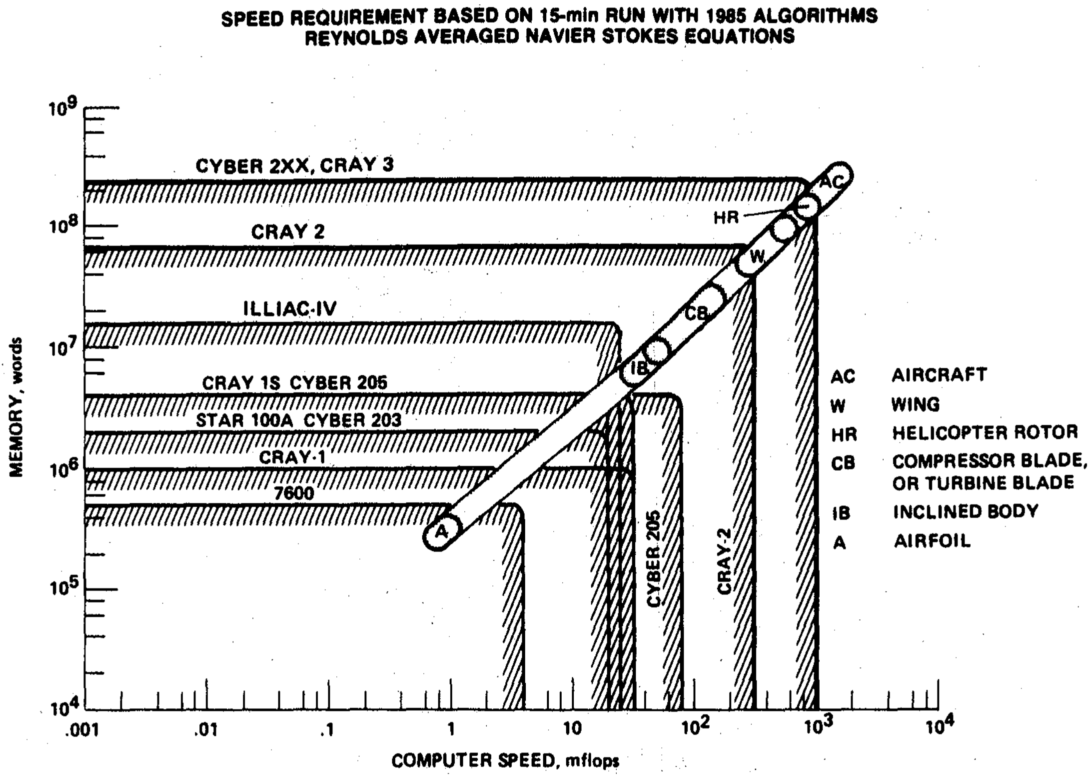

Figure 6. Computer Speed and Memory Requirements Compared With Computer Capabilities

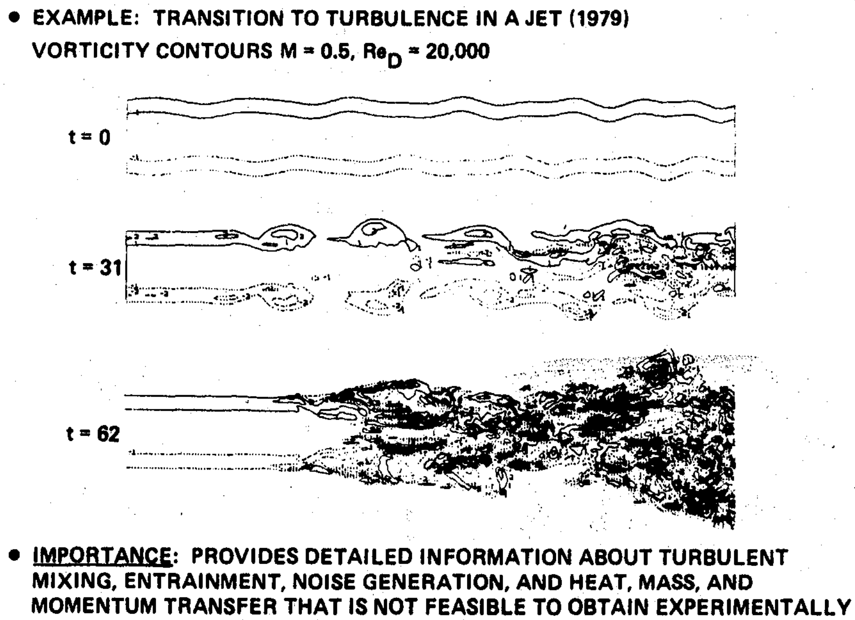

Figure 7. Provide Increased Understanding of Fluid Flow Physics

Specific performance goals of the NAS System are:

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

The capability envisioned for the NAS System in 1987 will permit three-dimensional viscous flows, based on the Stage III approximation, about multiple-component (wings, fuselage, tail, nacelles, etc.) aircraft configurations. Furthermore, this capability will permit Stage IV computations in a research mode that will lead to a better understanding of turbulence and, hence, to the better turbulence models essential to full realization of the Stage III approximation.

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 workstations, 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 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. Figure 8 illustrates the continuing development of the NPSN functional and performance capabilities through successive introduction of advanced high-speed processors into the network. The introduction of each new high-speed processor involves the 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.

As shown in Figure 9, 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 Communication Subsystem (LHCS)
7. High-Speed Data-Network Subsystem (HSDN)
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.

The High-Speed Processor Subsystem (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 applications at a sustained rate of 250 mflops within a minimum 64×10^6 word memory capacity. The second (HSP-2), planned for integration in 1987, will increase these values to 1 gflops and 256×10^6 word memory.

Whereas the HSP is the ultra-high-speed, large-scale computer resource serving the global user community, the Workstation Subsystem (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.

The Support Processing Subsystem (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.

The Graphics Subsystem (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

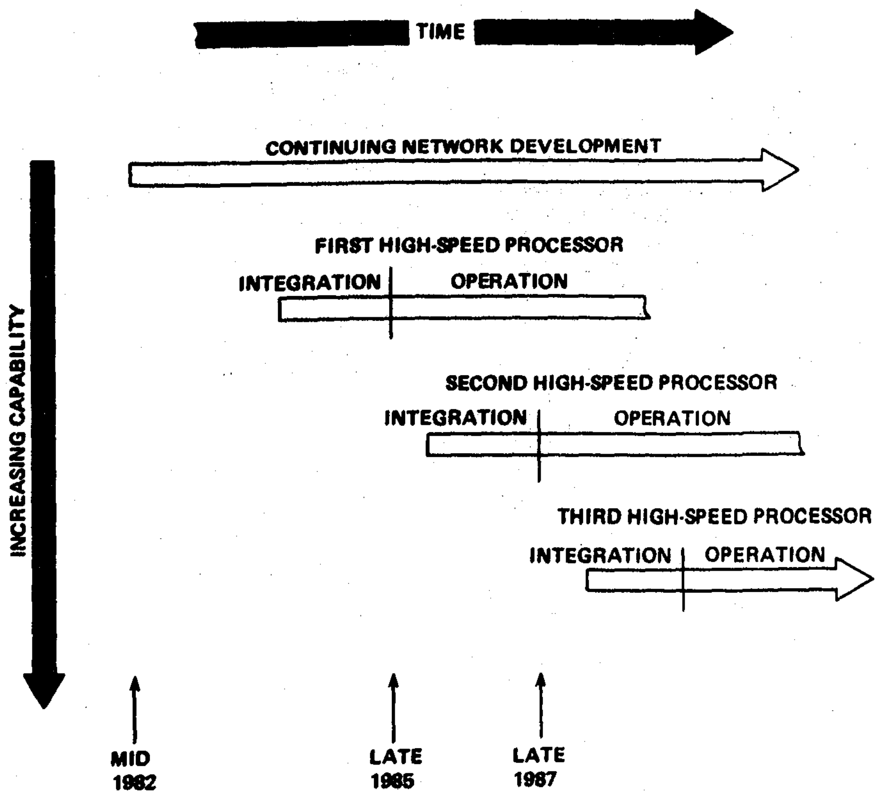

Figure 8. Evolution of NAS System

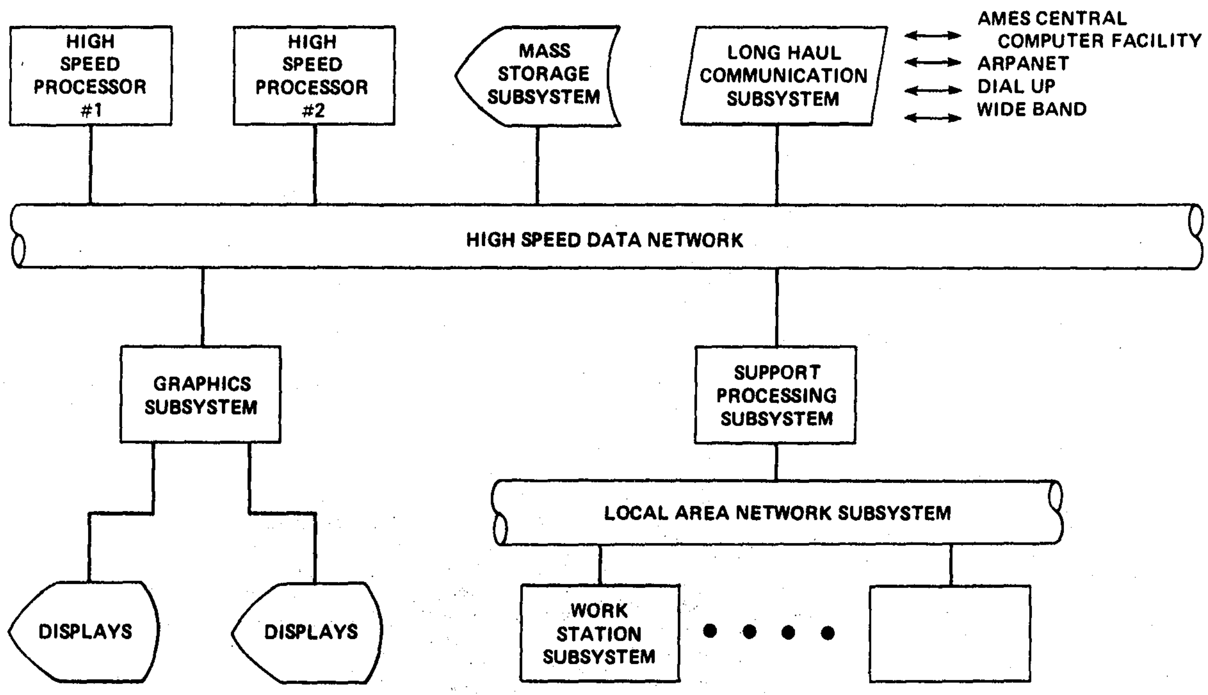

Figure 9. Numerical Aerodynamic Simulation Processing System Network

beyond that provided by workstations, and will be shared by first-level user organizations.

The Mass Storage Subsystem (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. Current requirements call for on-line storage of 200 gigabytes in 1985 and 800 gigabytes in 1987.

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.

The Long-Haul Communication Subsystem (LHCS) will provide the data-communication interface between the NPSN and 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).

The High-Speed Data-Network Subsystem (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.

The Local-Area Network Subsystem (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.

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.

The NPSN software 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. A common and consistent user environment 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™ or UNIX™-look-alike operating system on the user 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™, vendor UNIX™ look-alikes, and virtual operating system approaches using Software Tools. The UNIX™ 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. Vendor independence and high performance are aided by the implementation of a highly functional and efficient network protocol, 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.

7. Supporting Computer Science Research at NASA Ames

In mid-1983 the Research Institute for Advanced Computer Science (RIACS) was formed at NASA Ames and is operated under contract by the University Space Research Association. The intent was to bring additional computer science expertise to NASA and to initiate a number of research ventures involving personnel from NASA, RIACS, industry, and academia. The long range research theme is to automate the process of scientific investigation from problem formulation to results dissemination. Principal areas of focus include: concurrent processing, computer system networking, knowledge-based expert systems, and fault tolerant computing. It is envisioned that appropriate research advances from these ventures will lead to future upgrades in the NAS System in concert with the NAS Program pathfinding effort to provide a leading-edge computational system to a national community of aerodynamic and fluid dynamic users.

8. Concluding Remarks

Remarkable progress was made during the decade of the 1970s in advancing the discipline of computational aerodynamics. This progress was sparked by increases in available computer power and was focused primarily on the development of inviscid flow-field simulation techniques. These techniques have been incorporated into computer codes that are now vital tools in the aerospace industry for design and analysis.

Today, research is focused primarily on extending simulation capability to the treatment of more complex aerodynamic configurations and more complete treatment of complex aerodynamic phenomena. A major goal is the simulation of the three-dimensional viscous flow field about a complete aircraft. Such a capability will substantially reduce the level of uncertainty in current simulations. This will permit early detection of design deficiencies, thereby avoiding severe penalties in cost, schedule, and performance in many aerospace system developments.

To achieve this objective, and for the rapid rate of progress in computational aerodynamics to continue, further substantial increases in available computer power will be required. The Numerical Aerodynamic Simulation (NAS) Program, a NASA FY '84 new start in aeronautics, is designed to meet this challenge. The goal of this effort is to provide an advanced capability by the mid- to late-1980s that will be available for use by government, industry, and academia. The NAS System will be continually upgraded beyond that date as computer technology advances.

Acknowledgments

The author acknowledges the substantial contributions to the development of the NAS concept, as described here, by his friends and colleagues, F. R. Bailey and V. L. Peterson of NASA Ames, and D. R. Chapman of Stanford University.

References

1. D. R. Chapman, "Computational Aerodynamics Development and Outlook," AIAA Journal, Vol. 17, No. 12, pp. 1293-1313, Dec. 1979.

2. A. L. da Costa, "Application of Computational Aerodynamic Methods to the Design and Analysis of Transport Aircraft," ICAS Paper, 11th Congress of the International Council of the Aeronautical Sciences, Lisbon, Portugal, Sept. 10-16, 1978.

3. J. A. Jupp, "Wings of Success," World Airnews, pp. 4-8, Dec. 1980.

4. A. A. Wray, "Study of Transitional Phenomena in a Flat Plate Boundary Layer and Effects of Small Scale Turbulence on Large Scale Motions," DCW Industries Report DCW-R-22-01, Dec. 1978.

5. V. L. Peterson, W. F. Ballhaus, Jr., and F. R. Bailey, "Numerical Aerodynamic Simulation," NASA TM-84386, July 1983. To appear in the Proceedings of the Conference on Large Scale Scientific Computation, Academic Press (to be published).

1. Report No. NASA TM-85887	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle COMPUTATIONAL AERODYNAMICS AND SUPERCOMPUTERS		5. Report Date January 1984
		6. Performing Organization Code
7. Author(s) W. F. Ballhaus, Jr.		8. Performing Organization Report No. A-9583
9. Performing Organization Name and Address Ames Research Center Moffett Field, CA 94035		10. Work Unit No. K-1757
		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 505-31-01

15. Supplementary Notes

Point of Contact: W. F. Ballhaus, Ames Research Center, MS 200-4, Moffett Field, CA 94035 FTS 448-5065 or (415) 965-5065

16. Abstract

This paper reviews some of the progress in computational aerodynamics over the last decade and describes what remains to be done in the years ahead. The discussion includes a description of the Numerical Aerodynamic Simulation Program objectives, computational goals, and implementation plans.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
Computational fluid dynamics Numerical aerodynamic simulation NAS Program	Unlimited Subject Category - 02

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price*
Unclassified	Unclassified	14	A02

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