CR-152058

PRELIMINARY STUDY

FOR A

NUMERICAL AERODYNAMIC SIMULATION FACILITY

SUMMARY REPORT

BY: N. R. Lincoln

OCTOBER, 1977

Prepared under Contract No. NAS2-9457 by:

CONTROL DATA CORPORATION
2800 Old Shakopee Road
Minneapolis, Minnesota 55440

for

AMES RESEARCH CENTER
NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

(NASA-CR-152058) PRELIMINARY STUDY FOR A NUMERICAL AERODYNAMIC SIMULATION FACILITY (Control Data Corp.) 17 p

N77-85509

Unclas

00/09 52523

NUMERICAL AERODYNAMIC SIMULATION FACILITY
SUMMARY REPORT

For the past 6 months the Research and Advanced Design Laboratory of Control Data Corporation has been conducting a joint study in cooperation with Ames Research Center of the National Aeronautics and Space Administration (NASA). The objective of this study was to determine the methodology and feasibility of construction of a Numerical Aerodynamic Simulation Facility (NASF). This facility would be utilized by NASA as an integral component of a complete service to the aerodynamic design and evaluation community represented by industry and government engineering organizations alike. These services would include the open availability of the NASF, physical wind tunnels of all sizes, and the vast expertise possessed by NASA engineers, physicists, and mathematicians.

The study began with several assumptions. First, no existing computational ensemble could provide the necessary solutions to three-dimensional Navier-Stokes equation systems representing aerodynamic shapes in all speeds of airflow. The second assumption was that such a facility would find its most critical needs arising about 1982. This date was itself a compromise between the desire for a high performance computational capability to meet immediate needs and the known state of the computer art in 1977 which is not capable of meeting even the most modest objectives set for the NASF. The third assumption was that no more than two computational approaches would be viable for the NASF, and that work at Ames in development of the program was sufficiently mature to permit actual codes to be used in the study.

The Control Data approach to the study was then to make a quick, early assessment of the probability of achieving computational performances in excess of 100 times the CDC 7600 speeds being realized by the existing Ames installation. At the outset it was felt that with technologies already in hand and architectural principles already demonstrated, achieving the performance goals by 1982 was a certainty. At the direction of Ames personnel, however, Control Data proceeded to examine the state-of-the-art of relevant technologies, the state-of-the-art of systems and processor architectures, and the measurable computational requirements of the two Navier-Stokes solution programs then in existence. The purpose of this phase of the study was to provide NASA with sufficient information so that its staff members could make an independent evaluation of the best approach for construction of the facility. At the same time Control Data would attempt to develop a system design to meet the objectives.

The general technical approach to the system design was to use, wherever possible in the design, standard parts and components to reduce development costs and risks for those components. This resulted in the identification of two main components in the NASF, the front-end or support processing system, composed of commercially available equipment and software, and the back-end or Navier-Stokes Solver (NSS), which must utilize special design, special technology, and special software to meet the speed requirements of the facility. Initially, it was felt that a derivative of the STAR-100 architecture and design could be used for the NSS. This would further reduce the development costs and project risks, as well as manufacturing costs due to volume ordering of common components. Since a member of the STAR family, the 100C, appeared to possess a basic computational speed on which to build a specialized processor, the concept of commonality appeared quite appealing.

About two thirds of the way through this study effort, however, it was found that some radical departures from STAR architecture and design had to be taken to meet the goals of the NASF. It did appear, however, that certain of the technological achievements in LSI technology and system organization of subcomponents could be borrowed from the STAR-100C project to reduce design time and risk of completion of the NASF.

SIGNIFICANT RESULTS OF STUDY

Given this initial orientation, the study yielded significant results that are summarized below.

TECHNOLOGY

- The basic memory unit for an NSS is still best constructed of bipolar memory parts of the emitter coupled logic (ECL) family or a family with similar speeds. For a system of this generation, memory access speeds in the 30-to 40-nanosecond range for up to 8 million words of data are attainable.

  A lower range of memory speeds is available with current technology, with attendant cost and power reductions over the high performance ECL memory. To meet the needs of the three-dimensional Navier-Stokes codes, there are a known number of occasions when memory must be accessed in an unstructured, random manner. To reduce the delays accompanying sequences of random accesses, the memory can be built into a multitude of banks such that the probability of two successive references can be almost eliminated. There are times, however, when all computation must pause while a required operand is retrieved from the memory system. In such cases, the access time delay for a single operand becomes important. Thus, to ensure that no facet of the Navier-Stokes solution becomes a bottleneck, the memory must exhibit the combination of properties of high bandwidth, fast access, and multiple banking. It is felt that the fastest, reliable technology available today is the correct choice for memory technology.

- The basic logic element for a processor of this type will be based on high-speed, large scale integration (LSI) devices with switching speeds in the 500-picosecond range. Exotic elements such as Gallium Arsenide and Josephson devices have not progressed sufficiently in initial research to be used in a manufacturing environment in 1980 to 1982.

  Studies of various technology families and architectural alternatives have revealed that it is more cost-effective and more reliable to build a superprocessor from a minimum number of parallel units implemented with the fastest technology available than to attempt to meet the same level of performance with a large number of parallel, but individually slower speed processors. The NSS should therefore be constructed of the best technology available in the 1977 to 1982 time frame. Of course, the performance, manufacturing, and cost advantages of LSI dictate the use of the highest integration possible. For ECL speeds, the number of gates possible today per LSI component is between 150 and 200. Expectations for an LSI component with 400 to 500 gates to be available for construction of the NSS are reasonable, though not without some risk.

Lower speed components of the MECL variety will necessarily be employed where circuit speeds are not as important as power dissipation, cooling, and cost. For example, the I/O system, trunks, and some peripheral subsystems will be constructed from existing technologies, both MOS and lower-powered ECL. If at all possible, all components should be built with industry standard parts, even the LSI portions. Membership in a larger family ensures some long-term longevity for spare parts and support from a variety of semiconductor vendors.

- Slower speed memories will be fabricated with charge coupled devices (CCD) for NSS applications, since the state of development of electron-beam memories (EBAM) and magnetic bubble memories cannot yield components of the desired bandwidth or reliability.

  Million-word (64 bits) systems of CCD memories are being built to practical specifications today with 65K circuits. There is a realistic chance that operational CCD parts containing 265 kbits will be available for prototype system implementations in 1978. If the analysis of the NSS memory requirements is sustained by later studies, a 256-million-word system will be needed by 1982. Within the limitations of packaging, cooling, and reliability, it therefore appears quite practical to anticipate a 256-million-word system to be available for an operational NASF in 1981 to 1982. The programmatic study of the specimen flow model codes shows that a brute-force swapping technique can be employed between the main memory and the auxiliary storage medium. If this technique greatly simplifies hardware and software control, it must also possess data bandwidths of at least 1.6 billion bits per second (each way) to achieve the sustained processing rates desired for the NSS.

  Although million-bit bubble memories are now available for prototype experimentation, the bandwidths of such chips are limited to the 400 to 500 kilohertz range. In addition, the access time for data blocks is quite a bit higher than for the corresponding CCD technology. The ability of bubble memories to retain data in the event of power failures is desirable, but if the total run time for which data must be retained is less than 20 minutes, the loss of bandwidth and access time is not worth the cost. For example, existing million-bit chips would have to be arranged in parallel, with 4000 chips simultaneously transferring data, to achieve the 1.6-gigahertz data rate. Bubble memories of smaller size will most likely be found in some of the peripheral subsystems as replacements for small disks and fast-access drums that now hold directories and store-and-forward message buffers.

- Rotating magnetic media will remain the primary form of mass storage and archival storage for a system built in the early 1980's.

  Extensions of existing knowledge and technologies involved in rotating magnetic memory are readily projected for the next 4 years. There remain only the solutions to several nagging engineering questions before another improvement in density and transfer rates can be seen. The most probable direction will be in the form of sealed (almost hermetic) units containing disks, positioners, and head groups. These units will employ plated disk (rather than oxide-coated disks) to reduce film thickness and thus improve resolution. Factors of 4 to 16 times the existing storage densities will be achieved in the NSS timeframe. As an example, an 819-size unit (one single disk unit) will be able to house from 40 to 50 billion bits.

Laser and photostorage devices are not yet in the same ballpark with rotating mass storage for reliability and system availability. In the case of the NASF, the predicted on-line storage requirements can be met with the next forseeable generation of disk storage devices.

Archival storage is an area still undergoing great upheaval and experimentation. Although the IBM and CDC mass storage subsystems represent today an imperfect engineering approach to archiving, offshoots of them will probably still engage magnetic tape technology and random selection systems being pioneered by them. For this reason, site requirements were based on existing units such as the 38500 mass archival storage.

PROBLEM ANALYSIS

A large portion of time was spent in the analysis of two-dimensional specimen codes provided by NASA/Ames personnel. These were the explicit code being evolved by Bob MacCormack, and the implicit code under development by Steger, Pulliam and Lomax. Both codes were first run in their original FORTRAN form on the STAR-100 where the STAR instrumentation could be used to sample the key elements of the code operation. Both codes were then vectorized for the STAR-100 as a first step in the process of developing parallel algorithms to match the NSS, and as guidance for the creation of a unique NSS processor.

Finally, as the NSS structure took shape, the implicit code was restructured to match the new architecture and a set of rough estimates made as to the behavior of that code on the proposed NSS.

A summary of some of the results of this phase follows:

1. The explicit code required 7 minutes of 7600 time to compute a particular solution for the Garabedian-Korn airfoil to 256 time steps. The original scalar version of this code with no vectorization or optimization required 16 minutes of STAR-100 time reflecting the state of the compiler development, as well as the 80-nanosecond scalar issue rate of the STAR-100. A partially vectorized version of this code (one of the split operators) was run at 4.5 minutes. A fully vectorized version was not completed due to the diversion of attention to the implicit code. The explicit code was operating at an average rate of two megaflops for the total run on the 7600.

2. The implicit code, processing basically the same problem as the explicit code, was timed at about 12 minutes on the 7600 and 35 minutes in scalar FORTRAN on the STAR-100, while a first attempt at vectorization for the STAR-100 yielded a five-minute running time. The implicit code does not rely on special casing of computational regions and thus performs many more floating-point computations than does the explicit form. The implicit code operated at an average of around two megaflops on the 7600 also. The code developers are convinced that the three-dimensional form of this implicit program can be refined to reduce the computational requirements. This programming ploy is essential to the NASF meeting its system goals.

3. The implicit code was then singled out for restructuring for a hypothetical NSS. A method of processing slices of the data, similar to the scheme used by Lomax on the ILLIAC IV, was devised to permit a reduction in the size of the costly, high-performance main memory. A system of small, high-performance buffers, backed up by 8 million words of main memory, and that backed up by 256 million words of block transfer memory, can be effectively utilized by the slice mechanism. Depending on slice lengths the restructured implicit code was estimated to perform on the NSS between 660 and 940 megaflops in 64-bit mode and from 950 to 1910 megaflops in 32-bit mode.

4. A three-dimensional form of the implicit code can be sliced more efficiently and, by using 32-bit computation mode for a majority of calculations where accuracy permits, it is estimated that the NSS should run at an average rate in excess of 3000 megaflops, assuming a main computer clock of 10 nanoseconds.

SYSTEM

Figure S-1 gives a block diagram overview of the NASF system envisioned.

It can be seen from this figure that the NSS processor represents only a small portion of the equipment volume, as well as only about one-third the total system cost. The mass storage equipment and graphics subsystem needed to support the NASF are shown in rough outline form only, but represent the projected needs of an installation that will be operational in the 1982-1983 timeframe.

Some salient features of the displayed system are:

- A dual processor front-end configuration composed of computing equipment available in 1977 would provide sufficient power and reliability to meet the demands of a front-end system for the NASF. Computing equipment currently under development for standard sales in the 1980's promises even higher performance and reliability along with reduced cost, thus ensuring that the computational facility will have substantial power in the supporting subsystems.

  Experience with the STAR-100 system has shown that the development of even a minimal operating system to meet today's normal needs for system access and features is a monumental undertaking. From a manufacturer's point of view, when P and L statements become pursuasive inhibitions to grandiose plans, some means of reducing cost and schedules for putting a new computer architecture into production are absolutely essential. The computational facility concept was thus defined, wherein the STAR processor performed primarily calculations, and CDC CYBER processors performed all the data management functions, file and user security, access functions, and communications management functions necessary for a production system. This substantially reduced the resource and time requirements for STAR software. Further, it meant that a more stable operating system was available earlier in the production cycle.

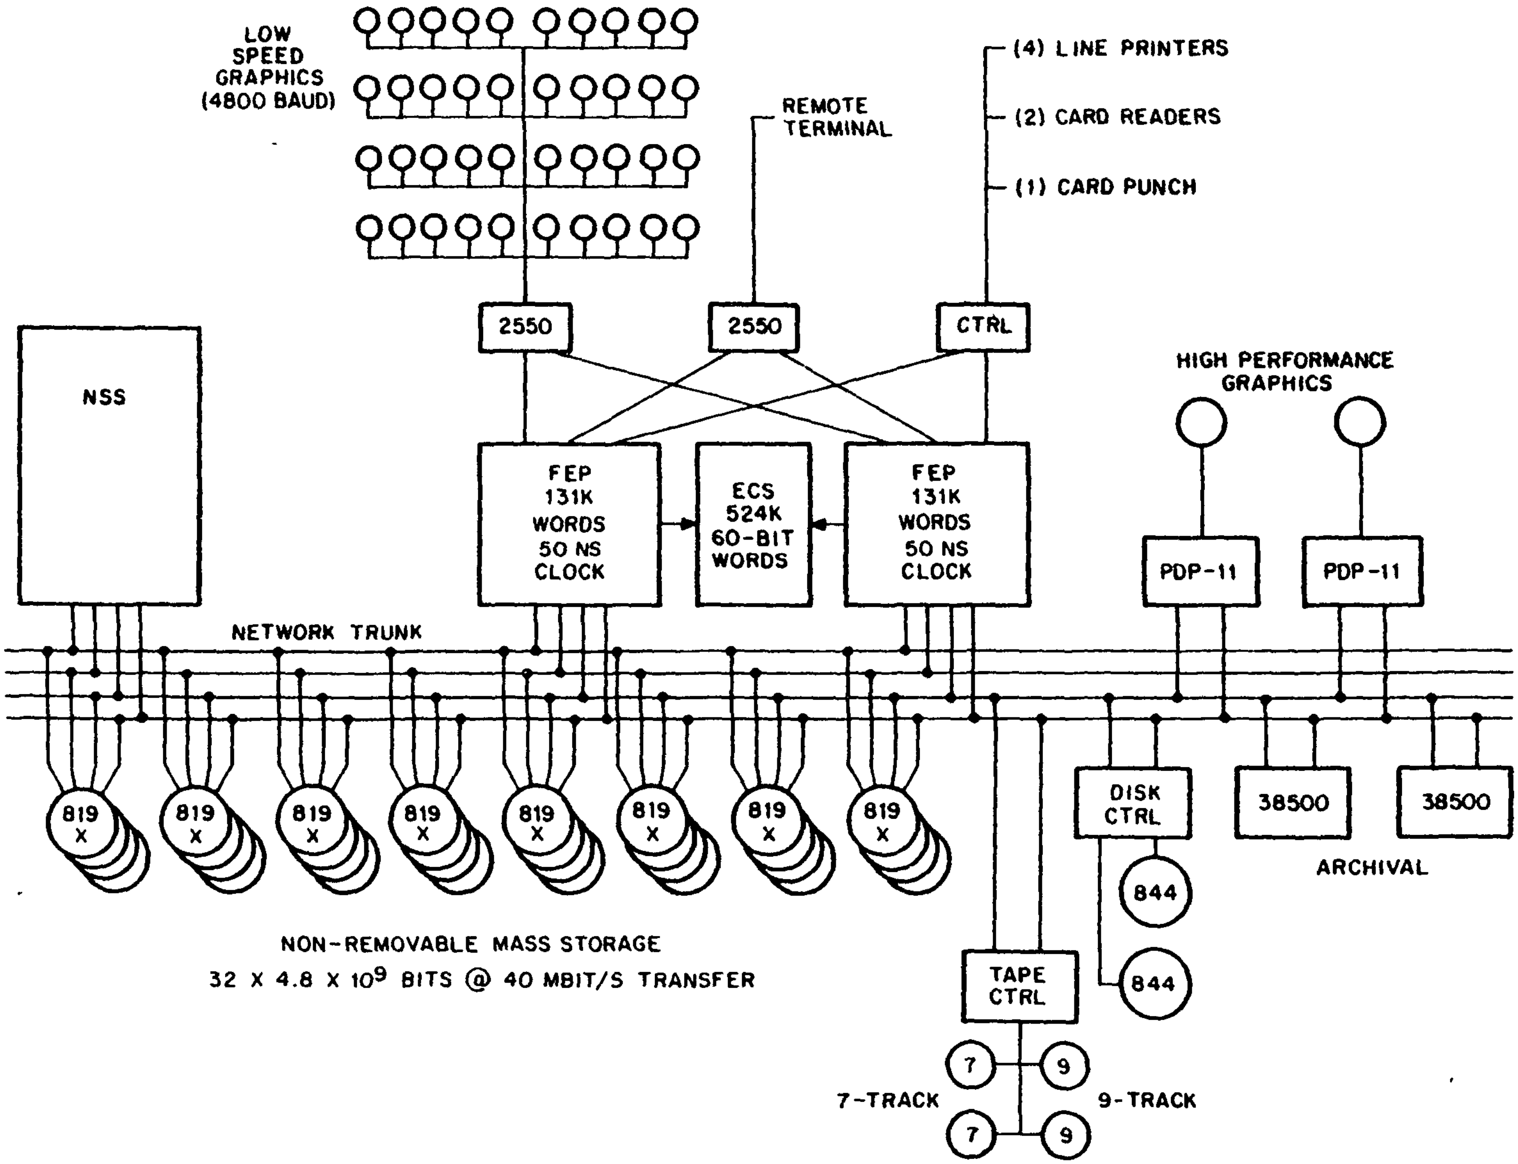

Figure S-1. NASF System Interconnection

By choosing a mature computer system for the front-end function, fully supported with the entire range of software available, NASA can be assured that the continuation of high levels of effort on performance, feature and stability aspects will yield a better system in 1982 than one designed specifically for Ames.

- A network trunk scheme of system interconnect would provide a more flexible means of harnessing all the equipment needed in the NASF. The distances which can be achieved, the number of connections to one trunk, and the sustainable bandwidths make this system quite appealing to meet the system requirements of the NASF.

  Network trunks with 50 million bits/second transmission capability and cable lengths of approximately 600 meters (2000 feet) are now operational. In addition to allowing peripheral devices and peripheral subsystems to be more remote from the attached computer, the trunk scheme is specifically designed to mate with alien equipment. This becomes a plus for users, such as NASA, permitting them to make the best choice of equipments to be attached (with the appropriate, moderate-cost adapter) to the trunk without concern for matching electronic channel and software protocol requirements.

  Such a network system allows the user to determine whether data can be transferred from one disk storage system to any attached processor without having to pass through a front-end machine. This can reduce bottlenecks due to demands for processor attention, as well as ensuring that the fastest I/O channels can be matched with available trunk bandwidth.

- Graphics hardware and software which are generally available and not customized for a particular site still leave much to be desired when matched against NASF requirements. Most notable, terminal costs and reliability, as well as response times, for complex 3-D displays need substantial improvement.

  However, graphics systems are receiving considerable industry attention and are being increasingly recognized as effective design tools. Also, developers of graphics systems seem to be placing growing emphasis on reducing, or eliminating, application dependence and equipment dependence. While these factors are favorable for expectations of adequate graphics capability, technology advances (such as the advent of the microprocessor) are providing cost improvements and increased reliability.

- As recommended by Ames study team personnel at the outset, compiling and scheduling of the NSS back-end is best performed on the front-end computing system. This makes possible early development and checkout of those very complex software elements on existing processors, well in advance of the availability of the NSS. Although experience has shown that a compiler operating on the target machine is better able to optimize code for the target machine, the time scale for this project dictates an early start on the compiler that could best be supported by existing equipment.

It would appear at this time that the best approach for the language processor is to identify the front-end processor as soon as possible. Then an examination of the existing compiling system on the front-end processor could determine the feasibility of using the basic front-end compiler with vector extension modifications to compile for the NSS. As much as possible, new compiler design, programming, and documentation must be reduced to accommodate the schedules.

NSS PROCESSOR

Figure S-2 gives a broad overview of the proposed NSS processor. Each of the major blocks represents a separately designed, and somewhat modular, functional entity. The vector units, map unit, scalar unit and swap unit can operate concurrently with each other, and in many cases, independently of each other. The major architectural feature shown here, in addition to the massive memory and memory bandwidth, is the utilization of 'functional' parallelism. The process of extracting data from memory for processing, and putting it back again, is called mapping. Thus, the map unit can perform memory access operations for restructuring data, while the vector units are performing computations on a separate piece of data that is held in buffer registers within the vector units.

Correspondingly, the management of the memory hierarchy (the main memory and the backing storage unit) requires the addressing and transfer of large blocks of data. This operation can proceed at the same time as vector arithmetic and mapping. Finally, many setup and housekeeping chores are necessary in nature and can be performed concurrently with the swapping, mapping, and arithmetic.

The choice of 8 vector units was based on tradeoffs between the search for a higher performance logic family than exists today, the amount of trunking and data alignment required, and the maximum amount of hardware that appears feasible to assemble, from power, cooling, physical geometry, and reliability standpoints.

Some additional points to be considered in the design and utilization of the NSS are:

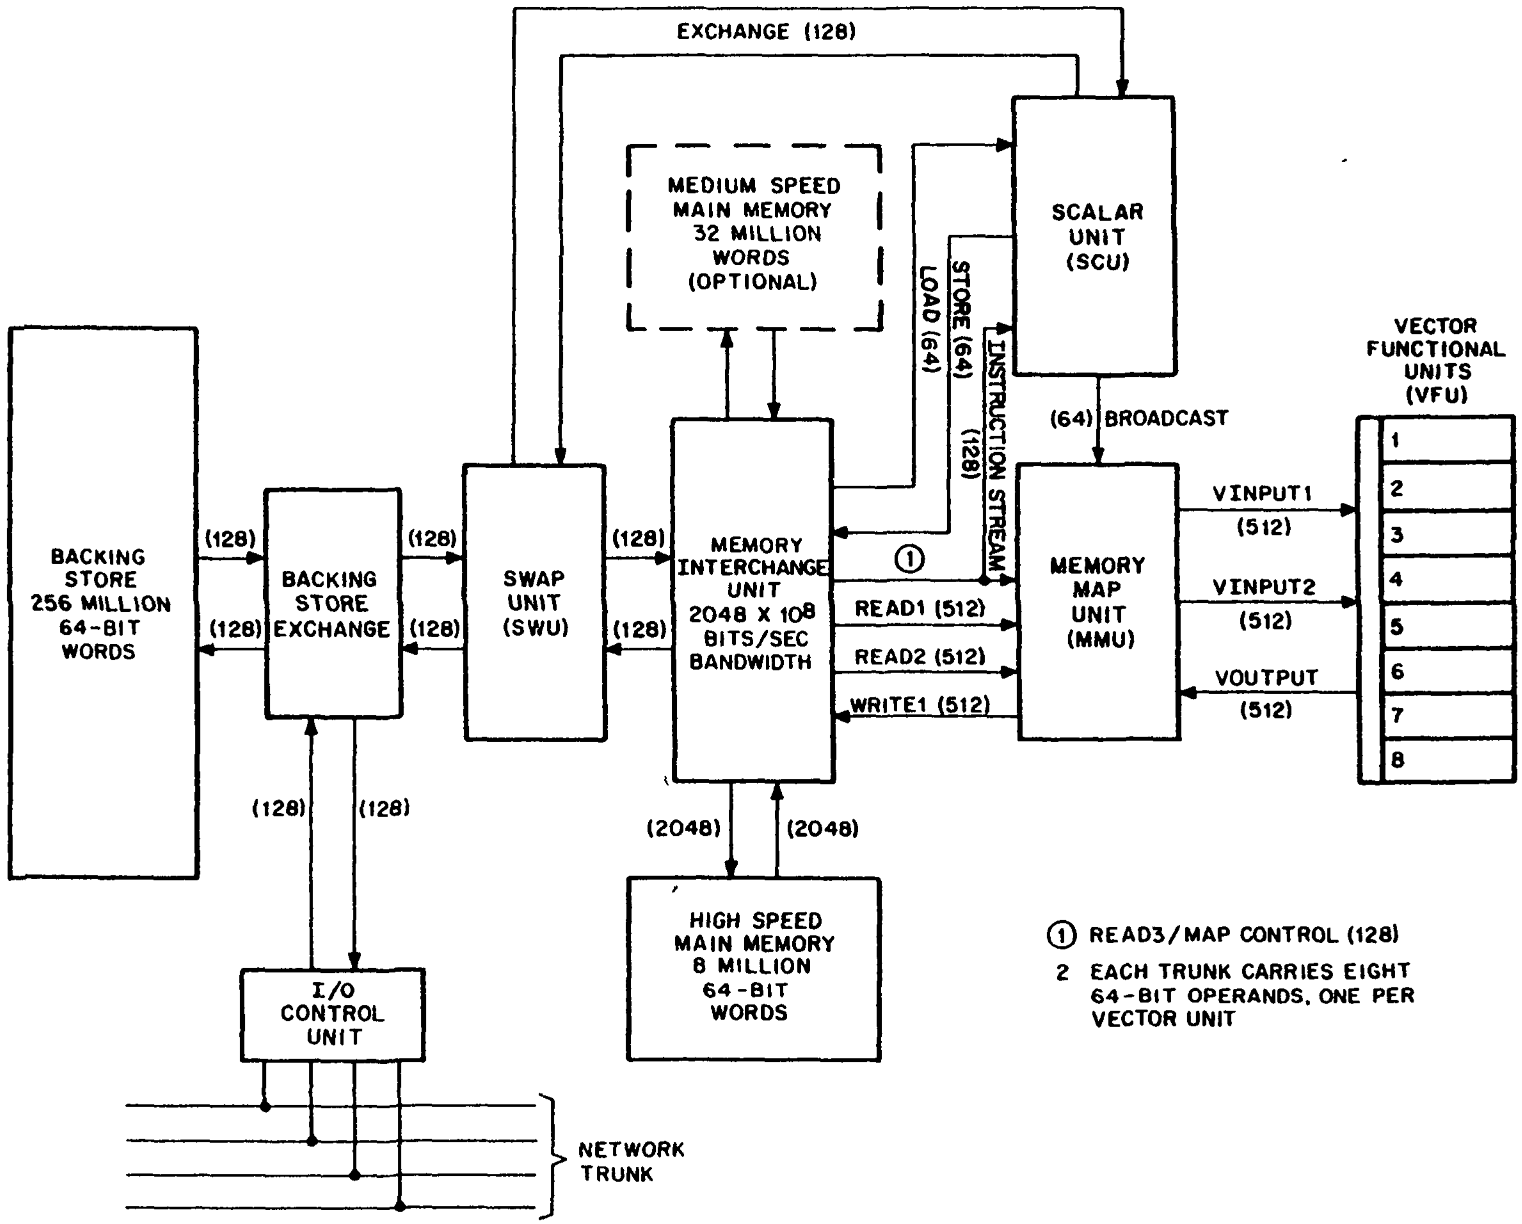

Figure S-2. Major NSS Components and Data Paths

- No existing computer system can perform the computations needed for 3-D Navier-Stokes solutions for flow field simulation. The NASF objective of complete solutions of these simulations in 7 to 15 minutes requires that such a processor achieve a sustained rate of computation between 1 to 2 gigaflops (1 to 2 billion floating-point operations per second). The fastest known machines today can attain peak rates for 64-bit computation slightly more than 100 million floating-point operations per second (100 megaflops), with sustained rates closer to 20 megaflops. This is a factor of 50 times slower than required.

- Given the projected technologies for the 1980's, no known existing computer architecture will yield the desired machine performance.

- With sufficient parallelism, such a machine is possible to design and build for operational employment in 1982.

- Key factors in achieving these goals are: the construction of sufficient memory to contain the entire problem on-line, without recourse to accessing slow speed mass storage devices; the ability to build a reliable collection of highly parallel hardware; and the programming and control of all the parallel hardware.

- Most of the data base (95 percent) can be maintained in 32-bit format, which reduces storage cost. Most of the computations (85 percent) can be performed in 32-bit form, with extended precision of at least 40 bits of coefficient required for a limited set of calculations. This makes possible the doubling of throughput of functional units when run in 32-bit mode instead of 64-bit mode.

- A processor containing a fast access memory of 8 million words of working storage and 256 million words of secondary storage can hold all projected problems. More importantly, such a memory can be made with known technologies and be made highly reliable through the use of error detection and correction techniques that are becoming commonplace in commercially available equipment.

- A processor with an 8-to 12-nanosecond clock and only eight separate functional units, each containing some localized parallelism, could achieve the 1-gigaflop threshold.
- The major problem to be solved in such an ensemble is that of sustaining the computing rate regardless of the manner in which memory is being accessed, linearly or randomly.
- Programmability and control of the necessary parallelism can be accomplished by melding together concepts taken from the STAR-100, the Texas Instruments ASC, and the ILLIAC IV.
- The most direct means for achieving programmability, reliability, and buildability is to begin with a single instruction stream, multiple data stream (SIMD) architecture.
- The NSS should be time-shared only in the most brute force manner, full rollout of the job in progress and the rollin of a new job, and then only in extraordinary circumstances. Otherwise, jobs should be permitted to go to completion.

RISKS

- Hardware risks anticipated for the proposed NASF range from negligible or minimal for front-end systems and network trunks, to moderate for graphics subsystems, to considerable for the NSS mainframe. The processors and peripheral devices with sufficient capability for front-end systems exist today and network trunks need little maturing to be sufficient. Graphics subsystems require some additional development of hardware and software as well as stabilization of approaches and techniques.
- To achieve the cost, performance, and reliability objectives established for this project, the NSS should be built with a second-generation, high-speed LSI. This technology is not yet available, and only expert opinion is available to ensure that this new generation will be available in time for the NASF. Alternative approaches can be taken yielding various degrees of reduced performance but decreasing the risk. Rough estimates of some of these approaches are:
  - Use of the planned STAR-100C for a run time of 30 minutes at essentially no risk.
  - An eight-pipe NSS using existing technology should yield a 15-minute run time with a risk factor of 0.1.

- An eight-pipe NSS using double-density chips should yield a 10-minute run time with a risk factor of 0.35.

- An eight-pipe NSS with double-density chips, 400-ps gate, should yield a 5-minute run time with a risk factor of 0.6.

- Software development absorbs an incredible amount of resources for even simple, uniprocessor systems. With much of the software expected to be used off-the-shelf, this risk can be ameliorated somewhat; however, considerable elapsed time will be required to stabilize the NSS compiler to the point where it can be put into general use. Three years is generally a minimum for such activity, even with a well-known language such as FORTRAN.

- The evolution of better algorithms for solving a system of partial differential equations such as the Navier-Stokes system could yield programs that would diverge radically from the form of the performance metrics. Thus, a specially tuned NSS could perhaps not be optimally tuned for the new algorithms.

- Although it is felt that costs and performance objectives can be tightly controlled to meet NSS requirements, scheduling remains very significant. The time frame is short, the technology is not yet in hand, and the design and simulation labor is extensive to produce the hardware complex. The biggest schedule risk, however, comes from the software development. Some steps which can be taken to help minimize the risk due to scheduling are:

  - Earliest possible initiation of each program phase, and earliest possible definition and stabilization of requirements.

  - Early selection of the front-end processor and leasing of time from the vendor of the target processor for software checkout.

  - Early release of software without all features for broader use and exercise of the software.

REQUIREMENTS

Results of this preliminary study indicate that pursual of this program to an installed, operational NASF entails the following estimated requirements.

- Development cost — 3.8 million dollars
- Acquisition cost — 40.8 million dollars (excluding housing)
- Operating costs — 1.8 million dollars per year (after installation)
- Space — 8000 ft^2 (excluding remote terminals)
- Energy — 980 kVA
- Cooling — 700,000 kcal/hr (2,780,000 Btu/hr)

RECOMMENDED FUTURE WORK

The next logical step in the development of the NASF is to refine:

- The structure and architecture of the NSS computational engine
- The analysis of the various forms of 2-D Navier-Stokes solutions
- The 3-D versions of the Navier-Stokes codes
- The definition of the resulting 3-D version of the Navier-Stokes program as the performance metrics
- The preliminary Navier-Stokes programming for the proposed NSS
- The definition of the programming language
- The system structure, applying workload data for peak and average operating periods to demonstrate that the supporting system will be adequate
- The schedules for all remaining aspects of the program

The final work product of this effort should be a series of detailed specifications for every component, whether it be programs, hardware, or buildings to be used to direct the design and construction of the NASF as well as to measure progress throughout the project.

AFTERWORD

The phase I study of the NASF has been a worthwhile experience for Control Data Corporation, and in particular, the RADL study team. It should be apparent that the design of the system, and most specifically that of the NSS, has undergone revision and evolution. This came about through a process of give-and-take with the staff at Ames. With the openness and candor permitted by the cooperative nature of this study phase, it was possible, RADL believes, to arrive at a better solution for the NSS architecture than could have been arrived at solely by the best resources within Control Data or Ames.

The probability for success of this project will rely heavily on the continuation of this excellent contractual relationship between NASA and vendor study teams. Only by merging the strengths of hardware production experts and mathematical and programming specialists can the most optimum system be obtained with the least cost and risk to all.