THE IMPACT OF "FOURTH GENERATION" COMPUTERS

ON NASTRAN

James L. Rogers, Jr.
NASA Langley Research Center

INTRODUCTION

The NASTRAN computer program (ref.1) is currently capable of executing on three different "third generation" computers, the CDC 6000 series, the IBM 360/370 series, and the UNIVAC 1100 series. In the past, NASTRAN has proved to be adaptable to the new hardware and software developments for these computers. The NASTRAN Systems Management Office (NSMO), as part of NASA's research effort to identify desirable formats for future large general-purpose programs, funded studies on the impact of the STAR-100 (ref. 2) and ILLIAC IV (ref. 3) computers on NASTRAN.

The STAR-100 and ILLIAC IV are referred to as "fourth generation" or "4G" computers in this paper. "Fourth generation" is in quotes because the differences between generations of computers is not easily definable.

Many new improvements have been made to NASTRAN as it has evolved through the years. With each new release, there have been improved capabilities, efficiency improvements, and error corrections. The purpose of this paper is to shed light on the desired characteristics of future large programs, like NASTRAN, if designed for execution on "4G" machines.

Concentration will be placed on the following two areas:

1. Conversion to these new machines
2. Maintenance on these machines

The advantages of operating NASTRAN on a "4G" computer is also discussed.

BACKGROUND

Figure 1 shows an example of the system changes NSMO has dealt with in the past and of some changes presently being contended with. Minor changes had to be made to Level 15 of NASTRAN when IBM released their 3330 disk packs. The changes by CDC to a SCOPE 3.4 Operating System and by IBM to Virtual Storage systems are causing considerable modifications for the operation of NASTRAN.

25

The STAR-100 and ILLIAC IV computers both have significant hardware and software features to support their respective pipeline and parallel processing capabilities. Pipeline and parallel processors can result in significant increases in computation speed when used on vector-type operations.

A diagram depicting the pipeline operations of the STAR-100 is shown in figure 2. When operating on two vectors, A and B, the pipeline works in the following manner. Elements A(1) and B(1) are received into the pipeline. They then proceed to the next unit in the pipeline, which is sign control. At this time elements A(2) and B(2) are received into the pipeline. A(1) and B(1) move to the align unit, A(2) and B(2) move to the sign control unit, and A(3) and B(3) are received by the pipeline. Each pair of elements then proceeds down the pipeline, with a new pair of elements entering the pipeline at each transfer, until the result is calculated and placed in the result stream.

The conceptual design of the ILLIAC IV with its 64 processing elements (PE) is shown in figure 3. The parallel processors operate differently on vectors than a pipeline processor. With parallel processors, PE_0 operates on A(1) and B(1), PE_1 operates on A(2) and B(2), ..., and PE_{63} operates on A(64) and B(64).[1] All of these operations take place simultaneously.

The STAR-100 and ILLIAC IV studies were conducted to gain insight into the potential impact of major system changes on large finite element programs like NASTRAN. In each of these studies there was one main objective: to investigate the feasibility of modifying Level 16 of NASTRAN in order to make it execute efficiently on the subject computer. This objective was to be accomplished in the following four steps:

1. Identify and describe the areas in NASTRAN which (a) easily lend themselves to or (b) could cause problems in conversion to the subject computer.
2. Determine the areas of NASTRAN where (a) modifications are needed to improve efficiency, and (b) significant benefits could be expected from using new strategies or algorithms for the subject computers.
3. Determine whether or not the above changes can be made in a way that the efficiency of NASTRAN can be improved with little or no increase in the number of computer dependent subroutines.
4. Estimate the time and cost involved in designing, coding, and implementing each of the modifications identified above.

Many different aspects of NASTRAN were studied. These items include:

1. Linkage Editor
2. Input/Output

3. Paging Problems
4. Machine-Dependent Code
5. Matrix Operations
6. Checkpoint/Restart
7. Compilers

The details of these aspects are discussed subsequently.

The STAR-100 and ILLIAC IV are completely dissimilar in the method of operating on vectors. Because of this and other dissimilarities, finite element programs like NASTRAN may require distinctly different versions to function efficiently on each machine.

CONVERSION

This section is concerned with the effort required to convert an existing version of NASTRAN to execute efficiently on a "4G" computer. Two basic questions are answered in this section. (1) What is the scope of the required changes in terms of time and manpower? (2) Which areas of NASTRAN must be converted to exploit "4G" technology?

Scope

The conversion effort to a new computer may be conveniently divided into a two-step process. The first step involves converting the currently existing NASTRAN to execute on the "4G" computer. The second step takes the converted code and improves it so that NASTRAN will execute efficient-ly on that computer. Table 1 summarizes the total effort required to complete both steps on the STAR-100 and ILLIAC IV.

An effort of 67 man months (ref. 2) over 9 months is estimated to convert NASTRAN to execute on the STAR-100. This effort results in only a scalar version of NASTRAN, which does not exercise the vector processing capability, and results in almost no improvement over the CDC 6600. To exploit the vector processing capability of the STAR-100 would require another 30-60 man months over a 10-18 month time period. Of the 67 man months in the initial conversion step, only 12 man months are to be used in actual NASTRAN code conversion.

An effort of 60 man months over 18 months (ref. 3) is estimated to convert NASTRAN to execute on the ILLIAC IV. This effort would not make full use of the parallel processing capability, but it is estimated that this will give the user 37% faster NASTRAN execution than the same run on an IBM 370/165. To make efficient use of the ILLIAC IV would require another 50-80 man months over an 18-24 month period. This effort would

allow NASTRAN to execute an estimated 90% faster than on the IBM 370/165. Of the 60 man months in the initial conversion step, 43 man months were estimated for actual NASTRAN code conversion.

Required Changes to NASTRAN

The changes that would be needed in the NASTRAN system include the following:

1. The Linkage Editor - The modifications to the Linkage Editor depend on the particulars of the computer involved. Conversion to the ILLIAC IV could require the Linkage Editor to be completely rewritten (a formidable task), since the present Linkage Editor on the ILLIAC IV has no overlay capability (ref. 3). Whereas, conversion to the STAR-100 could involve dropping the Linkage Editor in favor of executing NASTRAN as a single controllee file (ref. 2).

2. Input/Output - There are several important differences between NASTRAN and STAR I/O techniques (ref. 2). NASTRAN has hundreds of data blocks allocated for over 50 files, while the STAR Operating System (OS) provides less than 15 files. NASTRAN has open-ended files, while STAR OS requires allocation of the file space at the time the file is opened. The NASTRAN GINO provides random access methods employing indexed-sequential files, while STAR OS employs a simple sequential record manager. On the ILLIAC IV, the NASTRAN I/O package must be optimized to handle the bulk of data transfers between the processing element memories and the ILLIAC IV disk memory. In either the STAR-100 or the ILLIAC IV computer, because of the increased computational speed, the I/O must be highly optimized so as not to decrease overall efficiency.

3. Paging Problems - In STAR, a Virtual Storage computer, paging is a method for associating virtual memory with real memory. Several major factors influence the page size determination in a scalar virtual machine, namely code organization, compression, transport time, and the page replacement algorithm. Additional factors influencing the page size are created with the introduction of the vector capability. These factors include the cost of halting a vector instruction to replace a page, the cost of restarting a vector instruction, and the vector length. Before any conversion could take place, all of these factors would have to be examined and an optimal page size determined.

4. Machine-Dependent Code - All machine-dependent subroutines would, of course, require complete recoding in a "4G" assembly language.

5. Matrix Operations - The matrix operation modules of NASTRAN should be highly optimized for a "4G" computer, in order to exploit the special advantages of these computers. It was suggested in reference 2 that the NASTRAN matrix file structure could be optimized for the STAR-100 by dividing the matrix files into two separate files. One file would contain all the control information such as column, row position, and members of coefficients. This often enables one to operate directly on the coefficients without intermediate reorganization of the coefficients that for efficient pipeline processing. It was suggested in reference 3 that matrix operation modules could be optimized on the ILLIAC IV by developing detail specifications before beginning implementation. These preliminary design criteria would consider definition of array storage conventions within the ILLIAC IV processing element memories, and specialized storage schemes and disk mapping criteria for internal file communications and external files used in intermodular communication.

6. Checkpoint/Restart - Indiscriminate checkpointing of data files is most undesirable on "4G" computers. The transfer rates to and from a disk and central memory are slow compared to the execution power of a "4G" computer. Often, therefore, the cost effective approach would be to recalculate rather than checkpoint and restart.

7. Compilers - The STAR-100 FORTRAN compiler encompasses the NASTRAN FORTRAN subset with one exeception: the use of the ampersand symbol (&) in a calling sequence to signify a non-standard return label. FORTRAN specifies that the symbol be a dollar sign ($). The ILLIAC IV has a compiler option which will convert standard FORTRAN to IVTRAN, the ILLIAC IV FORTRAN-based language. This option examines DO loops of standard FORTRAN programs and converts them into more efficient DO FOR ALL loops for use on the ILLIAC IV.

Although all of these aspects of conversion are important, both studies (ref. 2,3) concluded that the majority of time in any conversion effort would be spent in optimizing the matrix operations.

Single-Programming and Multi-Programming

The ILLIAC IV is a single-programming computer, i.e. it is dedicated to execution of only one job at a time. Whereas, the STAR-100 is eventually anticipated to operate in a multi-programming mode, i.e. it will execute many jobs simultaneously at any one time. Reference 2 concludes that the STAR-100 CPU would remain idle most of the time if NASTRAN were executed on the STAR-100 in a single-programming environment. This, of course, would be very inefficient. Because of its configuration, the ILLIAC IV cannot handle a multi-programming environment. Thus one must definitely take the configuration of the conversion computer into consideration before conversion begins.

"Front-End" and Complete Conversion

One of the major questions that arose during both the STAR-100 and the ILLIAC IV studies was, Is the preferred configuration to have NASTRAN execute in a host plus "4G" computer ("front-end") environment (e.g. let STAR do what STAR does best and leave the rest to the CDC 6000) of for NASTRAN to be completely converted to the "4G" computer? Both studies concluded that if conversion were contemplated, the preferable mode is for NASTRAN to be converted to do all of its executions on a "4G" computer. There are several reasons for recommending the complete conversion concept over the "front-end" concept.

1. The STAR-100 requires 180 msec to transfer one page of data from the CDC 6000 to the STAR-100.

2. Once the "front-end" concept was working, the remaining conversion effort, to get all of NASTRAN on a "4G" computer, while involving significant volumes of code, would not require the further system type extensions.

3. The cost of total conversion is estimated to be less than that of the "front-end" concept.

There are some differences when converting to a host plus "4G" computer and a total conversion effort.

The conversion of NASTRAN to a host plus "4G" computer involves only a subset of NASTRAN. Prime candidates for the conversions are the functional modules which have modest input requirements, heavy computer and/or internal I/O requirements, and modest output requirements. New code must be generated to pass data between the "4G" computer and its host. Further, new code would have to be developed so that when NASTRAN is running on the host computer it can either continue processing or go into RECALL until a needed file is received from the "4G" computer.

If the complete conversion takes place, the resulting NASTRAN code would be computer dependent. It would no longer be compatible with a "third generation" NASTRAN and probably not even compatible with another "4G" NASTRAN. This would complicate the maintenance of NASTRAN, a situation discussed in the next section.

MAINTENANCE

Once a large computer program has been developed or converted and released to users, the maintenance of that program becomes the primary concern. NASTRAN's maintenance effort centers around an archive version.

This version is continually being modified and contains all of the latest
error corrections and new capabilities. The CDC, IBM, and the UNIVAC
versions are generated from this archive version. Each of these versions
also has its own unique features which must be maintained separately.
These features include machine-dependent subroutines, special linkage
editor control cards, and subroutines with multiple entry points or
non-standard returns.

As figure 4 shows, the archive version is used to create a particular
test version. Demonstration problems are then run on this version. If an
error occurred in a machine-independent subroutine, then its correction
in the archive version probably results in a correction in all versions.
However, if the error occurred in a machine-dependent subroutine, then it
may or may not occur in other versions and further testing is required.
After the known errors are corrected, the next version is tested. The
looping of this procedure is continued until all three versions of NASTRAN
are ready for delivery to the public. The extensive machine-independent
code and other well developed relationships among the three versions are
fully utilized to minimize the testing effort required.

The "4G" computers involve radical departures from the "3G" machines
and strong variations among themselves requiring different special pro-
gramming language. Thus, for such machines, all code is essentially
"machine-dependent". The cost of maintenance efforts for different
machine versions cannot be minimized through extensive commonality of code,
as it is for the three existing NASTRAN codes.

ADVANTAGE OF CONVERTING NASTRAN TO
A "4G" COMPUTER

The primary advantage in converting NASTRAN to a "4G" computer is the
gain in computational speed, especially for vector-type operations. Tables
2, 3, and 4 show some timing comparisons for the ILLIAC IV, STAR-100, and
present "third generation" NASTRAN computers. From table 2 it can be seen
that the STAR-100 and ILLIAC IV are on the order of 5 to 10 times faster
than the fastest "third generation" computer when a large number of steps
are involved in the calculation. Table 3 compares the potential efficiency
of NASTRAN operations performed on the ILLIAC IV with the IBM 370/165 and
the CDC 6600 computers. For this comparison, the process of matrix decom-
position was selected as a representative operation involving large amounts
of both computation and input/output processing. The decomposition of the
10,000 degree-of-freedom matrix would take 100 hours on the IBM 370/165
or 150 hours on the CDC 6600 when spill occurs. This same job, however,
could be run in 4 hours on the ILLIAC IV. Table 4 shows a time comparison
between the STAR-100 (anticipated) and CDC 6600 computers for decomposing
a stiffness matrix. The algorithms used are Gauss elimination (in symmetric
form) or Cholesky decomposition (with or without square roots) (ref. 2).

The effective bandwidth depends upon the numerical algorithm used in implementing the mathematical algorithm. For this table, the effective bandwidth has been set to $4\sqrt{N}$, where N is the number of equations. It is also assumed that both computers have full machine utilization of CPU time. It can be seen from the table that for 20,000 equations, the Cholesky method on the STAR-100 is 30 times faster than the Cholesky (FORTRAN) method on the CDC 6600 (13 minutes on the STAR-100, 6 hours 38 minutes on the CDC 6600). For the above tables, it is obvious that "4G" computers have a speed advantage when performing the large vector-type operations that are so common in finite element programs.

CONCLUDING REMARKS

As a part of NASA's research toward identification of desirable forms for future large finite element programs, studies were made of the required scope and technical changes which would be necessary to make NASTRAN operate efficiently on two "4G" computers, the ILLIAC IV and the STAR-100. Conversion efforts for either of these two computers could conveniently be divided into two steps. The first step would result in a working, not efficient, version of NASTRAN. The second step would optimize the results of the first step and yield an efficient version of NASTRAN on a "4G" computer. The first step alone was found not worth the effort, since the resulting version of NASTRAN would show only small improvements in execution speeds over similar "3G" versions. The time frame to complete both steps and release a "4G" version of NASTRAN to the public would take a minimum of three years.

Numerous areas of NASTRAN would need modification to take advantage of the increased computational speed of a "4G" computer. Areas requiring changes include the Linkage Editor, input/output, machine-dependent code, matrix operation subroutines, and the checkpoint/restart capability. Most of the effort, however, would be spent optimizing the matrix operation subroutines to exploit the capabilities of "4G" computers. A total conversion to a "4G" computer appears to be preferable to using a host "4G" computer environment. However, the converted "4G" NASTRAN would not be cost effective. Moreover, required changes would yield essentially all machine-dependent code and greatly amplify the burden of maintenance.

There are no current plans for NASA to convert NASTRAN to a "4G" computer. There are, however, other projects to develop structural analysis codes for "4G" computers. These are the ILSA (ILLIAC IV Structural Analysis) project sponsored by the Advanced Research Projects Agency and supervised by the Defense Nuclear Agency and a project designated as FESS (Finite Element System for STAR-100) at Langley Research Center.

REFERENCES

1. McCormick, C. W. ed.: "NASTRAN Users' Manual, "NASA SP-222(01), June, 1971.

2. Control Data Corporation: "Study of the Modification Needed for Efficient Operation of NASTRAN on the Control Data Corporation STAR-100 Computer." NASA CR-132644, 1975.

3. Universal Analytics, Inc.: "Feasibility Study for the Implementation of NASTRAN on the ILLIAC IV Parallel Processor." NASA CR-132702, 1975.

TABLE 1

EFFORT REQUIRED FOR EFFICIENT CONVERSION OF NASTRAN TO EXECUTE ON STAR-100 AND ILLIAC IV COMPUTERS

COMPUTERS	STEP ONE CONVERSION		STEP TWO OPTIMIZATION		TOTAL CONVERSION	
	MAN MONTHS	MONTHS	MAN MONTHS	MONTHS	MAN MONTHS	MONTHS
STAR-100	67	9	30-60*	10-18	97-127	19-27
ILLIAC IV	60	18	50-80	18-24	110-140	36-50

*ONLY INCLUDES VECTOR PROCESSING CAPABILITY, NO OTHER CONVENIENCE OR PERFORMANCE FACTORS.

TABLE 2

COMPARATIVE SPEEDS* OF EXECUTION BETWEEN

"THIRD AND FOURTH GENERATION" COMPUTERS

[For the ILLIAC IV and STAR, the figures are preliminary and for illustrative purposes only. For the 360/195 and 7600, the numbers are sensitive to the way the smaller, faster memories are used.]

Operation	Steps per stage	IBM 360/75	IBM 360/195	CDC 7600	CDC STAR	ILLIAC IV
Addition	$N = \infty$	0.24	4.6	5.2	50	50
	$N = 1$	0.24	0.55	1.6	0.57	0.78
Multiplication	$N = \infty$	0.14	4.6	5.2	25	44
	$N = 1$	0.14	0.53	1.5	0.57	0.69
Division	$N = \infty$	0.096	1.7	2.0	12.5	17
	$N = 1$	0.096	0.43	0.93	0.56	0.27

*64-bit precision computation speeds (memory to memory) in millions of operations per second.

TABLE 3

DECOMPOSITION TIMING ESTIMATES FOR NASTRAN

[For the ILLIAC IV, the figures are preliminary and for illustrative purposes only]

	IBM 370/165	CDC 6600	ILLIAC IV (After Step 1)	ILLIAC IV (After Step 2)
Decomposition (no spill)				
Matrix order 10000 dof				
Semibandwidth 300 dof	0.6 hours	1.3 hours	0.7 hours	0.01 hours
Active columns 50 dof				
Decomposition (with spill)				
Matrix order 10000 dof				
Semibandwidth 600 dof				
Active columns 50 dof	100 hours	150 hours	9 hours	4 hours
Working storage 50000 words				

TABLE 4

TIMING ESTIMATES FOR THE DECOMPOSITION

OF A STIFFNESS MATRIX

Number of Equations	Time (seconds)			
	STAR (FORTRAN)		CDC 6600 (CHOLESKY)	
	GAUSS	CHOLESKY	FORTRAN	COMPASS
100	.08709	.07469	2.7502	1.5520
250	.35107	.32035	10.862	6.1366
500	1.3634	1.3345	31.601	18.251
750	2.6060	2.6346	59.856	35.331
1000	4.0498	4.2010	94.844	57.044
3000	24.392	28.236	743.38	546.28
5000	55.013	67.703	2035.2	1612.2
10000	170.49	230.09	6839.7	5639.7
20000	537.25	798.85	23914.	20524.

LEVEL OF NASTRAN	12	15	15.5	16	
YEAR OF RELEASE	1970	1972	1973	1976	

IBM 3330 DISKS

CDC SCOPE 3,4 OS

IBM VIRTUAL STORAGE

Figure 1 - Examples of Impacts of Hardware and Software Developments on NASTRAN Releases

Figure 2 - Data Movement within the STAR-100 Pipeline

PROCESSING ELEMENTS
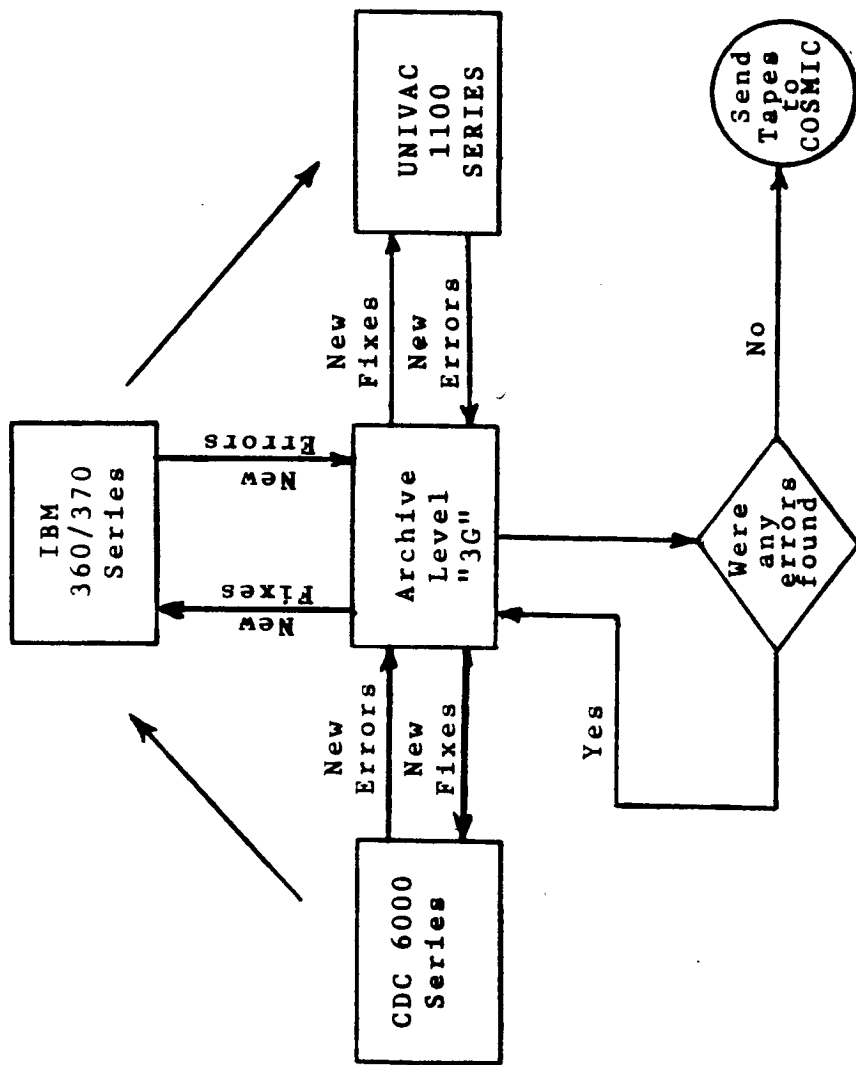

Figure 3 - Data Movement within the ILLIAC IV Processing Elements

Figure 4 - Computers and the NASTRAN Maintenance Effort