Experiences With MSC/NASTRAN On Various Hardware Platforms

By

Mulji Patel
FMC CORPORATION
Naval Systems Division
4800 East River Road
Minneapolis, MN 55421

Abstract

A number of production Finite Element (FE) analyses were run on different hardware platforms such as the Silicon Graphics workstation (380), IBM main frame (3090) and CRAY supercomputers (XMP and YMP) using MSC/NASTRAN Versions 65 and 66. The study was conducted to find:

1. Elapsed (Wall Clock) time for various hardware platforms using MSC/NASTRAN Version 65 and 66 using real world FE analysis problems in a production type environment.

2. Effect of various hardware platforms on overall FE analysis turnaround including pre- and post-processing time.

The study shows that the elapsed time is not directly proportional to the CPU rate (MFLOPS), rather it is significantly affected by other factors such as memory size, I/O operations, cache memory utilization, software constructs (V65 to V66), batch job queue, and general system parameters. The study also shows that significant time is required to interface software, hardware and communications media into a complete, usable system for pre and post processing and analysis as well as training the users when a change is made from one hardware platform to another.

1.0 Introduction

In recent years there has been a significant growth in computer technology as well as in Finite Element Analysis (FEA) software. Today these programs are available on a variety of hardware ranging from PCs to supercomputers to meet the demands of entry-level to experienced full time FEA users.

Though FEA software has been around for a long time, it takes several years for a user to become proficient, especially with a general purpose finite element program such as MSC/NASTRAN. Therefore, it is unlikely that an experienced FEA group such as the group at FMC Naval Systems Division will switch to another analysis program. However, the availability of high speed computing power at lower cost encourages FEA users to explore different hardware for their FEA work.

When a switch from one computer system to another is desired, the question of how to measure performance of FEA on the new hardware arises. The most common performance parameters used are MFLOPS (Millions of Floating Point Operations per Second), clock cycles in nsec (nanoseconds), MIPS (Millions of Instructions Per Second), etc. Supercomputers are vector processing machines that can generate many operations from one instruction which make MIPS ratings of these machines less important. MFLOPS is only a measure of the peak speed of the computer's arithmetic unit. It is hard to realize this peak speed with general purpose FEA codes because of intensive I/O (Input/Output) operations.

For FEA work, however, the most important factors are the WALL CLOCK (ELAPSED) time and the CPU time. WALL CLOCK time indicates the time a computer takes to complete an analysis, while the computer cost of doing FE analysis is based on the CPU time. In this study, these two factors have been evaluated on five different computers with respect to single user vs multi-user, batch job queues, MSC/NASTRAN V65C vs V66B, etc. Also, other factors associated with FEA such as remote printing and plotting, high speed communications, interfacing modules for pre- and post-processing, operating systems, etc. have been investigated for a complete, usable system.

2.0 Computer Systems Used

The computers used for the study are listed on Table I. The IBM 3090 located at FMC's Corporate Data Center is shared by all FMC's divisions. It is mainly used for business and accounting applications; technical computing is a small part of its use. Though it is a multi-user system, the queue time for processing is generally not long.

The NCSA (National Center for Supercomputing Applications) at the University of Illinois, Champaign-Urbana, Illinois has two supercomputers: CRAY-2 and CRAY Y-MP. They had a CRAY X-MP until the end of October 90; it has now been replaced by a CRAY Y-MP. MSC/NASTRAN V65C and V66B are available on Y-MP. Both computers are used by academic as well as industrial users from all over the United States. Thus, they are heavily used.

The CRI/YMP is a Cray Research Inc. supercomputer at their Cray Research Park facility in Eagan, Minnesota. It is mainly used for software code testing, benchmarking, etc.

The SGI is a workstation from Silicon Graphics Computer Systems, Inc. The MSC/NASTRAN benchmark runs were done at their facility in Arden Hills, Minnesota.

Table I: Computer System Specifications						
Computer	No of Procrs	Memory (MW)	Clock (ns)	Word (bits)	Operating System	I/O Device
1. IBM	2	32	14.5	32	MVS	–
2. NCSA/XMP	4	8	8.5	64	UNICOS	SSD
3. NCSA/YMP	4	64	6.1	64	UNICOS	SSD
4. CRI/YMP	4	64	6.1	64	UNICOS	SSD
5. SGI	1	14	30.3	32	UNIX	–

Notes:

1. IBM Computer: 3090/600E, located at FMC's Corporate Data Center, Dallas, TX. Multi-user.
2. NCSA/XMP and NCSA/YMP: CRAY supercomputers X-MP/48 and Y-MP4/464 at NCSA (National Center for Supercomputing Applications), University of Illinois, Champaign, IL. SSD-Solid State Device for I/O. Multi-user.
3. CRI/YMP: Cray Research Inc.s' supercomputer Y-MP SN/1033, located at Cray Research Park, Eagan, MN. Single-user.
4. SGI: Silicon Graphics' workstation model 4D/380S, located at Arden Hills, MN. Single-user.
5. Though multiprocessors are listed, only one processor/computer was used doing the analysis.

3.0 Finite Element Models

The four finite element models used for the study are shown on Figures 1 through 4. Table II lists their characteristics such as number of nodes, number and type of elements, MSC/NASTRAN Solutions and Versions. These are all production, real world finite element models comprised of many different elements. Though Static Solution (SOL24) is typically used quite often for such analyses, static comparisons are not included in this study because they generally take less CPU time and computer resources when compared to the common dynamic analyses done at FMC.

4.0 CPU Time

Table III and Graph 5 show the CPU time for all the four FE model runs on the five computers. Since we at FMC have been using the IBM computer system for a long time, the CPU and WALL CLOCK time comparison has been made with respect to the IBM time. The CPU time for the other computers does follow, in general, the computer's clock speed. It takes less CPU for the analysis as the clock gets faster. However, the CPU time varies a lot depending on the analysis and the environment. For example, the NCSA/XMP CPU ratios varies from 0.50 to 0.77 in comparison to the nsec ratio of 0.59. A similar spread in CPU ratio, 0.40 to 0.81 against the nsec ratio of 0.41, is found for the NCSA/YMP. However, with CRI/YMP the ratio is only 0.34 to 0.51 against the nsec ratio of 0.41. This demonstrates the effect of the computer environment on CPU time. The CRI/YMP used a single-user environment while the NCSA computers are configured for a multi-user environment. For the SGI workstation the CPU time includes the I/O time. No separate data is available for the I/O time on SGI. For the other computers in the study, the processing cost is based on the CPU time which has been separated from the I/O time. Therefore, it will be hard to draw any conclusion for the SGI regarding its actual CPU time.

Graph 7 with Table V and Graph 8 along with Table VI show the CPU time for the MSC/NASTRAN V65C versus the MSC/NASTRAN's new version V66B on the NCSA/XMP and the NCSA/YMP respectively. For the case of V66B, both the Unstructured and the Structured solutions were benchmarked for the CPU time. For some reason, the NCSA/XMP showed higher CPU time for the V66B solutions except for the FEM4 model. While in the case of the NCSA/YMP the reduction in the CPU time was realized for all the analysis runs; however, the reduction in CPU time varies from 0.62 to 0.90. Also, the Unstructured and the Structured solutions of V66B did not make much difference in CPU time except for the FEM4 analysis.

5.0 The WALL CLOCK Time

The WALL CLOCK times as shown in Table IV and Graph 6, vary significantly depending on the computer used as well as the environment. The WALL CLOCK time is generally inversely proportional to the computer's clock speed; the faster the clock, the less WALL CLOCK time is required for the analysis. Considering the IBM WALL CLOCK as a base time, the NCSA/XMP and the NCSA/YMP WALL CLOCK times vary 32% to 78% and 33% to 89% respectively. In the case of the CRI/YMP, the WALL CLOCK time has been reduced by a significant factor – as low as 0.15. Graph 9 and Table VII illustrate the environmental effect on WALL CLOCK time of the YMP, which is used as single-user at CRI and multi-user at NCSA facilities. This again illustrates how the multi-user environment slows down the analysis work even on high speed supercomputers. The SGI workstation shows better performance for the WALL CLOCK time than the IBM or the NCSA/XMP, however, it is only executing a single job at a time.

Another WALL CLOCK time factor for a multi-user system such as NCSA is the number of users on the system at a time. It varies with time in a given day as well as from day to day. Our experience with NCSA is that some days when the system is very busy it can take two to three times longer WALL CLOCK time than reported in this study.

The Graph 10 and Table VII show another system variable which affects WALL CLOCK time. It is the complex batch job queue system (NQS) at NCSA. This system is used by the industrial as well as the academic users. The queues are prioritized according to CPU time and memory size required to process a job. Also there are limitations on how many jobs a user is allowed per queue, the total number of jobs per user at any time, and the number of jobs per queue at any time. The system operator has an overriding power to set aside some jobs, especially the big jobs, if the system reaches its full capacity. However, NCSA redesigned this NQS system in July 1990 and the WALL CLOCK performance has been improved as shown in Graph 10.

Some other factors affecting the WALL CLOCK time are related to fine-tuning the hardware to application software such as MSC/NASTRAN. This includes optimum memory size, cache memory utilization, adequate size of the I/O disks, record length size for databases, post files, etc. On the IBM, for example, specifying the record length and a block size of 23472 bytes for a database reduces the WALL CLOCK time significantly. The 23472 bytes size corresponds to the half track size on an IBM disk. The NCSA supercomputers are not fine-tuned yet for such factors but they are working on such performance changes.

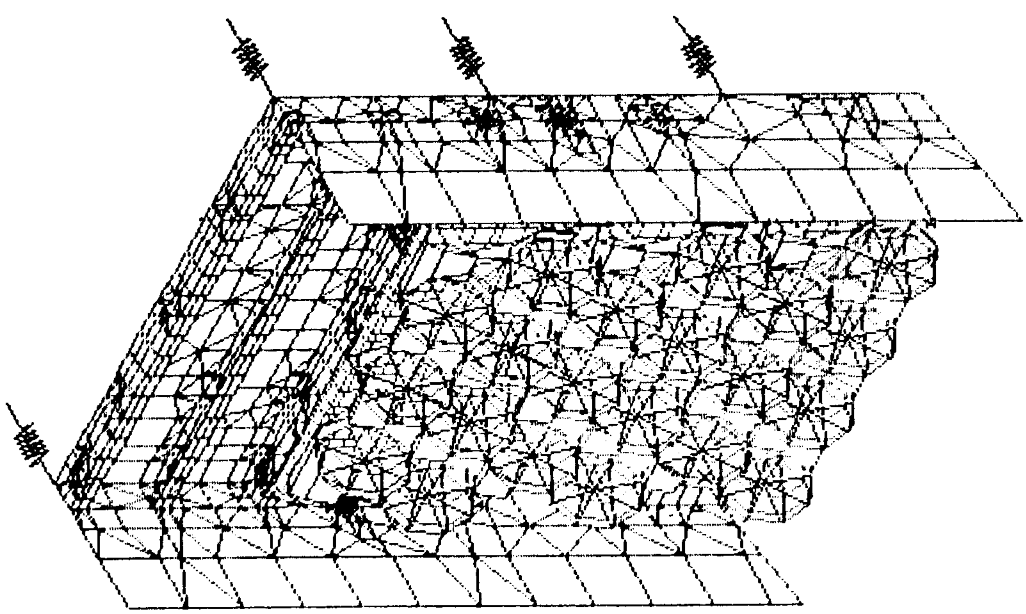

Figure 1 : FEM4

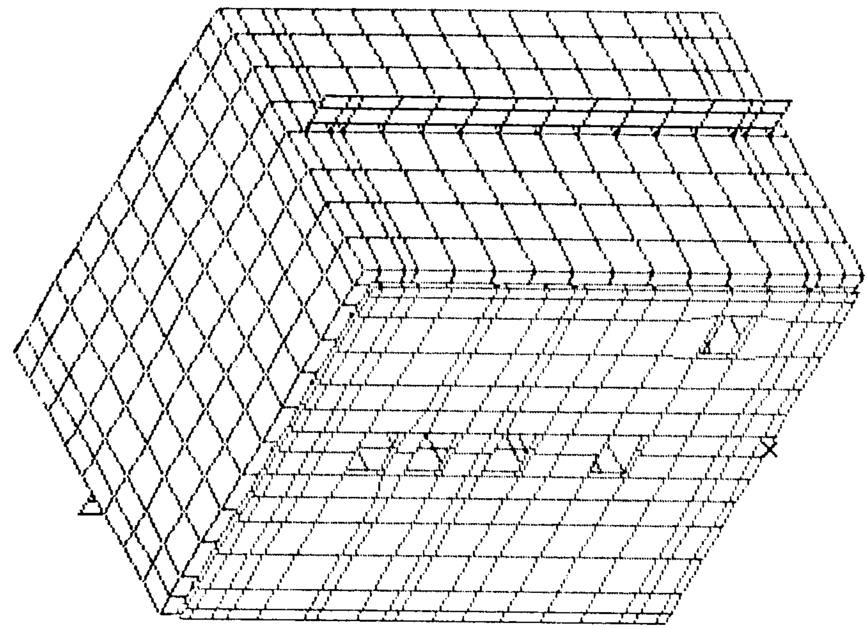

Figure 2 : FEM1

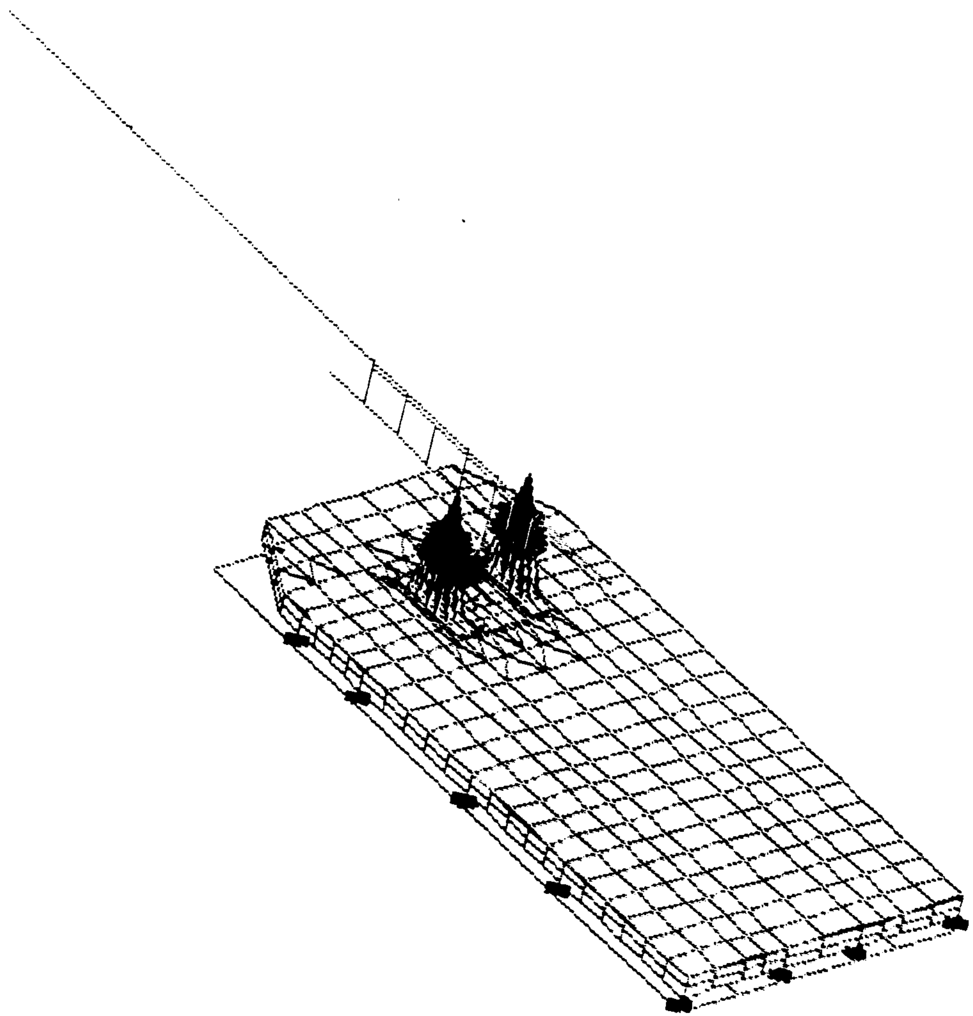

Figure 3 : FEM3

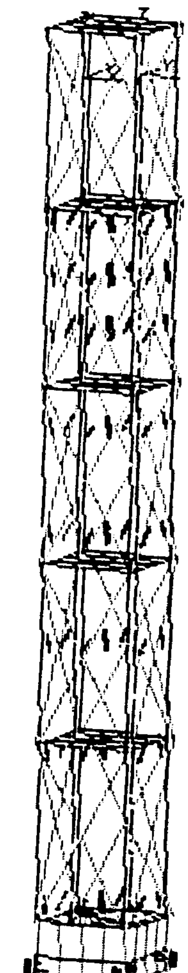

Figure 4 : FEM2

Table II: Finite Element Models

		Grids	Elements		Solution
1.	FEM1	1527	QUAD4 =	1177	V65C– SOL63, Modal Extraction
			BAR =	118	V66B – SOL63 Unstructured
			ELAS2 =	132	V66B – SOL103 Structured
				1427	
2.	FEM2	902	QUAD4 =	148	V65C – SOL69, Direct Transient
			TRIA3 =	28	V66B – SOL69 Unstructured
			BAR =	576	V66B – SOL109 Structured
			ELAS2 =	90	
				842	
3.	FEM3	1116	QUAD4 =	831	V65C – SOL69, Direct Transient
			TRIA3 =	58	V66B – SOL69 Unstructured
			BAR =	439	V66B – SOL109 Structured
			ELAS2 =	411	
				1739	
4.	FEM4	4931	QUAD4 =	1605	V65C – SOL3, Modal Extraction
			TRIA3 =	2108	V66B – SOL3, Unstructured
			BAR =	254	V66B – SOL103 Structured
			BEAM =	44	
			ELAS2 =	6	
			HEXA =	944	
			PENTA =	82	
			TETRA =	174	
				5217	

6.0 Other Factors

Though CPU and WALL CLOCK time are important factors, there are a number of other factors which directly affect FE analysis. Some of these factors are transmission line response for interactive sessions, large file transfer rate, remote printing and plotting, networking with local computers, terminals, printers, plotters, etc. with such remote computers, pre- and post-processing, security issues especially with the defense contractors such as FMC, technical support and training, long term file storage, operating system, etc.

High speed transmission line reliability is a must when dealing with the huge amount of data generated with a FEA code such as MSC/NASTRAN. It is very easy to generate 40 to 50 Megabytes of output or plot files with transient analysis. Generally, such processing is done at the supercomputer and only the result are brought back to a local computer for the post processing. Therefore, a quick transfer of these output files is necessary. There is a significant amount of investment and time required to install such a high speed line.

Also, the transmission line is an important factor for the local terminal response time for the interactive work with the remote computers. A high speed line gives almost an instantaneous response while editing a file or reviewing the results interactively. This saves a lot of analyst's time and facilitates quick turn-around time for an analysis.

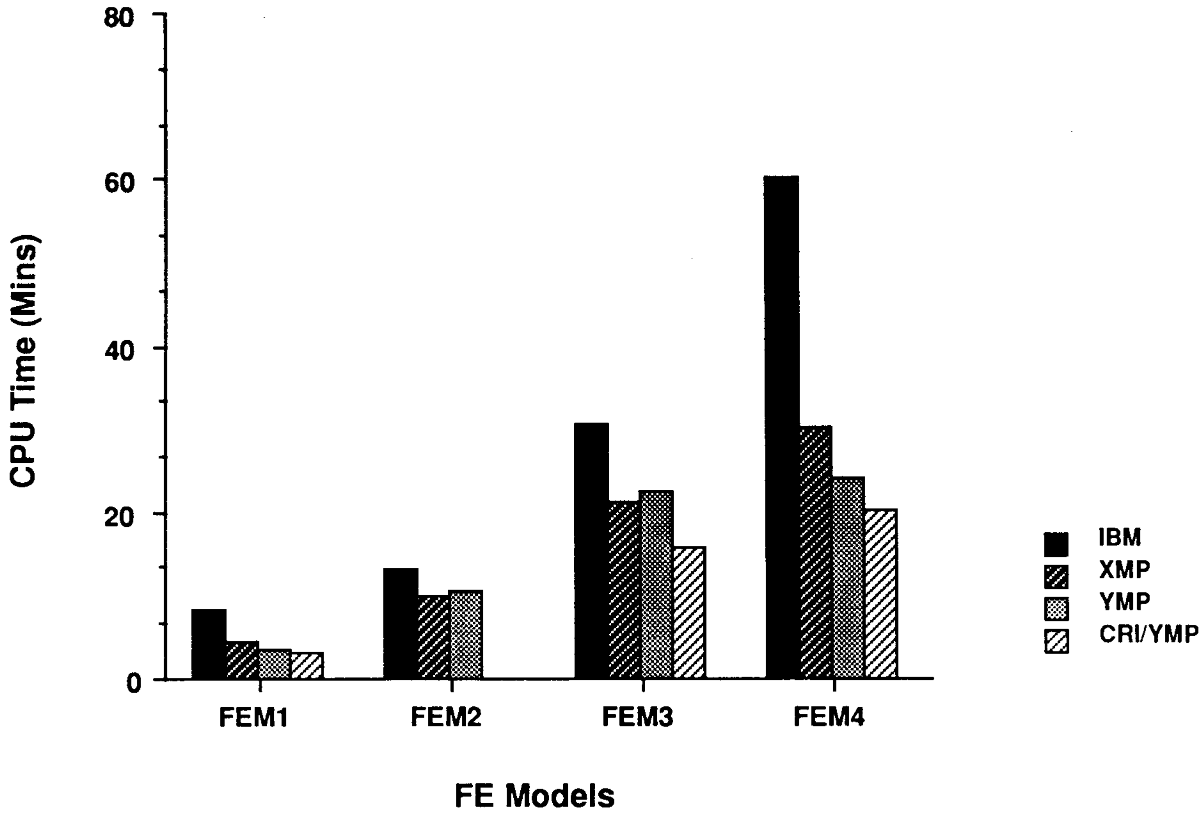

Graph 5: CPU Time

Table III: CPU Time – MSC/NASTRAN V65C				
FE Model	IBM HH:MM:SS	XMP HH:MM:SS	YMP HH:MM:SS	CRI/YMP HH:MM:SS
FEM1	8:13	4:37 (56)	3:25 (42)	3:19 (40)
FEM2	13:06	10:07 (77)	10:34 (81)	N/A
FEM3	30:40	21:20 (70)	22:20 (73)	15:38 (51)
FEM4	1:00:01	30:04 (50)	24:08 (40)	20:15 (34)
nsec	14.5	8.5 (59)	6.1 (41)	6.1 (41)

( )– % of IBM

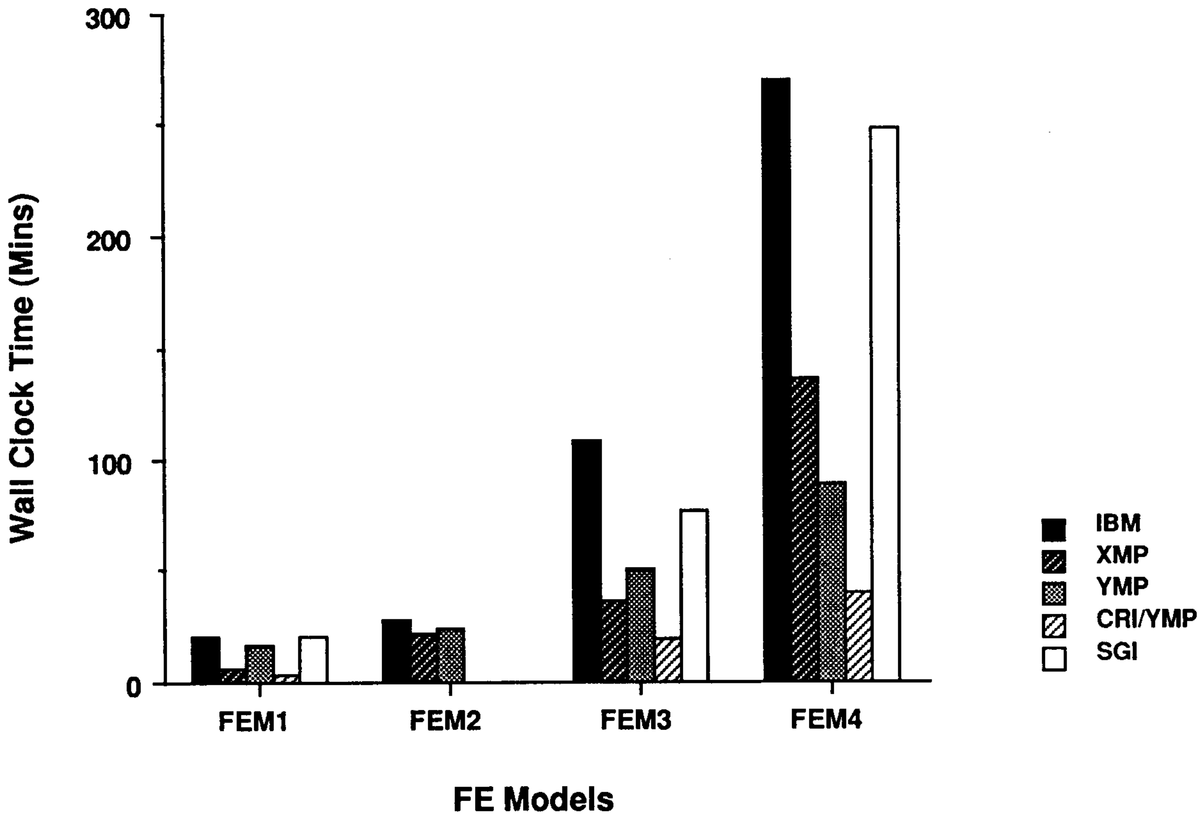

Graph 6: Wall Clock Time

Table IV: WALL CLOCK Time – MSC/NASTRAN V65C

FE Model	IBM HH:MM:SS	XMP HH:MM:SS	YMP HH:MM:SS	CRI/YMP HH:MM:SS	SGI HH:MM:SS
FEM1	20:20	6:27 (32)	17:07 (84)	3:33 (18)	21:03* (104)
FEM2	27:28	21:21 (78)	24:14 (89)	N/A	N/A
FEM3	1:48:04	35:46 (33)	50:33 (47)	19:04 (18)	1:17:16* (72)
FEM4	4:30:00	2:16:23 (51)	1:28:56 (33)	39:38 (15)	4:07:52* (92)

( )– % of IBM
* – MSC/NASTRAN V66A

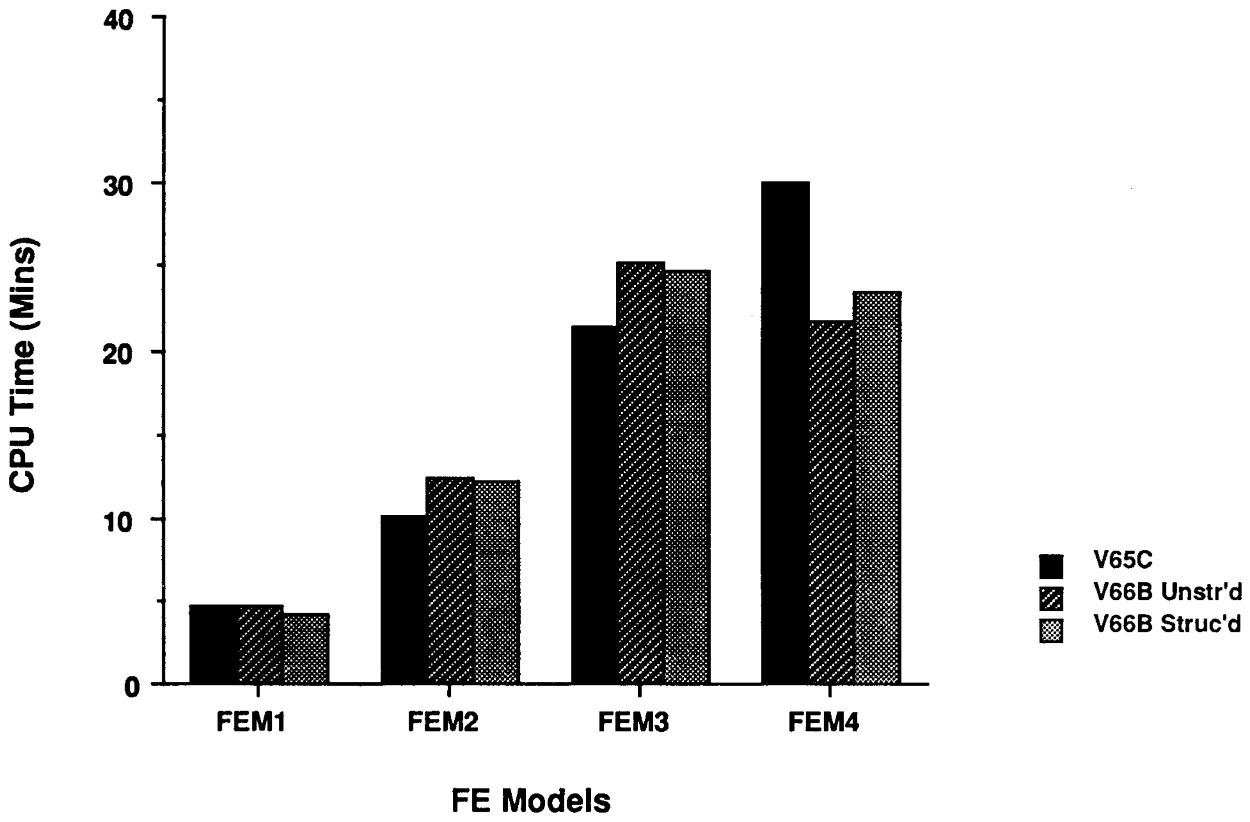

Graph 7: MSC/NASTRAN V65C vs V66B (NCSA/XMP)

Table V: MSC/NASTRAN V65C vs V66B
CPU Time – NCSA/XMP

FE Model	V65C HH:MM:SS	V66B/UNSTR HH:MM:SS	V66B/STRC HH:MM:SS
FEM1	4:37	4:37 (100)	4:15 (92)
FEM2	10:07	12:23 (122)	12:16 (121)
FEM3	21:20	25:14 (118)	24:49 (116)
FEM4	30:04	21:46 (72)	23:32 (78)

( ) – % of V65C

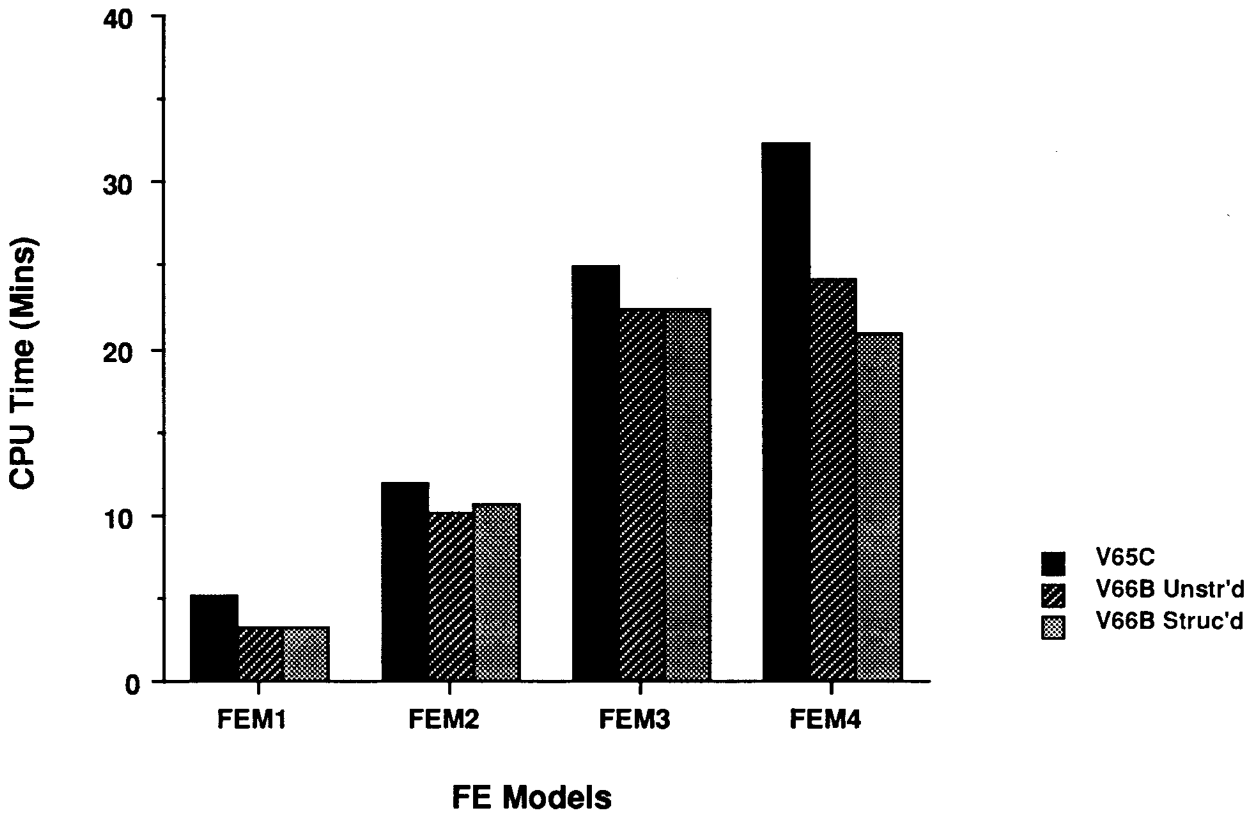

Graph 8: MSC/NASTRAN V65C vs V66B (NCSA/YMP)

Table VI: MSC/NASTRAN V65C vs V66B
CPU Time - NCSA/YMP

FE Model	V65C HH:MM:SS	V66B/UNSTR HH:MM:SS	V66B/STRC HH:MM:SS
FEM1	5:13	3:25 (62)	3:17 (63)
FEM2	11:49	10:07 (86)	10:34 (89)
FEM3	24:50	22:20 (90)	22:23 (90)
FEM4	32:21	24:08 (75)	20:57 (65)

( ) – % of V65C

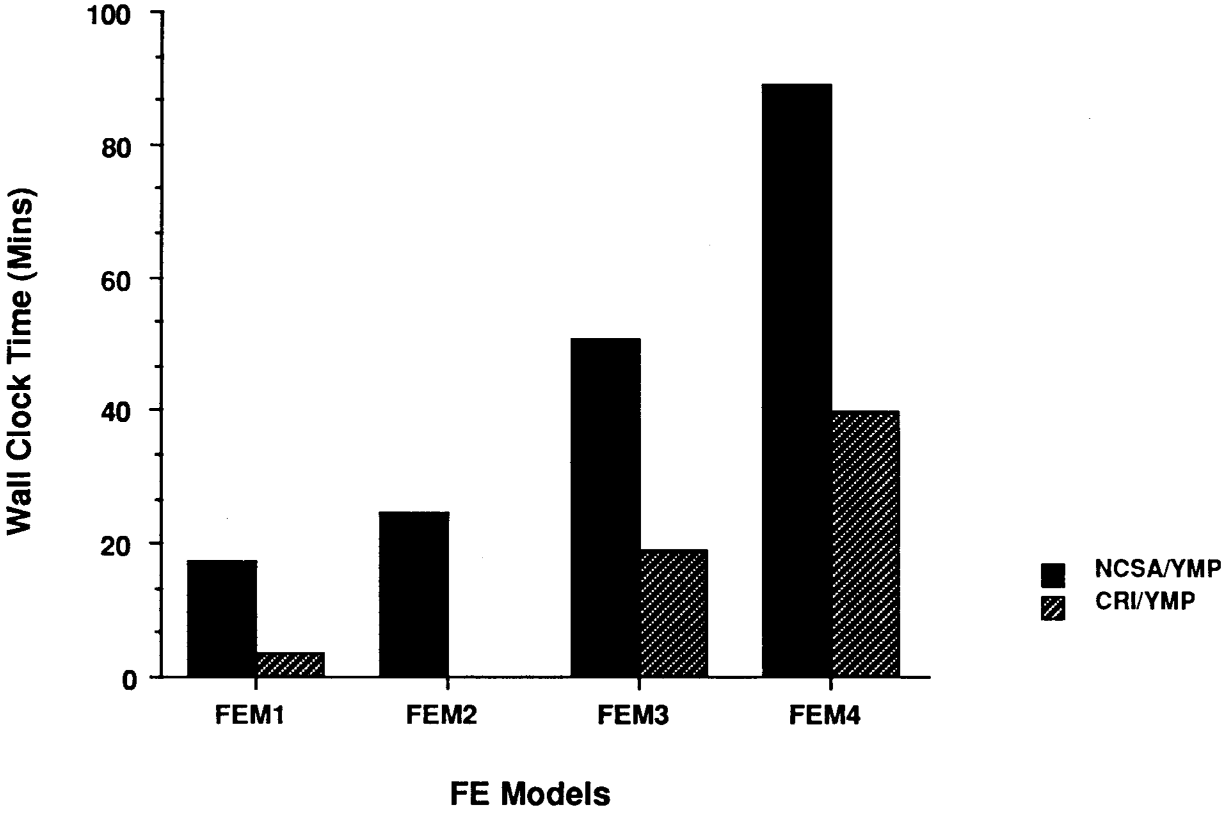

Graph 9: NCSA/YMP vs CRI/YMP

Table VII: Multiusers vs Single-user WALL CLOCK Time
NCSA/YMP vs CRI/YMP
MSC/NASTRAN V65C

FE Model	NCSA/YMP HH:MM:SS	CRI/YMP HH:MM:SS
FEM1	17:06	3:33 (21)
FEM2	24:14	N/A
FEM3	50:33	19:04 (38)
FEM4	1:28:56	39:38 (45)

( ) – % of NCSA/YMP

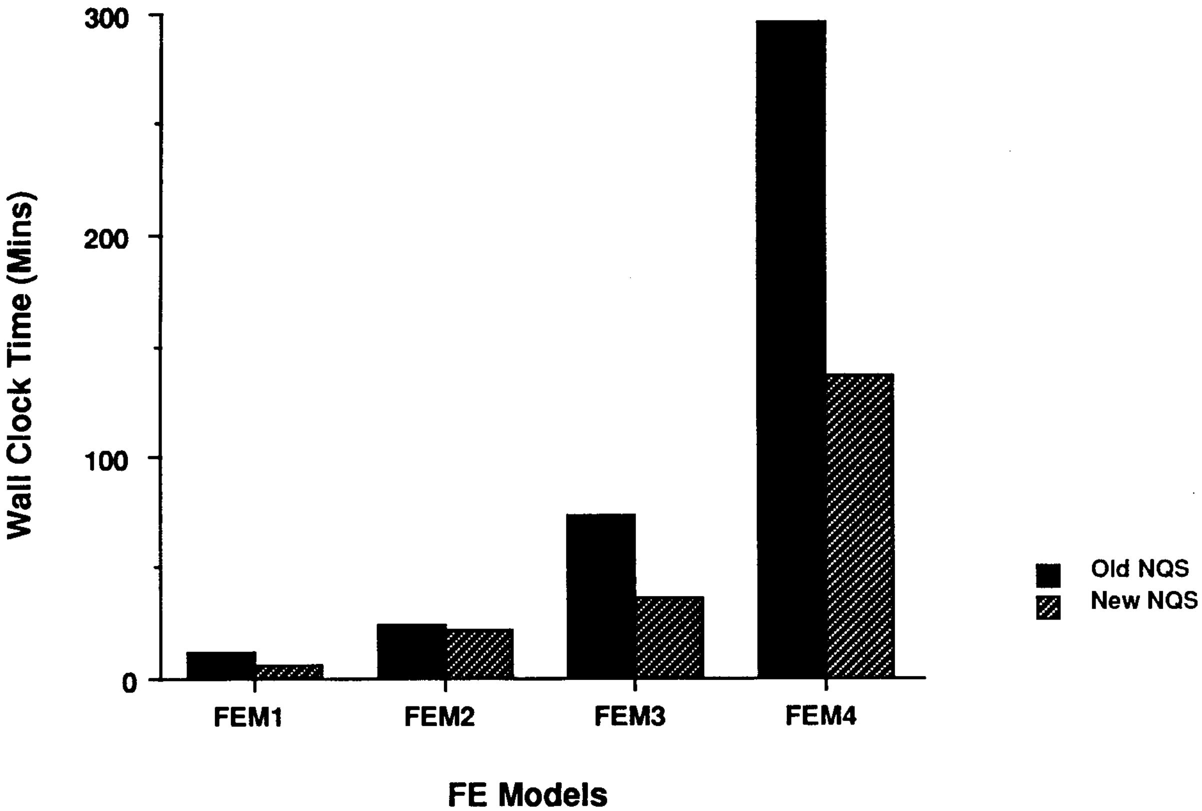

Graph 10: Old NQS vs New NQS (NCSA/XMP)

Table VIII: Old NQS vs New NQS (Batch Queue) NCSA/XMP - WALL CLOCK Time MSC/NASTRAN V65C

FE Model	Old NQS HH:MM:SS	New NQS HH:MM:SS
FEM1	12:10	6:27 (53)
FEM2	24:20	21:21 (88)
FEM3	1:13:02	35:46 (49)
FEM4	4:55:34	2:16:23 (46)

( ) – % of Old NQS

Remote printing of the output on paper or microfiche is necessary for documentation and record keeping. Also, since the transient analysis generates a huge amount of data, selective plots on screen or paper are required for efficient evaluation of such results. Therefore, the local printer/plotter has to be properly connected to the remote computer through high speed lines. Chances are high that the existing local hardware may not match the new remote computer environment and may require additional investment.

Besides the local hardware compatibility with a remote computer, there is a security requirement for the defense contractors that the communication has to be one way only. The remote computer cannot communicate directly with the local computer. Because of this requirement, the printing or plotting on local hardware cannot be initiated by the remote computer in batch mode. This means the files have to be transferred interactively through the local terminal, which ties up the local terminal until the file transfer is completed.

Getting a rapid as well as quality response from the technical support group at the remote computer is very important for the analysis in reducing the turn-around time. It takes a number of years to build a technical support group who can understand FE analysis requirements, the computer system and application software such as MSC/NASTRAN. The FE analyst does not have time and training to understand thoroughly the computer system and to interface the software with the hardware efficiently. It is the technical support group who puts all these pieces together for an efficient FE analysis system and keeps tuning the system as new requirements arise.

Also, when a change is made from a system like IBM to a system like NCSA which uses the Unix based operating system (UNICOS), it takes a significant portion of an analyst's time to learn the new operating system. Besides the operating system, there are other new things to learn such as the file editor, file handling system, CPU cost estimating, tracking and reporting, etc. Thus, training an FE analysis group on a new system can be a significant part of the cost and time.

7.0 Conclusions

Though MFLOPS, MIPS or nanosecond parameters are somewhat a measure of relative performance of computers/supercomputers, the real practical measures are CPU and WALL CLOCK time in a production type FE analyses. For a general purpose FE program such as MSC/NASTRAN, the CPU time depends on the problem size, number and type of elements, software constructs, the hardware configuration including available memory, cache memory utilization, I/O media, and whether the computer environment is for single or multi users. Since the computer cost and the software royalty charges are based on CPU time, comparing the CPU time in benchmark work is important. The study also shows that the WALL CLOCK time varies significantly from computer to computer, and very much depends on its environment. When a computer/supercomputer is moved from a single-user to a multi-user environment, the WALL CLOCK time may go up as much as 3 to 4 times as illustrated by the benchmarks done using the NCSA/YMP and the CRI/YMP.

Besides CPU and WALL CLOCK time consideration for a FE analysis group, there are other important factors that significantly affect the FE analysis turn-around time. Close attention and thorough investigation should be made when deciding to process FE analysis on a remote computer for such factors as high speed transmission line, remote printing and plotting, interfacing local computers, terminals, printers, plotters, security issues, technical support and training, long term file storage, etc. Significant time, manpower and capital investment are required to build a complete FE analysis system.

8.0 Acknowledgment

The benchmark work on CRI/YMP was done by Mr. Dawson Deuermeyer of Cray Research Inc., at Cray Research Park, Eagan, Minnesota on their supercomputer Y-MP SN/1033. Also, the SGI workstation benchmark timings were provided by Mr. Charles S. Larson of Silicon Graphics Computer Systems, Inc., Arden Hills, Minnesota. Both Mr. Deuermeyer and Mr. Larson were very helpful in providing their system configuration and information. Special thanks to Mr. Dave Wang of FMC's Dallas Data Center, and Mr. Ilhan Dilber, Dr. Fouad Ahmad and Mr. Ramesh Jayaraman of NCSA for their technical support in running these FE models at their computer centers.

9.0 References

(1) Ahmad, Fouad M., and Dilber, Ilhan, "Performance Evaluation of Issues of Finite Element Codes on Supercomputers", First NCSA Conference on Finite Element Applications In Computational Mechanics, May 1990.

(2) Taylor, Ian A., and Huseby, John, "MSC/NASTRAN In A High Performance Computer Environment", The MSC/NASTRAN 1987 World Users Conference Proceedings, Vol 1, 1987.

(3) "MSC/NASTRAN Application Manual, Vol II, V66B CRAY (UNICOS), The MacNeal-Schwendler Corporation, Los Angeles, CA, May 1990.