Siemens PLM Software

Satisfying the perpetual need for increased digital simulation performance

www.siemens.com/nx

white paper

Maximizing the business impact of technology requires far more than just buying the right point functionality and handing it off to the analysis department. Digital simulation needs to be at the core of every PLM business process because it enables management to make faster, more informed decisions. Overall, increasing use of digital simulation leads to better products that are more salable, have better performance and higher margins, all of which directly benefit the bottom line.

PLM Software

Answers for industry.

SIEMENS

▶ Satisfying the perpetual need for increased digital simulation performance

Table of contents

Executive summary

With the increasing acceptance of the value of digital simulation combined with seemingly inexhaustible advances in computer hardware, there is a correlating trend – the perpetual need for improved performance. Leading software solutions are based on geometric, frequency and hierarchic domain decomposition methods. The fundamentals of these methods and an industrial case study are presented in this paper.

Furthermore, it is important to develop technologies that befit the current hardware architectures. The proper pairing of software technologies with appropriate hardware environments and future directions are discussed as well.

Introduction

The superior importance of performance in today's implicit finite element analyses is due to very large (and ever increasing) model sizes and extremely wide frequency ranges of interest. To analyze the dynamic behavior of the structures under these circumstances, modal solution techniques are commonly used in the industry.

Modal solutions are based on the modal representation of the problems, and the cornerstone of this is the efficient calculation of the free, undamped vibration of the mechanical system – that is, the eigenvalue analysis problem of

$$(K_{aa} - \lambda M_{aa})\phi_a = 0$$

Here the eigenvalue λ represents a natural frequency, the eigenvector ϕ is a free vibration shape of the finite element model. The a subscript refers to the analysis partition of the finite element model. This is one of the most time-consuming computations of large-scale global analyses in the automobile and the aerospace industries. Hence, this is chosen to be the basis for discussion in the following sections.

Geometric domain decomposition technique

The geometric domain decomposition is executed by automatically subdividing the geometry based on the connectivity information obtained from the finite element model. Such a subdivision of a simple rectangular plate problem is shown in figure 1.

Here the o_i partitions refer to the interior of the domains and the t partition is the common boundary shared by the domains. There are automatic tools available to execute this task.[1] The automatic tools attempt to create components with similar interior size and minimal boundaries for computational efficiency. The geometric domain decomposition is important to reduce the very large problem sizes, but it does not affect the frequency spectrum. The geometric domain decomposition is mathematically represented with the following partitioning:

$$\begin{bmatrix} K_{oo}^1 - \lambda M_{oo}^1 & & K_{ot}^1 - \lambda M_{ot}^1 \\ & K_{oo}^i - \lambda M_{oo}^i & K_{ot}^i - \lambda M_{ot}^i \\ K_{ot}^{T,1} - \lambda M_{ot}^{T,1} & K_{ot}^{T,i} - \lambda M_{ot}^{T,i} & K_{tt} - \lambda M_{tt} \end{bmatrix} \begin{bmatrix} \phi_o^1 \\ \phi_o^i \\ \phi_t \end{bmatrix} = 0$$

This partitioned eigenvalue problem may be solved by a special formulation of the Lanczos method.[2] Note that the formulation allows for a computationally exact solution. The efficiency highly depends on the size of the boundary in relation to that of the interiors.

This technology is most efficiently executed on shared memory parallel (SMP) computer architectures. These computers have a common memory shared by the multiple processors, accessed via a memory bus. The individual processors may also have a cache memory distinct to them, accelerating some operations when the data could be kept resident in the local cache.

The main reason for geometry domain partitioning being applicable to this architecture is the communication cost related to the boundary between the geometric partitions. If that data component is kept on the shared memory accessible to all processors, the communication cost is reduced.

Note that this technology is also applicable to other types of analyses – linear statics, for example.

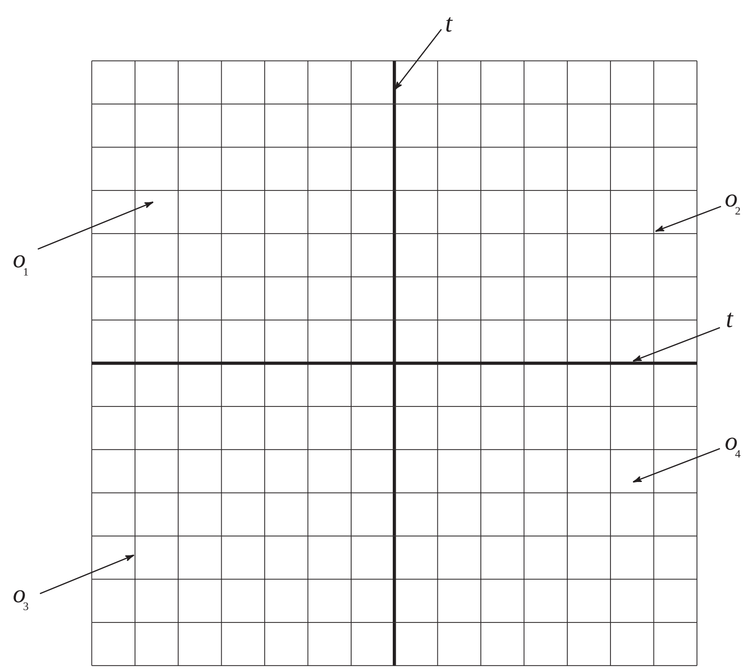

Figure 1: Geometry domain partitioning of finite element model.

Frequency domain decomposition technique

The frequency domain decomposition technique is a systematic introduction of a series of intermediate frequency boundaries in the user-given frequency range specified by $F_{\min}$; F_{max}. The modal density distribution is estimated and taken into consideration to produce frequency segments containing approximately the same number of modes for computational balance. Such decomposition of the frequency domain is ineffective with respect to the large problem size, but reduces the frequency range for a process. The frequency domain decomposition is represented by:

$$\begin{bmatrix} \textit{Segment} & \textit{Lower frequency} & \textit{Upper frequency} \\ 1 & F_{min} = f_o & f_1 \\ 2 & f_1 & f_2 \\ \dots & \dots & \dots \\ j & f_{j-1} & f_j \\ \dots & \dots & \dots \\ s & f_{s-1} & f_s = F_{max} \end{bmatrix}$$

In commercial finite element software tools, such as Nastran, the user also has the possibility of directly inputting the intermediate frequencies on data cards. That is especially advantageous in repeated analysis, which is commonplace in the industry.

To automatically subdivide the frequency range, a heuristic formula developed for Nastran is shown below. In this equation v_i is the eigenvalue corresponding to the f_i frequency as $f_i = \sqrt{v_i}/(2\pi)$ and F_{max}; F_{min} are similarly related to V_{max}; V_{min}. This is due to the fact that the equal distribution is really important in the eigenvalue spectrum, where the computations occur.

$$v_i = V_1 + (V_2 - V_1)\,\frac{1-\alpha^i}{1-\alpha^s};\ \alpha \neq 1.0;\ i = 1,..,s - 1$$

An α value chosen in the range of 1:25 to 1:5 provides a good distribution.

This technology clearly fits well the distributed memory parallel (DMP) computer architectures. In such architectures, the individual processors do not share a memory. The memory is distributed directly to the processors. Note that in this scheme the computation of the mode shapes in the individual frequency segments is independent of each other. Therefore, they may be executed in distributed fashion on each processor. The only communication needed is when gathering up the results for the master processor.

Note that this technology is also efficient in other types of analyses, such as frequency response analysis, for example. Naturally, this method also produces computationally exact eigenvalue solutions.

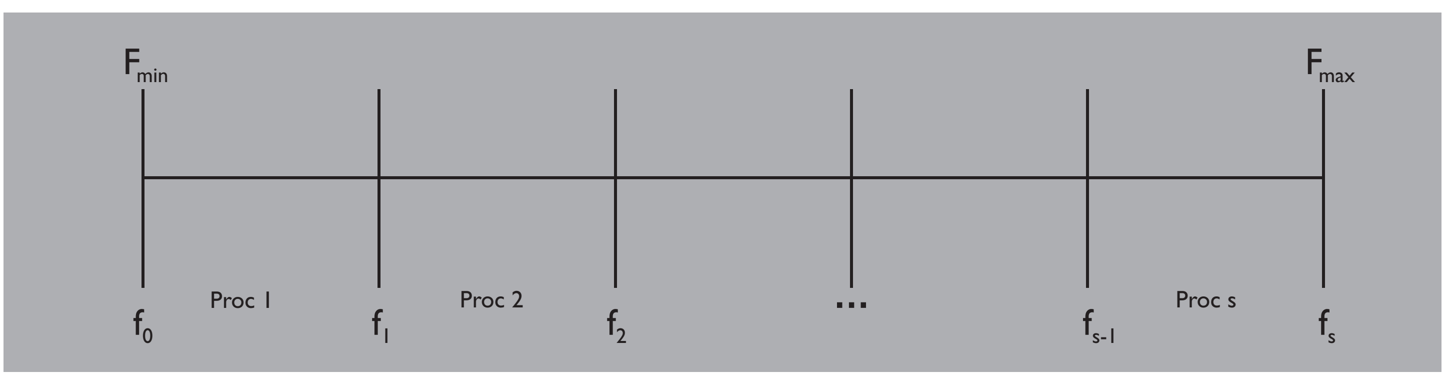

Figure 2: Frequency domain decomposition of frequency range.

The hierarchic computation scheme executes the two paradigms simultaneously.

It is important to emphasize that this version still solves the eigenvalue problem computationally – exactly as the two above methods. The significant contribution is in doing that while both the global matrices are geometrically partitioned and the frequency range is subdivided.

This computation is based on a very specific computational scheme and processor assignment shown in Table 1, where we assumed to have $p * s$ processors available. The $((j-1) * p + i)$-th processor computes the eigenvalues of the j-th frequency segment (Λ^j) and the i-th geometric partition of the corresponding eigenvectors ($\Phi^{i,j}$).

	Partition^1	Partition^i	Partition^p
Processor	1	i	p
Segment^1	$\Phi^{1,1}, \Lambda^1$	$\Phi^{i,1}, \Lambda^1$	$\Phi^{p,1}, \Lambda^1$
Processor	$(j\text{-}1) * p+1$	$(j\text{-}1) * p+i$	$j * p$
Segment^j	$\Phi^{1,j}, \Lambda^j$	$\Phi^{i,j}, \Lambda^j$	$\Phi^{p,j}, \Lambda^j$
Processor	$(s\text{-}1) * p+1$	$(s\text{-}1) * p+i$	$s * p$
Segment^s	$\Phi^{1,s}, \Lambda^s$	$\Phi^{i,s}, \Lambda^s$	$\Phi^{p,s}, \Lambda^s$

Table 1: Hierarchic domain decomposition concept.

The preferred hardware environment for this scenario is a cluster of multiprocessor workstations. Such workstations are usually tied together with either a hardware switch or an Ethernet network. The favored arrangement is to retain the geometric domain component of the hierarchic approach on the workstations and spread the frequency domain segments across the network.

Assuming s workstations of p processors each, the scheme of Table 1 should be executed with careful designation. The tasks of each row (the geometric partitions of a particular frequency segment) are residing on one workstation. The tasks of each column (the various frequency segments of a particular geometry partition) are spread across the cluster.

The next section will show that the hierarchic method scales well up to 64 processors and is also applicable to unequal number of subdivisions in the geometry and frequency domains.

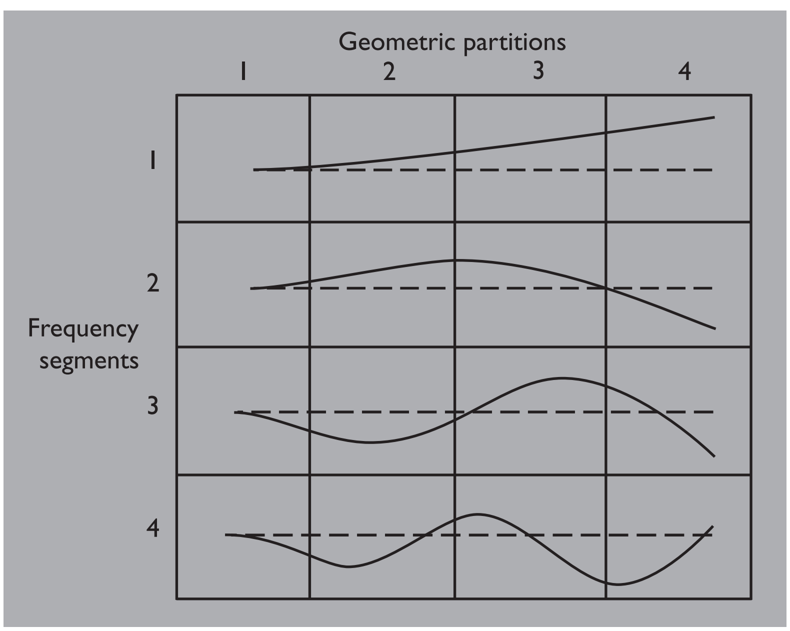

Figure 3: Example demonstrating the hierarchic domain decomposition concept.

To analyze the effect of the above technologies, a trimmed automobile model car body will be considered.[3] Such models have all major components of the car, such as wheels, engine, etc., incorporated. The statistics of Table 2 show the characteristics of such a finite element model.

Model data	Number of nodes	Number of shells	Number of solids	Number of rigids
	380,007	**361,249**	**3,762**	**9,056**
Sizes	g 2,280,042	n 2,223,139	f 2,223,109	a 1,937,282

Table 2: Model statistics of trimmed car body model.

The various sizes refer to the stages of finite element processing. Specifically, the g-size means all degrees of freedom in the system; the n-size is the dimension after all the constraints and rigid elements were processed. After applying the boundary conditions the size reduces to f, while the final analysis problem is size a with all singularities removed. As it is shown, such models usually contain a variety of elements and a large amount of constraints and rigid components.

Number of partitions	Maximum interior	Minimum interior	Maximum boundary	Minimum boundary
2	197,460	182,199	354	354
4	104,788	90,791	572	418
8	54,496	40,067	518	450

Table 3: Effect of geometry domain decomposition.

Table 3 demonstrates the effect of the automated geometry domain decomposition on this model by showing the number of interior and boundary nodes of the partitions. Several observations may be made. The range of the interior sizes depends on the quality of the automated partitioning. Naturally, the boundary size is increasing with larger number of partitions. Finally, the boundary size is at least two orders of magnitude smaller than the interiors, which is important to the computational efficiency.

The automated frequency domain decomposition results are shown in Table 4. In the frequency range of interest there were 840 modes of the model.

Number of segments	Minimum modes	Maximum modes
2	380	460
4	193	252
6	133	165
7	105	141
8	92	119

Table 4: Frequency domain decomposition statistics.

Note that the automated geometry domain partitioning techniques usually provide only even, and preferably binary, numbered domains. This is due to the fact that these techniques are mostly based on binary graph partitioning. This is not a restriction as the shared memory workstations tend to have even number of processors. On the other hand, in the frequency domain decomposition, odd numbers of segments are also allowed. This technology is insensitive to that issue and enables the use of odd numbered workstations via the hierarchic technology in workstation cluster environments.

The execution times of the case study example with the hierarchic domain technology are shown in Table 5. The analysis was executed on a cluster of eight workstations, each containing eight processors with 1.5 GHz clock cycle. The cluster had a 1-GB Ethernet network connection.

Number of processors	Number of partitions	Number of segments	I/O GByte	Elapsed min:sec	Elapsed speedup
1	1	1	1,028.4	523:58	1.00
4	2	2	431.4	167:15	3.13
8	4	2	266.9	83:41	6.26
16	8	2	191.1	45:16	11.57
32	8	4	98.3	34:07	15.35
48	8	6	77.8	27:14	19.23
56	8	7	67.1	24:41	21.22
64	8	8	61.4	27:00	19.40

Table 5: Execution times on workstation cluster.

The task of finding the natural frequencies and corresponding mode shapes of such a model is truly an enormous one. It is an overnight job with more than a tera-byte of I/O operations. The execution on a single processor is impractical considering the work environment and time schedule at automobile companies.

The elapsed time utilizing 32 processors is already a practical execution. The efficiency above that decreases, but the speedup is still increasing. It peaks at 56 processors, although a wider frequency range for this model may extend that peak to 64 or slightly beyond.

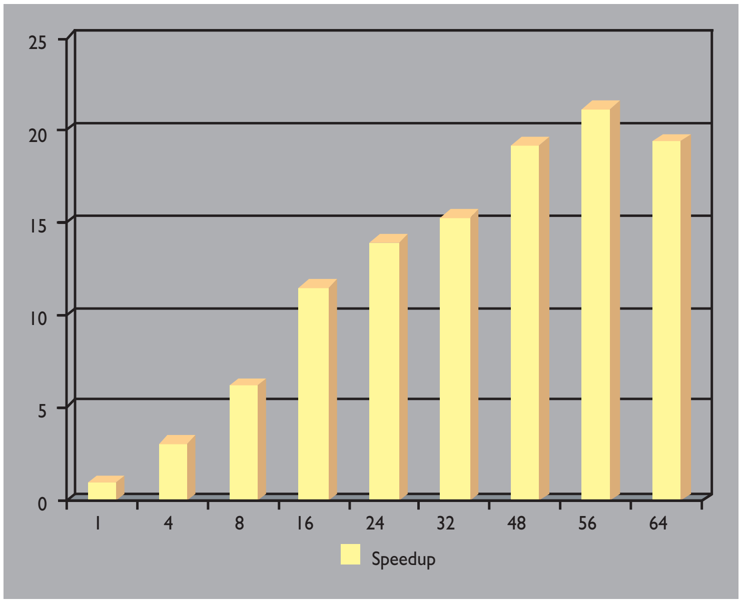

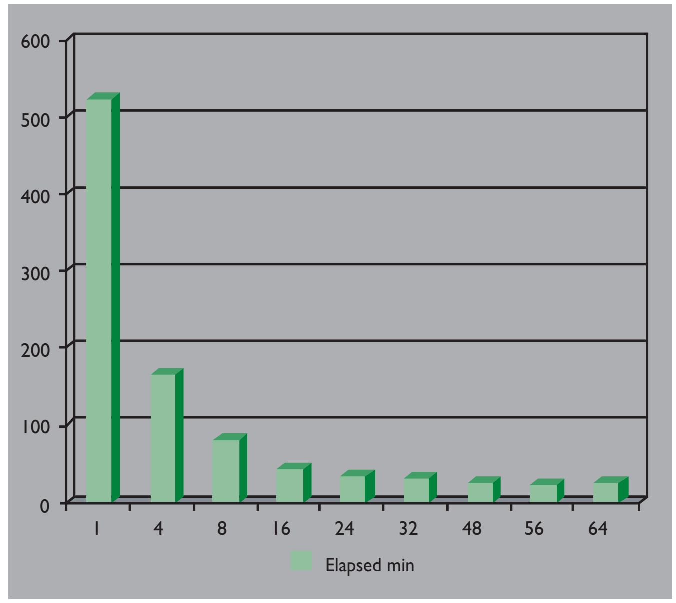

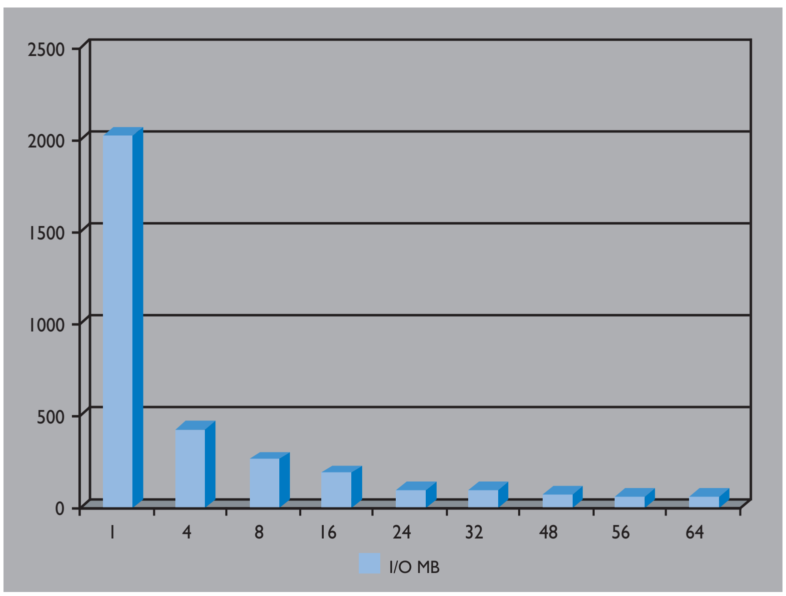

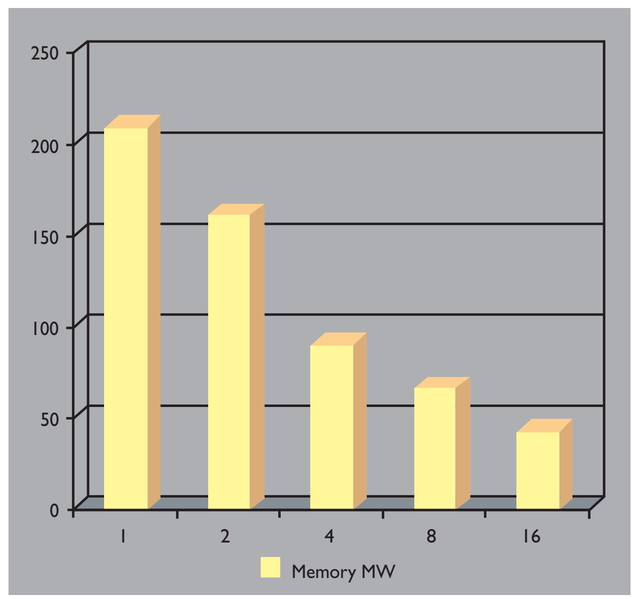

Figures 4, 5, 6 and 7: The speedup, elapsed time, disk use and memory requirements as functions of the number of processors indicated in the case study.

Conclusions and future directions

The new computational methods mentioned above are available in NX™ Nastran software, a new generation of finite element analysis tool aimed at solving a variety of product lifecycle management (PLM) analysis problems. The methods described have proven to be successful for a wide range of our customers in solving the global eigenvalue problem in the computational accuracy afforded by computer. These methods enable engineers to achieve superior performance on all major current industrial hardware utilizing up to 64 processors.

There is, however, an industrial desire to scale these computations to several hundred processors. The desire is fueled by the successful execution of explicit dynamic (crash) computations on hundreds of processors. The explicit technology is, of course, intrinsically better suited for that. Trying to address this need for implicit finite element analyses is much more challenging.

To utilize a much larger number of processors efficiently, the computational work related to the boundaries must be minimized. This is technologically possible by utilizing a multilevel geometric partitioning in connection with the classical component mode reduction concept.[4]

Of course, such method will only produce approximate global mode shapes and is prone to miss some natural frequencies, especially at the higher end of the frequency range. This is why such methods are not universally accepted yet, and the approach may always remain unacceptable in certain industries (such as aerospace).

The multilevel hierarchic computational technique may extend the scalability to several hundred processors. As the industry trend is likely toward even larger problems and wider frequency ranges in the future, such an approach is the subject of current development work.

References

[1] Karypis, L. and Kumar, V. – A fast and high quality multilevel scheme for partitioning irregular graphs, Technical Report TR 95-035, University of Minnesota, 1995

[2] Komzsik, L. – The Lanczos method: evolution and application, ISBN 0-89871-537-7, SIAM, 2003

[3] Komzsik, L. – Computational techniques of finite element analysis, Taylor and Francis, A CRC Press Book, 2005

[4] Hurty, W. C. – Dynamic analysis of structural systems using component modes, AIAA Journal, Vol. 3, No. 4, pp. 678-685, 1965

Siemens PLM Software

Siemens PLM Software, a business unit of the Siemens Industry Automation Division, is a leading global provider of product lifecycle management (PLM) software and services with 6.7 million licensed seats and more than 63,000 customers worldwide. Headquartered in Plano, Texas, Siemens PLM Software works collaboratively with companies to deliver open solutions that help them turn more ideas into successful products. For more information on Siemens PLM Software products and services, visit *www.siemens.com/plm*.

Siemens PLM Software

Headquarters
Granite Park One
5800 Granite Parkway
Suite 600
Plano, TX 75024
USA
972 987 3000
Fax 972 987 3398

Americas
Granite Park One
5800 Granite Parkway
Suite 600
Plano, TX 75024
USA
800 498 5351
Fax 972 987 3398

Europe
3 Knoll Road
Camberley
Surrey GU15 3SY
United Kingdom
44 (0) 1276 702000
Fax 44 (0) 1276 702130

Asia-Pacific
Suites 6804-8, 68/F
Central Plaza
18 Harbour Road
WanChai
Hong Kong
852 2230 3333
Fax 852 2230 3210

www.siemens.com/plm

W7 5647 5/10 C