NX Nastran performance

SIEMENS

White Paper

What Siemens PLM Software is doing to enhance NX Nastran

Siemens PLM Software's product development group enhances NX™ Nastran® software in three key areas: by providing discipline extensions, ease-of-use and process improvements, and better NX Nastran's performance. While each of these enhancements provides value in its own way, this white paper focuses on performance improvement – which makes NX Nastran the solution of choice for users who need to solve large problems.

Contents

Executive summary

Siemens PLM Software's product development group focuses its NX Nastran enhancements along three primary axes:

• By providing discipline extensions

• By improving ease of use and the solution's processes

• By delivering better performance

Discipline extensions

Siemens PLM Software's discipline extensions expand the reach and the different types of physics problems that Nastran can solve. For example, advanced nonlinear capabilities allow users to study nonlinear contact problems, to define complex materials like hyper-elastic materials that are used for door seals or similar applications, and even to solve very high speed physics applications where the time duration is short but the nonlinear effects are very large. Other examples include rotor dynamics, as well as integrated multi-disciplinary solutions including motion, coupled thermal and computational fluid dynamics (CFD) capabilities.

Ease-of-use and process improvement

Ease of use and process improvements make NX Nastran users more productive and include core NX Nastran enhancements, as well as better integration capabilities for the NX CAE pre-processing and post-processing environments. A good example is the way NX Nastran makes it easier to connect meshes. NX Nastran has a process called gluing that enables users to stick multiple components together into a finite element assembly. This process supports a range of connections, including linear contact found in a press fit where two items are tightly connected. New automatic penalty factor calculations that were introduced into NX Nastran's linear contact algorithms are looking very good today. Fewer contact iterations mean less run time. Improvements of 8:1 and higher have been realized in customer test cases.

Other process improvements reflect NX Nastran's ability to deal with interior acoustic problems, where a new approach to calculating the acoustic structural coupling matrices has resulted in speed improvements in excess of 100:1 over the original method. In addition, a new virtual mass solution method dramatically improves the performance for medium to large models that vibrate in fluid like fuel tanks and ships (see the accompanying figure).
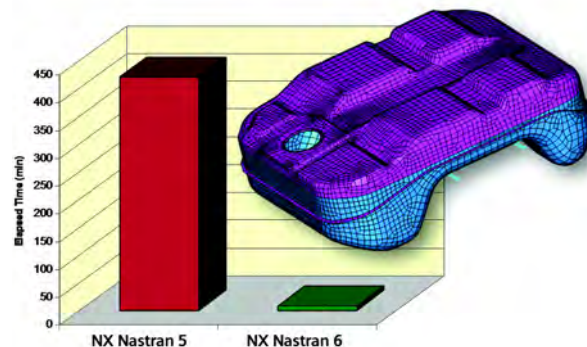
NX Nastran's virtual mass solution method.

Performance improvement

While NX Nastran's discipline extensions and ease-of-use/process improvements are highly valuable, performance improvement is at the heart of Siemens PLM Software's enhancements. Performance is what makes NX Nastran the solution of choice for users who need to solve today's increasingly large problems. Siemens PLM Software has worked really hard in key areas, starting with improving NX Nastran's performance of large scale dynamics problems that can be divided into several pieces and sent to multiple processors. NX Nastran now uses the distributed memory parallel process (DMP) to support this kind of problem solving in a very unique way.

Discipline extensions
Additional physics solutions Common modeling format
Ease-of use/process improvements
Easier modeling Greater accuracy
Performance improvements
Growth in model size State-of-the-art

The rest of this white paper discusses NX Nastran's performance enhancements in detail.

Perfomance improvements

NX Nastran users can divide problems up and almost linearly improve performance by leveraging multiple processors that are available in a DMP cluster environment. NX Nastran also has added extensions to memory addressability that enables users to solve huge problem sizes as evidenced by several benchmarks with 500M degrees of freedom and above. NX Nastran solves very large dynamics problems very quickly.

Processing large matrices

A major improvement relates to the way NX Nastran processes very large matrices, adding things like iterative solvers and related procedures to improve the time it takes to decompose these matrices and solve related problems when necessary. Siemens PLM Software has implemented an iterative solver in NX Nastran that enables highly complex geometry – typically castings and other types of structural applications – to be decomposed quickly, allowing users to interrogate results rapidly.

The iterative solver was an early project. Siemens PLM Software's development team updated performance and achieved a 7:1 performance advantage over the original capability. When combined with memory management enhancements, this opens up whole new opportunities to solve large problems that were impossible with competitive solvers of the time.

Leveraging new computer system architectures

After that, the challenge became how to enable NX Nastran to most effectively leverage today's new computer system architectures. In NX Nastran, shared memory processing (SMP) is used only for lower level operations, such as matrix decomposition and matrix multiplication for all solution sequences. Since every solution sequence involves matrix multiplications, SMP can be activated in all solution sequences for all analyses as long as the hardware supports SMP. In these cases, users should expect a significantly faster solution.

Using multiple processors in parallel can significantly reduce solution run times compared to more traditional serial solutions that use one processor. In addition, much higher levels of scalability can be obtained with distributed memory processing (DMP) compared to shared memory processing (SMP) – which is the preferred technology on multiprocessor nodes with shared memory or with processor nodes with multiple cores.

DMP solutions use a cluster with multiple processors and multiple I/O channels communicating over a network. The figure shown below represents a typical architecture of a hardware system that runs DMP. Each processor has its own memory and one or more disks.

DMP solutions are based on domain decomposition of the geometry (frequency) domain or the load domain. DMP methods achieve their solution speed by dividing the finite element model into smaller pieces that can be solved simultaneously. Users only need to specify the number of processors; the partitioning of the solution is done internally.
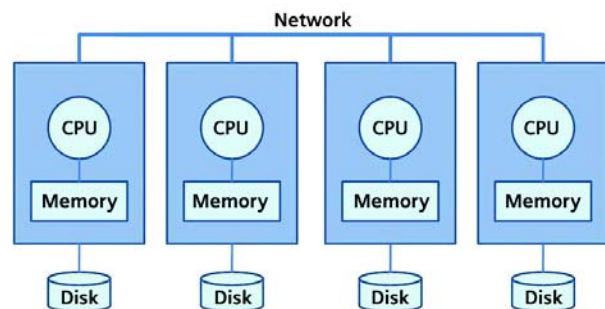
This division is performed with respect to geometry or frequency range individually or both at the same time. Although each processor is working on its own partition of the geometry or frequency range, it communicates with the others to share information. Once the solution is complete, the results are merged, creating a single results file.

NX Nastran has many options for partitioning of the solution domain, which are highlighted in the accompanying table.

Solution	Geometric	Frequency	Hierarchic	Recursive	Load
Static analysis (101)	●				●
Normal modes (103)	●	●	●	●	
Direct frequency (108)		●			
Modal frequency (111)	●	●	●	●	
Modal transient (112)	●	●	●		
Design optimization (200)	●	●	●		

Geometric domain partitioning is available for static and dynamic solutions. The physical model is automatically divided into geometry partitions that are solved on different processors. Conceptually, this is like an automated substructuring (super-element) approach; performance scales well with the number of processors available. This approach significantly reduces the amount of disk space and memory needed to solve problems, yielding a computationally exact solution for eigensolutions.

Frequency domain partitioning is available for dynamic solutions. The frequency range of interest is automatically partitioned into frequency range segments that are solved separately. Each processor solves the full model within its frequency segment. The only communication needed is when gathering the results for the master processor.

Hierarchic domain partitioning combines the geometric and frequency domain methods. This approach is used for modal solutions; it allows scalability to higher levels than could be obtained with either method individually. With this approach, a subset of processors or a cluster solves the eigenvalue problem for the local geometry while communicating with other subsets of processors or other clusters in order to consider the other fre-quency ranges. The preferred hardware environment for hierarchic domain is a cluster of multiprocessor workstations usually tied together by either a hardware switch or network.

Load domain partitioning is useful when there are a large number of load cases in a linear static analysis problem. Instead of partitioning the finite element model, the load matrix is partitioned among the processors as evenly as possible, and the linear solution is calculated within each of the respective processors for its own load cases. Like frequency domain partitioning, load domain partitioning, which does not need communication between processors, is nearly linearly scalable.

Recursive substructuring is the newest DMP method. It works for modal solutions and is the most scalable solution yet. It operates on multilevel partitions of the mass and stiffness matrices by performing automated matrix reduction. As a mathematical reduction method, the resulting solution is approximate when compared to the other DMP methods. However, experience has shown that the approximation is very good and has very little difference with an exact solution.

Good scalability for this method has been achieved out to 128 CPUs as shown in one of the following examples. Using the Lanczos method on a single processor as the baseline, the recursive DMP solution can solve more than 100 times faster. The recursive method also can be run on a single processor and, because of the effect of substructuring, runs 4 times faster than Lanczos.

# Illustrating use of the recursive domain solution

Two examples can be presented to illustrate the use of the recursive domain solution, starting with a normal modes (eigenvalue) analysis of a trimmed car body FE model. This model has around 6.5 million freedoms, including approximately 1.3 million grid points and approximately 1.2 million shell elements. This model was run with sparse eigenvector recovery, which is new in NX Nastran 7.0, turned on to determine scalability in terms of elapsed time vs. number of processors.
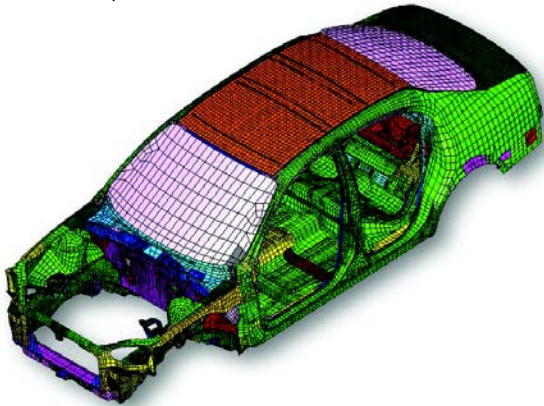
As can be seen in the accompanying chart, solution speed continues to scale well with the number of processors used; linearity of the speedup is very good with scalability appearing to depend on the size of the model and number of processors available.
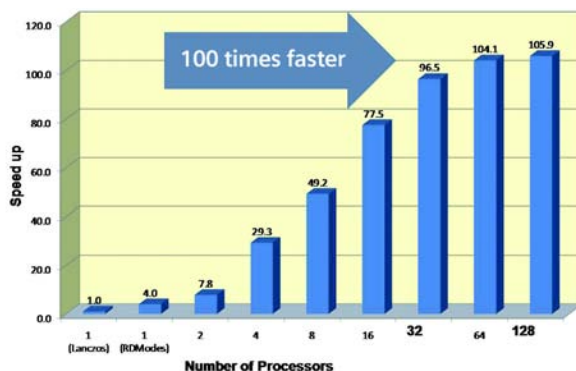
Approximately 1000 modes were computed up to 300 Hz in the initial run. However, the analysis was also used to test for any performance impact related to the number of modes requested. As can be seen in the accompanying charts, the solution appears relatively flat compared to the number of modes computed.
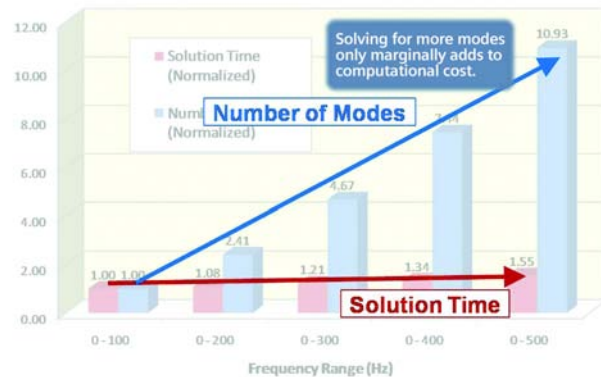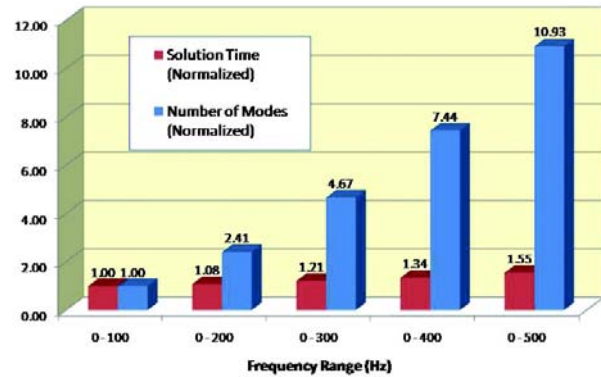

This case also highlights the value of sparse data recovery. This approach was implemented for large problems because, in many instances, a user is only interested in the solution at a few key locations. In cases like this one, the sparse eigenvector recovery method, which is automatically invoked, can significantly reduce the overall computation time and storage resource (5:1 in this case).
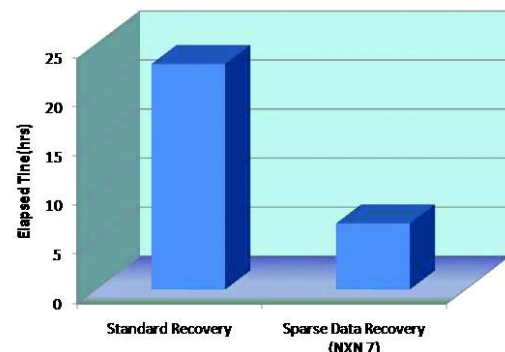
The second example involves a V8 engine model with 3.6 M grid points, 2.3 M parabolic test elements and 10.8 M degrees of freedom.
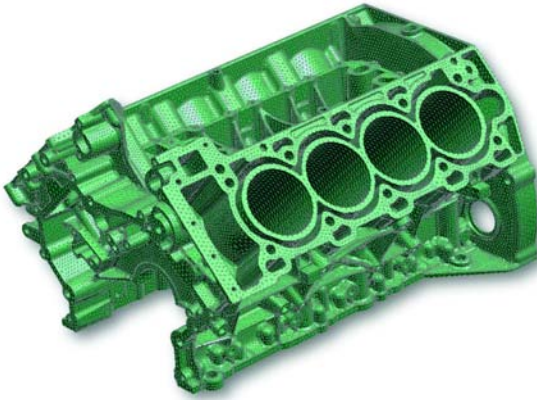
The problem was run on an HP Cluster with 64 CPUs running LinuxOS. As can be seen in the accompanying results, performance again scales well. It is always important to look at models with different types of mesh since the matrices being solved exhibit different forms and this changes the way solver algorithms need to be efficient to achieve targeted performance gains.
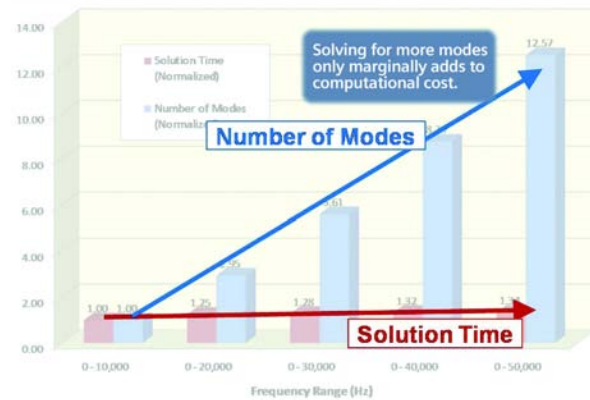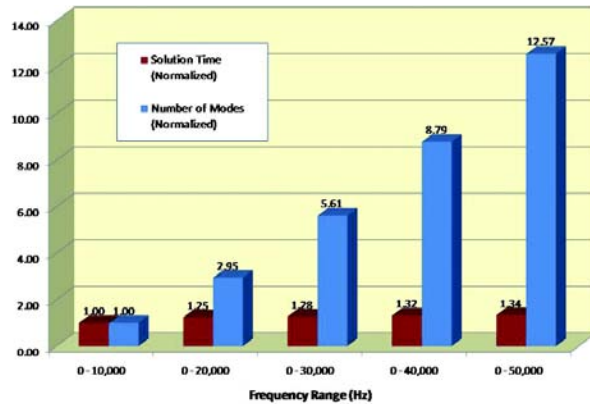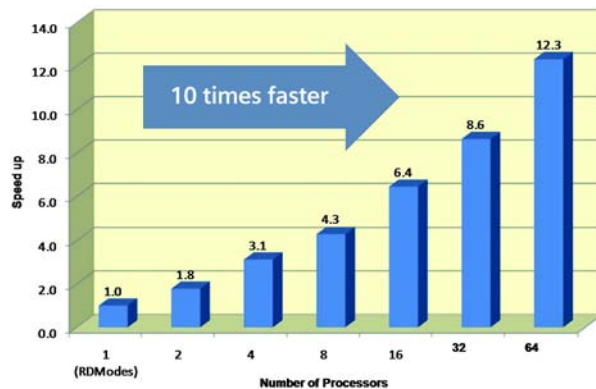
In the above case, approximately 250 modes were computed in the 0 – 10,000 Hz range. Testing the relationship between the number of modes requested and CPU time used again found that solving for more modes only marginally added to computation cost.

Conclusion

Complex models with millions of nodes and tens of millions of degrees of freedom are becoming the norm today. As analysis models continue to grow in terms of size and complexity, solver performance becomes a critical factor in a company's ability to produce results in time to impact key design decisions. Siemens PLM Software understands the importance of solver performance, which forms the heart of its NX Nastran development strategy.

About 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 69,500 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.

www.siemens.com/plm

X5 19887 5/11 C

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