

NX Nastran DMP

Distributed memory parallel processing with NX Nastran

fact sheet

Siemens PLM Software

www.siemens.com/nx

► Summary

NX™ Nastran DMP is an add-on module to the enterprise version of NX Nastran Basic software. The DMP license allows users to run NX Nastran with multiple processors on a cluster hardware system. Using multiple processors in parallel can significantly reduce solution run times compared to the more traditional serial solution that uses one processor. Also 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.

Benefits

Reduced solution time on large problems
 Reduced hardware requirements for individual nodes on a cluster
 Decreased turnaround time for each solution; allowing you to analyze more designs under more conditions in a shorter amount of time

Features

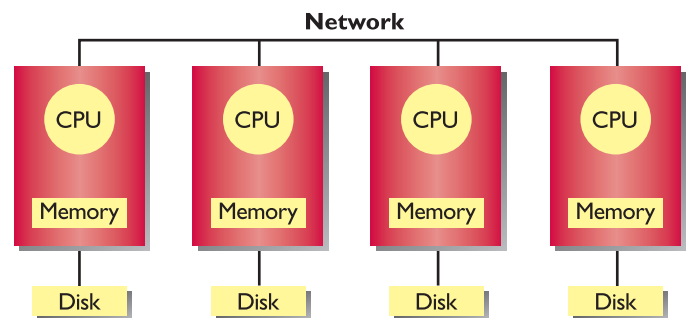
High scalability
 Static, modal, dynamic and optimization solutions
 One run uses one DMP license independent of the number of processors used
 Message passing interface (MPI) libraries are provided

Leveraging multiple processors

A distributed memory parallel (DMP) solution uses a cluster with multiple processors and multiple I/O channels communicating over a network. The figure below shows a typical architecture of a hardware system for running DMP. Each processor has its own memory and one or more disks.

DMP is a high level parallel processing methodology that works by partitioning the solution into domains that are distributed to different processors. The memory used for each domain is independent of the memory used by other domains. In contrast, SMP operates at a lower level

and does not divide the solution into domains but rather uses parallelization for efficient math operations such as matrix factorization and multiplication. The processors work independently but the memory is shared. The independence of the memory usage by processors in a DMP solution allows for much greater leveraging of parallelization thus providing a higher level of scalability. With NX Nastran DMP, scalability levels to 64 processors have been obtained with no loss of performance.



DMP Implementations

DMP solutions involve a physical subdivision of the model solution domain into partitions that are distributed and processed separately. The user only needs to specify the number of processors, and the partitioning of the solution is done internally.

NX Nastran has the following options for partitioning of the solution domain.

- Geometric domain
- Frequency domain
- Hierarchic domain
- Load domain
- Recursive substructuring domain

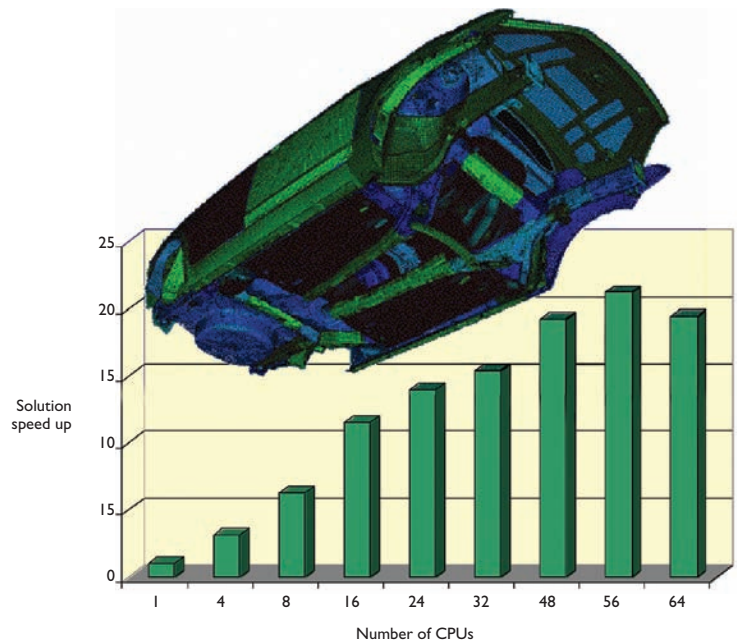
Geometric domain partitioning is available for static and dynamic solutions. The physical model is divided into geometry partitions that are solved on different processors. This table, in which N is the total number of physical DOFs and p is the number of processors, demonstrates the potential advantages for such decomposition.

| | <i>Not partitioned</i> | <i>Partitioned</i> |
|----------------|------------------------|--------------------|
| Problem Size | N | $n=N/p$ |
| Operation Cost | $N^3=p^3n^3$ | pn^3 |
| Remarks | $p^3 > p$ | |

Not included in the table is the effect of the connections between the geometry partitions. These connections require communication through the message passing interface (MPI) between the processors. Therefore the geometry should be partitioned such that there is minimal connectivity between the partitions to minimize the communication overhead. For this reason, NX Nastran DMP has several automated partitioning algorithms for determining an optimal set of geometric domains. In general, partitioning works better on shell dominant models where small domain boundaries can be found compared to "chunky" solid models that result in a large number of DOFs at the domain boundaries.

Frequency domain partitioning is available for dynamic solutions. The solution is partitioned into frequency range segments that are solved separately with no communication needed between the smaller solutions. For example, in an eigenvalue analysis for computing modes to 200 Hz, two domains can be used – one from 0-120 Hz, another from 120-200 Hz. The actual frequency boundaries are automatically computed by NX Nastran based on a modal density estimate – the goal to evenly balance the number of modes solved on each processor.

Hierarchic domain partitioning is a hybrid of geometry and frequency domain methods. It is used for modal solutions and allows for partitioning of geometry and the modal frequency domains at the same time. Using the two methods together allows scalability to higher levels than could be obtained with either method individually. The figure shows the scalability obtained by an hierarchic partition of a car body model out to 64 processors.

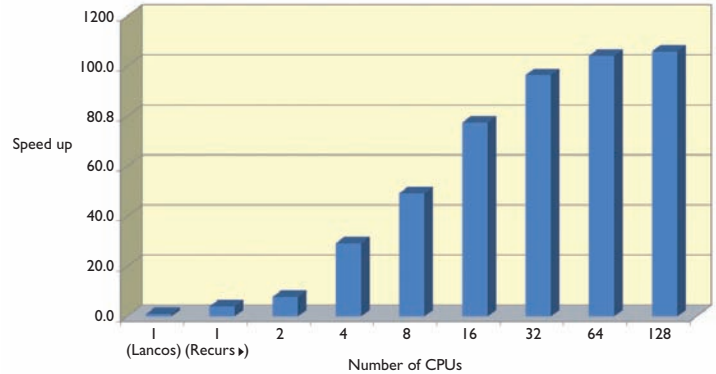


Load domain partitioning is available for use in static analyses and is useful

when there are many load cases (> 100) to solve. For such solutions the forward-backward substitution (FBS), needed for each load set, can take more time to solve than the one time numerically intensive

stiffness matrix decomposition. Partitioning the load domain leverages parallelization for the FBS by distributing load sets to multiple processors. Like frequency domain partitioning, load domain partitioning does not need communication between processors and 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 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 of this method has been achieved out to 128 CPUs for a car body model as shown by the figure. Using the Lanczos method on a single processor as the baseline, the recursive DMP solution can solve more than 100 times faster. It is noted that the recursive method can also be run on a single processor and because substructuring leverage runs 4 times faster than Lanczos.

Solutions and domain options

The table lists the NX Nastran solutions that have enabled DMP processing and the domain decompositions are available for each.

DMP Implementations in solution sequences

| Solution | Geometric | Frequency | Recursive | Hierarchic | Load |
|----------|-----------|-----------|-----------|------------|------|
| 101 | • | | | | • |
| 103 | • | • | • | • | |
| 108 | | • | | | |
| 111 | • | • | • | • | |
| 112 | • | • | | • | |
| 200 | • | • | | • | |

Availability

NX Nastran DMP is contained in the NX Nastran Advanced bundle or can be purchased as an add-on with any of the Nastran solution types outlined in the chart above. NX Nastran DMP is not available with the desktop versions of NX Nastran.

► **Contact**
 Siemens PLM Software
 Americas 800 498 5351
 Europe 44 (0) 1276 702000
 Asia-Pacific 852 2230 3333
www.siemens.com/plm

