Future Techniques for High Frequency NVH

Stan Posey, Automotive Industry Market Development
SGI Mtn View, CA, sposey@sgi.com, 650.933.1689

Cheng Liao, PhD, Principal Engineer, CAE Applications
SGI Mtn View, CA, liao@sgi.com, 650.933.3579

Christian Tanasescu, PhD, Manager, CAE Applications Engineering
SGI Munich, GR, christi@munich.sgi.com, +49.89.46108124

Acknowledgments

Abstract

The use of NVH analysis provides essential benefits towards designing vehicles for ride comfort and quietness, an increasingly competitive advantage in today's global automotive market. Requirements for NVH analysis at increasingly higher excitation frequencies is driving NVH modeling promoters beyond practical limits for conventional NVH methods. This paper examines details behind conventional NVH practice, NVH modeling directions for the future, and an alternative to conventional NVH that will allow future modeling targets to be achieved.

Introduction

Automotive manufacturers undergo increasing market pressures to satisfy consumer demand for vehicles designed with improved ride-comfort and quietness. Good management of a vehicle's NVH characteristics helps auto makers produce a more competitive product, especially in luxury vehicle markets. Sources of NVH are dynamic and acoustic response to typical mechanical loads -- those mainly applied from vehicle interaction with the road and operation of the powertrain. Significant NVH improvements have occurred recently such that vehicle sounds and vibrations previously masked by road noise are becoming a substantial noise source. As such, additional NVH reductions will only occur with design investigation at higher levels of fidelity and precision. Only the NVH laboratory offers such levels today since conventional NVH analysis still has several limitations.

Analysts continue to push NVH modeling to higher excitation frequencies in order to capture an increasingly larger audible range for additional NVH reductions. Subsequently this requires that NVH model parameters grow substantially larger than those in common modeling practice today do. Common global practice for NVH analysis on trimmed body-in-white (BIW) typically has upper bound limits on excitation frequencies to between 250Hz and 300Hz. However many automotive companies have desires to increase this to 600Hz and beyond during the next few years. These predictions for higher frequency modeling in the future has existing issues with conventional NVH methods, regarding both numerical accuracy and suitable job turn-around times.

The automotive industry has historically invested in vector systems to satisfy the high-performance computing (HPC) resource demands of CAE applications, but in particular for NVH analysis using MSC/NASTRAN. It is estimated that NVH analysis consumes from 20% to 25% of all HPC automotive cycles globally, which is second behind crash simulation at 55% to 60%. While crash simulation is the most CPU intensive automotive application (with little demand on other HPC resources), NVH requires high demand of virtually all HPC resources -- CPU, storage, memory bandwidth, and I/O rates on the order of TBs for a single NVH job. What's more, crash simulation benefits from moderate to high parallel scalability, whereas conventional NVH analysis techniques are limited in parallel scalability and usually restricted to a uniprocessor through-put environment.

Lately there has been growing concern over how the industry will address future NVH modeling requirements. This concern comes at a time when several automotive companies are migrating from vector to more cost-effective scalable RISC as their strategic HPC architecture. This shift is the result of new algorithms and methods recently implemented for nearly every automotive HPC application except for NVH. The vector-to-RISC migration began for the automotive industry during 1995 when a direct sparse solver was introduced in MSC/NASTRAN, and quickly became the choice over the expensive skyline solver. The sparse solver reduced CPU and storage requirements by one order of magnitude over the skyline solver and as such, structural analysis (statics) rapidly migrated from vector to RISC.

During early 1996, commercial computational fluid dynamics (CFD) software standardized on "unstructured-mesh" technology owing to its ease in automatic meshing and potential for high levels of parallelism. Later that year highly-scalable domain decomposition parallel techniques were introduced in applications like FLUENT and STAR-CD that could provide linear scaling to 64 processors for some CFD models, and commercial CFD migrated to RISC. Crash simulation applications like LS-DYNA, PAM-CRASH, and RADIOSS have gained benefit during that period from moderate parallel scalability with shared memory parallel implementations, but most recently, domain decomposition and highly-scalable parallel, RISC-based crash simulation has become a production environment for some automotive companies. This vector-to-RISC migration trend for crash simulation is likely to accelerate such that only large NVH remains without a clear vector migration path towards more cost-effective RISC systems.

Today's typical model sizes for current NVH analysis of an trimmed body requires a vector class system for effective job turn-around -- usually defined as overnight. With the current NVH techniques in place, even next generation vector systems will not deliver the performance required for conventional NVH modeling targets of the future. Model parameters will exceed the practical limits of these vector architectures. The conventional eigenvalue extraction and modal response methods (MSC/NASTRAN SOL 103 and 111) are widely considered to be too costly for vector, such that alternative methods are being proposed. This paper examines an alternative NVH analysis technique, direct frequency response

(MSC/NASTRAN SOL 108) as a consideration to achieve future modeling targets.

Conventional NVH

An NVH methodology that emerged as a conventional "industry-standard" was developed as rapid growth of finite element structural analysis began in the mid-1980's. During this period NVH analysis was migrating from component mode synthesis methods to application of MSC/NASTRAN and its efficient Lanczos algorithm on vector systems. The Lanczos algorithm is used to perform an eigenvalue analysis that computes the natural frequencies of a structure over a given frequency range of excitation. Dynamic response of the structure is then determined with a frequency response analysis on the generalized modal coordinates obtained from the eigenvalue analysis. This NVH analysis method of modal frequency response (MSC/NASTRAN SOL 103 and/or 111) is conventional practice for practically every automotive manufacturer and supplier world wide.

Today's current model sizes of a automotive body, in the range of 1.5M DOF with 1250 modes is most effectively executed on a vector class of system for overnight turn-around. For the current release of MSC/NASTRAN, release 70.5, parallel scaling is of little to no advantage and rarely used, such that NVH must rely on a fast single processor. Note that this applies to vehicle body NVH only, since this type of NVH analysis is dominated by time spent in the MSC/NASTRAN Lanczos eigensolver. Components and powertrain NVH are dominated by sparse solver decomposition and are well suited to RISC systems for models of even 3M DOF or more. A typical powertrain is not as flexible as a vehicle body and usually requires fewer than 150 modes for the largest of solutions. This compares with body NVH where models today average 1250 modes. Model sizes of body NVH at the high-end are approaching 3M to 5M DOF and more than 2500 modes.

An examination of the various solution paths of MSC/NASTRAN help to explain the demands required of a particular hardware architecture feature. Generally speaking, finite element sofware exhibits a range of compute behavior depending upon the kind of analysis being conducted and model size, such that a balanced hardware architecture is desired. For NVH modeling, parameters such as the size of the model, the type of geometry, the types of elements, and the excitation frequency of interest, all affect the MSC/NASTRAN execution behavior. A profile is provided in Table 1. that describes the behavior for certain MSC/NASTRAN tasks associated with NVH analysis.

Table 1. Compute Profiles for MSC/NASTRAN and NVH Analysis

Compute Task	Memory Cycles	CPU Cycles
Sparse Direct Solver	7%	93%
Lanczos Solver	60%	40%
Iterative Solver	83%	17%
I/O Activity	100%	0%

This profile highlights the importance of a balanced system since the sparse direct solver requires a fast processor for effective execution while the Lanczos solver requires high memory bandwidth speeds. The argument for memory bandwidth is even greater when considering I/O and use of the iterative solver rather than the sparse direct for matrix decomposition. The I/O requirement in particular is very critical to good elapsed time turn-around since Lanczos is highly dependent upon large amounts of I/O for models with a large modal density, such as those typical of body NVH modeling. vector architectures offer much higher rates of memory bandwidth than RISC systems, which is why they are the favored architecture for conventional NVH. For example, a typical SOL 103 body model that exceeds 1.5M DOF and 1000 modes executed on a single processor Cray T90 is roughly 5-fold faster in elapsed time than a single processor SGI Origin2000/250Mhz.

Still, the Origin2000 RISC architecture offers several advantages to conventional NVH. Any MSC/NASTRAN NVH job other than the large body models will turn-around in a matter of a few hours. It has also been observed that for many body NVH models in the range of up to 750K DOF and 1000 modes, half-day turn-around is consistently achieved. The SGI Origin2000 is a breakthrough implementation of the shared memory ccNUMA architecture. The motivation and direction towards ccNUMA evolved at SGI as traditional shared bus architectures like that of the CHALLENGE server began to exhibit high latency bottlenecks as processor counts were growing within a single system image. During this same time, non-coherent distributed memory architectures started to emerge, but the programming of applications for message passing in such an environment was considered too difficult for commercial success.

The Origin ccNUMA architecture exploited the latest design trend known as distributed shared memory: a cahce coherent but physically distributed memory parallel system that appears logically as a shared memory parallel system to the user. This offers the best features of popular contemporary architectures, meaning that by having memory distributed to individual processors, latencies that inhibit high bandwidth and scalability are greatly reduced. At the same time, the ability to globally address all distributed memory as a singular memory resource, simplifies the programming task substantially. Origin combines high-performance with ease of use in a single environment. Behind the high-performance is the Origin's unique directory based cache coherence and non-blocking interconnect design that delivers low latency, high bandwidth, and the highest SMP scalability in the industry with up to 256 processors in a single system image. This novel system design, coupled with extremely high IO bandwidth and expandability, makes Origin an exceptionally good throughput system for executing several CAE large-user applications, such as crash simulation and NVH concurrently.

Future Techniques

One well documented trend for the automotive industry is the wide spread application of implicit finite element analysis (FEA) for improved structural response and vehicle weight reductions. An application that just four years ago required vector HPC resources for any reasonably sized model, mainstream FEA is performed on desktop computer systems today. The advancement of RISC performance played a key role in this trend, but the most significant contribution came in the way of new software algorithm technology. During recent years, direct sparse solvers have been implemented in practically every commercial FEA package, and provide on average, a 10-fold performance improvement over the previous generation of profile solvers they replaced. Sparse solvers also greatly reduce storage requirements by roughly an order of magnitude, which also contributes to the ability of conducting meaningful FEA modeling on desktop systems. This breakthrough in sparse solvers fueled the rapid growth of commercial FEA software as an important tool in mechanical design for automotive and other manufacturing industries.

It is clear from an automotive industry perspective that a similar breakthrough is needed for NVH analysis to reach the same level of design benefit and pervasiveness. The ability to increase model sizes for high frequency resolution, while maintaining adequate solution turn-around that fits within design cycle times offers great commercial advantages to automotive manufacturers. Possible software breakthroughs include Lanczos algorithm performance improvements and parallelisation, or alternatives to Lanczos. Any alternative to Lanczos must provide a significant performance improvement with a minimum of equivalent numerical accuracy. During development of the next MSC/NASTRAN release, 70.7, MSC and SGI investigated all of these possibilities. Performance improvement to Lanczos has historically been an ongoing project between MSC and SGI, and results have shown substantial performance increases year over year. New for MSC/NASTRAN 70.7 is a distributed memory parallel capability implemented by MSC that shows good parallel scaling for conventional NVH modeling with its parallel Lanczos scheme. An example of the parallel scaling that is possible with SOL 103 is provided in Table 2. for the model *xxcmd* executed on an Origin2000/300Mhz. The model *xxcmd* is a BIW with 1.5M DOF and 1076 modes, with an upper bound on the excitation frequency of 200Hz. For this class of NVH model, Origin2000 shows suitable turn-around with as little as 2 processors at roughly 10 hours elapsed time.

Table 2. MSC/NASTRAN 70.7 SOL 103 MPI Parallel for *xxcmd*

Processors	Elapsed Seconds	Parallel Speed-up
1	61,595	1.0
2	36,734	1.7
4	24,501	2.5
8	19,474	3.2
16	14,660	4.2

Lately two potential alternatives to Lanczos have been identified for future modeling requirements. One is a shift from modal to direct frequency response, and the other is a new algorithm called Automated Multi-Level Substructuring [1] -- both of which are based upon MSC/NASTRAN. For MSC/NASTRAN 70.7, direct frequency response provides a potential for greatly improved turn-around times over conventional modal response NVH, owing to the highly parallel nature of the algorithm. Parallelization is implemented for the independent frequency steps which each perform the same amount of work, and therefore provides good load balance.

Results for direct frequency response are given for two vehicle bodies. The first demonstrates the potential for high parallel scalability, in this example up to 32 processors on an SGI Origin2000. Table 3. shows results for the model xlifr -- a vehicle body with 536K DOF with 96 frequency steps. It is observed that parallel efficiency is very high even to 32 processors.

The second example compares a vehicle body using direct frequency response on an SGI Origin2000, to conventional modal response on a Cray T90. The model contains 525K DOF and 2714 modes for an analysis of 96 frequency steps. Table 4. shows results that demonstrate roughly equal performance between SOL 111 on a single processor Cray T90 with a four processor SOL 108 on an Origin2000/300Mhz. This model contains a higher modal density than what is typically analyzed today, and is perhaps more representative of future modeling practice. These results are very encouraging based on both performance, and price-performance – and for a method that provides improved solution accuracy.

Table 3. MSC/NASTRAN 70.7 SOL 108 MPI Parallel for *xlifr*

Processors	Elapsed Seconds	Parallel Speed-up
1	114,264	1.0
2	57,282	2.0
4	28,887	4.0
8	14,883	7.8
16	8,070	14.2
32	5,047	22.6

Table 4. Origin2000/SOL 108 Comparison with Cray T90/SOL 111

Processors	Elapsed Seconds	Parallel Speed-up
1	120,720	1.0

2	61,680	2.0
4	32,160	3.8
8	17,387	6.9

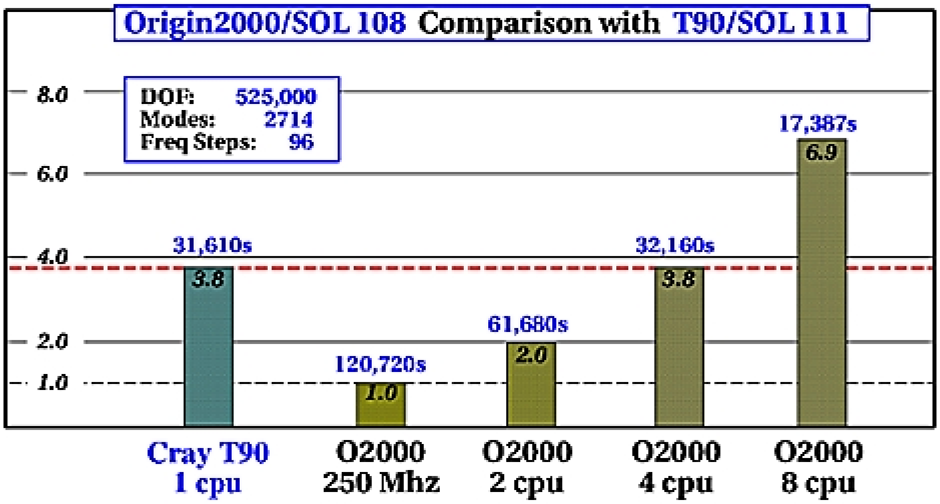

Conclusions

Future modeling requirements for trimmed body NVH are not practical with the conventional eigenvalue extraction and modal response method currently in practice. It has been demonstrated that a highly parallel direct frequency response method is a viable alternative that offers practical job turn-around time, and with equivalent-or-better numerical accuracy. Use of the direct frequency response method also offers the automotive industry a migration path from vector to more cost-effective RISC since it is highly parallel. The low cost of RISC computing will also enable rapid growth of design optimization and even makes multi-discipline optimization within practical reach. The significance of a highly scalable NVH solution like direct frequency response is that engineers will be able to model at increasingly critical frequency levels within a wider hearing range of occupants.

Modeling at increasingly higher frequency levels for acoustic response is conducted today in the automotive industry, but new capability with highly scalable direct frequency response will eventually enable aerospace, turbomachinery and other industries to consider acoustic improvements to their designs. This has been impractical in the past for aircraft design, owing to the excitation frequencies of interest combined with the large geometric scale of the models required. The authors believe that this highly scalable, cost-effective NVH methodology using direct frequency response has the potential to shift modeling practices for manufacturing on an industry-wide and global basis.

Acknowledgments

The authors would like to thank Mr. Jeff Konz, SGI Detroit, Mr. Takahiko Tomuro, SGI Japan, and Mr. Paul Conti, MSC.Software Detroit, for their contributions to this paper.

References

1. Bennighof, J. K. and Kaplan, M. F., "Frequency Window Implementation of Adaptive Multi-Level Substructuring," *Journal of Vibration and Acoustics*, Vol. 120, No. 2, pp. 409-418, April 1998.

2. Duncan, A., Su, F., and Wolf, W., "Understanding NVH Basics." Proceedings of the International Body Engineering Conference. Body Design and Engineering, Vol. 20, pp. 111-116, 1996.

3. Komzsik, L., MSC/NASTRAN Numerical Methods, Version 70, 1998.