Speeding-up acoustic predictions

M Tournour, J-P Rossion, L Bricteux, C McCulloch
LMS International, Leuven, Belgium, Email: michel.tournour@lms.be

Introduction

Optimising a product by 'testing' virtual prototypes requires multiple calculations of functional performance, including vibro-acoustics. During the past two decades, Finite Element Methods (FEM) and Boundary Element Methods (BEM) have been extensively used, made possible by advances in computer performance. Nevertheless, the analyst wants answers about his design within hours, even minutes, to be able to steer the design: calculations only for verification purposes may take days or weeks, but that is not acceptable in virtual prototyping processes. Furthermore, there is a clear need for advanced tools in the mid-frequency range. Unlike some high-frequency methods, the useful frequency ranges for FEM and BEM are mostly not limited theoretically, but by the capabilities of the computer and the solution time.

The technologies presented make accurate acoustic predictions timely and effective: speed-ups can be one or even two orders of magnitude. They tackle a wide range of applications ranging from engine acoustics to interior acoustics and enable users to design practical solutions and effectively reduce time-to-market and development costs.

The commercial package LMS SYSNOISE was used to compute the numerical examples which are shown.

Acoustic Transfer Vectors

ATVs are input-output relations between the normal structural velocity of the radiating surface and the sound pressure level at a specific field point. ATVs only depend on the configuration of the acoustic domain, i.e. geometry and properties (sound velocity and mass density), the acoustic surface treatment (local impedance), the frequency and the field point location. They do not depend on the loading. The calculation effort for a single ATV is about the same as for a single load case response calculation.

ATV for multi-load-case forced response

Because the calculation cost is relatively low, and ATVs are (by definition) independent of the acoustic loading conditions, ATVs can be used efficiently in multi-load-case acoustic response analyses. There are also interesting extensions to Panel Acoustic Contribution Analysis and Inverse Acoustic Numerical Analysis. A very-useful data reduction can be had from using Modal ATVs, the modal counterpart of the ATVs, expressing the acoustic transfer from the radiating structure to a field point in modal coordinates. A further speed-up can be had from an interpolation scheme for the ATV coefficients, using master and slave frequencies, applicable to exterior problems (where strong resonances do not occur). It can be shown that a safe choice of master frequencies is such that $\Delta f < c/4r$, where c is the speed of sound and r is the maximum distance between the mesh nodes and the field points.

ATV Example

An acoustic BEM mesh of 7504 nodes (Figure 1) is used for an engine. Both the radiated acoustic power and the pressure at 19 specific locations were computed (Figure 2).

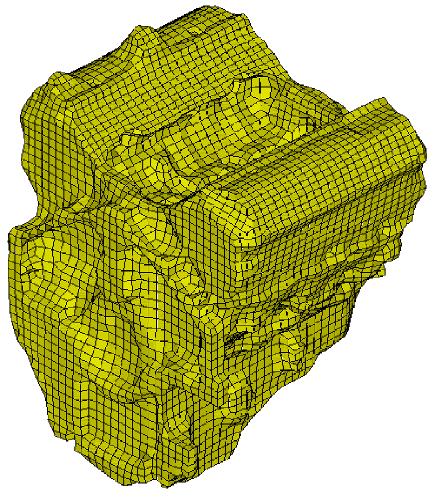

Figure 1: V6 engine mesh

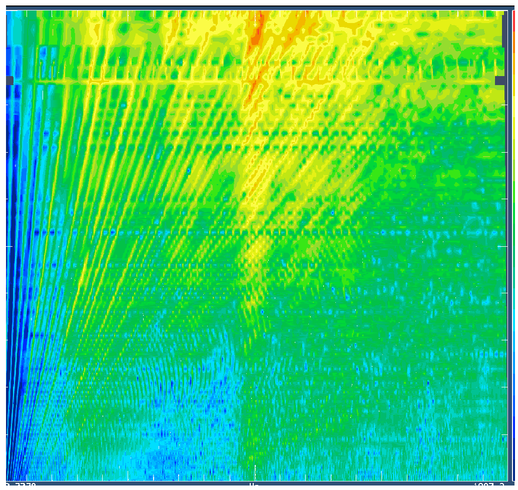

Figure 2: Sound pressure level above engine

The response was computed from 1000 to 6000 RPM, with a step of 50 RPM, and from 0 to 2000Hz with a frequency step ranging from 4.2Hz at 1000 RPM to 25Hz at 6000 RPM (21400 frequency solutions in total). Different frequency steps were used at different RPMs, to identify the orders. The complete run took approximately 13.5 hours. A conventional BEM approach would need about 223 days to perform the same computation - an impossible task!

Padé expansions

The aim of Padé expansion is to solve the Helmholtz integral equation for a complete frequency band, using factorization of the matrix at selected frequencies (called *master* frequencies). For large problems, the computation time is dominated by the factorization time of the matrix. Therefore, avoiding multiple factorizations gives large time savings. This is done by approximating the frequency response function by a Padé Approximation. Using Padé Approximation rather than more-classical Taylor expansions is justified by the fact that the coefficients can have singular points (poles or eigen frequencies) and is usually not holomorphic but only meromorphic, so that its Taylor series does not converge everywhere. The calculation requires the knowledge of the successive derivatives of the acoustic quantities with respect to frequency, evaluated at a 'central' frequency. A recurrence scheme results, having the important property that the calculation of the successive derivatives requires the factorization of a single matrix. In a given frequency range, we compute several matrices for different frequencies. The higher-order derivatives are computed from these different evaluations, and for increased performance, it can be assumed that there is a smooth variation with frequency.

Practical Issues

The practical value of the procedure depends on accuracy, speed (speed-up can be 10 times) and memory and disk requirements (several matrices need to be stored). The approach is limited to free field radiation without resonances. Master frequencies and the order of derivatives are selected automatically based on an accuracy level specified by the user.

Padé expansion example

An acoustic response was calculated for the engine example given before, 80 frequencies from 700Hz to 1500Hz on HP C3600. The computation using a conventional BEM approach took 11 hours 36 minutes, whereas with Padé expansion it took 1 hour 12 minutes: a speed-up of 9.6. A comparison of the pressure above the engine is so close that the frequency-function curves appear to be superimposed.

(The Padé Approximation is implemented in the SYSNOISE solver in partnership with CADOE SA, France.)

Krylov Iterative Solvers (FEM)

For FEM problems, as model size increases the total computation time is dominated by the solution of the linear system. An iterative solver circumvents this and also reduces memory requirements. The iterative solver is based on two Krylov subspace iterative methods, the restarted generalized minimal residual (GMRES) method and the quasi-minimal residual (QMR) method. These two methods are known to be robust.

The GMRES method is the most robust but also most expensive, as it requires the storage of the whole sequence of vectors to be orthogonalized. In practice, restarted or truncated versions are used to alleviate this drawback. The QMR method is less expensive in terms of computation time and memory requirements but also less robust. Therefore, the iterative solver we use is based on the QMR method with an automatic shift to GMRES in case of breakdowns.

The approximate factorization technique

The convergence of Krylov subspace methods is strongly influenced by matrix conditioning. Pre-conditioning transforms the original system into a better-conditioned equivalent one by a pre-multiplication.

A parallel version of the iterative solver can be used, for additional time gains. The strategy is based on the concept of pseudo-overlapped sub-domains.

Iterative solver example

An air intake is shown in Figure 3 and contains 117608 elements and 37593 nodes. Resources for a single frequency on SGI Origin 3000 were:

Solver	Memory (MB)	CPU time (sec)
Direct	**538**	**633**
QMR	**15**	**17 (117 iterations)**

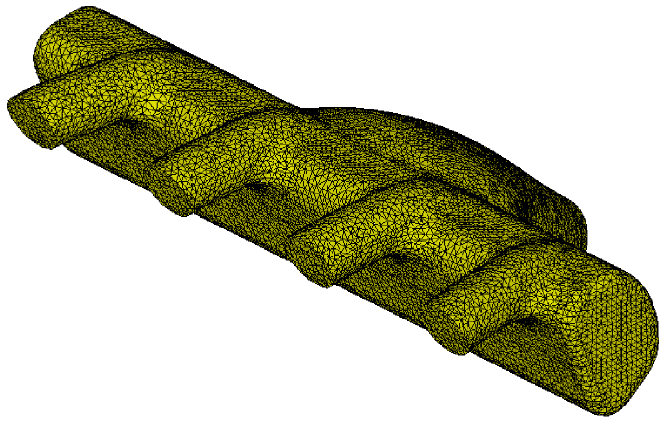

Figure 3: Air intake FE mesh

Domain Decomposition

Large problems can be tackled by division into sub-domains. Two methods have been used*: Finite Element Tearing and Interconnecting* (FETI) and a parallel form of the AFT introduced in the previous section. The FETI approach is based on a direct resolution for each sub-domain but an iterative one to solve the interface problem between domains; whereas the parallel AFT is fully-iterative, and we have used a concept of pseudo-overlapped domains with continuity between domains forced at the end of each Krylov subspace iteration. The AFT method therefore gives a speed-up of 30 to 40 times compared to the FETI approach.

Network solvers

Parallel processing can be executed at several 'levels'. At the **Frequency Level**, several processors (eg on a network) can each be given a sub-set of the complete frequency range to solve: the speed-up is nearly linearly-proportional to number of processors. At the **Matrix Level**, the complete system for one frequency is partitioned between processors: it enables extremely-large BEM problems to be solved, but the speed-up is less than linear. At the **Thread Level**, the operations within computations are assigned to several processors: the speed-up is less than linear, and shared-memory parallel hardware is needed, with compilers accepting OpenMP directives. Thread Level and Frequency Level can be combined. Further details will be given in the presentation.