Transactions on the Built Environment vol 69, © 2003 WIT Press, www.witpress.com, ISSN 1743-3509

Accelerating FEM and BEM acoustic solutions

M. Tournour, J. -P. Rossion, L. Bricteux, C. McCulloch & I McGann
LMS International, Belgium

Abstract

Reducing time-to-market, reducing development and production costs and increasing acoustic comfort, means less physical prototyping and more predictive modeling. However, predictive acoustic and vibro-acoustic models using FEM and BEM methods have often had limited usefulness, partly due to the long times taken to get the substantial amounts of results needed for engineering design optimization.

In this paper, recently-developed technologies are presented which accelerate acoustic solutions. An array of approaches is presented, using finite, infinite and boundary element methods, which are based on re-usable Modal Acoustic Transfer Vectors, Padé methods for rapid frequency-sweep solutions, domain decomposition, finite element iterative solvers and multi-processor 'netsolvers'. These technologies make for timely and effective acoustic predictions, which are also accurate. They tackle a wide range of applications, such as engine acoustics and other machinery noise radiation, interior vehicle acoustics and component vibro-acoustics. They enable the design of practical solutions and are effective in reducing time-to-market and development costs. The fundamentals of various methods, their deployment in software and a selection of practical applications, timing benchmarks and case studies are presented.

1 Introduction

The economic drivers of reduced time to market, improved product quality and reduced risks and costs, mean that virtual prototyping has become essential in many industries. 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

Transactions on the Built Environment vol 69, © 2003 WIT Press, www.witpress.com, ISSN 1743-3509

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 not limited theoretically, but by the capabilities of the computer and the solution time.

Successful deployment of FEM and BEM for vibro-acoustic simulation thus demands faster and more effective solutions. This paper presents state-of-the-art speed-up methods for FEM and BEM:

- ATV technology for multi-frequency and multi-RHS BEM computations
- Padé expansion for multi-frequency BEM computations
- Iterative Solvers for acoustic FEM computations
- Domain Decomposition for acoustic FEM computations
- Network Solvers for acoustic FEM and BEM computations (coupled or uncoupled).

The commercial packages LMS SYSNOISE Rev 5.5 [1] and Rev 5.6 (beta release) have been used to compute the numerical examples which are shown.

2 ATV™ Technology (BEM and FEM)

2.1 Acoustic Transfer Vectors

The Acoustic Transfer Vector (ATV) method, and its extension to Modal Acoustic Transfer Vectors, provides a huge speed-up for vibro-acoustic computations when there are multiple frequencies and multiple structural vibration load-cases (such as from an engine run-up). Furthermore, once computed, the ATVs or MATVs can be re-used efficiently with revised vibration load-cases or for structural optimization.

Acoustic Transfer Vectors are input-output relations between the normal structural velocity of the radiating surface and the sound pressure level at a specific field point. They can be interpreted as an ensemble of Acoustic Transfer Functions from the surface nodes to a single 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 or admittance), the frequency and the field point location. They do not depend on the loading. A highly-efficient formulation and implementation in SYSNOISE Rev 5.5 [1] limits the calculation effort for evaluating a single ATV to the computational cost of a single load case response calculation.

Transactions on the Built Environment vol 69, © 2003 WIT Press, www.witpress.com, ISSN 1743-3509

2.2 ATV for multi-load-case acoustic 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. Basic theory and an industrial application are in [2] and the use of ATVs in Panel Acoustic Contribution Analysis in [3] and in Inverse Acoustic Numerical Analysis in [4] and [5]. 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.

2.3 Interpolation Technique

In the case of sound wave propagation in an .open space, the fluid domain around the radiating object usually exhibits no resonances. Therefore, ATVs are rather smooth functions of the frequency, and coefficients can be accurately evaluated at any intermediate (*slave*) frequency, using a mathematical interpolation scheme based on a discrete number of *master* frequencies. (Note that the structural normal velocities cannot be similarly interpolated, since these directly depend upon the highly-resonant dynamics of the structure).

For large problems, computation time is dominated by the factorization time of the matrix. ATV interpolation does not required a factorization at the slave frequency, so frequency interpolation in an ATV-based forced response gives a huge time gain. Additional gains can come from using Netsolvers (Section 5.5).

It is shown elsewhere [6] that a safe choice for the 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.

2.4 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).

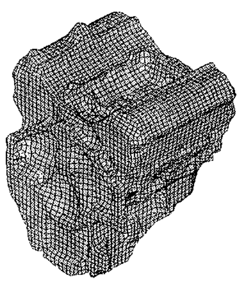

Figure 1: V6 Engine

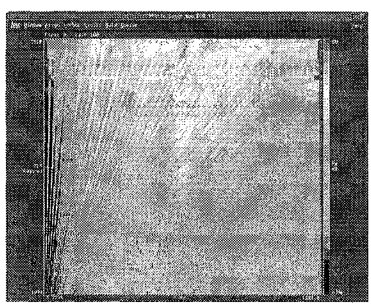

Figure 2: Sound pressure level above engine.

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!

3 Padé expansions

3.1 Background

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.

Conventionally, a linear (BEM) system has to be solved for each frequency:

$$\mathbf{A}(f)\boldsymbol{\mu}(f) = \mathbf{b}(f) \tag{1}$$

The coefficients of the matrix $\mathbf{A}$ depend on the frequency, so $\mathbf{A}$ must be calculated and factorized for each frequency of interest. An alternative is to approximate the frequency response function by a Padé Approximation (see [7,8]). Using Padé Approximation rather than more-classical Taylor expansions is justified by the fact that $\boldsymbol{\mu}(f)$ 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 of a Padé Approximation requires the knowledge of the successive derivatives of the acoustic quantities with respect to frequency, evaluated at the central frequency f_0. The first derivative of the surface potential $\boldsymbol{\mu}$ is the solution of:

$$\mathbf{A}\frac{\partial\boldsymbol{\mu}}{\partial f} = \frac{\partial\mathbf{b}}{\partial f} - \frac{\partial\mathbf{A}}{\partial f}\boldsymbol{\mu} \tag{2}$$

The derivative of $\boldsymbol{\mu}$ is obtained by solving the same system of equations that was solved to obtain the potentials $\boldsymbol{\mu}$ but with a different right-hand side involving the derivative of the matrix $\mathbf{A}$, the derivative of the original right-hand side vector $\mathbf{b}$ and the potentials $\boldsymbol{\mu}$.

The second- and higher-order derivatives of $\boldsymbol{\mu}$ with respect to f are the solution of:

$$\mathbf{A}\frac{\partial^n\boldsymbol{\mu}}{\partial f^n} = \frac{\partial^n\mathbf{b}}{\partial f^n} - \sum_{i=1}^{n} C_n^i \frac{\partial^i\mathbf{A}}{\partial f^i}\frac{\partial^{n-i}\boldsymbol{\mu}}{\partial f^{n-i}} \tag{3}$$

Transactions on the Built Environment vol 69, © 2003 WIT Press, www.witpress.com, ISSN 1743-3509

The most important property of this recurrence scheme is that the calculation of the successive derivatives of $\boldsymbol{\mu}$ requires the factorization of a single matrix $\mathbf{A}$. The right-hand side vectors associated with the higher-order derivatives of $\boldsymbol{\mu}$ are based on the successive derivatives of $\mathbf{A}$.

The coefficients of $\mathbf{A}$ are explicit functions of the frequency f which appears only as a parameter of the Green function G. In a given frequency range, we compute several matrices $\mathbf{A}(f_i)$ for different frequencies f_i. The higher-order derivatives of $\mathbf{A}$ are computed from these different evaluations of $\mathbf{A}$. For increased performance, it is assumed here that $\mathbf{A}$ varies smoothly with frequency.

3.2 Practical Issues

The practical value of the methodology depends on three issues: accuracy, speed (the speed-up can be 10 times) and memory and disk requirements (several matrices $\mathbf{A}$ need to be stored). The approach is also 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.

3.3 Example

The engine example of Section 2 can also show the benefit of the Padé expansion. The acoustic response was computed at 80 frequencies from 700Hz to 1500Hz on HP C3600. The computation using a conventional BEM approach took 11 hours 36 minutes, whereas the computation using Padé expansion took 1 hour 12 minutes: a speed-up factor of 9.6. A comparison of the pressure above the engine is so close that the curves appear completely superimposed.

4 Iterative Solvers (FEM)

4.1 Krylov subspace iterative methods

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 [9] and the quasi-minimal residual (QMR) method [10]. These two methods are known to be robust: theoretical and numerical evidence of their efficiency can be found in [11-13].

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

Transactions on the Built Environment vol 69, © 2003 WIT Press, www.witpress.com, ISSN 1743-3509

on the QMR method with an automatic shift to GMRES in case of breakdowns. (In fact, this automatic shift never occurred during the validation and testing).

4.2 The approximate factorization technique

The convergence of Krylov subspace methods is strongly influenced by the conditioning of the matrix. It is therefore recommended to use a preconditioning that transforms the original system into a better-conditioned equivalent one by a pre-multiplication. Theoretical and numerical descriptions of the approximate factorization technique are given in [13].

Finally, a parallel version of this iterative solver can be used, for additional time gains. (See Domain Decomposition, below). The parallelization strategy is based on the concept of pseudo-overlapped sub-domains described in [14].

4.3 Examples

Two examples are presented with velocity and impedance boundary conditions.

A cube is shown in Figure 3 and contains 91126 elements and 97336 nodes. The test looks trivial, but is often used in validations: it gives the densest FEM matrix. Comparison of memory and CPU time using the direct and the iterative solvers is given in Table 1 for one frequency on SGI Origin 3000 computer.

An air intake is shown in Figure 4 and contains 117608 elements and 37593 nodes. Table 2 shows results for a single frequency on the SGI Origin 3000.

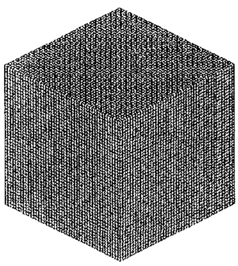

Figure 3: Cube (97336 nodes)

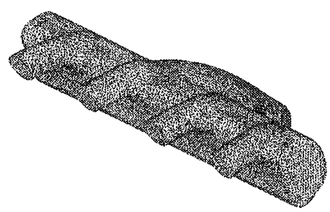

Figure 4 : Air Intake (37593 nodes)

Solver	Memory (MB)	CPU time (sec)
Direct	3400	3563
QMR	200	59 (52 iterations)

Table 1: Direct and iterative solver for the cube.

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

Table 2: Direct and iterative solver for Air Intake

Transactions on the Built Environment vol 69, © 2003 WIT Press, www.witpress.com, ISSN 1743-3509

5 Domain Decomposition and Network Solvers

5.1 General principles of domain decomposition

A domain decomposition method decomposes the entire domain into sub-domains and solves the problem on each sub-domain separately. Because of the de-coupling of the sub-domains, domain decomposition is well-suited for parallel computing, when each sub-domain is allocated to a single processor. In order to restore the connection between the sub-domains, continuity conditions are imposed along the interfaces between the sub-domains. This leads to a so-called *interface problem* that describes the coupling of the sub-domains. The solution of the interface problem readily produces the solution of the global domain. Whereas the solution of the problem on each sub-domain happens without communication, the solution of the interface problem requires communication between neighbouring sub-domains. The amount of communication depends mostly on the size of the interface between the sub-domains. In this paper, we consider two specific domain decomposition methods, called *the Finite Element Tearing and Interconnecting* method for the *Helmholtz* equation (FETI-H method) and the parallel *Approximate Factorization Technique* introduced in the previous section.

5.2 The Finite Element Tearing and Interconnecting method

The FETI-H method is a non-overlapping domain decomposition method for solving linear systems arising from the finite element discretization of the Helmholtz equation on a bounded domain.

The FETI method introduced in [15] is a domain decomposition technique that uses a direct method on the sub-domains but an iterative one for solving the interface problem. It aims at combining the robustness of a direct method with the ease of the parallel implementation of an iterative procedure. The main features of the methods from both the numerical analysis and the distributed parallel implementation points of view are presented in [15].

5.3 The Parallel Approximation Factorization Technique

Whereas the FETI approach is based on a direct resolution for each subdomain, the *Parallel Approximation Factorization Technique* is a fully iterative approach. The parallelization of the *Approximate Factorization Technique* introduced in Section 4.2 is based on the concept of pseudo-overlapped subdomains described in reference [14] where the continuity at the domain interface is enforced at the end of each Krylov subspace iteration.

5.4 Example of domain decompostion

Table 3 compares the performance of the FETI and AFT methods on SGI Origin 2000 computer for the air intake shown in Figure 4.

Transactions on the Built Environment vol 69, © 2003 WIT Press, www.witpress.com, ISSN 1743-3509

Number of CPU's	FETI(sec)	AFT(sec)	Ratio FETI/AFT
1	633	17	37
4	162	4	39
8	70	2	33
16	54	1.6	33

Table 3 : Computation time for the air intake using FETI and AFT approaches.

5.5 Network solvers

5.5.1 General principles A "Netsolver" aims to solve very large jobs by means of parallel processing (splitting tasks over multiple processors). As well as using Domain Decomposition, three 'netsolving' techniques are possible: Frequency Level, Matrix Level and Thread Level Distribution. The Netsolver supports various types of parallel platform architecture: Shared or distributed memory parallel computers, homogeneous or inhomogeneous clusters of workstations. Three standards for data exchange are used: Public Domain MPI (Message Passing Interface) for data exchange between homogeneous or heterogeneous computers linked through a LAN network or between CPUs of parallel solvers, native MPI for data exchange between the CPUs of parallel solvers, and OpenMP for shared-memory parallel solvers.

5.5.2 Frequency Level Distribution assigns to each processor a subset of the total set of frequencies to be computed. Each processor runs its own executable and performs its own independent analysis. A BEM example (similar to the engine shown earlier, but with 4532 nodes) was solved at 48 frequencies on four machines, defining subsets of frequencies and merging the final results database automatically. The speed-up is about linear (2 machines, 2 times; 4 machines, 4 times): the overheads for merging etc are negligible.

5.5.3 Matrix Level Distribution assigns to each processor a subset of the global matrix to be computed (partitioning in blocks). Since each node processes only a sub-part of the matrix, the memory requirement is divided by the number of nodes. This is an advantage when compared to the Frequency Level Distribution, for large problems, but the speed-up is not linear: 4 processors, 2.9 times (on an SGI Origin 3000).

5.5.4 Thread Level Distribution assigns the operations within computations, to several processors. This requires shared-memory parallel computers and compilers accepting OpenMP directives. The memory required is the same as a sequential version. On the same SGI Origin 3000, with parallel operations in both assembling and solving, the speed-up was 3.2 times.

5.5.5 Combined Levels Frequency and thread levels can be combined: The previous example was solved for 100 frequencies using 80 CPUs. 20 CPUs were used for the frequency level (5 frequencies per CPU) and for each frequency 4 threads (slave CPUs) were used. The sequential run took 9h 43min on SGI Origin 3000, the combined MPI-OpenMP parallel solution took 9min 15 sec, a speed-up of 63 times.

6 Conclusions

There are several new speed-up technologies - modal acoustic transfer vectors, Padé methods, domain decomposition, iterative solvers and netsolvers - for FEM/BEM acoustic solutions. Domain decomposition and iterative solvers are very efficient for acoustic FEM; ATVs are very useful for multi-frequency and multi-RHS BEM; and Padé expansion is effective for multi-frequency BEM. Netsolvers are very efficient for all computations (Acoustic FEM and BEM) and can be combined with other technologies.

Acknowledgements

The integration of Padé expansions in SYSNOISE Rev 5.6 is the result of a collaboration between LMS International and CADOE S.A., France.

The integration of the domain decomposition techniques and the iterative solver in SYSNOISE Rev 5.6 has been done within the ESPRIT project DOMINOS and the IWT project TRICARMO.

The introduction of Thread Level parallel processing and use of OpenMP in SYSNOISE Rev 5.6 has been done within the IST project INTONE.

References

[1] SYSNOISE Rev 5.5 User's Manual. © LMS International, 2000.

[2] F. Gerard, M. Tournour, N. El Masri, L. Cremers, M. Felice, A. Selmane, "Numerical Modeling of Engine Noise Radiation Through the Use of Acoustic Transfer Vectors - A Case Study", *Proc and exposition report, SAE Noise and Vibration Conference.* April 30 – May 3, 2001, Traverse City, Michigan, USA

[3] L. Cremers, M. Tournour and C.F. McCulloch, "Panel Acoustic Contribution Analysis based on Acoustic Transfer Vectors", *Proc. of The 2001 International Congress and Exhibition on Noise Control Engineering.* The Hague, The Netherlands, 2001 August 27-30.

[4] M. Tournour, L. Cremers, P. Guisset, F. Augusztinovicz, F. Marki, "Inverse numerical acoustics based on acoustic transfer vectors", *Proc. ICSV-7*, Garmisch-Partenkirchen, Germany, Jul. 4-7 2000.

[5] F. Augusztinovicz and M. Tournour, "Reconstruction of source strength distribution by inverting the boundary element method", In "*Boundary*

Elements in Acoustics. Advances and Applications", ed. O. von Estorff, WIT press, 2000.

[6] M. Tournour, S. Dessart and C.F. McCulloch, "On the accuracy of vibro-acoustic solutions using the ATV method", *Proc. ISMA2002*, Leuven, Belgium, 16-18 Sept 2002.

[7] Brezinski C. "Padé type approximation and general orthogonal polynomials". ISNM Vol. 50, Birkhauser Verlag, Basel, 1980

[8] ADOC 3.0 User's Manual. CADOE S.A., France.

[9] Y. Saad and N.H. Schultz. "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems". *SIAM J. Sci. Statist. Comput.,* 7:856-869, 1986.

[10] R.W. Freund. "Conjugate gradient-type methods for linear systems with complex symmetric coefficient matrices". *SIAM J. Sci. Statist. Comput.,* 13:425-448, 1992.

[11] R.W. Freund and N.M. Nachtigal. "An implementation of the QMR method based on coupled two-terms recurrence". *SIAM J. Sci. Statist. Comput.,* 15:313-337, 1994.

[12] M. Magolu monga Made. "Preconditioning of discrete Helmholtz operators perturbed by a diagonal complex matrix". *Comm. Num. Methods Eng.,* 16:801-817, 2000.

[13] M. Magolu monga Made. "Incomplete factorization-based preconditioning for solving the Helmholtz equation". *Int. J. Num. Methods Eng.*, 50:1077-1101, 2001.

[14] M. Magolu monga Made and H.A. van der Vorst. "Parallel incomplete factorizations with pseudo-overlapped subdomains". *Parallel Computing,* 27:989-1008, 2001.

[15] A. de La Bourdonnaye, C. Farhat, A. Macedo, F. Magoulès and F.-X. Roux. "A non-overlapping domain decomposition method for the exterior Helmholtz problem", *Contemporary Mathematics*, 218:42-66, 1998.