PACAN-D

ESPRIT Project Number 24871

PAC

SER Systems Ltd

Celestion

MIRA

Public Final Report of ESPIRT HPCN PST Activity PACAN-D, Parallel Acoustic Analysis Demonstrator.

PAC/PACAN-D/D0.1 Version 1

T. Cooper, N. K. Allsopp, Parallel Applications Centre
J. R Wright, Celestion
S. Roberts, MIRA
P. Macey, SER Systems

10 June 1999

Abstract

The PACAN-D (Parallel Acoustic Analysis Demonstrator) project intended to show how a parallel acoustic code could be deployed on low cost hardware (both NT and Unix) and thereby innovate the deployment of numerical simulation in two different business models.

Synopsis

This report is the final report of the PACAN-D Project (ESPRIT). The project involves the following partners: PAC (UK), SER Systems UK (UK), Celestion(UK) and MIRA (UK). The project was co-ordinated by PAC, and further information may be obtained from

Tim Cooper,
Tel: +44 (0)1703 760834,
Fax: +44 (0)1703 760833
mailto: tpc@pac.soton.ac.uk

Virtual prototyping, through the use of numerical simulation, offers a way for companies to reduce the time to market for new components. However numerical simulation software is not easy to use and requires staff with a wide range of skills. For many organisations this is costly. The PACAN-D project has demonstrated two novel ways in which a parallel simulation code can be deployed to reduce this cost.

The Motor Industry Research Association acts as a centre of excellence in the UK for automotive research. Part of its business is the provision of consultancy service to manufacturing companies (both OEMs and component manufacturers), and it uses numerical analysis to support this activity. At the start of the project, the turnaround time to analyse a component was approximately 3 months. Most of this time was spent in the preparation of the component mesh (a numerical representation of the component geometry required by all numerical analysis packages). Although packages exist which will create these meshes semi-automatically in a matter of days, if not hours, the resulting mesh is not of the quality required for the complex geometries studied by MIRA, and if used would take up to a week to solve. This is not acceptable, given the computational resources currently available to MIRA, and purchasing extra resources rapidly becomes too expensive to justify. However when these automatically prepared meshes are run with a parallel code, the execution time is reduced to only a day. This means that the whole process time is reduced from 3 months to 1 month or less. As the cost of the consultancy is determined by the effort, the saving equates to an equivalent saving in cost meaning that MIRA are able to sell the service to more organisations.

Celestion manufacture loudspeakers, and have invested heavily in numerical simulation over the years. Whilst they have reaped the benefit of this, in a market place which is ever more competitive, there is constant pressure to lower costs and get more and more value from investments. Celestion are therefore looking at ways in which they can make simulation more accessible to designers (as opposed to the analysts who currently do the work). They have identified that once the analyst perfects a simulation scheme it remains stable for a range of parameters. By automating the analysis process using in-house software they hope to be able to give the simulation to designers to use, thereby removing the potential bottleneck on the analyst. However this extra use of the software means that they potentially need more hardware on which to run it, and the cost of this is hard to justify. In the PACAN-D project we have shown how a parallel code can be run on an existing network of office PCs without impacting on the everyday work of those using them.

Executive Summary

The use of parallel simulation codes has been widely demonstrated in industry. However their take up by smaller organisations has often been limited by the cost of the hardware needed to run the codes. In the PACAN-D project we set out to show how a code which had been ported in part to run on a massively parallel machine could be fully ported to run on low cost hardware, namely a cluster of NT workstations and a low cost Unix SMP box. The aim of the project was not just to show that the port is technically possible, but also to investigate the business models by which the parallel code could be used.

The approach adopted in the project was to take a code that had previously been partly parallelised for an old architecture (the Intel Paragon), and port it to a code based on MPI. Since MPI is available on most architectures, this was felt to be the least risk option. The next stage was to deploy the code on a cluster of NT workstations and assess its usability. The feedback from the end user was then used to develop models for the deployment of parallel simulation on a cluster of NT machines. Finally, the code was deployed on the Unix SMP box. The problem to be run in this case was much larger than the NT problem, and again a business model was developed for how the parallel code could be used.

The Motor Industry Research Association acts as a centre of excellence in the UK for automotive research. Part of its business is the provision of consultancy service to manufacturing companies (both OEMs and component manufacturers), and it uses numerical analysis to support this activity. At the start of the project, the turnaround time to analyse a component was approximately 3 months. Most of this time was spent in the preparation of the component mesh (a numerical representation of the component geometry required by all numerical analysis packages). Although packages exist which will create these meshes semi-automatically in a matter of days, if not hours, the resulting mesh is not of the quality required for the complex geometries studied by MIRA, and if used would take up to a week to solve. This is not acceptable, given the computational resources currently available to MIRA, and purchasing extra resources rapidly becomes too expensive to justify. However when these automatically prepared meshes are run with a parallel code, the execution time is reduced to only a day. This means that the whole process time is reduced from 3 months to 1 month or less. As the cost of the consultancy is determined by the effort, the saving equates to an equivalent saving in cost meaning that MIRA are able to sell the service to more organisations.

Celestion manufacture loudspeakers, and have invested heavily in numerical simulation over the years. Whilst they have reaped the benefit of this, in a market place which is ever more competitive, there is constant pressure to lower costs and get more and more value from investments. Celestion are therefore looking at ways in which they can make simulation more accessible to designers (as opposed to the analysts who currently do the work). They have identified that once the analyst perfects a simulation scheme it remains stable for a range of parameters. By automating the analysis process using in-house software they hope to be able to give the simulation to designers to use, thereby removing the potential bottleneck on the analyst. However this extra use of the software means that they potentially need more hardware on which to run it, and the cost of this is hard to justify. In the project we have shown how a parallel code can be run on an existing network of office PCs without impacting on the everyday work of those using them.

Contents

1 State of the Art

1.1 The numerical acoustic analysis of loudspeakers

Acoustic analysis is well established and is well understood. Numerical simulations form a part of this and are used by many organisations, particularly in the automotive sector. However the numerical acoustic analysis of loudspeakers is at present limited. The problem is not that the physical processes are not well understood, although there are many complex fluid/structure interactions that must be captured if believable results are to be obtained. The problem is rather that to capture them the analyst must build large complex models. Most of the organisations interested in this field are SMEs who cannot afford the large computing resources needed to analyse these large problems.

At the start of the project, Celestion were able to model the coupled interactions between the loudspeaker cone and the air. The model of the code was realistic, but there were limitations in the model of the loudspeaker box. The main problems, as far as Celestion were concerned, were that to analyse larger problems took too much memory and too much time to execute. The requirements of Celestion were

1. to be able to run larger problems in parallel in the same elapsed time that it took to run their existing problems,
2. to be able to run their existing problems is around 1/3 time and
3. to be able to do the above on their existing cluster of 4 NT workstations. These workstations are used for office work, and an important requirement was that the impact on the users of the machines be minimised.

1.2 The numerical analysis of the acoustic properties at MIRA

The numerical analysis of the acoustic properties of automotive components is well established. For most cases, there is minimal loading from the air on the surface, so the analysis is relatively simple. However this is not always the case, and in these instances a more specialised analysis procedure is called for.

The UK Motor Industry Research Association is a centre of excellence in automotive research, and provides consultancy services to the automotive industry. One of the areas where it provides consultancy is in the area of coupled fluid structure acoustic analysis. This is particularly a problem in engine air intakes, and the problem is further compounded by the fact that these structures often have convoluted geometries that are hard to mesh. When providing an item of consultancy, much of the cost is the time taken by MIRA to mesh these complex geometries. A typical air intake might take up to 2 months to mesh, followed by a further month to perform the analysis (a process which involves solving the mesh a number of times). Automatic mesh generators would reduce the 2 months mesh generation time down to a week or less, but the poor quality meshes generated would take up to a week to execute, thereby increasing the time perform the analysis by so much that all benefit is lost.

The requirement of MIRA, therefore, was to obtain a code that would be able to process an automatically generated code in around a day. This would allow them to reduce the time to analyse a component from 3-4 months down to approximately 1 month.

1.3 The use of simulation by designers

The extent to which designers use simulation varies considerably between sectors. In some hi-tech areas, for instance aerospace, there is minimal distinction between designer and analyst. However in many other sectors the analysis is carried out by a dedicated team, separate from the design process. If one goes back 30 years, the process of design was to go from paper drawings to physical tests and iterate until an acceptable design was reached. This is an expensive process, and the advent of numerical simulation allowed virtual prototypes to be built at a much lower cost. However the process of passing a design from the design team to the testing team has remained the same. What this process fails to exploit is that once the analysis process is defined it will remain stable for some set of variations from the original design. Therefore there is no reason, from the point of view of physical validity, why members of the design team could not run the analysis. Although this idea is simple to express, the obstacles to its implementation are many. Simulation codes are not easy to execute, the method of changing parameters of interest is not consistent, simple changes may have unexpected results and so on.

However at Celestion they have started to look at many of these problems. At the start of the project they were at the stage where, for a limited set of problems, the designers were able to change parameters of interest and run the simulation. The problem for the designers at Celestion was that the answers did not arrive back for a number of hours, and this way of working does not fill well with the established practice of design teams. Their conclusion was that if they were to be able to use the tool, answers would have to be available within an hour.

1.4 The maintenance of parallel simulation codes

One of the hindrances to the take up of parallel simulation is that parallel codes can be hard to maintain. The original parallel port of the PAFEC VibroAcoustic code was developed before the advent of standard message passing libraries and was specific to the Intel iPSC/860 and Paragon, machines to which SER Systems Ltd did not have direct access. Consequently the code was not maintained. The original port was highly tuned to the architecture of the iPSC/860 and Paragon. This made it hard for the original developers of the software to read and understand the ported code. Furthermore any development of the parallel version for another platform would be unlikely to respect the optimised code, meaning that the code would be both poorly optimised and hard to understand.

One of the intentions, therefore, of this parallelisation was to make it as portable as possible and not to tune it to any specific architecture. This of course represents no new thinking, and many initiatives in the years since the original port have developed technologies to increase the portability of parallel codes. However the creation of a maintainable code is a skill, and this requirement must be included from the start of the design.

2. Approach taken

The approach taken in the project was to parallelise the latest version of the serial PAFEC VibroAcoustic code using the MPI message passing standard. The code was then tested against two industrial benchmarks, a loudspeaker cone and cabinet from audio equipment manufacturer Celestion and a car engine air intake from MIRA. In this section we shall discuss the issues associated with the choice of MPI to implement the parallelisation, the motivation behind the two test cases, the design of the parallel code and the roll out of the code to the end users.

2.1 The choice of MPI

The choice of the message-passing library was determined by the requirements of SER systems Ltd, the owners of the PAFEC VibroAcoustic code. They need a library that is

- portable across many architectures (Unix and Windows),
- commercially supported and
- easy to understand.

MPI (Message Passing Interface) has rapidly become established as the de facto message-passing library for parallel codes. It is portable across a wide range of platforms. Implementations are supported either under the GNU agreement or by hardware vendors. It is easy to understand, and most codes can be parallelised using a small subset of the functions available. At the start of the project PAC believed the parallelisation could be carried out using at most 6 directives. For these reasons it appeared to be the natural choice for implementing the parallelisation of the code.

However at the start of the project, although there were many implementations for Unix machines, there were no robust implementations for NT. A number of initiatives were under way to port the library, but only beta releases were available. In order to be certain that MPI for NT was the right way to go PAC organised a one-day conference on PC clustering to look at the various technologies available.

The project workshop attended by members of the HPCN-TTN and the outcome documented in the report 'PC Clustering: a Report on Current Best Practice, T Cooper, Oct 1997'. The workshop looked not only at MPI for NT, but also at other libraries to support parallelism such as threads, and the hardware to support clustering.

As a result of the workshop it became clear that there were a number of initiatives which were developing implementations of MPI on NT, most notably the WINPAR project, co-ordinated by Genias and the MPICH-NT initiative at Argonne in the USA. Other initiatives which looked at building clusters of PC machines were felt to be too specialist or too academic. For SER Systems Ltd to be able to support the code commercially it was important that a third party should support the parallel libraries. Therefore the final choice was to develop the code using the beta release from the WINPAR project. When final versions from both initiatives became available, a benchmarking test was performed to ensure that the best version for the PAFEC VibroAcoustic code was used.

2.2 The industrial test cases

Two industrial test cases were used in this project, an air intake from MIRA and a half model of a loudspeaker cone within a cabinet.

2.2.1 The MIRA test case

The MIRA test case is a model of the air intake for an automobile. The mesh was created using semi-automatic methods, resulting in a mesh which would normally take too long to execute.

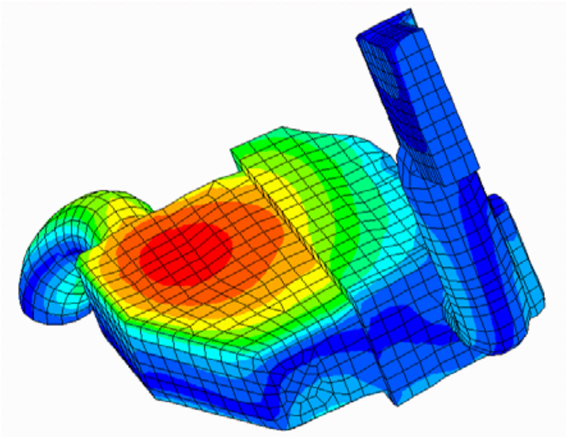

The MIRA test case, a car intake system, consisted of 27538 structural degrees of freedom, 4758 acoustic F.E. freedoms (internal) and 1992 acoustic boundary element freedoms

2.2.2 The Celestion test case

The use of FE techniques in the design process at Celestion is worth discussing as an example of the take up of simulation technology by a small manufacturing enterprise. As might be expected, Celestion started their FE work with simple models for which the issue was transmission through the member. This type of model is fairly easy to solve, in terms of the physical processes to be captured, and does not typically need large models.

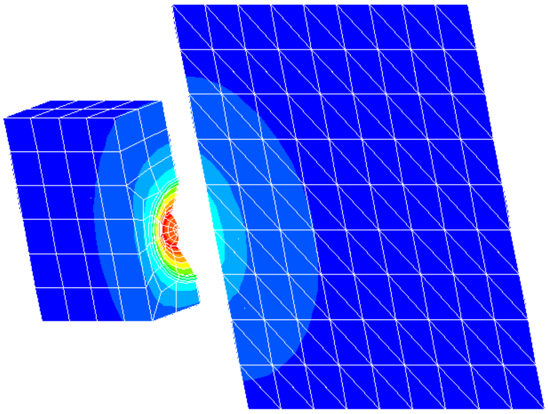

Test case with pressure waves radiating from the code of the speaker box which used 140 quadratic shell elements and a mass element containing 1917 elements with a front size of 112 to model a dustcap, surround, cone and box of the speaker cabinet.

As their experience in the techniques increased, Celestion looked to build more complex models of the radiating acoustic surfaces. These problems are harder to model as the effects of the interfaces between components becomes critical, and as a consequence require much larger models. At this stage Celestion found that full models were hard to run as the memory requirements became prohibitive. The Celestion test case represents the largest problem that they could run at the start of the project on a single workstation.

2.3 The design of the parallel code

Traditionally a parallel system consists of a master processor and a large number of slave processors. The master shares out the work to the slaves and collects the results after the computational phase of the analysis. Since we were dealing with a small cluster of NT workstations the parallel algorithm was written such that the master acted as a

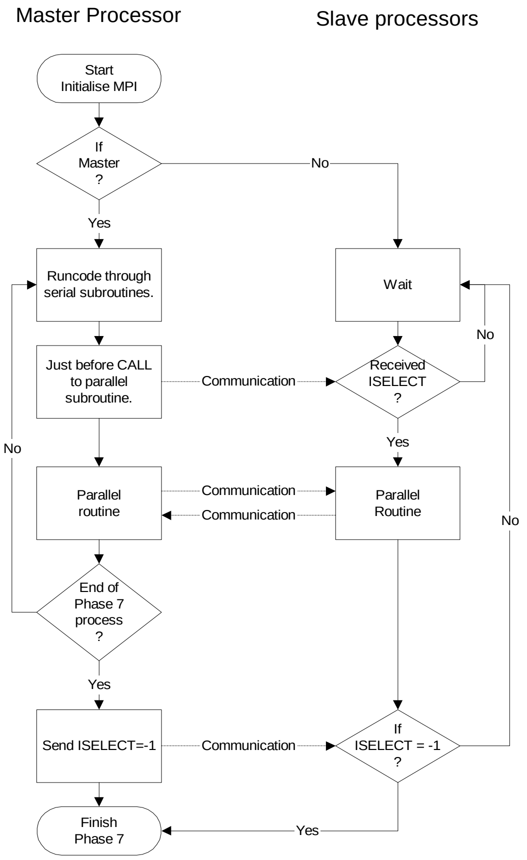

Flow diagram of overall methodology of serial code parallelism.

slave during the numerically intensive sections of code. There is usually a loop on frequencies analysed, which includes virtually all the computation. One strategy would have been to share out frequencies between processors. This would have been fairly easy to implement since there would be virtually no communication between processes. However this result can be achieved simply by concatenating the output files from some serial analyses run over sections of the frequency range on unconnected processors. A more ambitious approach was chosen, parallelizing at a lower level, so that all processors are used in the computation of each frequency. This means that the combined resources of cache/memory/disk are available to tackle each frequency, which is beneficial for very large problems, exceeding the physical memory requirements of a single processor, but there is a greater need for communication between processors, which may delay the computational procedure. From profiling it was known that the numerically intensive routines formed only a small part of the total number of lines of code. Hence the technique used was to let the master follow the original serial route through the code, but when a computationally demanding section is reached all of the processors are given a similar amount of work. This means that the system is load balanced within the parallel sections of the code, whilst the main I/O is dealt with by the master processor only. This is demonstrated in the figure below, which shows the underlying methodology of the parallelism of the analysis routines.

The PAFEC VibroAcoustics system comprises a suite of programs called *phases*. There are 10 phases in total. Phases 1-6 perform pre-processing tasks. Phase 7 performs the main solution of the equations. Phases 8-10 perform post-processing tasks.

The underlying ethos to the parallelism of phase 7 is that the master processor proceeds through the code in the same way as the serial version, as shown in flow diagram above. When the master processor is about to enter a numerically intensive routine a message is sent to the slave processors. All processors perform an equal part of the required calculation. When the routine is finished the master processor continues to progress through the serial code whilst the slave processors wait for the next numerically intensive section to be reached by the master processor.

The equations for a fully coupled vibroacoustic solution using an acoustic BE mesh coupled to a structural FE mesh, are

$$\begin{bmatrix} [S]+i\omega[C]-\omega^2[M] & [T]^T \\ -\omega^2\rho[G][E]^T & [H] \end{bmatrix} \begin{Bmatrix} \{u\} \\ \{p\} \end{Bmatrix} = \begin{Bmatrix} \{F\} \\ \{p_I\} \end{Bmatrix}$$

where {u} is a vector of displacements on the structural mesh, {p} is a vector of pressures on the BE mesh, [S], [C] and [M] are structural stiffness, damping and mass matrices and are large and sparse. [H] and [G] are small dense matrices derived from the BE formulation. [T] and [E] are coupling matrices. Sometimes a simpler modal structural representation is used, but that does not permit variation of properties with frequency, as occurs with the cone and surround of a loudspeaker. The current work was based on a full solution of equation (1) using the 4 stages below.

- **Stage 1 - FE merging/reduction**. Merging contributions from individual finite elements forms the dynamic stiffness matrix and coupling matrix [T]. A frontal solution is used to eliminate degrees of freedom during the merge. The matrices are shared between processors and the elimination is done in parallel.
- **Stage 2 - forming the BE matrices**. For each collocation point on the BE surface it is necessary to integrate over the surface to form a row in the BE matrices. Parallelization is achieved by sharing these collocation points between the processors. Distributed BE matrices are formed.
- **Stage 3 - reducing the BE matrices**. The matrix $-\omega^2\rho[G][E]^T$ is formed and reduced using resolution with the structural elimination equations from stage 1. As above the matrices are distributed between processors.
- **Stage 4 - Gaussian elimination of final equations**. The resulting compact dense set of equations is solved using a column cyclic parallelized form of Gaussian elimination on the distributed matrix.

2.4 The installation of the code at the end user sites

PAC has a wide variety of computer hardware at its offices. These range from Unix boxes through to desktop NT workstations, both single and dual processor.

Celestion are using a cluster of NT workstations linked together using 100Mbit fast Ethernet whilst MIRA are using a dedicated multi-processor SGI PowerChallenge. To facilitate the parallelism on the two operating systems MPI was chosen. For the SGI using UNIX the message passing was obtained using MPICH whilst on the NT system the newly emerging WMPI was used. Implementing the message passing using MPI was therefore vindicated, as the high level coding was identical for both NT and UNIX operating systems.

One the issues that arose from the distributed cluster of NT workstations was that the slave processors were to be used as dual-purpose machines. That is the machines were to be used on people's desk for administration tasks as well as being used as part of the computing resource of the clustered machine. One option, which was investigated, was to use to dual processor NT workstations on the slave processor machines. This enabled the computational cluster to be hidden from the uses of the slave machines but there was no method available to select which processor was used on the dual processor platforms. It was discovered, however that this did not seem to impact on the slave processor users.

3 Results, Achievements and Benefits

The following benchmark test results were performed on a cluster of NT workstations and a shared memory multi-processor SGI machine.

3.1 The Celestion test case benchmarking results

The input data set was supplied by Celestion International, which represents a typical loudspeaker system. This test case was used on a cluster of Pentium workstations running Microsoft NT and a multi-processor shared memory Silicon Graphics (SGI) machine running UNIX. Three version of the Message Passing Interface (MPI) were used. Two versions (WMPI PaTENT version 4.09 and MPI/Pro version 1.2.3) were used on the Microsoft NT platforms and one on the UNIX platform (MPICH version 1.0.13).

The solution time, in seconds, is given for each frequency. The number of FORTRAN words in the array used for most of the program data storage is tabulated as:

	1 processor	2 processors	3 processors	4 processors
FORTRAN words	19085236	11963350	9158554	7657310

Table 1: Number of FORTRAN words used.

This reduction in memory usage demonstrates that using the distributed memory of the slave processors allows a larger problem to be calculated when using 4 processors than could be calculated on a single processor.

1 processor		2 processors		3 processors		4 processors	
MPIPro	*WMPI*	*MPIPro*	*WMPI*	*MPIPro*	*WMPI*	*MPIPro*	*WMPI*

Stage 1	17	17	26	26	33	35	34	36
Stage 2	30	30	16	16	11	11	8	8
Stage 3	161	161	128	94	74	78	127	64
Stage 4	34	34	67	65	100	109	104	106
Total Time	246	246	237	201	221	233	274	216

Table 2. Performance using 450MHz Pentium II processors, 128 Mbytes RAM, 100 Mbit Ethernet with MPIPro 1.2.3 (left-hand columns) and PaTENT WMPI 4.09 (right-hand columns).

The following results were obtained from the shared memory SGI machine. Although 8 processors are available only four were used to give a comparison between the UNIX implementation and the NT implementation of the code.

	1 processor	2 processors	3 processors	4 processors
Stage 1	73	74	73	75
Stage 2	165	86	58	41
Stage 3	975	484	315	227
Stage 4	132	95	85	104
Total Time	1345	739	531	447

Table 3: Execution time using a shared memory 75MHz SGI Power Challenge with MPICH

3.2 The MIRA test case results

This particular test case was too large to fit onto one processors of desktop machine running NT. Therefore all the following analysis of this test case was performed on an SGI Power Challenge only.

This particular test case, of a car intake system, consisted of 27538 structural degrees of freedom, 4758 acoustic F.E. freedoms (internal) and 1992 acoustic boundary element freedoms. The sysetm was so large that the analysis was broken into two parts, ie the structural reduction and the boundary element analysis was calculated in two separate jobs. The total time for a single frequency are shown in the two following tables.

	1 processor	2 processors	4 processors	6 processors
Total time	5292	3394	1932	1698

Table 4: Execution time using SGI Power Challenge for structural reduction only.

	1 processor	2 processors	4 processors	6 processors
Total time	8000	3775	1901	1302

Table 5: Execution time using SGI Power Challenge for boundary element calculation.

These results were measured at MIRA on a 12 processor 194MHz SGI Power Challenge.

3.3 Analysis of benchmarking results.

For the Celestion test case the front size was too small for stage one to execute in parallel. There is however a large global summation at the end of stage one and this accounts for the increase in execution time as the number of processors is increased.

Stages two and three of the solution process are parallelised very efficiently even though stage three does contain some necessary communications. The efficiency of the different MPI protocols shows that WMPI is implemented more effectively than MPI/Pro for this particular test case.

Stage four when run over the 100 Mbit ethernet is slower when run on more than one processor. This is because there is a relatively high ratio of communication to computation. However as the problem size increases the ratio of communication to computation decreases in all stages. Thus it is believed on larger problems the execution time will decrease. This is shown by the SGI version of the code where the processor speed to communication speed is relatively slower than that of the NT with Ethernet communication system.

The MIRA testcase demonstrates that when the problem size becomes large enough so that the calculation dominates the communication then even stage 1, which by necessity has a large amount of communication, the parallel code does scale relatively well.

3.4 Celestion Assessment of the code

3.5.1 Speed performance

There are four stages of the analysis which have been implemented in the parallel code. Benchmarking showed that Stages 2 and 3 show almost linear improvement with the increase in number of processors, whereas Stage 1 runs slightly slower and Stage 4 shows a small improvement. (For further details see 'High Performance Computing Network Techniques for Vibroacoustic Analysis' by Macey, Wright and Allsopp, Audio Eng. Soc. Preprint #4883, May 1999.)

Since Stages 2 and 3 are usually the most critical stages for a Celestion model, we can conclude that the project has been highly successful.

3.5.2 Memory performance

The memory requirement per process of the parallel code reduces by a factor of 3 on 4 machines. This means that Celestion can now start to analyse problems of a loudspeaker within a room, or problems where they are able to account for coupling between most of the components of the loudspeaker.

3.5.3 Usability of the code

One further improvement that would be of particular benefit to Celestion, where the PCs are also used for day-to-day office work, is the implementation of Dynamic Load Balancing. This would allow the code to run in the background without impacting on the work of designers. The numerical analysis performed in the PAFEC VIBROACOUSTIC code could not support this, and to rewrite the code would have taken more effort that was available in the project.

The core issue with dynamic load balancing is to be able to run the code without impacting on other users of the network. To try and address this, the code was also installed on a network of dual and single processor machines at PAC and tested in various configurations and under various background loads. It was noted that when the parallel code was set to run on a cluster of dual processor machines, with only one process per machine, then the impact on users of those machines was limited. These results were demonstrated to Celestion.

The issue for Celstion, therefore, was one of cost benefit. To upgrade to dual processor machines would mean replacing the motherboards as well, and the benefit of this needed to be matched against either leaving their machines in their current configuration or simply upgrading each single processor to one twice the speed. In either case Celestion felt that the benefits from the parallel code were sufficient to meet their needs without upgrading their cluster, and have made no further improvements.

It is likely that if Celestion does upgrade their machines it will be to faster processors to allow all users to benefit, including the users of PAFEC VIBROACOUSTIC, rather than the dual processor machines which would only benefit the PAFEC VIBROACOUSTIC users.

3.5.4 Conclusions

In summary the results of the project are that Celestion are now able to run models four times more complex than previously, or run the same models four times faster. It should be noted that this improvement can be increased still further by the addition of more PCs to the cluster, although relative improvement will start to diminish.

The speed improvement in particular has allowed Celestion to meet the goal of moving FEA from a 'Research Tool' to a 'Design Tool', the most important aim of the project.

3.5 MIRA assessment of the code

The solution of coupled structural acoustic models at MIRA is currently done using the PAFEC VIBROACOUSTIC structural acoustic software. This is used for both finite element modelling of the structure and enclosed cavities and the boundary element solution of externally radiated noise. Currently the solution times required for many classes of problem are high and this restricts the use of the code to projects where lead times are small or financial constraints are significant.

As part of the PACAN-D project to highlight the advantages of high performance computing, MIRA prepared a test case to demonstrate that the application of a parallel version of the PAFEC VibroAcoustic code could reduce solution times, so opening up the techniques to a wider group of users.

3.5.1 Objectives

The aim of the demonstration case was to test the parallel code against a model representative of those likely to be used by MIRA. This means that the model must be representative in terms of size and content and in terms of solution type and element library to allow a realistic assessment of the advantages of parallel codes to be made.

3.5.2 Current position

MIRA currently uses the serial PAFEC VibroAcoustic acoustic software on HP workstations and on a single processor of a multiprocessor Silicon Graphics Power Challenge. The main issues with this configuration arise from the long run times demanded by more complicated coupled acoustic/structural models, in particularly when using boundary elements to model external radiation. Because of this there is little opportunity to run multiple iterations of a model as the solution times and costs become prohibitive. This tends to restrict the use of coupled acoustic modelling to projects which are either trivial, have a long lead time or are well funded. As a consequence smaller customers, due to the high costs involved, seldom use these modelling techniques. To overcome these restrictions the model preparation costs have to be reduced without compromising the applications of the model.

3.5.3 Perceived advantages

The use of a significantly faster code should allow the use of modelling to be extended significantly. Initially the reduced solution times will allow models of a size currently used to run more often, allowing a greater number of design iterations to be considered thus improving the utilisation of existing models. The resultant lead times for the model results to feed into the product engineering process should be significantly reduced.

The second area where benefits should accrue is in the ability to use less optimised models to produce results in a useful time frame. This should allow faster meshing times to be achieved and significantly reduce the costs of generating models. It is this aspect which should appeal to the less frequent users of modelling as the overall costs will reduce although the computational cost is expected to remain similar as the CPU utilisation will remain largely unaffected.

3.5.4 Test case

The above discussion highlights the current issues and aspirations for the coupled structural and acoustic models. The requirements for the test case were therefore drawn up to reflect both the current issues and expected uses of the code should the benefits be as significant as expected. The model size was chosen to be typical of a small automotive component or a limited model of a vehicle and the mesh was generated using methods which were less than optimum, producing a larger model generated in a a significantly reduced time.

The test case chosen was a vehicle intake system consisting of two ducts and a large central air cleaner box. The interior of the system was modelled using 3D acoustic finite elements with the structure modelled using shell and beam elements. The exterior of the structure was covered with a boundary element as the purpose of the model was to calculate the radiated noise, both from the surface and the intake orifice. In general the boundary element patches had one to one equivalence to the surface shell elements as this represents a rapid, though less efficient way, of producing the mesh. The 3D acoustic elements were linear as the meshing is faster for a given geometry and the final model size is reduced.

In total the model consists of :-

3680 3D linear acoustic FE elements

2100 2D quadratic shell elements

2000 Quadratic boundary element patches

The model was run in 2 stages. The first stage was a coupled solution of the structure and it's interior acoustic volume. The second stage was a solution of the BE problem using the structural response from the previous step. This allowed an assessment of the parallel code to be made for the FE and BE methods seperately.

3.5.5 Results

The improvements in solution time for the FE and BE stages are shown in figures 1 and 2 respectively. These are based on the times required to solve the system equations rather than the total run times and so represent only the parallelised section of the process.

As can be seen from fig.1 the speed up for the FE solution was limited to a factor of approximately three, with a maximum efficiency being achieved with 3 processors. A single processor required 5492 seconds to run through the equation solution phase while 3 processors reduced this to 2269 seconds; a reduction of 59%. This phase constituted only 49% of the total run time however, as the initial stages, required to preprocess the model, were also significant. The improvement in the total run time was therefore limited to 22%.

Greater efficiencies were achieved for the BE problem, however, with reductions of 85% when using 8 processors. The most effective solution was achieved with 6 processors and a reduction of 82% in the solution time; on a single processor the equation solution required 22240 seconds reducing to 3900 seconds when 6 processors were used. Unlike the FE model this improvement was carried through to the total run time as the additional, serial processing requirements were minimal. There was also a 28% reduction in the size of the intermediate data storage file, something that did not occur with the FE stage.

For the complete model, where we have both the FE and BE stages, the total run time has reduced from 34000 seconds to 12437 seconds; a reduction of 64%. This is due to the dominance of the BE solution times in the total run time and the large improvement obtained in this stage.

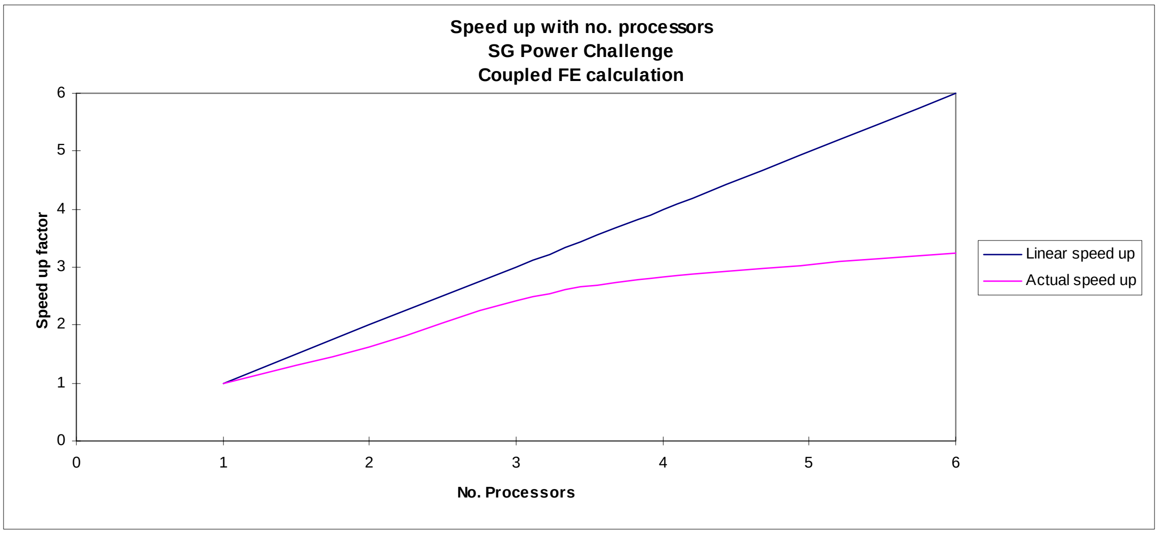

Fig 1

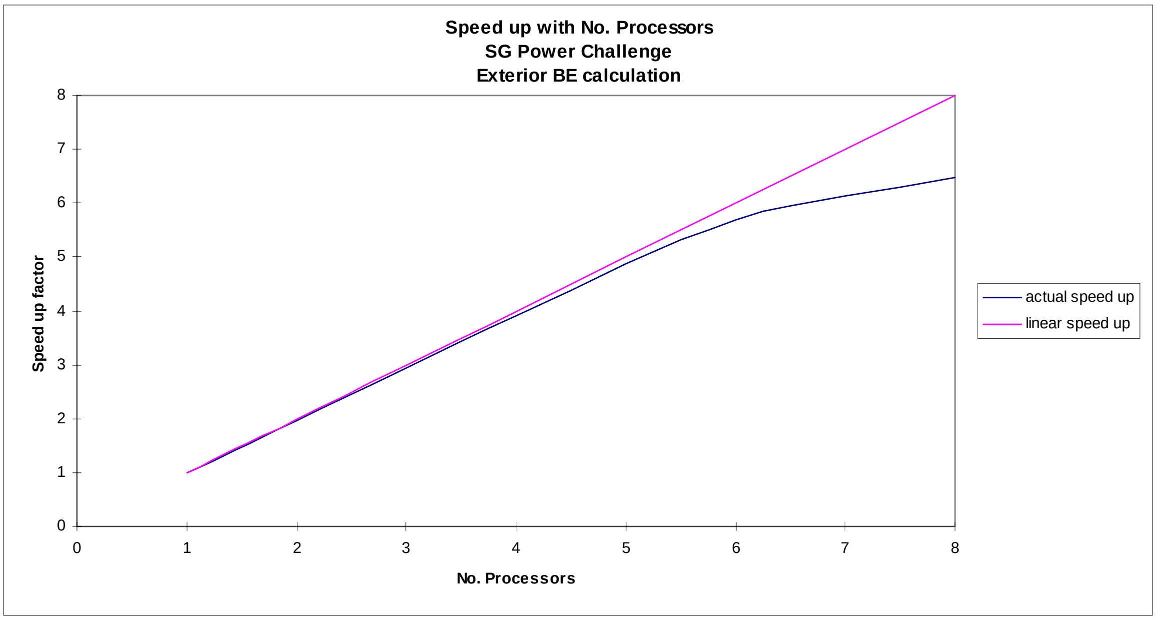

Fig 2

3.5.6 Conclusions

The increase in speed achieved with the parallel code is sufficient to allow large BE models to be a more practical proposition. The increase achieved for the FE model is less significant and, when considering the full run time, there is only a small benefit with the parallel code. Although this could still make the procedure more attractive it is likely to have limited impact on the utilisation of the code.

The benefit to the BE method, however, is large and the method becomes far more attractive. The reduction in solution time suggests that an auto-meshed surface with a higher number of elements is still a practical tool, and the use of more processors can easily compensate for the less efficient mesh.

The overall improvement in the run time of the combined FE/BE model, approximately 65% with 6 processors, means that far larger models can be run within the same time frame as the current models. Thus the use of automeshing becomes acceptable as the run time penalty is removed. This should allow a significant reduction in model preparation costs and allow modelling techniques to used on many projects where it is currently too expensive.

4 Dissemination

4.1 Conference / Journal Publications and the WWW

There are project WWW pages which were administrated by the PAC. They can be found at http://www.pac.soton.ac.uk/pacan-d

The following are a list of publications where results from the PACAN-D project can be found.

- Porting Legacy Engineering Applications onto Distributed NT Systems
  Allsopp_NK, Cooper_TC, Ftakas_P, Macey_PC
  3rd USENIX Windows NT Symposium, July 12-14 Seattle. This paper was accepted for presentation. D.A. Nicole will present the results.

- Parallel Computing on Windows NT Clusters
  Takeda_K, Allsopp_NK, Hardwick_JC, Macey_PC, Caton_MJ, Nicole_DA, Cox_SJ, Lancaster_DJ.
  3rd USENIX Windows NT Symposium, July 12-14 Seattle. This was accepted as a poster.

- Applications of High Performance Computing Network Techniques to Analysis of Submerged Structures.
  Macey_PC, Allsopp_NK, Gill_AS
  Undersea Defence Technology Europe 99, Nice 29th June-1st July 9th Paper accepted as a poster and will be presented by P.C. Macey.

- An Assessment of MPI Environments for Windows NT
  Takeda_K, Allsopp_NK, Hardwick_JC, Macey_PC, Nicole_DA, Cox_SJ, Lancaster_DJ.
  International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'99), 28 June - 1 July 99 Las Vagas. The paper will be presented by S.J. Cox.

- Applications of High Performance Computing Network Techniques to Sonar Design.
  Macey_PC, Allsopp_NK, Gill_AS
  Sonar Transducers 99, Birmingham 19-21st April 1999
  P.C. Macey gave the paper presentation.

- High Performance Computing Network Techniques for Vibroacoustic Analysis: Audio Applications.
  Macey_PC, Wright_JR, Allsopp_NK
  106th Audio Engineering Society Convention, Munich 8-11th May 1999
  P.C. Macey gave the paper presentation.

- High Performance Computing Network Techniques for Vibroacoustic Analysis.
  Macey_PC, Allsopp_NK, Gill_AS
  Proceedings of the Institute of Acoustics Vol 20 Part 5 1998 "Reproduced Sound 14 Windermere". ISBN 1 901656 13 6, ISSN 0309 – 8117. This presentation was given by P.C. Macey.

A success story article has been published within the journal Engineering April 1999 which gives an overview of the PACAN-D project. As a result of this article, one follow-up contact has been made with an industrial partner.

4.2 Automotive Sector Group

A durable full colour poster, laminated in plastic was produced. This will be shown at two automotive sector group events. These are:

- ISATA '99 International Symposium on Automotive Technology and Automation in Vienna 14th – 18th July 1999. ISATA stands for the NETAPDEX are having a stand at this event and the PACAN-D poster will be present.
- EAEC European Automotive Congress "Vehicle systems Technology for the Next Century" 30th June – 2nd July 1999 in Barcelona. CEBRA are having a stand and the PACAN-D poster will be presented.

4.3 NAFEMS Workshop

The PAC held a NAFEMS workshop entitled 'Changing the way we use simulation' where the results of PACAN-D were discussed. The target audience for this event was FE practitioners in industry and FE software developers. Results from the PACAN-D project were presented by P.C. Macey and Simon Roberts.

Full details of the event have not yet been received from NAFEMS, so the following information is provisional only

Number of Attendees	20
Industry/Academia	Mixed
Number of follow-ups	PAC has made 2 follow-up contacts with attendees
Attendee feedback	NAFEMS is co-ordinating this

4.4 Exploitation

PAC and SER Systems have reached an agreement whereby PAC will continue to provide support for the parallel code after the end of the project.

PAC and SER together have made one visit to a customer site during the life of the project, which looks likely to lead to a sale of the parallel code.

5. Conclusions

The PACAN-D project has demonstrated how low cost hardware can be used to run parallel FE/BE analysis. It has demonstrated how the savings in time can be used to leverage simulation into areas of the business where it would not usually be used, and thereby deliver savings and increased performance benefits to industry.

5.1 Simulation for designers

Celestion are committed to a program of making simulation a tool used by designers and not just analysts. Key to this is the use of codes that can deliver answers in a timeframe acceptable to designers, and which take account of all physical processes that are relevant to the problem.

The parallel PAFEC VibroAcoustic code has been demonstrated in the offices at Celestion on a cluster of NT workstations. The code was tested against a problem that, at the start of the project, represented the largest problem that Celestion could run. The code solved the problem in a time scale that allows designers to test a design over lunch, a significant improvement over the current state of affairs.

In addition to this, the memory saving per machine of the code suggests that Celestion will be able to build much larger models than they can at present. This will allow designer to test the impact of their changes much more thoroughly than they can at present by taking into account a wider range of physical processes. In addition they can test their designs within the environment where they will operate.

5.2 Reducing the cost of the simulation process

MIRA is a much larger organisation than Celestion. They already use simulation extensively within their design processes. However, acoustic analysis is still relatively expensive to perform. The time taken to perform an acoustic analysis is much more dependent on the number of elements in the underlying mesh than most other numerical techniques, and for this reason it is important to minimise these or else the execution time becomes prohibitive. The dominant cost in performing an acoustic analysis is therefore the time taken to prepare the mesh. There are methods of rapidly generating meshes, but when these are used the time to execute the problem become quicker. Only by addressing both the execution time and the mesh generation time can MIRA reduce the cost of acoustic analysis.

The parallel PAFEC VibroAcoustic code has been tested on a Silicon Graphics Power Challenge at MIRA. The code was tested against a problem generated using semi-automatic methods that produced a mesh that normally would have taken too long to execute. The code reduced the execution time from 9hours to 3hours in one case and from 6 hours to 1 hour in another case. This saving reduces the total process time of an acoustic analysis from 3 or 4 months down to around 1 month, and means that MIRA can make use of its expertise more widely both within the organisation and in its consultancy services.

5.3 Maintenance of parallel codes

A parallel numerical code typically requires one person to maintain it. For a small software house this can be too expensive, and so a parallel version of a code is not maintained unless a significant number of users require it. The parallelisation implemented in the PAFEC VibroAcoustic code is based upon a limited set of routines within a commercially supported library of parallel routines. The parallelisation is not tuned to any specific architecture, and has been shown to work on both a shared memory machine and a distributed cluster of machines. The method of parallelising the code was to select the numerically most intensive routines and to implement minimal changes throughout the rest of the code. This resulted in a code that was easier for the developers to understand.

The PAFEC VibroAcoustic code supports different types of analysis. As the project progressed, those which were needed by the different test cases were parallelised, allowing the code owners to gain confidence in the parallelisation strategies that were used.

5.4 Exploitation

To enable SER to quickly respond to requests for the parallel code, SER and PAC have negotiated a support deal. This has already resulted in one sale of the parallel code.

6. Contact Details

Company	contact	URL	Role
Parallel Applications Centre	Tim Cooper	http://www.pac.soton.ac.uk/pacan-d	Project Leader
SER	Neil Rigby	http://www.seruk.com	Software Provider
CELESTION	Julian Wright	http://www.celestion.co.uk	End-User
MIRA	Simon Roberts	http://www.mira.co.uk	End-User

The project ran for 24 months starting in May 1997 and ending May 1999. The European Commission contributed to the project 278,000 Euro.