before submitting their proposal to a Research Council. Secondly, applicants who have not used CSAR or HPCx before will be asked to submit a draft case for support, together with code and test data, to both the HPCx and CSAR services. The services will run the test code and will report back to the applicant on the suitability of their service and will provide the technical assessment. As necessary, discussions can follow between the applicant and the centres to determine which service (or possibly services) is best for their research and what resources are required. The applicant will then decide which service they wish to use and finalise their research proposal.

Experienced HPC users, confident of the service they require, will not be required to submit test code to both services, but it is hoped that obtaining a technical assessment of their proposal before submission will open up a dialogue between the service and applicant and allow fine tuning of the proposal.

Parallel Finite Element Analysis

Lee Margetts, Mike Pettipher, Ian Smith
University of Manchester

Why are some of the world's leading experts in parallel finite element analysis (FEA) coming to Manchester this summer? Over the past year, there has been growing interest in Manchester's parallel FEA work. So much so that some of the leading names from America, Europe and Japan are going to be teaching here at a week long Summer School jointly organised by the National Science Foundation (NSF) of America and the University of Manchester.

The story starts around 1995 when Ian Smith of the Manchester School of Engineering first approached Mike Pettipher about parallelising one of the programs from his book 'Programming the Finite Element Method'. By this time, the Domain Decomposition techniques popularised in the late 1980's by Farhat with his 'Greedy Algorithm', had grown in maturity. Despite this, the Domain Decomposition approach to parallel FEA never seemed to hit the mainstream, remaining until this day a largely specialist activity. Perhaps the algorithms, basically centred on matrix condensation or tearing the finite element mesh apart and distributing the pieces, were too difficult to master.

Ian Smith focused his attention on an alternative solution strategy whereby the pieces to be distributed are the finite elements themselves. The 'element-by-element' or 'mesh free' approach can be solved by purely iterative strategies. No mesh is ever assembled and consequently does not require tearing apart. The technique is basically the same as 'explicit' methods which have long been considered 'embarrassingly' parallel.

In 1998, EPSRC agreed to fund a joint research project between Manchester Computing and the Manchester School of Engineering. At this time, Lee Margetts joined the team to study for his PhD under the supervision of Ian Smith. The objective was to develop a parallelisation strategy that could easily be applied by a non-specialist to any general finite element problem.

By the end of 2002, the parallelisation strategy first implemented by Mike Pettipher had been successfully generalised and all the MPI coding was hidden away into a library of FORTRAN callable subroutines. This library was used to create a suite of ten example programs covering the three main types of problem found in Engineering: Static equilibrium; dynamics (or time dependent problems) and eigenvalues.

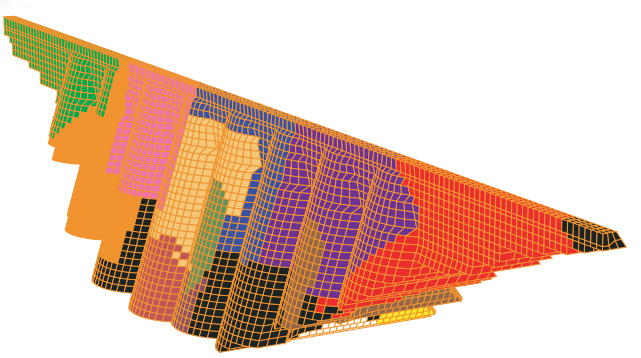

Figure 1: Finite element mesh of a buttress dam used to evaluate the structural response to earthquakes. The mesh is distributed across processors as illustrated by the colouring.

In March this year, a UKHEC workshop introduced the work to non-specialists in parallel computing. As the parallel programs retain the same structure and style of the original serial programs taken from Ian Smith's book, a researcher already familiar with serial FORTRAN programming and finite element analysis should be able to develop his or her own parallel codes. This was the philosophy behind the teaching and practical sessions at the workshop.

An article about parallel computing would not be complete without some performance figures! All the programs scale extremely well over large numbers of processors and for the results presented here, CSAR's 512 processor Origin3000, Green, was used. To give an impression of the scalability, our program for the direct numerical solution of the Navier Stokes equations boasts a speed up of 256 on 256 processors, whilst sustaining an impressive 30% of the machine's peak performance (see table 1). Similar performance is achieved for an elastoplasticity problem as illustrated in table 2. In this case, efficient use of up to 500 processors is clearly demonstrated. To name another example, an 8,000,000 equation eigenvalue analysis recently run on 256 processors found the first 100 eigenvectors in 460 seconds. Programs written to solve problems in heat conduction, dynamical systems and coupled processes such as magnetohydrodynamics show similar results.

Reynolds Number	256 processors	Serial	% Peak
10	20 minutes	2-3 days	29
100	47 minutes	8-9 days	29
1000	180 minutes	> 1 month	29

Table 1: Direct Numerical Solution of the Navier-Stokes equations – 4,500,000 equations

Processors	PCG(s)	MFlops	% Peak
4	2786	999	31
256	43.4	64190	31
500	23.9	116500	29

Table 2: Elastoplasticity – 6,000,000 unknowns

Why is this method so successful? To cover the three problem types mentioned earlier, three different iterative solvers are used: PCG (preconditioned conjugant gradients) for symmetric positive definite problems; BiCGStab(l) for non-symmetric systems and Lanczos for eigenvalue problems. These all have at their core matrix-vector products and other simple vector operations. Most of the communication and computation may be overlapped. The only possible downside is the unavoidable presence of a few global communications, such as dot products, which on some systems are known to limit scalability. If the problems are large enough, the communication to computation ratio is so low that the technique also works well on networks of PCs connected by Ethernet. This has been recently demonstrated by some collaborative work with the Civil Engineering department at the Universidad Politecnica de Madrid.

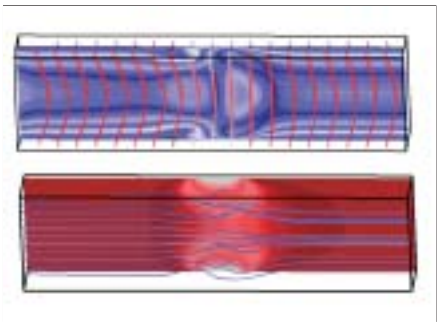

Figure 2: An example of magnetohydrodynamics: The flow of an electrically conducting fluid through an insulated rectangular duct under the influence of an externally applied magnetic field.

Although the parallel programs enable the solution of 3D problems with millions of degrees of freedom, understanding the results can itself present further challenges. Collaboration between the project and staff of the Manchester Visualisation Centre (MVC), in particular Joanna Leng, enabled the development of a powerful and convenient way of presenting and interpreting the results of the simulations. Using the AVS Express toolkit, Joanna produced an application that allows the user to view and manipulate the results of the simulations in stereo, in a virtual immersive environment such as an SGI Reality Centre. To interactively manipulate the largest data sets, ~ 10,000,000 variables, the multi-pipe edition of AVS developed by MVC was required. At Manchester this was run on an SGI Onyx300 with 6 dedicated graphics pipes.

These are exciting times. With the efficient use of powerful HPC resources and advanced visualization tools, scientists and engineers are not only able to investigate more complex models and systems, they are also able to explore their models intuitively and collaboratively through virtual reality visualization. In the future, we will highlight the merits of parallel FEA from another perspective. The same solution strategies can be applied to solve smaller problems very quickly – so fast that the engineer or scientist may be able to interact with their model in real time. This is the aim of the Advanced Virtual Prototyping Research Centre's Virtual Prototyping project in which Manchester are contributing their parallel finite element expertise.

Finally, if you would like to know more, please feel free to contact the authors or better still, enrol on the Summer School - it's free to CSAR users!

This paper presents parallel 3-D finite element analysis for distributed memory multiprocessors. Traditionally, finite element analysis has been performed on sequential computers. Current research in high performance finite element analysis shows considerable promise for fast, efficient implementation on MIMD and SIMD computers. This paper demonstrates the use of a standard, banded Cholesky method for solving the finite element system of equations. The uniformity of the underlying data distribution ensures high performance due to load balance.