



Visualizing Large Parallel Finite Element Simulations

J Leng¹, L Margetts²

¹ *Manchester Visualization Centre, University of Manchester*

² *Manchester Computing, University of Manchester
(currently at the Universidad Politecnica de Madrid, Spain)*

December 2002

Synopsis

The Centre for Civil and Construction Engineering, UMIST and Manchester Computing have been collaborating on a project to develop a suite of parallel finite element programs for solving a range of large engineering problems. With huge data sets, visualizing and interpreting the results of the simulations required some extensible software.

AVS/Express and VTK were likely candidates. However, only AVS/Express supported the specialist data types being produced. Even so, the data could not be easily imported. A number of visualization modules and applications were developed to make the visual analysis of these FE engineering problems possible. This report also draws attention to the need for High End visualization capability.

1. Introduction

Traditionally, engineering has been considered an established user of High End Computing facilities. However, it is generally not noted that the majority of these engineers work in the aerospace or fluid mechanics sector. Other subdisciplines such as Civil Engineering and Structural Engineering rarely use such facilities. The reader can get some idea of the type of user by reviewing the consortia currently using UK HPC resources.

Many engineering researchers use finite element models as a research tool, mainly carrying out analyses serially on single processor machines. In the UK, parallel computing facilities exist that allow larger, more complex analyses to be solved. In order to use these facilities, there is one major hurdle – there are no commercially available parallel finite element packages so users must develop a parallel program themselves. This is a difficult and time consuming task.

In recent years, the University of Manchester has been funded by EPSRC to investigate a new method of parallelizing finite element programs, using iterative solvers and element

by element or mesh free techniques [1]. This work has led to the development of a library of routines, callable from FORTRAN90 that can be used relatively easily to convert a serial program to run in parallel. These routines have been used to parallelize a suite of programs [2] covering 10 different algorithms taken from Smith and Griffiths [3], a well known teaching text for writing finite element programs. These cover a range of problem types including static and dynamic material behavior in structural mechanics and geomechanics, as well as fluid and heat flow. These programs have now been successfully converted for parallel execution, whilst retaining the style and clarity of the original serial programs. The codes have proven to be scalable up to at least 256 processors as will be demonstrated later in the text.

The parallel finite element project focused on the equation solution and parallel algorithms. No time was earmarked for visualizing and processing the results. With the ability to solve large analyses in a reasonable time, the next task was to find a way of validating and interpreting such large models. That is the subject of this case study, which describes the work carried out at Manchester Visualization Centre (MVC) for UKHEC.

MVC has supported users of graphics and visualization since the mid 70's. Users now tend to have larger 3D datasets that require more intensive analysis. These types of data often benefit from custom made tools for the visual analysis and are more efficiently handled in a virtual environment. These two requirements of users have driven research at MVC into two independent but related areas:

- Work with scientist in their domain area to produce appropriate toolkits.
- Develop interfaces with virtual environments so that toolkits can benefit from more intuitive user interfaces.

With parallel finite element analysis at their grasp, more complicated models can be investigated involving three dimensions, time dependency and the coupling of different physical processes. This has opened up a number of areas of research outside the engineering domain where a collaborative research effort is needed:

- Visualization/VR tools are needed to validate/'debug' such codes.
- Visualization/VR tools are needed to help explore and understand the physics in the FEM.
- Visualization/VR tools are needed in the engineering design process.

**In this report, we first describe the early stages of the work, the development of a finite element data reader. This is followed by an outline of the capability and merits of using AVS Express. In the remainder of the report, several examples are described which show the type of visualization work carried out so far. These include a fluid dynamics test case, a magnetohydrodynamic (MHD) problem, the use of stochastics in finite element

problems and a straightforward geotechnical analysis. Particular aspects of the work are highlighted where ‘High End’ capability such as immersive stereoscopic projection is arguably necessary. In other cases, simple modules are described which had to be developed by the authors, adding to the functionality of the visualization software.

1.1. Visualization Issues

Visualization shows features in a data set. It is easy to show the expected ones but more difficult to show the unexpected. To find the unexpected features a large number of interactive and possibly application specific tools are needed. This not only requires the use of high-end interactive visualization equipment but also the application of appropriate visualization techniques.

1.2. Application Toolkits

AVS/Express is a modular visualization system with a high level GUI with its own visual programming environment. Visualization systems generally have a high learning curve for new users. A new user will probably have to understand at least some components of all these subjects: the domain science; principals of visualization; specific visualization algorithms; graphic design; elements of cognition; and hardware configuration. The high level GUI is designed to reduce the learning curve. It allows users to experiment with many different settings and see the results instantly.

Advanced users of AVS/Express can use the available high quality libraries as well as developing their own libraries of modules. All modules whether native or not are accessed through the GUI. New libraries build up toolkits of functions required by specific users. When necessary these can be integrated into a stand alone application where the user is presented with a user interface especially developed for them.

These features of AVS/Express make it suitable for use by research groups from various domains including engineering. There is one basic tool but it can be customized into a number of tools which can then be used depending on the skills of the user and the work they are doing.

1.3. Interfaces to Virtual Environments

Common virtual environments include workbenches, such as ImmersaDesks, three-wall large screen environments such as RealityCenters and immersive environments such as CAVEs. A limitation of AVS/Express is that it was not designed for use in a virtual environment. It does provide active stereo support but this was mainly used on the desktop. There is no support for the multi-projector display environments often found in virtual environments. The AVS/Express Multipipe Edition (MPE) [4] project has extended the rendering engine within AVS/Express so that it can be used in these virtual environments.

The AVS/Express renderer has been extended using the SGI Multipipe SDK software. This allows MPE to take advantage of the multiple graphics pipes found in high end visualization supercomputers. The output from each graphics pipe can then be projected

in the display environment. The Multipipe Edition allows the user to be immersed in and interact with their visualization.

1.4 A General FE Reader

No visualization is possible if the data can not be read into the visualization system. The FE data is in a specialist format and is difficult to import into any turn key system.

The first stage in the work was the development of a specialist data reader which would enable FE data to be read into AVS/Express. The reader is designed to read AVS/Express cell data using a format similar to the AVS/Express field file descriptor. A simple text file is used to describe the data: its dimensions, number of nodes, data type, whether it is scalar or how many vector components it has and number and type of cells. After this, at the end of the file, pointers to the data files and position of the data are given.

The benefit of this type of reader is that the text descriptor file is easy to write and edit. The reader can read coordinates, cell information and node data from the files the engineers use. No new large data files need to be written or stored. Displacement data, important to this group of users, is a special case of vector data that can be marked and handled specially within AVS/Express. The format can be extended as the engineers require further functionality. Developing this reader was the key to the engineering case studies discussed here.

2. Engineering Case Studies

In the following section, four examples are presented. The first, the lid-driven cavity problem required visualization expertise to help validate the code through visualizing the results. The second involves a more complex physical problem, magnetohydrodynamics (MHD). In this case, visualization was used as an investigative tool to understand the complex physics being simulated. In the third case, involving stochastic finite element analysis, problem specific tools have been written to help the engineers at UMIST validate new codes that are under development. Finally, the visualization of a geotechnical problem, the Cern Cavern excavation is described.

2.1. The Lid-driven Cavity

The lid-driven cavity problem is a well documented test case for computational fluid dynamics (CFD) algorithms. A cubic cavity contains a fluid that is initially at rest. The top surface of the cavity or 'lid' is driven at a constant velocity. A steady state solution is then sought for the motion of the fluid inside the cavity.

The domain was subdivided into a quarter of a million finite elements, giving rise to 1 million grid points where values for the pressure and velocity field were to be calculated. With approximately 4 unknowns at each grid point, the computational task was to solve a system of 4 million non-linear simultaneous equations.

There are various ways of solving CFD problems that typically involve numerical techniques such as finite differences or finite elements. Various formulations are in use

that simplifies the Navier Stokes equations by taking mathematical short cuts. For simple test problems, these methods produce similar answers with reasonable computation times. However, for complicated geometries, these simplified methods may not give the correct answer or give any answer at all.

The closest we can get to the correct solution is to solve the full system of simultaneous equations with no simplifications. This is often referred to as direct numerical simulation or DNS. Being computationally very expensive, DNS is not typically used for everyday CFD problems. However, solving large complicated problems using DNS benefits the developers of simplified industrial algorithms by providing accurate solutions for validation work.

Table 1 shows the performance data for the problem described, which was analyzed using different values of Reynolds Number. The parallel time is for 256 processors, and the serial time is estimated from the time per iteration. It should be noted not only that the parallel version runs 256 times faster than the serial, but also that the actual performance relative to the peak performance of the processors is also very good.

The Reynolds Number is a dimensionless parameter that characterizes the flow. In simple terms, one can imagine a low Reynolds flow being laminar, like treacle and a very high Reynolds flow being turbulent, like smoke rising from a cigarette. As the Reynolds number increases, the computational cost increases and at a certain point a solution is unobtainable. Increasing the resolution of the analysis (number of grid points) enables the computation to proceed to a higher Reynolds number. To resolve fully turbulent flows using DNS, it has been suggested that around 1 billion computational grid points would be needed. We're a long way off at 1 million!

Problem	Reynolds Number	Parallel Time Mins	Serial Time Days	%Peak Performance	Gflops
Half Cubic Cavity10	20	20	2-3	29	59
	100	47	8-9	29	59
	1000	180	>1 month	29	59

Table 1. Performance data versus Reynolds number

As the Navier Stokes test problem shows, parallel computation can lead to significant improvements in simulation time. This not only increases the complexity of engineering model that can be investigated, it also means that more sophisticated tools are required to interpret the results.

The initial visualization application was developed using a low resolution data set (roughly 5000 grid points) on an SGI O2 machine. When the visualization functionality was defined and the application developed the high resolution data could be examined. The high resolution data had over 200 times the number of grid points. The application would not run at all on the O2 but was moved to a 32-processor (MIPS R14000) SGI

Onyx 300 visualization system with 6 graphics pipes. Initially the application was run across the network which meant the rendering was done locally on an O2. Some simple exploration could be done but it was difficult to explore the data for example it was difficult rotating the model, moving the orthoslices and probes.

The final visualization application used a texture based rendering method to show the complexity of the flow profile through the data. Cut planes were placed orthogonal to each of the principal axes of the data, although these could be moved by a user who wished to explore the flow profile. The magnitude of the data flow was assigned a texture coordinate across the diagonal of an image, in this case a blue shaded mandrill. The resulting image does not show the strength of the flow at any point but shows how it varies between points.

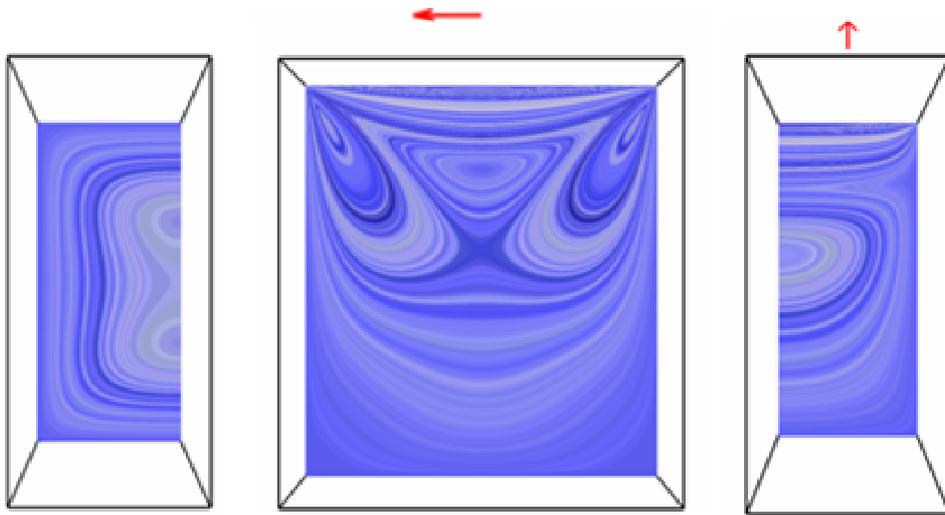


Figure 2. Flow profile for the central orthogonal slice for each of the three principal axes

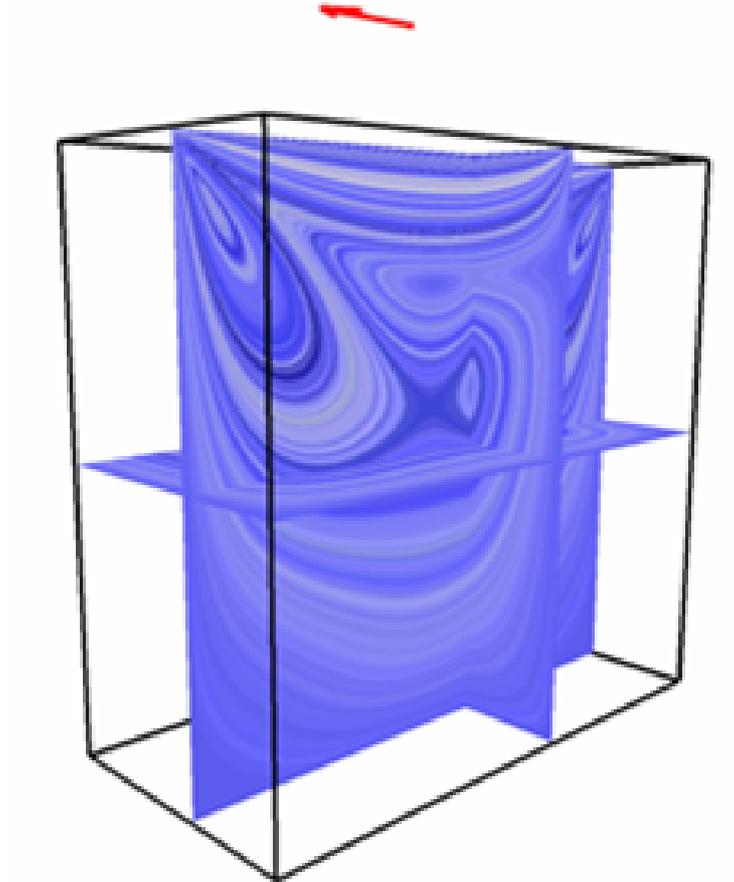


Figure 3. Flow profile of the three orthogonal slices in figure 2 but shown together.

The flow profile does not show where or how fast particles would move in this data. A probe is used to select a point or series of points within the data. These points are used as seed points which show what the movement of particles released from these points would be. The flow paths are shown as streamlines, the speed of the flow is shown as colour. The flow is fast where it is coloured red, against the lid. The flow is slow where the colour is blue, away from the lid. Flow is a time dependent phenomenon and as is nicely demonstrated by a well known animated technique called particle advection. An arrow is used to show in which direction the lid of the box is moving.

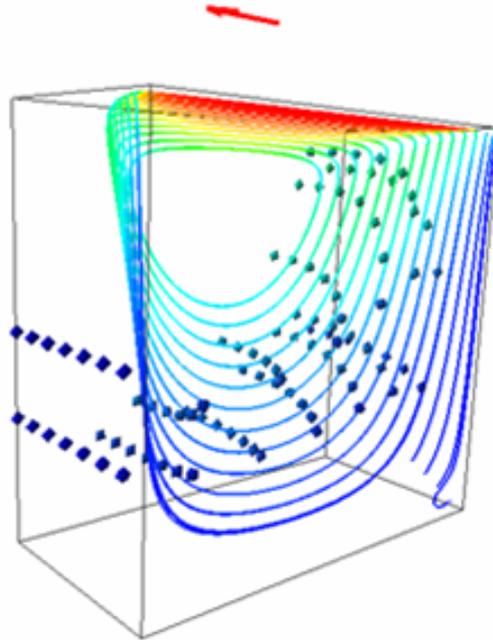


Figure 4. Streamlines and practical advection show the flow in 3D.

These images were viewed in Manchester's Virtual Immersive Projection Laboratory (VIPL) using the MPE version of AVS/Express - taking full advantage of the 'High End' graphics hardware. By moving into this immersive virtual reality environment it was possible to see the results in stereo and work easily with others. Data that had been difficult to validate was validated to the satisfaction of the engineer and associates within hours.

2.2. Magnetohydrodynamics

The parallel algorithms used to model the Navier Stokes equations were modified to simulate a more complicated flow problem [5] – the flow of an electrically conducting fluid in the presence of a magnetic field. In MHD flows, the two distinct physical processes, electromagnetism and fluid flow, are coupled. The fluid flow affects the magnetic field and the magnetic field affects the fluid flow.

A simple test problem is described here where a fluid flows into a duct with a prescribed parabolic velocity profile and uniform pressure applied across the inlet face. A uniform pressure gradient is applied at the walls, along the length of the duct. The interaction of the two physical processes is then investigated by visualizing the results of a parallel FE analysis. There is no simple analytical solution to this problem and the validation of the model can only be undertaken by comparing the features of the simulation with physical experiments.

Interpreting the results of a simulation where two or more physical processes interact is particularly difficult. The engineers were originally using conventional visualization tools that output to a two-dimensional medium – a computer screen or hardcopy. This was very limited and the physical processes were virtually impossible to understand.

This simulation used a very similar visualization application to the one used for the lid-driven cavity simulation. However, in order to make it possible to show both fields (magnetic and fluid flow) together, the previously described visualization application needed to be extended. Figure 5 shows a typical view given by the extended application. There are two instances of the duct. In the first, the magnetic field profile is shown together with streamlines indicating the fluid flow. Above this, the second instance of the duct shows the fluid flow profile with magnetic field lines. Colour was used to convey information about the physical processes. In both instances of the duct, the fluid flow was depicted using a blue colour scheme and the magnetic field was represented by red.

In the lower instance of the duct in figure 5, the position of the externally applied magnetic field is picked out by the distortion of the flow lines which twist both in the plane of the paper and perpendicular to it. This is impossible to make out visualizing in only two dimensions. To be able to identify and examine this feature, stereoscopic projection in an immersive environment proved an invaluable, if not necessary, analytical tool.

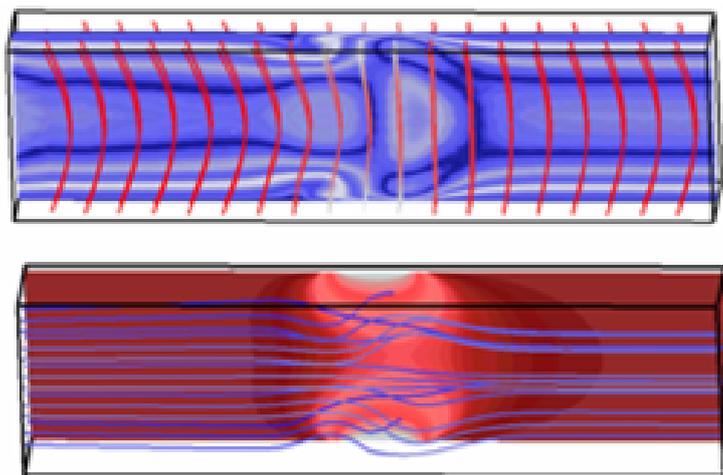


Figure 5. MHD duct flow

2.3. Validating Stochastic Finite Element Models

In geotechnical engineering, contemporary design is often deterministic with prescriptive engineering design codes providing specific values for certain engineering parameters. In a deterministic design, characteristic values of material properties are determined for the appropriate soil layers. These characteristic values are conservative and do not take into account the natural variability of the soil. With prescribed engineering parameters and conservative material properties, a design calculation will test whether a limit state is satisfied, i.e. failure of the structure will or will not occur.

In commercial finite element packages, simple models of material behaviour are typically used and the materials are assumed to be homogeneous. Soil properties are far from

homogeneous and can have a wide range of possible structures and physical properties. They may comprise sands or clays, which are layered or massive, compacted or loose, saturated or dry. These properties will vary laterally and with depth and perhaps also seasonally.

In 1994, the first draft of a European design code, ‘Eurocode 7 – Geotechnical Design’ was published [6]. In the codes, provision was made for engineers to determine their own values for certain engineering parameters so long as their use could be justified in terms of reliability. To facilitate this methodology, new concepts such as evaluating the probability of failure were introduced. With this flexibility, some engineers began to investigate the use of stochastic methods in engineering design. For example, Hicks and Samy [7] have been using stochastics in finite element analysis to model material heterogeneity in a number of simple geotechnical engineering scenarios, comparing their work with deterministic solutions.

Stochastic geotechnical design involves undertaking Monte Carlo type simulations in which the same finite element analysis is repeated over and over again using a different set of material properties for each analysis or realisation until a statistically stable result is achieved. A stochastic random field generator is used to generate the values of the material properties for each realisation. The material values generated have a random statistical distribution, which can be based on the horizontal and vertical variation in material property determined in the field or laboratory. Two examples of random fields are shown in Figure 6 [8]. Each finite element has a single material value which is plotted using a single colour or greyscale shade. The granularity of the shading represents the fineness of the finite element mesh.

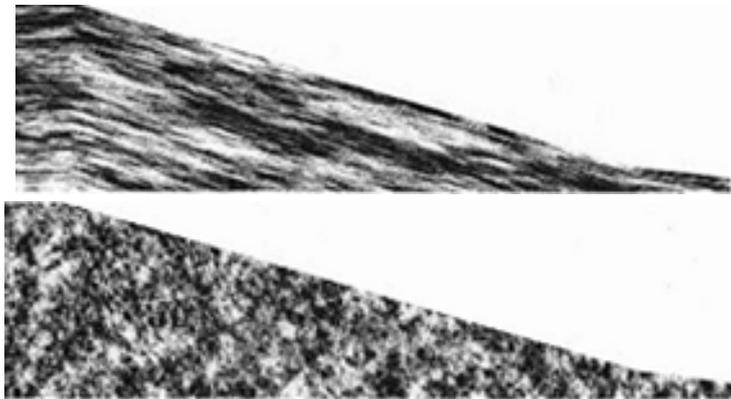


Figure 6. Material properties generated using a stochastic random field

Current research focuses on two dimensional analyses. Although the computational times are lengthy using serial computers, this is acceptable in the context of a research environment. However, in three dimensions the computational demands for stochastic finite element modeling will prove prohibitive unless a parallel strategy is adopted. If the memory requirements do not limit the size of an analysis that can be undertaken, time constraints will.

Parallel computing not only increases the size and complexity of the problems that can be solved, it also opens the door to new difficulties, particularly in visualization and interpretation. For example, the two dimensional stochastic random field data is validated qualitatively using visualization. In Figure 6, the values of the material properties are plotted using a discrete grayscale colour map and from the visualization, the analyst can quickly check whether the random field is as they expect, isotropic or anisotropic for example.

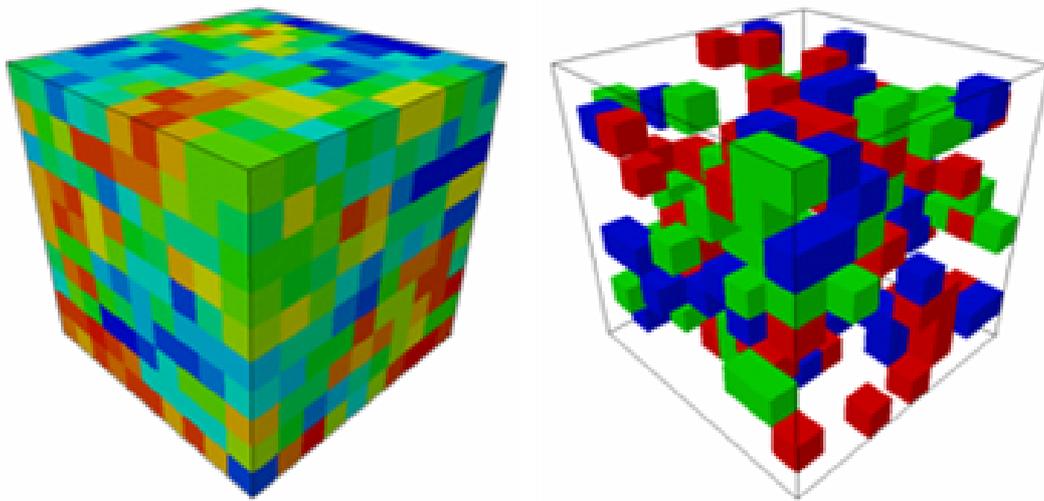


Figure 7. Visualization of a 3D random field

The interpretation of three dimensional random field data is considerably more difficult and validation requires visualizing the internal structure of the field. A preliminary study into visualizing this type of data was investigated as part of this case study. The same colour map idea is used for the three dimensional data. To view the internal structure, a thresholding module was written which allows the user to set a range of values for which finite elements are rendered visible or invisible. An example 3D field and a visualization using the thresholding module are shown in Figure 7. A variety of library modules are available for AVS Express such as transparency and immersive stereoscopic projection, that can further aid the visualization of this type of data.

2.4. Cern Cavern Excavation

The final example concerns a real geotechnical problem, the excavation of a new underground cavern for the CERN particle accelerator. A complex of tunnels and other structures is centered on the UX15 cavern which is 50m below ground and approximately 34m span, 60m long and 40m deep. Two large shafts, 16m and 9m diameter respectively, connect the cavern roof vault to the surface buildings, and a number of smaller caverns and tunnels lead off the cavern walls. The excavation of the UX15 cavern had to be carried out while another project was still running and that project was

not to be interrupted or influenced by excavation. A detailed 3-dimensional finite element analysis was used to predict the ground movements during excavation.

The finite element domain is a continuous solid 3D model in which the existing structures are holes and the surrounding soil is a deformable medium. A view of the soil displacement, using arrows, is shown in Figure 8. The position of the cavern is shown in grey. Of interest to the analyst is the movement of the existing structure.

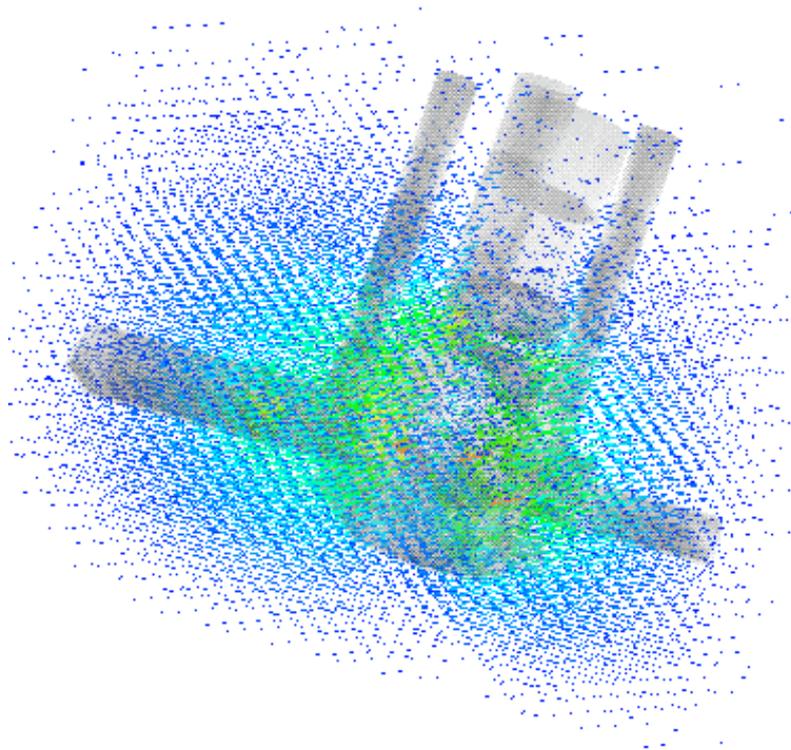


Figure 8. Entire dataset of displacement vectors

Although the finite element analysis must compute the displacement for the entire domain, visualizing all the available data obscures the actual information that is wanted. All the engineer wishes to see is the displacement of the actual structure. In order to do this, a filtering module was written. With the dataset encoded into different subsets, some representing the soil and others representing the underground structure, the filtering module could remove unwanted information as shown in Figures 9 and 10.

One of the items on the engineers' wish list was the ability to be immersed in the visualization and fly through the tunnel, viewing the displacement contours from the inside of the structure. This is only now possible with the latest version of AVS/Express 6.1.

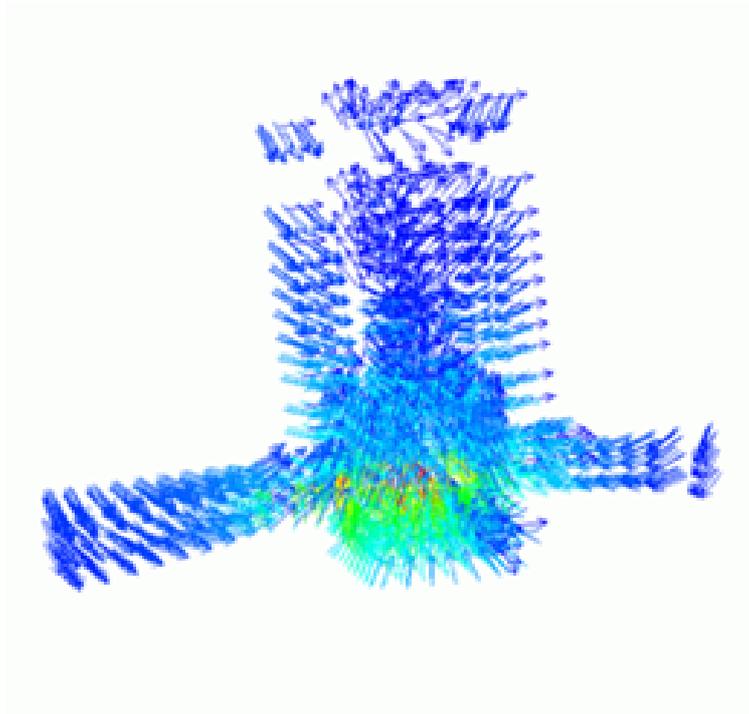


Figure 9. Filtered dataset showing displacement vectors on the structure

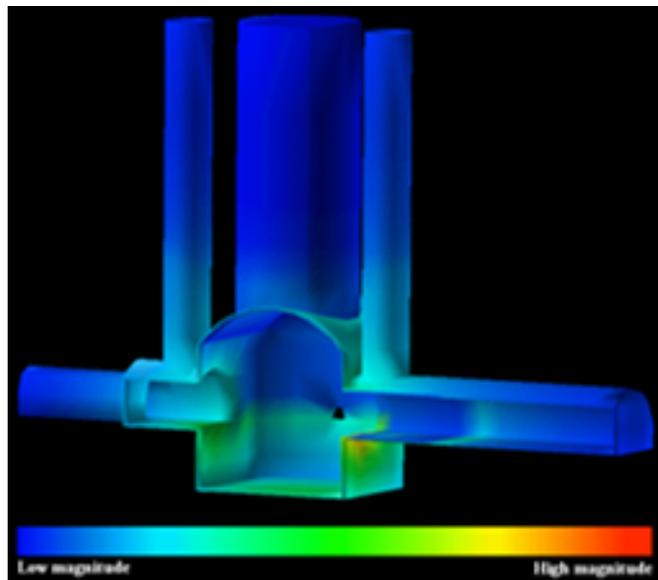


Figure 10. Filtered dataset showing displacement contours on the surface of the structure

3. Conclusions and Future Work

We have shown that immersive visualization has increased the understanding of these complex engineering models. The comments received from end users have been very favourable and they are keen to exploit this approach further. Although we have used MVC's immersive projection theatre, the applications developed could also be used without modification in any other immersive environment such as a CAVE or workbench.

Further visualization modules are being added to AVS/Express. In particular modules suitable for geotechnical engineering are being developed. New 3D models are being developed that require validation and “debugging”. A future aim of the engineers at UMIST is to model time dependent dispersal of a pollutant through a three dimensional geological domain, visualizing the results as the analysis proceeds. This too will require new visualization functionality.

The key point is that with the more complex 3D models that can be analyzed in a reasonable time through parallel computing, the development of visualization tools and the provision of ‘High End’ visualization facilities are becoming increasingly important in the wider discipline of engineering.

This work is being communicated to the wider academic community through a series of events based in Manchester, such as a UKHEC workshop in Parallel Finite Element Analysis and Visualization in March 2003 and an International Summer School following the same themes in September 2003.

4. References

- [1] L. Margetts, “Parallel Finite Element Analysis”, PhD Thesis, University of Manchester, 2002.
- [2] I.M. Smith and L. Margetts, “Portable Parallel Processing for Non-Linear Problems”, to appear in VII International Conference on Computational Plasticity, Barcelona, 2003.
- [3] I.M. Smith and D.V.Griffiths, “Programming the Finite Element Method”, Third Edition, Wiley, 1998
- [4] P.G. Lever, G.W. Leaver, I. Curington, J.S. Perrin, A. Dodd, N.W. John, W.T. Hewitt, "Design Issues in the AVS/Express Multi-Pipe Edition" IEEE Visualization 2000 Conference Works In Progress, Oct. 2000
- [5] I.M. Smith, L. Margetts, “Parallel Finite Element Analysis of Coupled Problems, Numerical Methods in Geomechanics”, Editors Pande G.N., Pietuszcak S., Balkema 2002.
- [6] ENV 1997-1. “Eurocode 7: Geotechnical design – Part 1: General rules” CEN, European Committee for Standardisation, 1994
- [7] M.A. Hicks and K. Samy, K. “Influence of heterogeneity on undrained clay slope stability” Quarterly Journal of Engineering Geology and Hydrology, 39, 41-49, 2002.
- [8] C. Onisiphorou. “Stochastic analysis of saturated soils using finite elements”. PhD thesis, University of Manchester, UK, 2000.