PERMAS-D

ESPRIT Project Number 24472

PAC

INTES GmbH

Dynalis

Structural Engineering A.S.

Public Final Report of ESPIRT HPCN PST Activity PERMAS-D, Parallel PERMAS Deployment.

PAC/PERMAS-D/FR4 Version 1

N. K. Allsopp, Paul Gordon
Parallel Applications Centre

24 May 1999

Abstract

This report has two main aims:

- to demonstrate the PERMAS FE code running effectively on shared memory (SMP) platforms with test cases of real concern to industry;
- to identify a migration strategy for potential users of parallel PERMAS to facilitate their take-up.

We will show that SME's will benefit from using Parallel PERMAS on a Shared Memory System. Significantly improved performance on a low or medium cost hardware platform will lead to different design and numerical simulations strategies. This ultimately leads to the market profile of the company being raised as they will be seen as on the leading edge of technology take-up.

Synopsis

This report is the final report of the PERMAS-D Project (ESPRIT no 24472). The project involves the following partners: PAC (UK), INTES GmbH (Germany), Structural Engineering A.S. (Norway) and Dynalis (France). The project was co-ordinated by PAC, and further information may be obtained from

Paul Gordon,
Tel: +44 (0)1703 760834,
Fax: +44 (0)1703 760833
E-mail to: plg@pac.soton.ac.uk

End users of the project, Structural Engineering and Dynalis, are facing an ever-increasing growth in the size of the finite element models they have to use. They are also being asked to perform more analyses in a shorter time. One solution is the purchase of a shared memory machine (SMP).

The addition of the SMP machine not only frees up desktop machines for interactive use by engineers but also reduces administration overheads by centralising back up procedures and licensing issues. Being able to run simulations in hours rather than days increases the transparency of the design process by allowing more client interactivity.

This report details how the costs and benefits of using the PERMAS finite element software on an SMP platform were assessed, analysed and demonstrated.

Executive Summary

The business objective of this project was to transfer a proven HPCN application technology into engineering consulting businesses, through deployment on low-cost platforms. This provided a model for other companies to follow.

To meet these objectives it was necessary to achieve the following technical milestones:

- specification of benchmark case studies from the two end-user partners;
- demonstration of performance levels needed to realise business benefits;
- analysis of cost-benefits across the range of relevant low-cost HPCN platforms;
- dissemination based on these results and lessons learned, to encourage replication of HPCN technology deployment by other consulting end-users.

The demonstrators addressed typical end-user problems from Dynalis and Structural Engineering AS (SE). Both companies currently use HP, DEC and SGI workstations, running problems such as dynamic modal analysis of ship structures, and car components analysis. These problems possess up to 200,000 degrees of freedom and take 10 to 20 hours to run. The goal of Dynalis was to assess both shared memory and cluster approaches with a view to reducing the run-time for such problems by a factor of 2-3 through the use of low-cost HPCN systems.

DYNALIS is doing engineering work in different industrial branches. Being a sub contracting party, the response time and the reactivity demanded by clients is very high. Very often, the initial design and scaling has to be improved and optimised. Moreover, the higher and higher accuracy required pushes to perform sensitivities to the analysis in order to take into account imperfections and/or uncertainties.

So, in the shortest possible time, several analyses have to be performed on a generic model. Moreover, in the car industry, finite element calculations are used for static and dynamic calculations and for crash analysis. The mesh used for crash analysis (explicit method) is generally finer than meshes used for implicit methods. Considering the modelisation time, using the same model for implicit and explicit methods will save significantly manpower, but increase computer time.

The main activities at SE are the design or verification of offshore structures and the development of procedures for marine operations. In this context SE uses the finite element method to calculate the behaviour of the structures subjected to different loads. The results of the finite element analysis are used as input to in-house written post-processors to check if the design fulfils the necessary requirements (rules).

The definition of the orthotropic material properties and the transfer of the hydrodynamic loads to the finite element model requires a large effort in terms of manpower to create and to check the model and is prone to error. The correction of the resulting stresses after the analysis is also prone to error. Generating three different models (with some thousand unknowns in the overlapping regions) is not especially effective.

SE wanted to improve the analysis procedures in different ways :

- use the real plate thickness, real stiffener cross section properties and real stiffener spacing, not orthotropic material properties,
- make only one model consisting of the whole ship,

The new strategy led to less effort in terms of manpower and fewer errors but to very large models. This was not possible at the start of the project because of limitations in the available computer power of SE. It was expected that parallel computing could definitely change the picture and have a high impact on the cost, the project time and the reliability of the results.

When an SME decides to upgrade to HPCN, there are two alternatives, namely shared memory system or a cluster (of legacy or low cost workstation). Which option is the best for a particular company depends on its core business and day-to-day method of operation. A company that relies on computational resources as its core business, i.e. SE or Dynalis, will probably find that SMP is the desired route, while a company that uses computation as a support tool for the designers will find that the cluster approach is more desirable.

Shared memory systems are medium cost (22000 to 45000 Euro) workstations that are capable of high performance computation. If the answer to the following question: "Are my engineers using their workstations to full capacity with no CPU cycles available for other applications?" is 'Yes' then the suggested route is to move to an SMP platform.

Low cost systems (e.g. Intel PC's or legacy workstations) operating in a cluster can be an effective way of increasing the computational resources of the company by using spare CPU cycles on idle machines. If the answer to the following question: "Are many of the machines in my organisation not being used at full capacity (e.g. for report writing or office tools)?" is 'Yes' then the suggested route is to move to an cluster solution.

Contents

1 Introduction

The main results of the project are the identification of a viable HPC migration strategy for small users of the parallel PERMAS code, a slightly tuned parallel PERMAS code suited to this hardware migration strategy, plus case study experiences and dissemination material.

The code owner INTES are planning to exploit the enhanced software and the other project results in their marketing of the parallel PERMAS product, especially to small customers similar to Dynalis and SE. The end-users will exploit the project results to enhance their design consultancy businesses, and to guide their future computing hardware procurement strategy.

This project will provide a valuable model for the adoption and maintenance of HPC technology and exploitation of its capabilities by small consultancy businesses, and other small manufacturers such as component suppliers. The PAC will exploit this to target promotion of HPCN technology transfer over the next 2-3 years on these companies, most of whom have yet to exploit HPCN capabilities.

2 Approach taken

The two main areas covered within this project were the development of the parallel PERMAS code onto a shared memory platform and the identification of a viable migration strategy for the SME end users.

These two areas are described in the following sections.

2.1 Parallelism of PERMAS on SMP Platform

The main areas for software enhancement were to be:

- selection of load distribution algorithms to suit the low-cost, shared memory system architectures; and
- incorporation of improved, partially asynchronous I/O from the sequential code into the parallel code.

The best load distribution methods for shared-memory systems are likely to be similar to those used on scaleable parallel systems, subject to the constraint that the algorithms can be specialised to small numbers of processors.

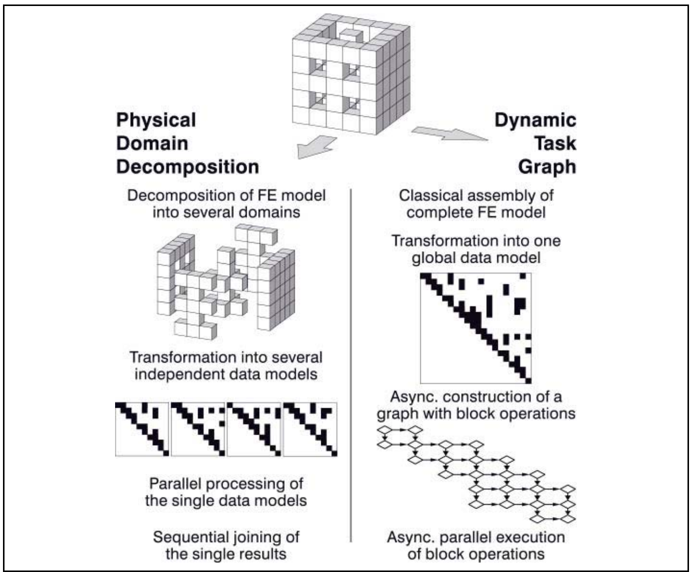

Parallelisation Strategy

Two load distribution algorithms were implemented in the SMP PERMAS code, and the user can select between them:

- DIRECT, this tries to minimise communication.
- BALANCED, which approaches load balancing as the main aim.

A further parameter that can be used to control the load distribution is the reordering algorithm.

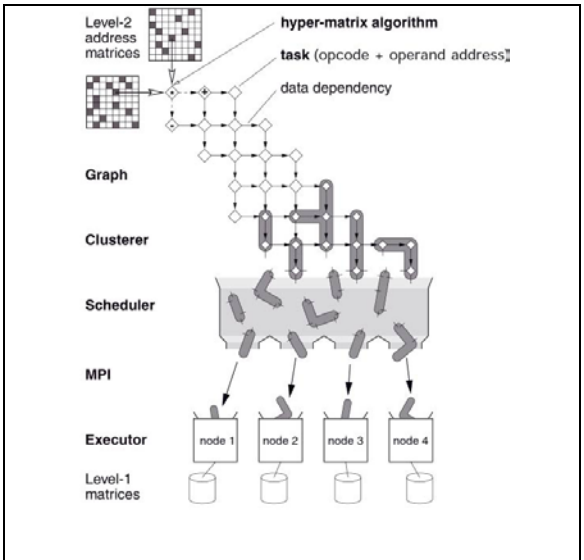

Dynamic Task Graph

The task graph model (TGM) execution model is applicable for shared memory and distributed memory architectures. The PERMAS-D project utilised a mulit-threaded design. The overall aim was to increase asynchronity by a process for disk I/O especially for a single processor. The graph below shows the full multi-threaded version without message passing that was implemented within the project.

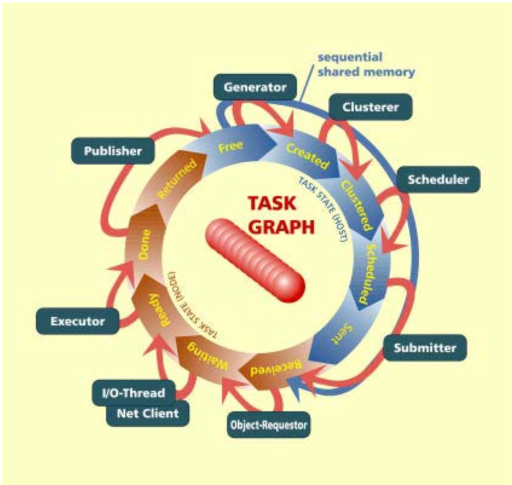

Task graph manager execution model

The difference between a sequential and this multi-threaded execution is the improved I/O handling. The same I/O strategy is employed for the sequential and parallel versions within the TGM, where data locality is the major issue.

2.2 Cost benefit analysis of migration

The hardware configuration at SE for solving engineering problems consisted of two HP 9000/735 workstations, one HP 900/720 workstations and X-terminals (mostly PC's with adequate software) in a switched 100-Mbit-LAN.

There are five engineers more or less constantly (temporarily sometimes more) working with analyses. SE uses FEMGEN as a pre-processor to generate the finite element model, and in-house written post-processors to calculate beam-, plate- and panel-buckling and fatigue life following the valid rules.

The definition of the orthotropic material properties and the transfer of the hydrodynamics loads to the finite element model requires a large effort in terms of manpower to create and to check the model and is prone to error. The correction of the resulting stresses after the analysis is also prone to error. Generating three different models (with some thousand unknowns in the overlapping regions) is not especially effective.

SE wanted to improve the analysis procedures in different ways :

- use the real plate thickness, real stiffener cross section properties and real stiffener spacing, not orthotropic material properties,

- make only one model consisting of the whole ship,
- eventually run the wave response analysis on the same model (if that is possible).

Their intention was that the new strategy would lead to less effort in terms of manpower and fewer errors but to very large models. This was not possible before PERMAS-D because of limitations in the available computer power of SE. SE thought that parallel computing could definitely change the picture and have a high impact on the cost, the project time and the reliability of the results.

DYNALIS is doing engineering work in different industrial branches. Being a sub contracting party, the response time and the reactivity demanded by clients is very high. Very often, the initial design and scaling has to be improved and optimised.

So, in the shortest possible time, several analyses have to be performed on a generic model. Moreover, in the car industry, finite element calculations are used for static and dynamic calculations and for crash analysis. The mesh used for crash analysis (explicit method) is generally finer than meshes used for implicit methods. Considering the modelisation time, using the same model for implicit and explicit methods will save significantly manpower, but increase computer time.

On a single processor workstation (like an SGI R4400), only 1 dynamic modal analysis may be performed each day (elapsed time > 12 hours). The expectation of DYNALIS users was to perform 2 or 3 analyses each day.

3 Results, Achievements and Benefits

The two main results of the project were:

- demonstration that SMP PERMAS gave improved performance; and
- demonstration that this performance was cost-effective to deliver and gave real business benefit.

The following show a typical benchmark result obtained from test cases provided by the end-users. The cost benefit analysis and the migration strategy for the purchase of an SMP machine is also discussed.

3.1 Test cases provided

SE and Dynalis supplied the following test cases for the purposes of this project. Each partner has supplied two test cases which are representative of the types of models which are frequently analysed by each company.

The test cases were specified by the PERMAS-D end-users as being representative of the types of problems they solve regularly. Each test case tests the limits of current sequential PERMAS installations, but with the addition of the more complex analyses required by the users, these test cases will required high-performance computing to solve in a viable time period.

3.1.1 Structural Engineering

The suction anchor is a relatively new technique for the mooring systems of drilling rigs, semi-submersible production platforms, floating storage and offloading vessels is the use of suction anchors. These are large cylinders (in this case 12 m high and 5m diameter), which are closed at the top and usually stiffened with webs inside (not in this case). They are installed at the sea bottom by pumping out the water inside the cylinders and so sucking them into the sea bottom. The anchor chains are connected to the anchor by the pad eyes.

SE provided an existing model of a pad eye (coarse, uniform mesh, no shackle) and their ideal model (fine, optimised mesh, shackle with contact).

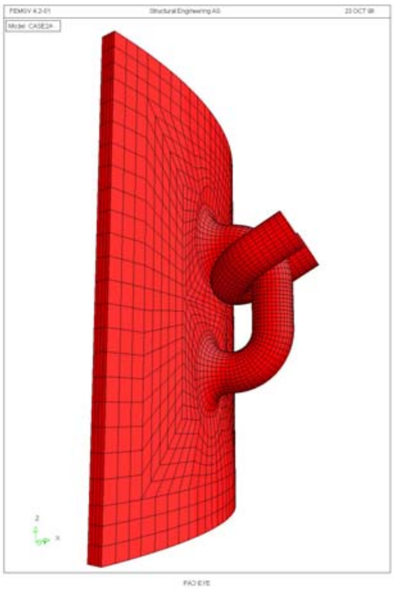

An advanced model of a pad eye, expanded with shackle to allow different coarseness of mesh, for contact and eventually shape optimisation.

3.1.2 Dynalis

Dynalis provided two models - a shell model of a car body (on which they would normally perform a model analysis to identify vibration problems) and a ski lift component.

For confidentiality reasons, the car model below is not an existing car but an old model modified (refined mesh and some modification of the design) to obtain a size of model similar to the size of models used today in the car industry.

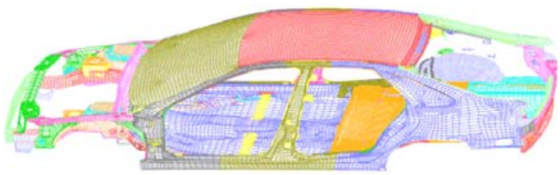

Mesh of a Car.

```
Nodes    : 125752
Unknowns : 754278
Elements : 134200
Elements used:
              680   BECOS
             8536   TRIA3
           124984   QUAD4
```

Analysis : 20 eigenmodes

The following 3D model is part of the equipment for a ski lift. For this model, the expectation is a reduction of the elapsed time for a refined mesh and/or for an analysis with many loading cases (for fatigue).

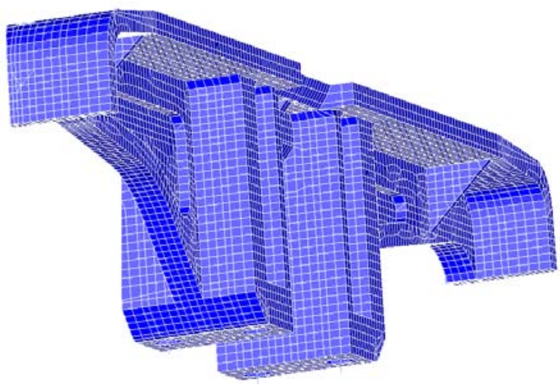	The 3D model is a part of the equipment for a ski lift.
	`Nodes    : 31666`
	`Unknowns : 94710`
	`Local guiding freedoms: 24`
	`Suppressed freedoms:    24`
	`Dependent freedoms:    576`
	`Elements : 24910`
	`Elements used:`
	`2634  PENTA6`
	`22272  HEXE8`
	`4  BECOS`
	`1062 Relative Contact DOFs with`
	`354 Normal CA-DOFs`
	
	`708 Frictional CA-DOFs`
	
	`Analysis: Static with contact.`

3.2 Benchmark results

To prove that the new version of the PERMAS code could deliver the required performance, the code was benchmarked on a number of different machines under different conditions:

- SE: HP9000/735 1 99MHz PA-RISC 7100 processor, main memory 176 MB
- SE: SGI Origin 200, 2 180 MHz R10000 processors, main memory 512 MB
- PAC: SGI Power challenge, 8 75MHz R8000 processors, main memory 1024 MB.
- INTES: HP 9000/800,4 processors, main memory 2 GB
- INTES: DEC-Alpha 4100, 4 processors, main memory 1 GB

To give a flavour of the results obtained the following benchmarks are tabulated from the test cases run at INTES

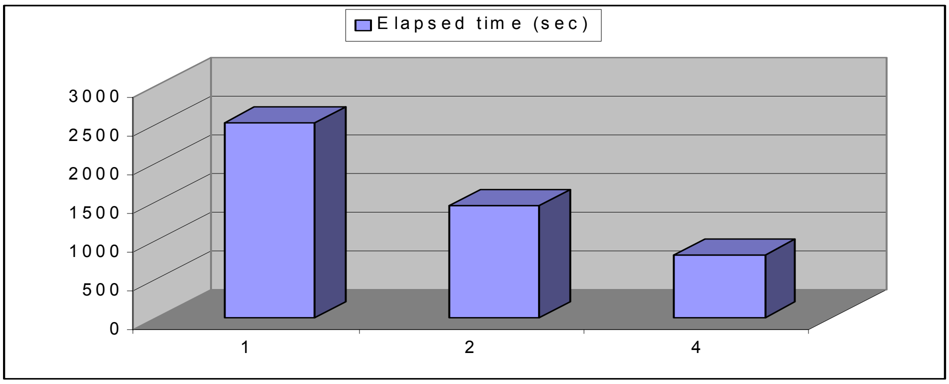

PAD Eye.

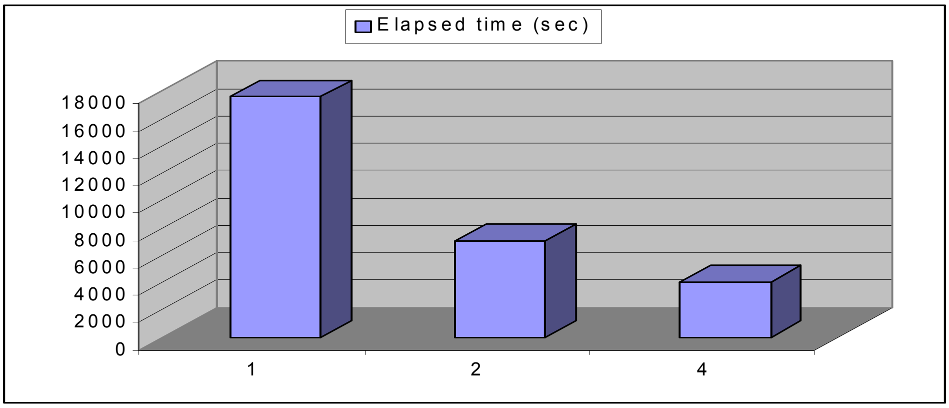

Car Body

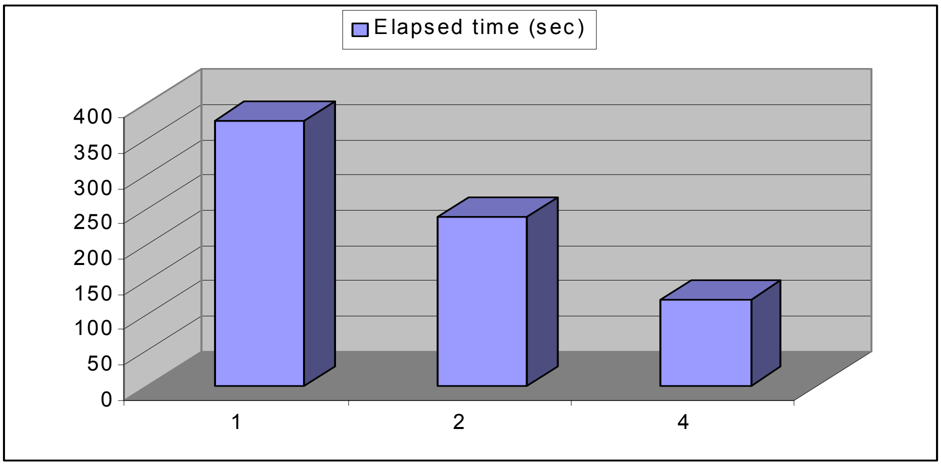

Sky Lift component

The benchmark suite on INTES machines was run on a loaded machine. That means that in the sequential case each benchmark was run as four jobs simultaneously, for the two CPU case each benchmark was run as two jobs simultaneously and for the four CPU case each benchmark was run as a single job. Only by doing so it could be guaranteed that all CPUs were similarly occupied. To ensure that the memory usage also was the same in all three cases, the memory for the two CPU jobs was doubled and for the four CPU jobs was quadrupled compared to the sequential jobs. It has to be said that the disk space needed and so the necessary IO differs between the three cases: the sequential jobs need four times more disk space and the two CPU job two times more disk space than the four CPU jobs. This has strong effects for the IO.

Suppose that you are working on a loosely coupled problem. Since for tightly coupled problems the (100Mbit fast Ethernet) network power is not sufficient to reach satisfying speedups. The limit of the SMP solution is the IO bottleneck, whereas the limit of the cluster solution is network power. The cluster solution has a limited "bandwidth" of possible applications that have to be evaluated by typical client applications.

3.3 Cost – Benefit Analysis

Having demonstrated that the required performance was achievable with SMP PERMAS, the project then set out to show it was viable to use SMP platforms. This was done through a cost-benefits analysis.

The cost-benefits analysis answers the question: *given an existing level of investment in workstations, interconnect and human resources, what is the most cost-effective approach to realising benefits from low-cost HPCN platforms?*

- **Costs**

There is a procurement process overhead consisting of contacting manufacturers of SMP's and testing of the possible platforms. Once a machine has been chosen there is a cost involved with installation. This includes time spent setting up the network, configuration, software and learning to use the machine as well as physically installing the equipment.

The perceived strategy for moving from an existing to an SMP based system can be seen as being done in a series of small, low-cost steps. The existing system is likely to be made up of a mix of desktop workstations and PC's that are being used for simulation, pre and post processing and report writing.

- **Benefits**

Increasing simulation tasks impacts on engineer's interactive use of their workstations and finding spare CPU cycles on another machine takes time. Additionally, spare CPU cycles may not be available.

Installing an SMP machine to be used as a calculation engine will free up desktop machines for more interactive tasks. This will also reduce system overheads by centralising backup and licensing issues. The main benefit will therefore be increased engineer's productivity and efficiency.

- **Migration**

When an SME decides to upgrade to HPCN, there are two alternatives, namely a shared memory system or a cluster (of legacy or low cost workstation). Which option is the best for a particular company depends on its core business and day-to-day method of operation. A company that relies on computational resources as its core business will probably find that SMP is the desired route, while a company that uses computation as a support tool for the designers will find that the cluster approach is more desirable.

Shared memory systems are medium cost workstations that are capable of high performance computation. If the answer to the following question: "Are my engineers using their workstations to full capacity with no CPU cycles available for other applications?" is 'Yes' then the suggested route is to move to an SMP platform.

Low cost systems (e.g. Intel PC's or legacy workstations) operating in a cluster can be an effective way of increasing the computational resources of the company by using spare CPU cycles on idle machines. If the answer to the following question: "Are many of the machines in my organisation not being used at full capacity (e.g. for report writing or office tools)?" is 'Yes' then the suggested route is to move to an cluster solution.

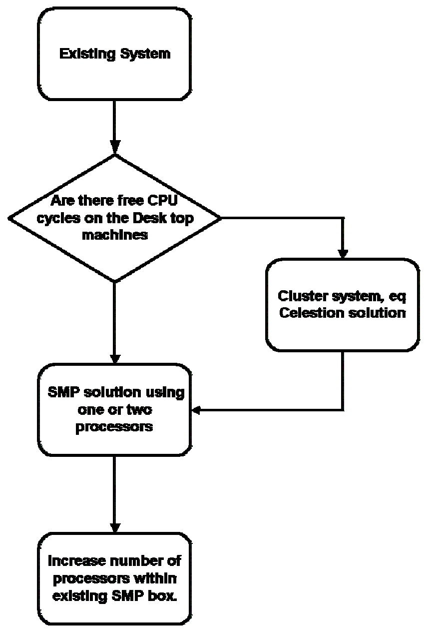

Migration Strategy to SMP system

The purchase price for a typical SMP machine is in the region of 22000 to 45000 Euro for a 2 CPU machine. There is a procurement process overhead consisting of contacting manufacturers of SMP's and testing of these possible platforms. Once a machine has been chosen there is a cost involved with installation. This includes time spent setting up the network, configuration, software and learning to use the machine as well as physically installing the equipment.

Installing an SMP machine to be used as a calculation engine will free up desktop machines for more interactive tasks. This will also reduce system overheads by centralising backup and licensing issues. Future upgrades to the SMP platform are easily performed by addition of CPU, memory and disk.

For a more detailed explanation of the cost benefit analysis the reader should refer to D4.2

4 Dissemination

A major activity of PERMAS-D was the dissemination of the benefits that HPCN can bring to engineering consultancies.

To aid with the dissemination of the PERMAS-D project three posters were produced during the lifetime of the project. Copies of the posters are shown below. Poster one gives an overview of the project and was designed to stand alone. Poster two demonstrated the possible speedup attainable on an SMP platform using test cases supplied by both SE and Dynalis. Poster three, the final poster, was designed to stand along side the first two and create a storyboard effect of the whole project. The posters turned out to be 790mm x 1180mm in size.

Full-sized copies of the posters can be downloaded from the PERMAS-D Web site at http://www.pac.soton.ac.uk/permasd

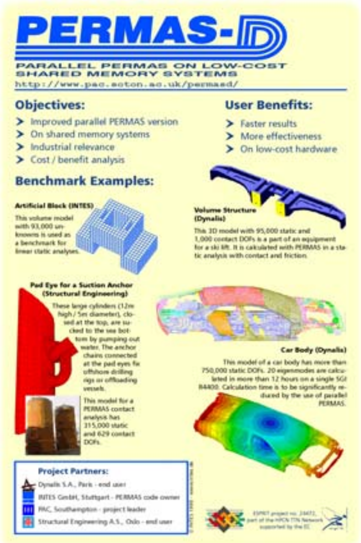

Poster One: Background

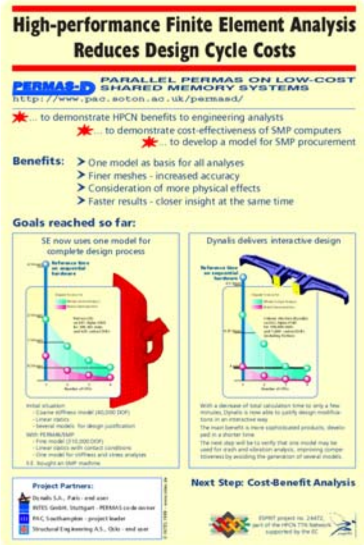

Poster Two: Benchmark Results

Poster Three: Cost-benefit & Migration Strategy

4.1 Publications

(See the PERMAS-D Web site for copies of the material produced)

4.1.1 PERMAS newsletter

This is a technical newsletter, which was sent to all the INTES clients in December 1998. Enclosed inside were the PERMAS-D flyers with descriptions of the test case models and first preliminary benchmark results.

4.1.2 SE company flyer

As part of on-going marketing and company promotion SE created and circulated a one-page flyer. A copy of the flyer is attached to the end of this document.

4.1.3 HARVEST

A success type story was also submitted for publication. To date no publication date has been given as Harvest is only circulated once every two months. This publication is read in many French "Bureaux d'Etudes" by FEM users.

Article to appear in a publication such as Harvest

PERMAS-D: A project to demonstrate the increase in competitiveness by moving to a low-cost shared memory system.

A consortium of organisations in Europe have ported and applied the PERMAS code for use on parallel SMP platforms, drastically reducing run-time for small companies that rely on this code for their core business.

Two companies have been evaluating the new parallel PERMAS code on existing problems looking for performance increases.

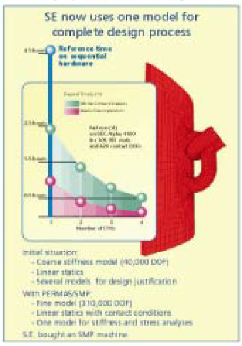

Description of the pad-eye problem
These large cylinders (12m high/ 5m diameter), closed at the top, are sucked to the sea bottom by pumping out water. The anchor chains connected at the pad eyes fix offshore drilling rigs or offloading vessels. This model for a PERMAS contact analysis has 315,000 static and 629 DOFs. Structural Engineering (SE) in Norway now run this analysis on a parallel SMP platform with a factor of 10 increase in performance. SE can now run analysis in a few hours rather than days.

Description of the ski-lift problem
This 3D model has 95,000 static and 1,000 contact DOF's and represents part of an equipment for a ski lift. It is calculated with PERMAS in a statics analysis with contact and friction.

In general, SME's will benefit from this new way of using PERMAS. They will be able to calculate more complex problems in the same time with greater accuracy using a finer mesh and more degrees of freedom. The addition of the SMP machine not only frees up desktop machines for interactive use by engineers but also reduces administration overheads by centralising back up procedures and licensing issues. Being able to run simulations in hours rather than days increases the transparency of the design process by allowing more client interactivity. The market profile of the company is therefore increased as they will be seen as on the leading edge of technology take-up.

The purchase price for a typical SMP machine is in the region of 22000 to 45000 Euro for a 2 CPU machine.

There is a procurement process overhead consisting of contacting manufacturers of SMP's and testing of these possible platforms. Once a machine has been chosen there is a cost involved with installation. This includes time spent setting up the network, configuration, software and learning to use the machine as well as physically installing the equipment.

The perceived strategy from moving from an existing to an SMP based system can be seen as being done in a series of small, low-cost steps. The existing system is likely to be made up of a mix of desktop workstations and PC's that are being used for simulation, pre and post processing and report writing. Increasing simulation tasks impacts on engineer's interactive use of their workstations and finding spare CPU cycles on another machine takes time. Additionally, spare CPU cycles may not be available.

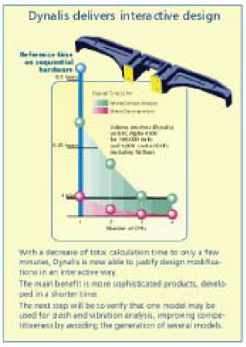

Installing an SMP machine to be used as calculation engine will free up desktop machines for more interactive tasks. This will also reduce system overheads by centralising backup and licensing issues. Future upgrades to the SMP platform are easily performed by addition of CPU, memory and disk.

Conclusion
There can be seen to be two major benefits for moving PERMAS code to a parallel SMP platform. For a small capital outlay, an SME can greatly increase its computational resources. Migration to a centralised SMP network means more efficient use of the company's existing machines.

The initiative is the result of collaboration by a group of organisations from across Europe under the European ***ESPRIT*** scheme, project number 24472. Led by the Parallel Applications Centre (UK), the consortium includes Dynalis (France) and Structural Engineering (Norway) representing the end-users and Intes (Germany) as the software providers.

More information can be obtained from: Paul Gordon at the Parallel Applications Centre on *plg@pac.soton.ac.uk* or visit the project Web site at *http://www.pac.soton.ac.uk/permasd/*

Article submitted to Harvest.

4.1.4 Other publications

The success story, which has been submitted to Harvest, has also been sent to a freelance journalist who frequently publishes within engineering publications. To date no article has been published.

4.2 TTN Network Industry Sector Groups

The main target industry sectors for PERMAS-D were those of the end-users, namely automotive (Dynalis) and Offshore (SE). Thus the most relevant TTN Network Sector Group is Automotive.

The PERMAS-D dissemination material was sent to the Automotive Sector Group co-ordinators. The project consortium has remained in touch with the Automotive Sector Group activities and the following two automotive sector group activities were targeted.

4.2.1 ISATA '99

International Symposium on Automotive Technology and Automation, in Vienna 14th – 18th June 1999. The NETAPDEX TTN are having a stand at this event and all three PERMAS-D posters will be present.

4.2.2 EAEC

European automotive congress "Vehicle systems Technology for the Next Century" 30th June – 2nd July 1999 in Barcelona. CEPBA-TTN are having a stand and all three posters will be presented.

4.3 Conferences Attended

4.3.1 German Finite Element Congress

INTES hosted a booth at this event, which took place in Baden Baden, Germany on 16/17 November 1998, as part of their normal marketing activity. This event had a large, mainly German audience (the language of the congress was German) of users of finite element analysis. These included analysts from engineering consultancies like SE and Dynalis. Poster 1 was presented as highlighting the expected benefits of the PERMAS-D project. The results of this exercise were that out of 100 participants 25 flyers were distributed and 4 new contacts were made.

4.3.2 MICAD

This conference was held in France, 9-12 February 1999. Dynalis hosted a booth as part of their normal marketing activity. The Automotive Sector Group of the TTN Network also hosted a stand. The MICAD audience was made up of designers, engineers, analysts, CAD users etc. and was directed towards a French speaking audience. The language of the conference/exhibition was French.

There was PERMAS-D material on both stands at MICAD and the material led visitors from one stand to the other. A second poster was produced detailing some of the actual benefits of SMP PERMAS as realised in the benchmarking phases of the project. This poster was displayed along with the first to give a complete PERMAS-D success story up to but not including the cost-benefit analysis which was to be contained within PERMAS-D poster 3.

The first two posters were presented at this event. There were 200 participants of which Dynalis targeted 50 individuals for follow-up after the event.

4.3.3 NAFEMS workshop

The PAC held a NAFEMS workshop, 25th May 1999, that discussed 'ubiquitous simulation' where the results of the PACAN-D and PARACOMP PSTs were discussed. The target audience for this event were FE practitioners. Rolf Fischer from INTES attended and presented a talk titled: "reducing design cycles by parallel FEA simulations with PERMAS."

4.3.4 Web site

The PAC produced and hosted a Web site for the project. This contains overview information, downloadable versions of the dissemination material produced and contact details for all the partners. This site will remain live until at least 31 March 2000.

See http://www.pac.soton.ac.uk/permasd

5 Conclusion

The results of PERMAS-D strengthen INTES position to supply a viable SMP platform and enables them to exploit the project by

- extending the available case-study material and
- making a limited extension to some of the software features.

These extensions will enable INTES to better address small users including engineering consultants and also component suppliers.

Code enhancements during the project were expected to be slight, but directed towards the needs of smaller end-users. Performance tuning measures implemented or identified during the project were incorporated into their PERMAS product software, to increase the potential benefits from HPCN technology for this group of users.

Case study results and deployment experiences from the project have also been used to enhance INTES marketing strategy for the parallel capabilities of their software, with respect to smaller users.

Finally, the existence of new users using low cost parallel systems will provide INTES with ongoing feedback necessary for maintenance of the parallel software, whilst the cost-benefit analysis results will provide input for strategic investment in future software improvements to address low-cost platforms.

The end users will exploit the enhanced analysis capabilities from using the parallel code to speed up their own design processes. They will also exploit the results from the technology assessment phase of this project to plan a coherent policy for future hardware procurement, incorporating upgrades to their parallel computing facilities to ensure lasting benefits from migration of their design processes. In fact SE have already migrated to an SMP platform and Dynalis plan to do so soon.

All partners will exploit the experiences and know-how gained during the project to enhance their business activities. In particular, the PAC will use this know-how and dissemination material from the project to promote and enable the application of HPC technology by industry throughout Europe.

This project has shown how the successful migration to HPC technology may be made using legacy systems as a first step. It has provided a useful model for other small manufacturers and consultancies.

There can be seen to be two major benefits for moving PERMAS code to a parallel SMP platform. For a small capital outlay, an SME can greatly increase its computational resources. Migration to a centralised SMP network means more efficient use of the company's existing machines.

6 Contact Details

Company	contact	URL	Role
	Paul Gordon	http://www.pac.soton.ac.uk/ plg@pac.soton.ac.uk	Project Leader
	Rolf Fischer	http://www.intes.de/~fischer/ fischer@intes.de	Software Provider
	Jean Michel Terrier	Dynalis@dynalis.fr	End-User
	Meinolf Muller	http://www.se-as.no meinolf.muller@se-as.no	End-User

The project ran for 12 months starting in April 1998 and ending April 1999. The European Commission contributed to the project 120,000 Euro.