

Industrial Stochastic Simulations on a European Meta-Computer

Ken Meacham, Nick Floros, and Mike Surridge

Parallel Applications Centre, 2 Venture Road,
Chilworth, Southampton, SO17 7NP, UK.
`{kem,nf,ms}@pac.soton.ac.uk`
<http://www.pac.soton.ac.uk/>

Abstract. This paper outlines the experiences of running a large stochastic multi-body simulation across a pan-European meta-computer, to demonstrate the use of the PROMENVIR tool within such a large-scale WAN environment. We describe the meta-application management approach developed by PAC and discuss the technical issues raised by this experiment.

1 Introduction

PROMENVIR (PRObabilistic MEchanical desigN enVIronment) is an advanced meta-computing tool for performing stochastic analysis of generic physical systems, which has been developed within the PROMENVIR ESPRIT project (No. 20189) [1, 2]. The tool provides a user with the framework for running a stochastic Monte Carlo (MC) simulation, using any preferred deterministic solver (e.g. NASTRAN), then provides sophisticated statistical analysis tools to analyse the results.

PROMENVIR facilitates investigations into the reliability of components, by analysing how uncertainties in their manufacture, deployment or environment affect key mechanical properties such as stresses. PROMENVIR does this by generating many hundreds of analysis “shots”, in which the uncertain parameters are generated from statistical distributions selected by the user. Key output values are then extracted from each analysis shot, and can be analysed to determine the sensitivity of the component with respect to uncertain parameters, correlations between physical properties and behaviour, and clustering characteristics (e.g. failure set distributions). PROMENVIR is an open environment, and can generate analyses for essentially any solver.

Since solver runs are independent, the overall simulation is intrinsically parallel and is therefore an ideal candidate for exploiting parallel HPC resources, which may include heterogeneous clusters of workstations, MPP or SMP platforms. These resources may reside within a company Local Area Network (LAN), or may be accessible over a geographically-distributed Wide Area Network (WAN).

The only restriction on the use of resources within a PROMENVIR simulation is the availability of hosts and solver licenses. The implication of this is that, by co-operating with partner sites across a WAN, a PROMENVIR stochastic simulation of many hundreds of solver runs may be executed within a few hours, rather than many days, hence providing a rapid turnaround time for analysis of new designs.

One of the main challenges within the PROMENVIR project was therefore to demonstrate the effectiveness of the PROMENVIR package within a pan-European meta-computing (WAN) environment, to solve a large stochastic problem of industrial significance. The Parallel Applications Centre (PAC) has successfully set up and run a series of “WAN experiments”, involving partner sites within a European consortium. This paper outlines the experiments which took place, the technical issues which arose and summarizes the results.

2 Testcase for WAN Experiment

The most widely available solver within the PROMENVIR consortium was the Multi-Body Simulation code SIMAID, developed by CEIT in San Sebastian [3], so it was decided that a suitable demonstration would be set up, using this code as the basis for a large stochastic simulation.

The testcase chosen was a simulation of a satellite antenna deployment, during which the antenna unfolds to give a planar rim (see Figs 1 - 3). The antenna was modelled using a set of beams, with springs and dampers at the joints. A single deterministic simulation had already been carried out, which predicted that the outer rim of the antenna would be essentially flat. However, the antenna simulation model was constructed using components with identical properties, and did not take into account the manufacturing tolerances inherent within the components and their assembly.

The PROMENVIR Monte Carlo simulation consisted of 500-1000 shots of the SIMAID solver, for the antenna testcase (at least 500 shots were required for convergence of results). For each shot, values for the critical component parameters (e.g. spring stiffness) were chosen from a normal distribution, which modelled the known tolerances. The results of the stochastic simulation could then be analysed to see the effect of manufacturing tolerances on the planarity of the deployed antenna, and hence the effectiveness of the antenna.

3 The European Meta-Computer

A total of 102 CPUs were provided for the experiment, from within the consortium (see Fig. 4). These were restricted to SGI architectures, since SIMAID was only available for IRIX or IRIX64 operating systems. However, many different types of SGI machines were made available, ranging from older Indigo R3000

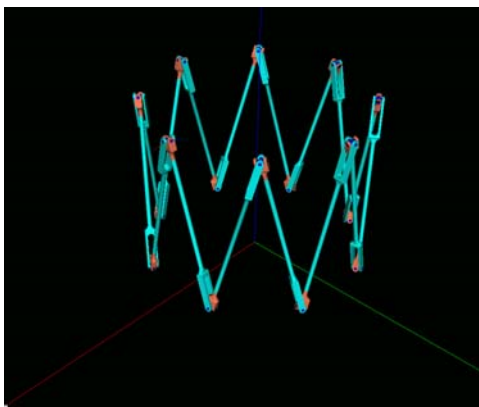


Fig. 1. SIMAID antenna model (folded)




Fig. 2. SIMAID antenna model (partially deployed)

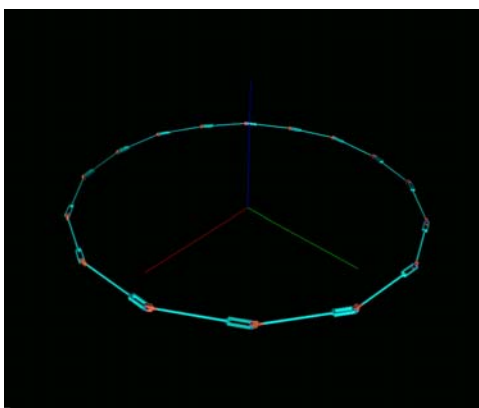


Fig. 3. SIMAID antenna model (fully deployed)

workstations up to a 64-cpu Origin (at UPC, Barcelona). The topology of this hardware configuration is shown geographically in Fig 4. The master host was running at PAC, while slave hosts were used in UK, Spain, Germany and Italy, at various partner sites. We refer to such a meta-computer as a Parallel Virtual Computer (PVC).

SIMAID was installed and licensed on each machine, and access permissions set up to allow remote access from the master SGI Indy at PAC, on which the PROMENVIR front end and resource manager were running.

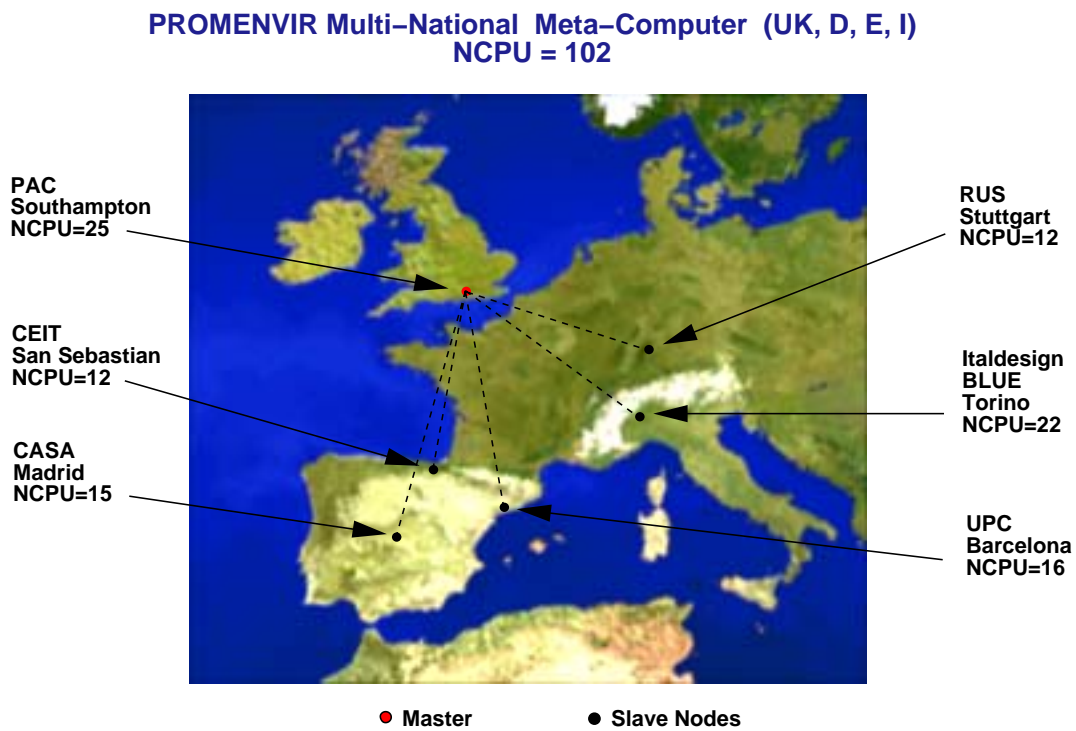


Fig. 4. Meta-computer (PVC) topology for WAN experiment

4 Technical Issues

4.1 Meta-Application Manager

A fundamental component of the package is the Advanced Parallel Scheduler (APS), which has been developed at the Parallel Applications Centre (PAC).

This has been designed as a meta-application manager, which orchestrates the use of the PVC by an application such as PROMENVIR. The APS is capable of making intelligent resourcing decisions, based on performance models (developed at PAC) of the application and resources. These models are used to predict resource usage for a solver in terms of CPU, memory, disk space and I/O traffic, and hence allows capacity planning of the available resources, before submission of any jobs.

A hosts (PVC) database is set up, using the System Definition tool (SDM), which automatically obtains characteristics for each required host, within the LAN or WAN. Performance characteristics for hosts may also be entered here, to be used by the APS.

The APS creates daemons on remote hosts, which are responsible for submitting jobs, copying files and communicating load information back to the master workstation. It is capable of initiating UNIX processes directly, but can also be used to control and submit meta-applications via conventional load-sharing software, such as LSF from Platform Computing [4].

4.2 Remote Host Access and Security

A major part of the work involved in setting up the WAN experiment was ensuring remote host accessibility. For example, some partners were able to set up, quite quickly, remote access to their hosts via `rsh` (remote shell); these sites generally consisted of academic institutions, where access is not normally too restricted. However, several partners (including PAC) had security features including firewall machines. This was a major obstacle, since `rsh` access is normally restricted and, even where available, a firewall machine generally allows access in one direction only. This makes it impossible to set up two-way `rsh` access between two secure sites, through two firewalls (i.e. one at each end).

A further requirement on the meta-computer was to make any file transfers between remote sites both robust and secure. For these reasons, we decided to explore the use of SSH/SCP protocols, which turned out to solve most of our problems and requirements.

SSH (Secure Shell) and SCP (Secure Remote Copy), developed by Data Fellows [5, 6], provide secure TCP/IP connections between trusted sites. Features include:

- strong authentication (via RSA keys)
- automatic data encryption and compression
- access through firewalls (via SSH officially registered port 22)

We were also able to exploit an advanced SSH facility called port forwarding, which enables TCP traffic to be forwarded to a remote host via the secure

connection.

Various partner sites already had an SSH facility, while others needed to install it. By using SSH and SCP protocols for all TCP connections, and by configuring partner firewalls to allow SSH access, we were then able to submit tasks transparently to the PVC, via the APS, and receive back results from solver runs, for collection by the PROMENVIR application. This demonstrated sharing of resources between sites for a single meta-application, even between sites which both have firewall security.

5 Running the Experiment

Several smaller-scale tests of PROMENVIR were carried out, firstly using machines that were available and easily accessible (i.e. without firewalls). It was soon found that the unreliability of the network was the major factor in any lost or failed solver runs (shots). Due to the statistical nature of the PROMENVIR environment, it is not critical if some shots are lost, since results will still converge, as long as enough shots are completed in total.

However, we were finding that up to 10-20% of the remote tasks could fail, due to slow and unreliable connections to certain sites (particularly to Italy). This figure was unacceptable, particularly if the WAN was intended for use with other (non-stochastic) types of application.

We therefore made several improvements to the APS, to enhance robustness and optimise data transfer. For example, rather than attempting single file transfers, we set up the APS to perform multiple retries. If a host could not be contacted after this, it was assumed to be down, and no further tasks were sent to it.

As more and more machines were made available, we were able to increase the number of CPUs used for the experiments. Our target was to reach 100 CPUs across the meta-computer.

6 WAN Experiment Findings

6.1 Results

A summary of the most successful experiment is shown in Table 1. This shows a list of the partner sites, and their contributions to the meta-computer in terms of CPUs provided (as SGI workstations or SMP platforms).

Of the 8 sites represented in this table, 5 had firewall security.

Table 1. PVC host usage statistics during WAN experiment

	CPUs	Availability			Shot Statistics		
Partner	Nproc	In PVC	Access	Used	Failed	Successful	Total
PAC	15	15	15	14	1	150	151
So'ton University	10	10	9	6	0	40	40
UPC	16	16	16	16	1	275	276
RUS	12	12	11	9	0	104	104
CASA	15	15	15	14	15	184	199
CEIT	12	12	12	8	0	98	98
ItalDesign	11	11	6	5	2	63	65
Blue	11	11	11	7	6	61	67
Grand Totals	102	102	95	79	25	975	1000
Total cpus installed	102	Elapsed Execution Time: 4:39:16 Approx Single CPU Time: 250 hrs					
Total cpus defined in PVC	102						
Total cpus available	95						
Total cpus used in WAN	79						

6.2 Availability and Reliability

Of course, the fact that remote hosts are set up and installed does not necessarily imply that they will be available at the time of running the experiment. Of the 102 CPUs defined in the PVC, only 95 could be contacted at the time of running the experiment. We found that this number fluctuated during the day; this was mainly due to the network load, but was sometimes due to machines being down, or turned off.

Of the 95 CPUs available to PROMENVIR via the APS (i.e. with daemons running), 79 were actually used during the simulation to run solver tasks. The reason for this lies in how the APS decides to allocate tasks to hosts. The CPU load is measured on each remote host (via the UNIX `uptime` command), and sent back to the APS. If this is below a certain threshold and the performance models indicate that it is advantageous to do so, a task will be submitted to that host; otherwise the task will be scheduled elsewhere.

We found that jobs were already running on certain hosts (generally those which had not been provided exclusively for the WAN experiment), so these were not used by the APS to run PROMENVIR tasks. However, we also found that some machines, which were apparently not running any jobs, were still not used by the APS; these hosts still appeared overloaded. We eventually traced this problem to certain SGI workstations which were running screen lock programs, which fully loaded the machines CPU resources, when no other (higher priority) jobs were running.

A total of 1000 SIMAID solver shots were submitted to the meta-computer, of which 975 completed successfully, and the total elapsed time of the PROMENVIR simulation was only 4h:39m:16s. This compares with an approximate time of 250 hours on a typical SGI workstation (a single SIMAID run taking around 15 mins). The 25 shot failures were still mainly due to network problems, which will always be present to some extent, within a large meta-computer, unless very fast and reliable links are used (e.g. ATM). A handful of shots were lost, due to solver problems. In certain extreme cases, combinations of random input parameters can cause the solver to crash, though this is rare.

6.3 Analysis of Results

The detailed analysis of the results is beyond the scope of this paper, and will be published elsewhere, in conjunction with results from further simulations which are being carried out using differing ranges of manufacturing tolerances.

However, preliminary results show that, when a stochastic approach is used, there is a broader variation in the deployed geometry, and hence a poorer planarity in the dish than predicted with the deterministic simulation.

By using the stochastic approach, further insights may be obtained than were previously available. For example, by examining the planarity variation according to the standard deviation of manufacturing tolerances, it is possible to work back and ask questions such as “what tolerances must be ensured in the manufacturing process, to produce an antenna of certain planarity”. These types of questions are highly relevant to industry.

7 Conclusions

PROMENVIR has been demonstrated as a highly useful tool for running industrial stochastic simulations across a European meta-computer, showing that corporate-scale meta-computing resources really can be used to solve large problems for industry. The APS, developed at PAC, has facilitated management of meta-applications and meta-computing resources (including firewall-protected systems), and is also being used with other meta-computing applications [7, 8].

Results of the WAN experiment have shown that, once a meta-computer has been set up, further experiments can be carried out routinely. However, the overheads of setting up remote sites can be quite large. Furthermore, it has been found to be difficult to obtain full (exclusive) access to machines, during a large-scale experiment, without the full cooperation from contributing partners.

Our experience suggests that, due to a combination of background loads, network performance and human factors (e.g. failure to log out of a machine),

around 80% resource availability is the maximum which can realistically be expected for large-scale meta-applications.

Finally, there will always be some inherent unreliability within a large-scale meta-computing resource, mainly due to network problems. These must be eliminated as much as possible and, where uncontrollable problems arise, the system and application must be made as robust as possible, to ensure that any simulation does not suffer fatally from individual task failures.

Acknowledgements

The work reported here was part of the ESPRIT PROMENVIR project. We are grateful to our partners (CASA, UPC, CEIT, RUS, Atos, ItalDesign and Blue Engineering) for a most fruitful collaboration.

References

1. Marczyk, J.: Meta-Computing and Computational Stochastic Mechanics. Computational Stochastic Mechanics in a Meta-Computing Perspective (1997) 1–18 Ed. J. Marczyk
2. PROMENVIR product web page: <http://www.atos-group.de/cae/index-promenvir.htm>
3. CEIT home page: <http://www.ceit.es/>
4. Platform Computing (LSF) home page: <http://www.platform.com>
5. Data Fellows home page: <http://www.datafellows.com/>
6. Further information on SSH may be found at: <http://www.cs.hut.fi/ssh/>
7. Upstill C.: Multinational Corporate Metacomputing. Proceedings of the ACM/IEEE SC97 Conference, San Jose, CA, November 15-21, 1997
8. Hey, A.J.G., Scott, C.J., Surridge, M. and Upstill, C.: Integrating Computation and Information Resources - An MPP Perspective. 3rd Working Conference on Massively Parallel Programming Models (MPPM-97), London, November 12-14, 1997