SAND99-2083C

AIAA

AIAA 2000–1651

SALINAS - AN IMPLICIT FINITE ELEMENT STRUCTURAL DYNAMICS CODE DEVELOPED FOR MASSIVELY PARALLEL PLATFORMS

Manoj Bhardwaj, Garth Reese, Brian Driessen,
Kenneth Alvin, and David Day
Sandia National Laboratories, P.O. Box 5800,
Albuquerque, NM 87185-0847.

41st AIAA/ASME/ASCE/AHS/ASC SDM
April 3-6, 2000/Atlanta, GA

DISCLAIMER

SALINAS - AN IMPLICIT FINITE ELEMENT STRUCTURAL DYNAMICS CODE DEVELOPED FOR MASSIVELY PARALLEL PLATFORMS

Manoj Bhardwaj, *
Garth Reese, †
Brian Driessen, ‡
Kenneth Alvin, §
and David Day ¶
Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87185-0847.

Abstract

As computational needs for structural finite element analysis increase, a robust implicit structural dynamics code is needed which can handle millions of degrees of freedom in the model and produce results with quick turn around time. A parallel code is needed to avoid limitations of serial platforms. Salinas is an implicit structural dynamics code specifically designed for massively parallel platforms. It computes the structural response of very large complex structures and provide solutions faster than any existing serial machine. This paper gives a current status of Salinas and uses demonstration problems to show Salinas' performance.

Background

The United States' commitment to ending underground testing, constraints on nonnuclear testing, and loss of production capability call for new means of verifying the security of the United States' nuclear stockpile.[1] Therefore, in 1996, the US Department of Energy announced its Accelerated Strategic Computing Initiative (ASCI) aimed at creating predictive simulation and virtual prototyping capabilities, and accelerating the development of high-performance computing far beyond what might be achieved in the absence of a focused initiative. More specifically, ASCI's vision is to shift promptly from test-based methods to computational-based methods of ensuring the safety, reliability, and performance of the US nuclear weapons stockpile. An initial result of this initiative was the installation in 1997 at the Sandia National Laboratories of an Intel 1.8-Teraflops (trillion floating-point operations per second) peak massively parallel system known as the ASCI Option Red supercomputer. Two additional multi-Teraflop systems known as the ASCI Blue Pacific and ASCI Blue Mountain machines were subsequently installed at the Livermore and Los Alamos National Laboratories, respectively. Harnessing the power of these ASCI machines and exploiting their full potential requires the development of scalable numerical algorithms, which for many applications is a significant challenge.

Part of the ASCI initiative is the development at Sandia of Salinas,[2] a massively parallel structural dynamics code aimed at providing a scalable computational workhorse for highly accurate structural dynamics models. Such large-scale finite element models require significant computational effort, but provide important information including vibrational and shock loads for components within larger systems, design optimization, frequency response information for guidance and space systems, modal data necessary for active vibration control, and characterization data for structural health monitoring.

Salinas

Salinas, a mostly C/C++ code, is an implicit structural dynamics code designed for optimal performance on massively parallel platforms. It's primary purpose is to be the analyst's tool of choice for structural dynamics solutions of finite element models with millions of degrees of freedom. Scalability on distributed memory systems is required to reduce analysis time from weeks to days or hours. Salinas uses the Message Passing Interface[3] (MPI) standard and is therefore portable to other parallel platforms including shared memory machines. Though Salinas is still under development, it has many capabilities which make it a viable tool, presently, to obtain structural dynamics solutions.

---

*Member AIAA, Senior Member of Technical Staff
†Principal Member of Technical Staff
‡Limited Term
§Senior Member AIAA, Principal Member of Technical Staff
¶Senior Member of Technical Staff

Capabilities

Salinas is capable of obtaining static, eigen,[4] direct transient, shock spectra, modal frequency response, modal transient, and modal shock solutions. Sensitivity analysis[5] capability and multi-point constraints are also available in Salinas.

Salinas' element library consists of the following elements: 8-noded hexahedral, 20-noded hexahedral, 15-noded wedge, 6-noded wedge, 4-noded tetrahedral, 10-noded tetrahedral, 3-noded triangular shell, 6-noded triangular shell, 4-noded quadrilateral, 8-noded quadrilateral, and 1-d beam elements. Salinas also has rigid bars, concentrated masses, and springs.

Parallel Solvers

At the heart of Salinas is a parallel solver, FETI[6] (Finite Element Tearing and Interconnecting). FETI is a domain decomposition based iterative method with Lagrange multipliers. In its simplest form, it is also known as the one-level FETI method, and can be described as a two-step preconditioned conjugate gradient (PCG) algorithm where subdomain problems with Dirichlet (displacement) boundary conditions are solved in the preconditioning step, and related subdomain problems with Neumann (traction) boundary conditions are solved in a second step. The one-level FETI method incorporates a relatively small size auxiliary problem that is based on subdomain rigid body modes. This coarse problem propagates the error globally during the PCG iterations and accelerates convergence. The FETI solver was chosen because of its underlying mechanical concepts, as well as its potential for delivering scalable performance.

For heterogeneous structures such as the reentry vehicle discussed below, FETI employs a Q projector.[7] This provides consistent scalability across subdomain boundaries for structures where the stiffness terms may vary by many orders of magnitude.

Running Salinas in Parallel

Parallel execution of Salinas has overhead associated with it which does not exist when running in serial. For example, the finite element model has to be "decomposed" into subdomains. These may be spread to redundant arrays of independent disks (RAIDs) (18 RAIDS on the Intel Teraflop) to improve parallel input/output (I/O). The mesh decomposition and optimization tools[8–11] will not be discussed here. With a model as large as one million degrees of freedom (d.o.f.), the time is minimal. The timings shown in this paper will show only the time associated with running Salinas in parallel. Salinas will be supported for all ASCI platforms, but all of the demonstration problems' results are reported for the Intel Teraflop, also known as the ASCI Option Red Supercomputer.

The ASCI Option Red Supercomputer

The ASCI initiative supports the ASCI Option Red Supercomputer, a massively parallel machine with a distributed memory multiple instruction and multiple data (MIMD) architecture, as well as the ASCI Option Blue Mountain and ASCI Option Blue Pacific supercomputers. The ASCI Option Red and Blue Mountain systems run MP LINPACK, one of the industry's standard speed tests for large systems, at 2.12 and 1.6 Teraflops respectively.[12] Peak speed of the ASCI Option Red Supercomputer is 3.0969 TFlops.

The ASCI Option Red supercomputer, also known as the Intel Teraflops machine, is the first large-scale supercomputer built entirely of commodity, commercial, off-the-shelf (COTS) components. It has 4,536 compute and 72 service nodes each with 2 Pentium Pro processors. The system has over 1 terabyte of real memory, and two independent 1-Terabyte disk systems. It occupies 1600 sq. ft. of floor-space. The system's 9216 Pentium Pro processors are connected by a 38x32x2 mesh.

The Pentium Pro processor runs at 333 MHz and has a peak floating-point rate of 333 Mflops. It has separate on-chip data and instruction L1 caches of 8 Kbytes each. It also has an L2 cache of 256 Kbytes packaged with the CPU in a single dual-cavity pin grid array (PGA) package. All cache lines are 32 bytes wide. The system was delivered with 128 Mbytes of memory per node, but has been upgrade to 256 Mbytes of memory per node. The two processors on each node support two on-board peripheral component interconnect (PCI) interfaces; each of these interfaces provides 133 Mbytes/sec I/O bandwidth. The memory subsystem is structured as four rows of four independently controlled and sequentially interleaved banks of DRAM to produce up to 533 Mbytes/sec of data throughput. Each memory bank is 72 bits wide. The router supports bi-directional bandwidths of up to 800 Mbytes/sec over each of the 6 ports. As many as four message streams can pass on any given port and at any given time.

Two UNIX-based operating systems collectively called the Teraflops OS run on the ASCI Option Red Supercomputer and present a single system image to the user. Compute nodes run an efficient small operating system called Cougar.[13–15] Service nodes run POSIX 1003.1 and XPG3, and AT&T System V.3 and 4.3 BSD Reno VFS.[16] The file system is concentrated on a small set of specialized nodes that process I/O requests. Symbios RM20 RAIDs are used for secondary storage. A Symbios RM20 RAID has two bays of ten drives each and two controllers. The disk drives are Seagate 4-Gbyte Barracudas with a 3.5" form-factor (Ref. 16).

Demonstration Problems

To demonstrate the scalability of Salinas, Benchmark Problems (described in next section) are run and timings obtained on varying number of processors between 8 and 1000. Here, scalability is the ability to solve an *n*-times larger problem on *n*-times larger number of processors in a nearly constant CPU time.[17]

A mock-up reentry vehicle is chosen to demonstrate the capability of Salinas for statics on a highly heterogeneous problem. The model is chosen to include many of the computational challenges inherent in models of complex structures. By doubling the mesh of the this RV model, a model with eight-times as many elements is used to show Salinas' capability in obtaining a structural dynamics solution on very large, complex finite element models.

An Electronics Package model, another heterogeneous model composed of 10-noded tetrahedral and six-noded triangular elements, is used to demonstrate the parallel eigensolver capability of Salinas.

A beam structure modeled using hexahedral elements is chosen to show the size of problems Salinas' can solve routinely in a relatively short time. Each of these models is described below.

BP1 and BP2

To determine the scalability of Salinas, model problems are needed which can grow proportionately as the number of processors increases. Two benchmark models are chosen for this paper. For both problems each processor has one and only one subdomain. Each subdomain is a cube of 12x12x12 8-noded hexahedral elements. Benchmark problem 1 (BP1) (Figure 1) is a cube which grows with parameter *n*, the number of subcubes along each side of the problem. Therefore for n=1, the global model is a cube of 12x12x12 8-noded hexahedral elements. For n=3, the global problem is a cube with 12*3 hexahedral elements along each edge. Since each processor only has one subcube of 12x12x12 elements, 27 processors are needed for this problem.

Benchmark problem 2 (BP2) (Figure 2) is a beam-like structure with the cross section being a square modeled by 4 subcubes. However, as opposed to growing like a cube like BP1, BP2 grows only in one direction along the length of the beam. Again each processor has a subcube of 12x12x12 elements. For n=3, the cross-section remains the same, but the length of the beam has changed so that there are 12*3 elements along the length of the beam. The total number of subcubes is 12, therefore 12 processors are needed.

BP1 and BP2 are run with the number of processors varying between 8 and 1000 on the ASCI Option Red Supercomputer. The problem size and processor size are both increased by a factor of *n* and the CPU time is recorded. These CPU times indicate the scalability of the code to obtain static solutions of benchmark problems.

Reentry Vehicle

As another demonstration problem, a mock-up reentry vehicle (RV) (Figure 3) is used to demonstrate the static analysis capability of Salinas on the ASCI Option Red Supercomputer. An RV can be expected to experience different loadings in normal and hostile environments. Its structural response during vibration is usually predicted by modal analysis, while its shock response is usually simulated by a direct transient analysis. The predictive computation of responses at component levels requires a detailed finite element model of the full body as well as the individual components.

The mock-up RV is composed of 330,300 elements and 334,759 nodes. With slightly more than one million degrees of freedom, this model requires significant computational power, and provides a reasonable benchmark for massively parallel computational platforms. All elements of the mesh are either 8-noded brick or 6-noded wedge elements. Decomposing this mesh into subdomains with good aspect ratios and without mechanisms is a difficult task because the RV aeroshell is composed of thin walls.

In addition to the decomposition complexities, there are eight different materials that are scattered within the RV model. The colors in Figure 3 represent the various materials. The Young's moduli vary from 10^2 psi to 3×10^7 psi. Hence, this RV structure is highly heterogeneous and can be expected to challenge any iterative solver.

A static analysis on this RV, decomposed into approximately 250 and 500 subdomains, is run on the ASCI Option Red Supercomputer. CPU times are obtained for various decompositions of the RV. The three decompositions are the following: (i) partition the mesh as is, with particular attention to the subdomain aspect ration, (ii) partition the mesh along its material boundaries, and (iii) re-organize all the material groups of the RV finite element model into two clusters and partition each cluster separately.

Extra Large Reentry Vehicle

To test model size limits, the mock-up reentry vehicle used above is refined such that the model contains eight times as many elements. This model, the extra large RV (XLRV), contains approximately 8 million degrees of freedom.

A static analysis on this XLRV, decomposed into approximately 1000 subdomains, is run on the ASCI Option Red Supercomputer. CPU times are reported for the arbitrary decomposition of the XLRV.

Electronics Package Model

To demonstrate the eigen solution capability of Salinas, an Electronics Package (EP) model is chosen. The model (Figures 4,5) consists of 0.5 million degrees of freedom and is decomposed into 256 subdomains. The first four flexible eigenvalues and eigenvectors are calculated. The model is composed of 100,000 10-noded tetrahedral and 12,000 six-noded triangular elements. The solution is compared to MSC/NASTRAN's solution.

An eigensolver was selected for Salinas based on robustness, accuracy, scalability and efficiency, with the underlying linear solver being FETI. Eigen solver methods based on Lanczos and subspace iteration were evaluated. The PARPACK Lanczos-based solver was selected due to its minimal memory storage requirements, reliability, and because the number of linear systems solved per mode is nearly minimized.

Large Beam Model

Since very large finite element models are difficult to find or create, a uniform beam is chosen as the structure to demonstrate Salinas' capability to obtain structural dynamics solutions quickly. The beam is modeled using 6.75 million 8-noded hexahedral elements. The finite element model is decomposed into 2000 subdomains. A static solution is obtained using Salinas on 2000 processors. This model contains more than 20 million equations.

Results

BP1 and BP2 problems

Results for the scalability analysis for models BP1 and BP2 are shown in Tables 1 and 2, respectively, and plotted in Figure 6. Figure 7 shows the solver time and total time for BP2. Tables 1 and 2 show the problem size, the number of subdomains (# of processors = # of subdomains), and the time to obtain a static solution. Problem size varied between 46,000 degrees of freedom and 5.6 million degrees of freedom.

The ideal result is a constant time for all the various cases shown in Tables 1 and 2. Since the problem size is the same per processor, a static solution for a 5.6 million degree of freedom cube takes 360 sec on 1000 processors while a small cube with approximately 50,000 degrees of freedom takes 336 seconds on 8 processors. This shows that the solution algorithm scales very well for both model problems.

Reentry Vehicle

Timings for the RV for various number of subdomains and partition types are shown in Table 3. The arbitrary decomposition, i.e. partition type (i), performed the best with the least amount of time for both 250 and 500 subdomain problems. Partitions (ii) and (iii) did not perform as well which indicates the degrading performance obtainable due to "bad" partitioning.

For the BP1 and BP2 problems, the problem size was kept fixed as the number of processors was increased. In contrast, the RV finite element model was kept fixed while the number of processors was increased from 250 to 513. The total time was cut in half as the number of processors doubled, which demonstrates the scalability of Salinas for complex structures.

Extra Large Reentry Vehicle

A static solution for the XLRV decomposed into 1000 subdomains is obtained. The total time for the solution is 11 minutes 21 seconds. This compares well with the 474 s (8 min) solution time for the RV with 250 subdomains. With many parallel codes, the communication time starts to become a larger factor in the total time and decreases the efficiency of the code. The XLRV demonstrates that Salinas can calculate a static solution of a complex structure with almost 10 million equations efficiently on a large number of processors.

Electronics Package

The first ten frequencies are calculated for the EP using MSC/NASTRAN and Salinas' parallel eigensolver. MSC/NASTRAN is run on an 18 node Sun Sparc server with 6 Gigabytes of RAM, while Salinas is run on 256 processors on the Intel ASCI Option Red Supercomputer.

Since the first 6 modes are rigid body modes (zero frequency), Table 4 compares the first four flexible frequencies obtained by MSC/NASTRAN and Salinas. The results indicate excellent agreement for the four modes between Salinas and MSC/NASTRAN. The time to obtain the 10 modes on a serial platform required more than 5 days, while on 256 processors, the solution was required approximately 30 minutes. Clearly there are significant differences in platforms, algorithms, and capabilities, but the potential for speed-up using parallel approaches cannot be denied.

Large Beam Model

The long hex beam is chosen to demonstrate the capability to handle large finite element models for obtaining structural dynamics solutions. A static solution of the 20.25 million degree of freedom beam model is obtained on 2000 processors. The total solution time is 21 minutes and 3 seconds. This example shows Salinas capability to obtain a static solution on a larger finite element model on a large number of processors quickly. The 2000 processors are not a "bottleneck" and indeed, Salinas scales well to thousands of processors.

Conclusions

Results for BP1 and BP2 problems show the scalability of Salinas by obtaining static solutions of a 46,000 d.o.f. model and a 5.6 million d.o.f. model in about the same time. The BP1 and BP2 problem sizes per processor were kept fixed as the number of processors increased.

The XLRV and RV models showed that Salinas performed well even when the model is highly heterogeneous. The material properties varied by 5 orders of magnitude. As the number of processors was increased from about 250 to about 500, the total time was reduce by one-half. The XLRV shows the capability of Salinas to obtain a static solution of a 10 million degree of freedom problem quickly.

The EP model demonstrates that Salinas' parallel eigensolver obtains structural dynamics solutions quickly and efficiently. The EP model also shows the comparison between Salinas and MSC/NASTRAN to be excellent when comparing the first four flexible frequencies. This model demonstrates the potential of using parallel approaches by reducing the required solution time from 5 days to 30 minutes.

The beam modeled using 6.75 million 8-noded hexahedral elements, shows Salinas' capacity to be efficient on 2000 processors. A static solution of this 20 million degree of freedom model is obtained in less than one hour. Many parallel codes would pay a heavy communication expense, but Salinas' scalability is shown in that a there is not a large communication "hit" even when using 2000 processors.

n	# of subdomains	# of degress of freedom	Time
2	8	46,875	336 sec
3	27	151,959	346 sec
4	64	352,947	355 sec
5	125	680,943	358 sec
6	216	1,167,051	365 sec
7	343	1,842,375	367 sec
8	512	2,738,019	380 sec
9	729	3,885,097	405 sec
10	1000	5,314,683	413 sec

Table 1 BP1 Timings

n	# of subdomains	# of degress of freedom	Time
2	8	46,875	336 sec
7	28	159,375	343 sec
16	64	316,875	354 sec
31	124	699,375	347 sec
54	216	1,216,875	350 sec
86	344	1,936,875	351 sec
128	512	2,881,875	354 sec
182	728	4,096,875	355 sec
250	1000	5,626,875	360 sec

Table 2 BP2 Timings

Acknowledgements

The authors would like to thank Charbel Farhat, Kendall Pierson, and Michel Lesoinne, from the University of Colorado at Boulder, for their help with the parallel solver. The authors would also like to thank David Martinez, James Peery, Dan Segalman, and Sam Key for their support, contributions, and insightful discussions. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

References

[1] "ASCI at Sandia Web Site," http://www.sandia.gov/ASCI.

[2] Reese, G., Bhardwaj, M., Segalman, D., Alvin, K., and Driessen, B., "Salinas - User's Notes," Tech. Rep. SAND99-2801, Sandia National Laboratories, 1999.

[3] Forum, M. P. I., "MPI: A Message-Passing Interface Standard," Tech. rep., 1994.

[4] Reese, G. and Day, D., "A Massively Parallel Sparse Eigensolver for Structural Dynamics Finite Element Analysis," Tech. Rep. SAND99-1158, Sandia National Laboratories, 1999.

[5] Alvin, K., Reese, G., Day, D., and Bhardwaj, M., "Incorporation of Sensitivity Analysis Into a Scalable Massively Parallel Structural Dynamics FEM Code," *5th U.S. Congress on Computational Mechanics*, August 1999, Boulder, CO.

[6] Farhat, C. and Roux, F. X., "A Method of Finite Element Tearing and Interconnecting and Its Parallel Solution Algorithm," *International Journal for Numerical Methods in Engineering*, Vol. 32, 1991, pp. 1205–1227.

[7] Rixen, D. and Farhat, C., "A simple and efficient extension of a class of substructure based preconditioners to heterogeneous structural mechanics problems," *International Journal for Numerical Methods in Engineering*, Vol. 44, 1999, pp. 489–516.

[8] Farhat, C., Lanteri, S., and Simon, H. D., "TOP/DOMDEC, A Software Tool for Mesh Partitioning and Parallel Processing," *Journal of Comput. Sys. Engrg.*, Vol. 6, 1995, pp. 13–26.

[9] Hendrickson, B. and Leland, R., "The Chaco User's Guide: Version 2.0," Tech. Rep. SAND94-2692, Sandia National Laboratories, 1994.

[10] Farhat, C., Maman, N., and Brown, G., "Mesh Partitioning for Implicit Computations Via Iterative Domain Decomposition: Impact and Optimization of the Subdomain Aspect Ratio," *International Journal for Numerical Methods in Engineering*, Vol. 38, 1995, pp. 989–1000.

[11] Vanderstraeten, D., Farhat, C., Chen, P. S., Kenuings, R., and Zone, O., "A Retrofit and Contraction Based Methodology for the Fast Generation and Optimization of Mesh Partitions: Beyond the Minimum Interface Size Criterion," *Computational Methods in Applied Mechanical Engineering*, Vol. 133, 1996, pp. 25–45.

[12] "ASCI Blue Mountain Web Site," http://www.lanl.gov/asci/bluemtn.

[13] Lillevik, S. L., "The Touchstone 30 Gigaflop DELTA prototype," *Sixth Distributed Memory Computer Conference*, IEEE Computer Society Press, 1991, p. 671.

[14] Traversat, B., Bitzberg, B., and Fineberg, S., "Experience with SUNMOS on the Paragon XP/S-15," *Intel Supercomputer Users*, 1994, San Diego, CA.

[15] Schuler, L., Riesen, R., Jong, C., van Dresser, D., and Maccabe, A., "The Puma Operating System for Massively Parallel Computers," *Intel Supercomputer North America Users*, 1995.

[16] Garg, S., Godley, R., Griffiths, R., Pfiffer, A., Robboy, D., Smith, S., Stallcup, T. M., and Zeisset, S., "Achieving Large Scale Parallelism Through Operating System Resource Management on the Intel TFLOPS Supercomputer," *Intel Technology Journal*, 1998, Q1'98 issue.

[17] Bhardwaj, M., Day, D., Farhat, C., Lesoinne, M., Pierson, K., and Rixen, D., "Application of the FETI Method to ASCI Problems: Scalability Results on One-Thousand Processors and Discussion of Highly Heterogeneous Problems," *International Journal for Numerical Methods in Engineering*, Vol. 47, 2000, pp. 513–536.

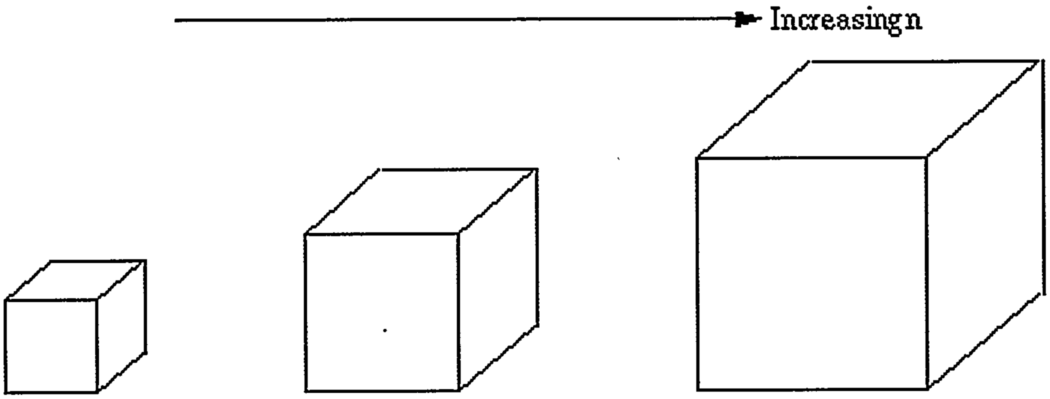

Fig. 1 Benchmark Problem 1

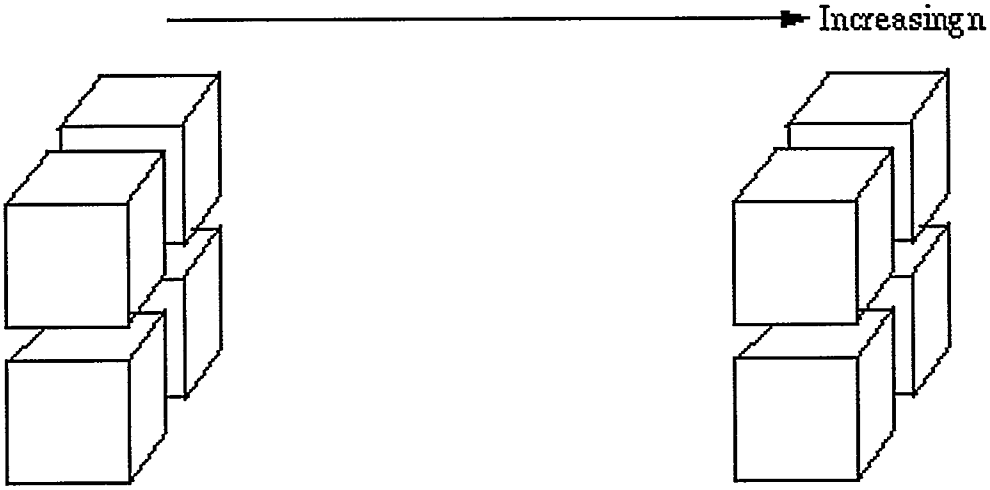

Fig. 2 Benchmark Problem 2

# of subdomains	Partition Type	Total Time
250	i	474 sec
241	ii	657 sec
257	iii	350 sec
513	i	219 sec
505	ii	502 sec
517	iii	218 sec

Table 3 Reentry Vehicle Timings

Mode No.	Salinas	MSC/NASTRAN
1	437.0 Hz	434.3 Hz
2	629.1 Hz	627.4 Hz
3	659.2 Hz	657.2 Hz
4	793.2 Hz	793.6 Hz

Table 4 Mode Comparison between Salinas and MSC/NASTRAN

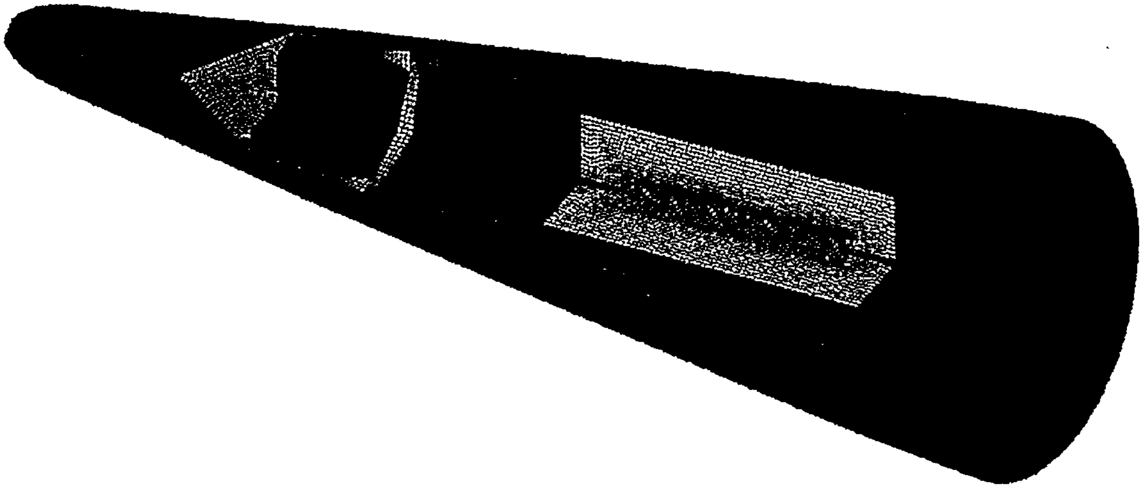

Fig. 3 Cut-away View of Reentry Vehicle Model

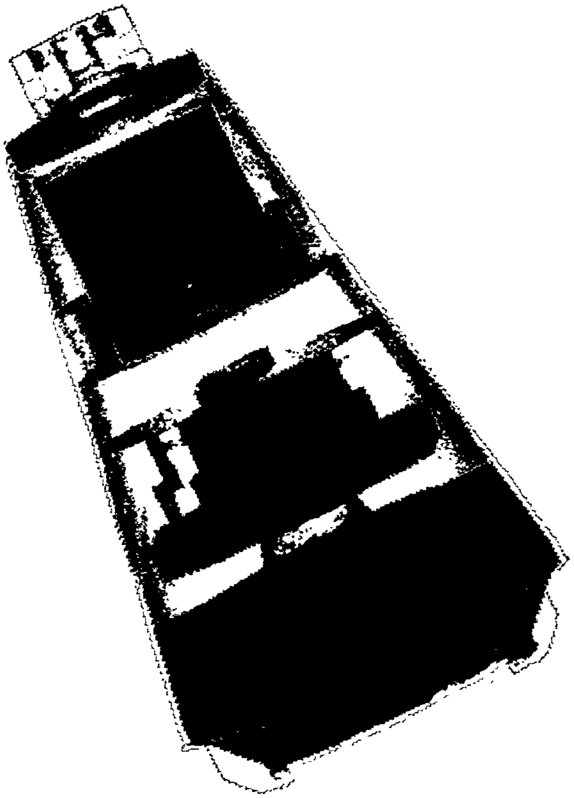

Fig. 4 Electronics Package Model (EPM)

Fig. 5 More Detailed Electronics Package Model(EPM)

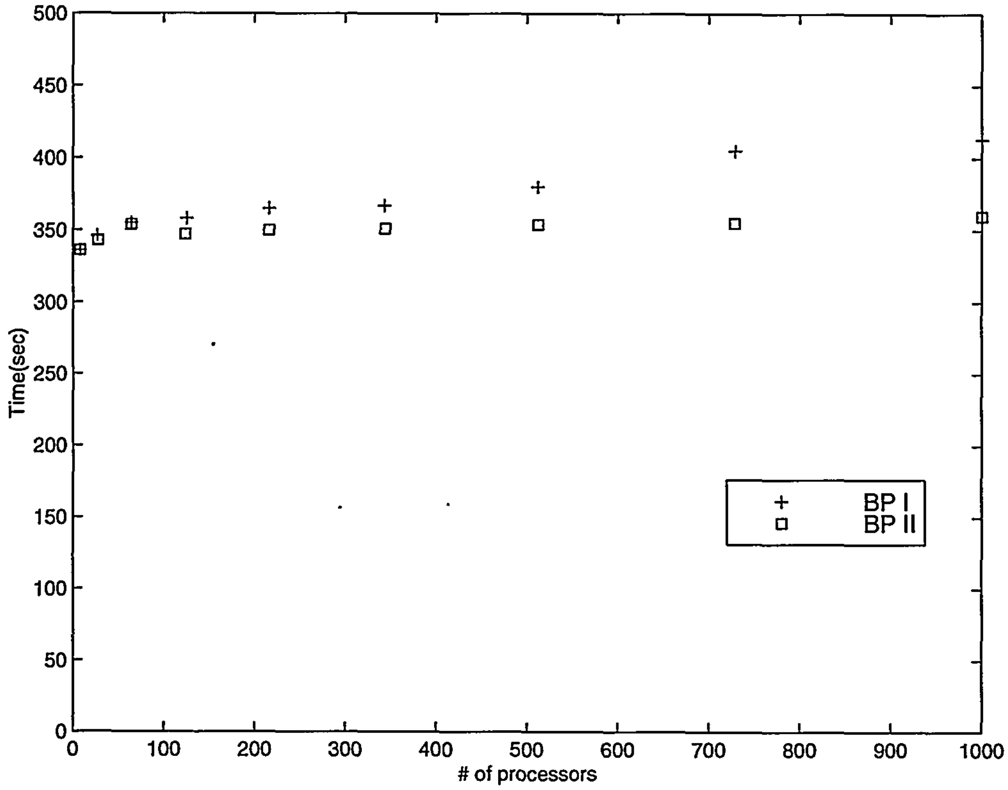

Fig. 6 Scalability of Benchmark Problems I and II to 1000 processors

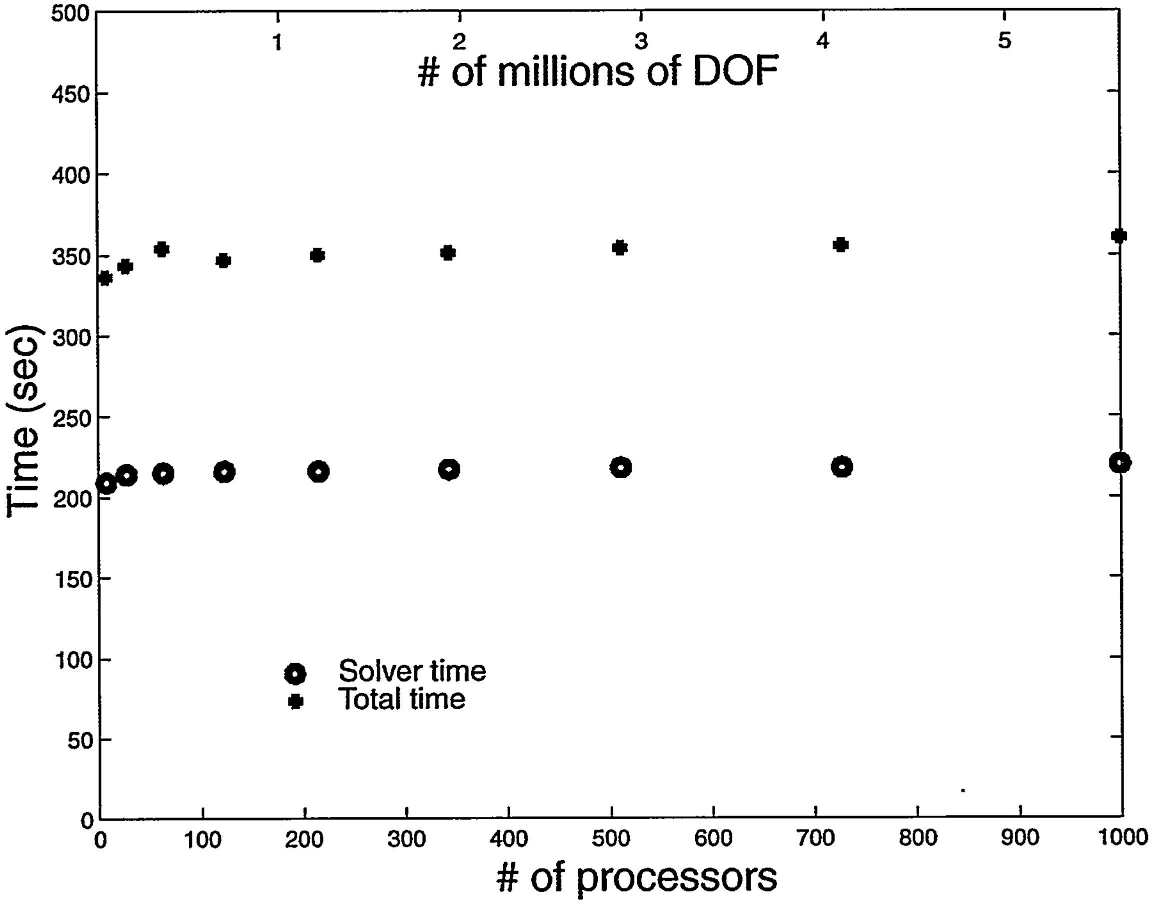

Fig. 7 Scalability of Benchmark Problem II of Solver and Total Time to 1000 processors