



Scalable Simulations of Penetration Mechanics on the SGI Origin 3800 and the IBM SP Power3 Computer Systems

by Stephen J. Schraml, Kent D. Kimsey,
and Thomas M. Kendall

ARL-TR-2537

June 2001

Approved for public release; distribution is unlimited.

20010723 168

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5066

ARL-TR-2537

June 2001

Scalable Simulations of Penetration Mechanics on the SGI Origin 3800 and the IBM SP Power3 Computer Systems

Stephen J. Schraml and Kent D. Kimsey

Weapons and Materials Research Directorate, ARL

Thomas M. Kendall

Computational and Information Sciences Directorate, ARL

Approved for public release; distribution is unlimited.

Abstract

This report presents an overview of an explicit message-passing paradigm for an Eulerian finite volume method for modeling solid dynamics problems involving shock wave propagation, multiple materials, and large deformations. Three-dimensional simulations of high velocity impact were conducted on the SGI Origin 3800 and the IBM SP Power3 computer systems. The scalability of the message-passing code on these architectures is presented and compared to the ideal linear multiple processor performance.

Contents

List of Figures	v
List of Tables	vii
1. Introduction	1
2. Scalable Paradigm	1
3. Parallel Architectures	3
4. Scalable Simulations	5
5. Scalability Results	7
6. Summary	11
7. References	13
Distribution List	15
Report Documentation Page	19

INTENTIONALLY LEFT BLANK.

List of Figures

Figure 1. CTH mesh decomposition with explicit message passing.....	3
Figure 2. Initial conditions for combined yaw and obliquity impact simulation.	5
Figure 3. Scalability of CTH on the SGI Origin 3800.....	8
Figure 4. Scalability of CTH on the IBM SP Power3.....	9
Figure 5. Comparison of CTH scalability results.....	10

INTENTIONALLY LEFT BLANK.

List of Tables

Table 1. Computational grids used in scalability study.....	7
---	---

INTENTIONALLY LEFT BLANK.

1. Introduction

The mechanics of penetration and perforation of solids have long been of interest for military applications in terminal ballistics. Kinetic energy penetration phenomena are also germane to applications involving high mass and high velocity debris attributable to accidents or high rate energy release, the transportation safety of hazardous materials, the safety of nuclear reactor containment vessels, the design of lightweight body armor, the erosion and fracture of solids because of repeated impacts by liquid or solid particles, and the protection of spacecraft from meteoroid impact.

Three-dimensional (3-D) continuum mechanics simulations of high velocity impact phenomena delineate the high performance computing resources for Army applications in terminal ballistics. Current applications in high velocity impact phenomena require the simulation time to increase from the microsecond to millisecond regime, and complex geometries dictate a finer mesh resolution. For a given computational domain, the number of cells in the domain scales inversely with the cube of the zone size. Reducing the zone size by a factor of N in each dimension will increase the number of zones (and thus the memory requirements) by a factor of N^3 . The explicit time integration scheme requires the time step to be proportional to the zone size to satisfy stability requirements. Thus, the number of integration cycles will increase as the zone size is decreased. The combined increase in number of zones and integration cycles resulting from grid refinement dictates that the processor requirements scale to the fourth power as the mesh is refined with smaller zones. These factors are strong stimuli for exploiting scalable architectures and algorithms.

In previous efforts, scalable penetration mechanics simulations were performed using a variety of commercially available parallel computer systems to evaluate the parallel performance of these architectures (Schraml and Kimsey 2000; Kimsey et al. 1998). The simulations in those studies were performed with the CTH hydrodynamics code (McGlaun and Thompson 1990). The scalability study was recently continued to evaluate the performance of two new architectures, the SGI Origin 3800 and the IBM SP Power3. This report describes the findings of the scalability studies of these systems.

2. Scalable Paradigm

CTH is an Eulerian finite volume code for modeling solid dynamics problems involving shock wave propagation, multiple materials, and large deformations in

one, two, or three dimensions. CTH is widely used across the defense research and development community to model problems in shock wave propagation. CTH employs a two-step solution scheme: a Lagrangian step followed by a remap step. The conservation equations are replaced by explicit finite volume equations that are solved in the Lagrangian step. The remap step uses operator splitting techniques to replace multidimensional equations with a set of one-dimensional (1-D) equations. The remap or advection step is based on a second order accurate method. To minimize material dispersion, several high resolution material interface trackers are available. Both analytical and tabular equations of state are available to model the hydrodynamic behavior of materials. Models for elastic-plastic behavior and high explosive detonation are also available.

Robinson et al. (1992) developed the algorithmic framework for conducting scalable Eulerian finite volume simulations for modeling problems in solid dynamics, based on object-oriented programming. Robinson demonstrated that the structured mesh of the Eulerian finite volume method is well suited for scalable paradigms employing message passing between computational subdomains.

One computing technique that can be employed on scalable computer architectures is referred to as single program multiple data (SPMD). Under the SPMD paradigm, the same executable code runs on each computational node, but each executable code works on a different set of data. Algorithms that depend on a fixed, logically connected mesh are readily adapted to the SPMD paradigm. The technique used for SPMD parallelism in CTH is similar to the formulation developed by Robinson et al. (1992) in that the entire problem domain is divided into subdomains that reside on individual computational nodes.

The use of "ghost" cells is a common technique for applying boundary conditions to finite difference and finite volume schemes, making the internal differencing computations independent of edges and corners in the Eulerian mesh. To adapt CTH to the SPMD paradigm, these ghost cells are used for passing messages between nodes. This practice of explicit message passing between subdomains allows each of the individual subdomains to have access to its neighboring subdomain's boundary cell data. Where a subdomain boundary is an external boundary of the overall computational domain, the ghost cell data are based on the appropriate boundary condition approximation. A simple example of this approach to mesh decomposition with explicit message passing is provided in Figure 1. For simplicity, the ghost cells are not shown in the primary computational domain or the subdomains of Figure 1. A thorough description of the distributed finite volume algorithm and message communication between subdomains is provided by Kimsey et al. (1998).

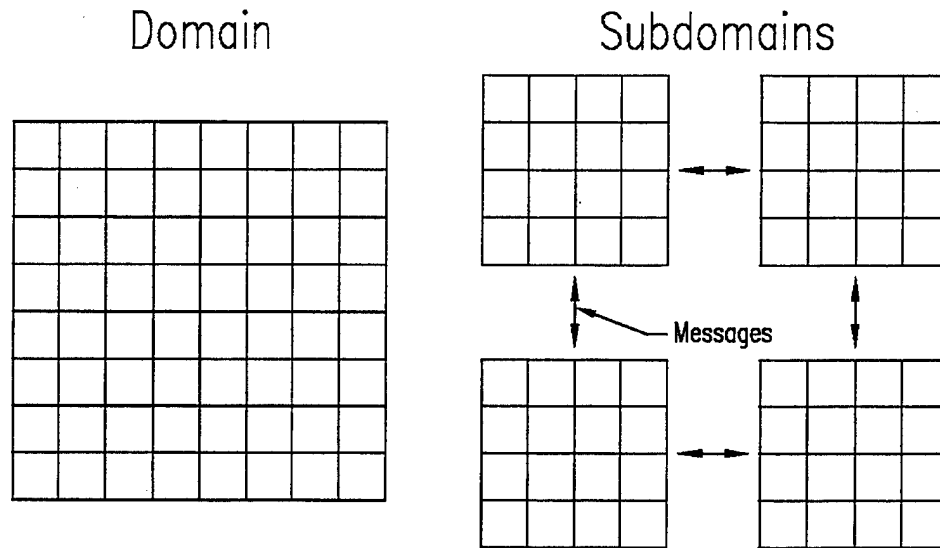


Figure 1. CTH mesh decomposition with explicit message passing.

3. Parallel Architectures

The SGI Origin 3800 is the largest in the SGI 3000 series of third generation cache-coherent non-uniform memory access (cc-NUMA) systems. The cc-NUMA architecture supports full access to the entire system memory from any processor. Both shared memory parallelism (e.g., OpenMP) and distributed memory parallelism (e.g., the message-passing interface [MPI]) models are supported. The system is modular in design and is comprised of (1) computing bricks that contain four processors and a maximum of 8 GB of memory, (2) router bricks that are used to interconnect computing bricks, and (3) input/output (I/O) bricks that provide interfaces for external connectivity (SGI Inc. 2000).

The computing brick uses MIPS R12000 processors with 8 MB of external level 2 data cache. It has a local memory subsystem with a bandwidth of 3.2 GB/s and a latency of 175 ns. A 3.2-GB/s link connects the computing brick to the router brick for access to memory on remote computing bricks. As the system is scaled to larger numbers of processors, the maximum number of router "hops" and the maximum latency increase. A 256-processor system, such as the one used for the current study, contains 64 computing bricks and 20 router bricks. The maximum number of router hops to a remote memory location is five, and the worst case latency is 485 ns – a little more than twice the latency of the memory within the local computing brick.

The current study employed 128-processor and 256-processor SGI Origin 3800 systems installed at the U.S. Army Research Laboratory (ARL) Major Shared Resource Center (MSRC). The systems at the ARL MSRC use 400-MHz processors, each containing a single floating point unit that supports a single cycle add-multiply instruction resulting in a peak floating point rate of 800 million floating point operations per second (MFLOPS). The 256-processor system has a peak floating point performance of 205 billion floating point operations per second (GFLOPS) and a peak memory bandwidth of 205 GB/s.

The IBM SP system consists of computing nodes that are interconnected via a specialized network switch. Each computing node is an independent system with its own operating system (OS), memory, and I/O devices. A single node may contain as many as 16 processors. The current investigation is focused on the 16-processor SMP node, also referred to as a "high" node. This architecture can support shared memory programming within a node and distributed memory programming within and/or between nodes. The two programming models can be mixed by employing shared memory programming within a node and distributed memory programming between nodes for multilevel parallelism. The current study employed only distributed memory programming with MPI.

The 16-processor SMP high node uses 375-MHz Power3-II processors (Amos et al. 2000). Each processor has two independent floating point units, each of which is capable of completing a multiply-add instruction every cycle resulting in a peak floating point performance of 1.5 GFLOPS per processor. The Power3-II high node memory subsystem includes (1) an on-chip 64-kB level 1 data cache and 32-kB instruction cache, (2) an off-chip 8-MB level 2 data cache, and (3) a cache-coherent switch-based memory that can be as large as 64 GB. The high node has a local memory system with a bandwidth of 16 GB/s and a latency of approximately 400 ns. Remote-to-local memory latency ratios are on the order of 100.

The IBM SP system installed at the ARL MSRC uses 32 Power3-II SMP high nodes, each with 16 GB of memory coupled with the SP switch2 interconnection system (Jennes 2000). Each node uses a single switch2 adapter, which has a peak transfer rate of 1 GB/s (500 MB/s in each direction) and a latency of approximately 50 μ s. The system topology is such that each node has a node-to-switch connection, and all switches are interconnected. A maximum of three hops is required to send a message from one node to another. The 512-processor system has a theoretical peak speed of 768 GFLOPS and a peak memory bandwidth of 512 GB/s.

4. Scalable Simulations

CTH with explicit message passing has been used to model a long rod projectile impacting an oblique steel plate on the two scalable architectures described previously. This problem was selected because of well-characterized experimental data and previous serial CTH simulations conducted by Hertel (1992). Fugelso and Taylor (1978) conducted a series of ballistic experiments to evaluate the effects of combined obliquity and yaw on high density long rod projectiles. Depleted uranium (DU) alloy long rod projectiles with little or no yaw were launched into an oblique, rolled homogeneous armor (RHA) plate that had been accelerated by an explosive charge, resulting in a yawed impact in the plate frame of reference. The DU alloy (DU 0.75%Ti) projectiles were right circular cylinders with a hemispherical nose, and the impact velocities ranged from 850 to 1650 m/s. Yaw and obliquity angles ranged from 0° to 70° and 10° to 0° , respectively, in the test series. The length and diameter of the projectile in Shot 58 of the test series are 7.67 cm and 0.767 cm, respectively, for a length-to-diameter ratio (L/D) of 10. The striking velocity was 1289 m/s, and the thickness of the RHA plate was 6.4 mm. In the laboratory frame of reference, the angle of obliquity was 73.5° , the plate velocity was 217 m/s, and the projectile velocity was 1210 m/s. In the plate frame of reference, the angle of obliquity was 64.2° , the projectile velocity was 1289 m/s, and the yaw angle was -9.3° . A schematic diagram of the initial conditions for Shot 58 is illustrated in Figure 2.




Figure 2. Initial conditions for combined yaw and obliquity impact simulation.

The scalability study was conducted with a nearly constant workload (i.e., number of computational cells on each processor for each of the simulations). This was done to keep the computation-to-communication ratio as close to constant as possible for simulations involving different numbers of processors. Maintaining a nearly constant computation-to-communication ratio and

eliminating disk access for intermediate plot and restart files during the time integration permitted the computational performance to be isolated and measured as a function of the number of processors used.

The single-processor baseline calculation used a Cartesian computational domain spanning 21.5 cm in the X direction, 3.0 cm in the Y direction, and 6.0 cm in the Z direction. The computational domain was discretized into uniform cubic zones 1 mm long, resulting in a 3-D grid of $215 \times 30 \times 60$. As the number of processors in a simulation increased, the number of zones in the model increased accordingly to maintain a nearly constant number of computational zones per processor.

All calculations were conducted for a simulated time of 40 μ s. The grid was incrementally refined by uniformly decreasing the characteristic zone length in each coordinate direction by a factor of $2^{-1/3}$. This approach doubles the total number of grid points with each successive mesh refinement. The characteristics of the grids used in the scalability study are summarized in Table 1. In this table, the columns NI, NJ, and NK refer to the number of Eulerian cells in the x , y , and z directions, respectively, and do not include ghost cells. The grid sizes listed in the table produce computational subdomains containing approximately 387,000 Eulerian cells each. For the 512-processor simulation, this results in a computational domain containing approximately 200 million Eulerian cells. An alternative to this mesh refinement technique would be to double the number of zones in one direction for one refinement, then double the number of zones in another direction for the next refinement, and so on. This approach would reduce the time step by a factor of two on the first refinement and would double the number of time integration cycles (i.e., computational cycles) to reach the desired simulation time of 40 μ s. The method of uniform zone size reduction resulted in a reduction of the time step by a factor of approximately $2^{-1/3}$ with each refinement. As a result, the number of computational cycles required to reach 40 μ s of simulated time increased only by a factor of approximately $2^{1/3}$ each time the number of processors doubled.

The scalable performance of the message-passing code is measured by the "grind time," which is the average processor time required for the code to revise all flow field variables for one computational cell in a given time increment (cycle). The grind time is expressed in units of μ s/(zone-cycle). In a case of ideal scalability, the grind time will decrease by a factor of two for every doubling of processors used if the ratio of computation to communication is held constant.

Table 1. Computational grids used in scalability study.

Number of Processors	NI	NJ	NK	Zone Length (mm)	Total Zones
1	215	30	60	1.000	387,000
2	271	38	75	0.794	772,350
4	341	48	95	0.630	1,554,960
8	430	60	120	0.500	3,096,000
16	541	76	151	0.397	6,208,516
32	683	95	191	0.315	12,393,035
64	860	120	240	0.250	24,768,000
128	1083	151	302	0.198	49,386,966
256	1366	190	382	0.157	99,144,280
512	1720	240	480	0.125	198,144,000

5. Scalability Results

All calculations in the scalability study were run to a simulated time of 40 μ s. The calculations run on the SGI Origin 3800 were run on power-of-two sets of processors between 1 and 256. The performance results from the simulations are presented in Figure 3. The measured grind times are represented by the circle symbols in the figure. The solid line represents the line of ideal scalability extrapolated from the single-processor simulation. Figure 3 shows that the measured results form a straight line, but the slope of that line does not follow the line of ideal scalability. Given the linear relationship of the measured grind time as a function of the number of processors, the actual scalability can be described by the equation

$$g_n = g_1/n^m \quad (1)$$

in which g_n is the predicted grind time for n processors, g_1 is the measured grind time from the single-processor simulation, and m is the parallel efficiency (the slope of the straight line formed by the measured results). A value of m equal to 1.0 represents the condition of ideal scalability. The actual value of m can be obtained by the application of a regression analysis to the measured data and results in a parallel efficiency of 0.878 for the SGI Origin 3800. A dashed line of this slope is plotted in Figure 3 and illustrates the divergence of the actual scalability from the ideal. The measured grind time in the single-processor calculation was 24.020 μ s/(zone-cycle). Using the parallel efficiency of 0.878 from the regression analysis and the scalability relationship in equation 1, the




Figure 3. Scalability of CTH on the SGI Origin 3800.

predicted grind time for 256 processors is $0.185 \mu\text{s}/(\text{zone-cycle})$. The measured grind time from the 256-processor simulation was $0.189 \mu\text{s}/(\text{zone-cycle})$, resulting in a performance ratio of 127 over the single-processor simulation.

The scalability simulations run on the IBM SP Power3 were run with power-of-two sets of processors between 1 and 512. Because this architecture is a collection of tightly coupled SMP nodes, there can be several different ways to configure a particular number of processors. For example, a 16-processor simulation can be run on one node using 16 processors, 16 nodes using one processor each, two nodes using eight processors each, eight nodes using two processors each, or four nodes using four processors each. In the study described here, at least one simulation was performed for every possible combination of nodes and processors per node to obtain the power-of-two processor sets.

The results from the scalability study on the IBM SP Power3 are presented in Figure 4. The measured results are represented by marker symbols in the plot. The measured results are organized by the number of processors per SMP node used in the simulations. Organizing the data in this manner helps to identify the effect of processor layout on the performance. The figure shows that the




Figure 4. Scalability of CTH on the IBM SP Power3.

measured grind time varies only slightly with the different combinations of nodes and processors per node used. This indicates that the data transfer rate between processors on different nodes is the same as that for processors on a common node.

As previously described, each SMP node in the system contains 16 processors and runs its own copy of the operating system. Before the scalability study was performed, it was considered that the cases using 16 processors per node might suffer a performance degradation as a result of contention with the OS. By using all 16 processors on the node, at least one CTH process might have to compete with the OS in its task of controlling the functions of the node. However, the results of simulations using 16 processors per node fall along the same straight line as the other results, indicating that contention with the OS is not a significant issue.

Figure 4 contains a line of ideal scalability that is extrapolated from the single-processor simulation. The measured results were combined to perform a regression analysis and resulted in a parallel efficiency, m , of 0.844, slightly less than the SGI Origin 3800. The grind time from the single-processor simulation on the IBM SP Power3 was $17.007 \mu\text{s}/(\text{zone-cycle})$, slightly faster than the single-processor simulation on the SGI Origin 3800. The 512-processor simulation on

the IBM resulted in a measured grind time of $0.082 \mu\text{s}/(\text{zone-cycle})$, resulting in a performance ratio of 208 over the single-processor case.

To compare the performance of both systems, the measured grind times and lines of linear scalability are plotted in Figure 5. The measured data from all processor configurations of the IBM have been consolidated into a single set of data and are represented by the circle marker symbol. The measured grind times from the SGI simulations are represented by the square marker symbols. The figure illustrates the faster single-processor performance of the IBM and the greater scalability of the SGI. Even though the IBM has a faster processor, the improved scalability of the SGI causes its performance to converge with that of the IBM as the number of processors is increased. The results of the 256-processor simulations for both systems overlap as the two linear scalability curves converge. For use in a production computing environment supporting large-scale continuum mechanics simulations, the performance difference between the two systems is practically negligible for processor configurations in the 256-512 range.

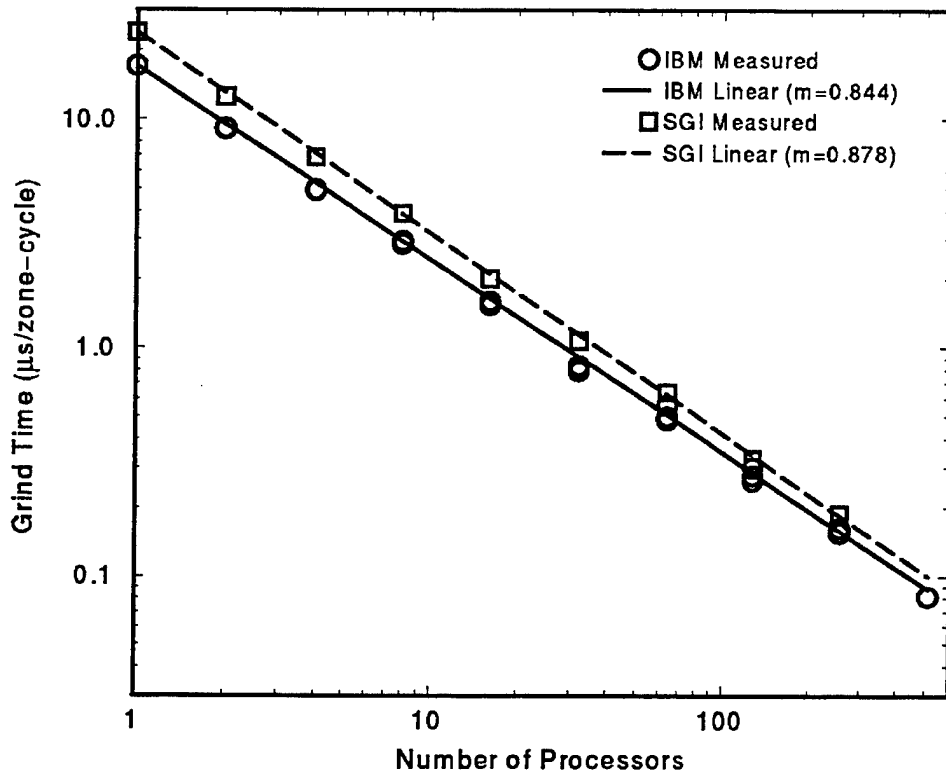


Figure 5. Comparison of CTH scalability results.

6. Summary

In a previous effort, the CTH hydrodynamics code was adapted to an SPMD programming paradigm to exploit large, scalable computer architectures. This paradigm involves the decomposition of the structured mesh into computational subdomains, with explicit message passing used to communicate data between the multiple processes used in solving the problem. This method has been previously demonstrated to scale linearly as the number of processors and corresponding problem size increased.

Two new entries into the scalable high performance computing community are the SGI Origin 3800 and the IBM SP Power3. A scalability analysis was performed on each system by running a series of 3-D parallel simulations on power-of-two sets of processors. The problem size was scaled with the number of processors to maintain a constant ratio of computation to communication. Both systems were found to scale linearly with parallel efficiencies of 0.844 for the IBM and 0.878 for the SGI. A parallel efficiency of 1.0 represents perfect scalability. Comparison of the performance of the two systems for large processor sets shows that they are both appropriate platforms for large-scale continuum mechanics analyses.

INTENTIONALLY LEFT BLANK.

7. References

- Amos, B., M. Mayfield, and F. O'Connell. "RS/6000 SP375 MHz Power3 SMP High Node." IBM Corporation, Austin, TX, 2000.
- Fugelso, E., and J. W. Taylor. "Evaluation of Combined Obliquity and Yaw for U-0.75% Ti Penetrators." LA-7402-MS, Los Alamos National Laboratory, Los Alamos, NM, 1978.
- Hertel, E. S. "A Comparison of the CTH Hydrodynamics Code With Experimental Data." SAND92-1879, Sandia National Laboratories, Albuquerque, NM, 1992.
- Jennes, P. "RS/6000 SP Switch2 Technical Whitepaper." IBM Corporation, Poughkeepsie, NY, 2000.
- Kimsey, K. D., S. Schraml, and E. S. Hertel. "Scalable Computations in Penetration Mechanics." *International Journal on Advances in Engineering Software Including Computing Systems in Engineering*, vol. 29, pp. 209-215, 1998.
- McGlaun, J. M., and S. L. Thompson. "CTH: A Three-Dimensional Shock Wave Physics Code." *International Journal of Impact Engineering*, vol. 10, pp. 351-360, 1990.
- Robinson, A. C., A. L. Ames, H. E. Fang, C. Pavlakos, C. T. Vaughan, and P. Campbell. "Massively Parallel Computing, C++, and Hydrocode Algorithms." *Proceeding of the Eighth Conference in Computing in Civil Engineering*, Dallas, TX, 1992.
- Schraml, S. J., and K. D. Kimsey. "Scalability of the CTH Hydrodynamics Code on the Sun HPC 10000 Architecture." ARL-TR-2173, U.S. Army Research Laboratory, Aberdeen Proving Ground, MD, March 2000.
- SGI Inc. *SGI 3000 Family Reference Guide*. Mountain View, CA, 2000.

INTENTIONALLY LEFT BLANK.

<u>NO. OF COPIES</u>	<u>ORGANIZATION</u>
2	DEFENSE TECHNICAL INFORMATION CENTER DTIC OCA 8725 JOHN J KINGMAN RD STE 0944 FT BELVOIR VA 22060-6218
1	HQDA DAMO FDT 400 ARMY PENTAGON WASHINGTON DC 20310-0460
1	OSD OUSD(A&T)/ODDR&E(R) DR R J TREW 3800 DEFENSE PENTAGON WASHINGTON DC 20301-3800
1	COMMANDING GENERAL US ARMY MATERIEL CMD AMCRDA TF 5001 EISENHOWER AVE ALEXANDRIA VA 22333-0001
1	INST FOR ADVNCD TCHNLGY THE UNIV OF TEXAS AT AUSTIN 3925 W BRAKER LN STE 400 AUSTIN TX 78759-5316
1	DARPA SPECIAL PROJECTS OFFICE J CARLINI 3701 N FAIRFAX DR ARLINGTON VA 22203-1714
1	US MILITARY ACADEMY MATH SCI CTR EXCELLENCE MADN MATH MAJ HUBER THAYER HALL WEST POINT NY 10996-1786
1	DIRECTOR US ARMY RESEARCH LAB AMSRL D DR D SMITH 2800 POWDER MILL RD ADELPHI MD 20783-1197

<u>NO. OF COPIES</u>	<u>ORGANIZATION</u>
1	DIRECTOR US ARMY RESEARCH LAB AMSRL CI AI R 2800 POWDER MILL RD ADELPHI MD 20783-1197
3	DIRECTOR US ARMY RESEARCH LAB AMSRL CI LL 2800 POWDER MILL RD ADELPHI MD 20783-1197
3	DIRECTOR US ARMY RESEARCH LAB AMSRL CI IS T 2800 POWDER MILL RD ADELPHI MD 20783-1197
	<u>ABERDEEN PROVING GROUND</u>
2	DIR USARL AMSRL CI LP (BLDG 305)

<u>NO. OF</u> <u>COPIES</u>	<u>ORGANIZATION</u>	<u>NO. OF</u> <u>COPIES</u>	<u>ORGANIZATION</u>
2	DOD HPCMO L DAVIS C HENRY 1010 N GLEBE RD SUITE 510 ARLINGTON VA 22201	1	DIR US ARO WASH AMXRO W KA BANNISTER ROOM 8N31 5001 EISENHOWER AVE ALEXANDRIA VA 22333-0001
1	US MILITARY ACADEMY MATH SCIENCES CTR OF EXC DEPT OF MATH SCIENCES MDN A MAJ DON ENGEN THAYER HALL WEST POINT NY 10996-1786	3	COMMANDER US ARMY RESEARCH OFFICE J CHANDRA K IYER TECH LIB PO BOX 12211 RESEARCH TRIANGLE PARK NC 27709-2211
4	COMMANDER US ARMY ARDEC AMSTA AR AEE WW E BAKER R FONG J WALSH C CHIN PICATINNY ARSENAL NJ 07806-5000	1	COMMANDER US ARMY CORPS OF ENGINEERS T BLEVINS 3909 HALL FERRY ROAD VICKSBURG MS 39180-6199
1	COMMANDER US ARMY ARDEC AMSTA AR AEE WW TECH LIBRARY PICATINNY ARSENAL NJ 07806-5000	3	DARPA J RICHARDSON B WILCOX TECH LIBRARY 3701 N FAIRFAX DR ARLINGTON VA 22203-1714
1	COMMANDER US ARMY ARDEC AMSTA AR AET M TECH LIBRARY PICATINNY ARSENAL NJ 07806-5000	2	CDR US ARMY TACOM AMSTA RSK J THOMPSON S GOODMAN WARREN MI 48397-5000
1	COMMANDER US ARMY ARDEC AMSTA AR FS E ANDRICOPOULOS PICATINNY ARSENAL NJ 07806-5000	2	DIR NRL J A NEMES A E WILLIAMS CODE 6684 4555 OVERLOOK AVE SW WASHINGTON DC 20375
3	COMMANDER US ARMY MERDEC AMSME RD ST WF L CRAFT D LOVELACE M SCHEXNAYDER REDSTONE ARSENAL AL 35898-5250	3	CDR NSWC W H HOLT W MOCK CODE G22 TECH LIBRARY 17320 DAHLGREN ROAD DAHLGREEN VA 22448-5000
		1	COMMANDER NWC T GILL CODE 3261 CHINA LAKE CA 93555-6001

<u>NO. OF COPIES</u>	<u>ORGANIZATION</u>
1	NAVAL POST GRAD SCHOOL J STERNBERG CODE 73 MONTEREY CA 93943
1	USAF PHILLIPS LAB PL WSCD F ALLAHDADI KIRTLAND AFB NM 87185
3	USAF WL MNMW W COOK M NIXON TECH LIBRARY EGLIN AFB FL 32542-5434
1	ALLIANT TECHSYSTEMS INC R STRYK 600 SECOND ST NE HOPKINS MN 55343
1	APPLIED RSCH ASSOCIATES INC J D YATTEAU 5941 S MIDDLEFIELD RD STE 100 LITTLETON CO 80123
11	DIR LOS ALAMOS NATL LAB J BOLSTAD CIC ACL T CARNEY EES 5 MS C305 E FERM DX 3 MS P940 G T GRAY MST 8 MS G755 K HOLIAN CIC 12 MS D413 L HULL DX 3 MS P940 J JOHNSON F663 L LIVERSKY X HM MS D413 D A MANDELL X CI MS F663 L SCHWALBE X HM MS D413 TECH LIBRARY PO BOX 1663 LOS ALAMOS NM 87545
6	DIR LAWRENCE LIVERMORE NATL LAB C HOOVER L125 D LASSILA L170 P RABOIN L125 J E REAUGH L282 R E TIPTON L170 TECH LIB PO BOX 1663 LIVERMORE CA 94550

<u>NO. OF COPIES</u>	<u>ORGANIZATION</u>
2	DIR SANDIA NATL LAB D BAMMANN MS 9405 M CHIESA MS 9042 PO BOX 969 LIVERMORE CA 94550
15	DIR SANDIA NATL LAB R L BELL MS 0820 R M BRANNON MS 0820 D E CARROL MS 0819 L C CHABILDAS MS 1181 M G ELRICK MS 0819 M FORRESTAL MS 0315 E S HERTEL JR MS 0819 M KIPP MS 0820 J M MCCLAUN MS 1202 J S PEERY MS 0819 S A SILLING MS 0820 R M SUMMERS MS 0819 P A TAYLOR MS 0820 T TRUCANO MS 0819 P YARRINGTON MS 0820 PO BOX 5800 ALBUQUERQUE NM 87185-0307
4	APPLIED RSCH ASSOCIATES INC J CREPEAU S HIKIDA C NEEDHAM R NEWELL 4300 SAN MATEO BLVD SE STE A220 ALBUQUERQUE NM 87110
1	COMPUTATIONAL MECH ASSN J A ZUKAS PO BOX 11314 BALTIMORE MD 21239-0314
1	IAT UNIV OF TEXAS AT AUSTIN PO BOX 202797 AUSTIN TX 78720-2797
1	IAT UNIV OF TEXAS AT AUSTIN D LITTLEFIELD 4030 2 W BRAKER LN AUSTIN TX 78759-5329

NO. OF
COPIES ORGANIZATION

- 2 MAXWELL TECHNOLOGIES
J BARTHEL
S ROGERS
8888 BALBOA AVE
SAN DIEGO CA 92123
- 1 ORLANDO TECH INC
D A MATUSKA
PO BOX 855
SHALIMAR FL 32579
- 2 SOUTHWEST RSRCH INSTITUTE
C ANDERSON
J WALKER
8500 CULEBRA ROAD
PO DRAWER 28510
SAN ANTONIO TX 78284

ABERDEEN PROVING GROUND

- 27 DIRECTOR
US ARMY RESEARCH LAB
AMSRL CI H
T KENDALL
R NAMBURU
P PAPADOS
AMSRL WM T
B BURNS
AMSRL TM TA
W BRUCHEY JR
Y HAUNG
P KINGMAN
D KLEPONIS
H MEYER JR
AMSRL WM TB
R FREY
J STARKENBERG
R LOTTERO
AMSRL WM TC
R COATES
K KIMSEY
D SCHEFFLER
S SCHRAML (5 CPS)
AMSRL WM TD
R L BITTING
A M DIETRICH JR
F GREGORY
D J GROVE
M RAFTENBERG
S SCHOENFELD
S SEGLETES

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 2001	3. REPORT TYPE AND DATES COVERED Final, August 2000–March 2001	
4. TITLE AND SUBTITLE Scalable Simulations of Penetration Mechanics on the SGI Origin 3800 and the IBM SP Power3 Computer Systems			5. FUNDING NUMBERS 1L162618AH80	
6. AUTHOR(S) Stephen J. Schraml, Kent D. Kimsey, and Thomas M. Kendall				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: AMSRL-WM-TC Aberdeen Proving Ground, MD 21005-5066			8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-2537	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) <p>This report presents an overview of an explicit message-passing paradigm for an Eulerian finite volume method for modeling solid dynamics problems involving shock wave propagation, multiple materials, and large deformations. Three-dimensional simulations of high velocity impact were conducted on the SGI Origin 3800 and the IBM SP Power3 computer systems. The scalability of the message-passing code on these architectures is presented and compared to the ideal linear multiple processor performance.</p>				
14. SUBJECT TERMS computer simulation, penetration mechanics, shock waves			15. NUMBER OF PAGES 22	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

INTENTIONALLY LEFT BLANK.

USER EVALUATION SHEET/CHANGE OF ADDRESS

This Laboratory undertakes a continuing effort to improve the quality of the reports it publishes. Your comments/answers to the items/questions below will aid us in our efforts.

1. ARL Report Number/Author ARL-TR-2537 (Schraml) Date of Report June 2001

2. Date Report Received _____

3. Does this report satisfy a need? (Comment on purpose, related project, or other area of interest for which the report will be used.) _____

4. Specifically, how is the report being used? (Information source, design data, procedure, source of ideas, etc.) _____

5. Has the information in this report led to any quantitative savings as far as man-hours or dollars saved, operating costs avoided, or efficiencies achieved, etc? If so, please elaborate. _____

6. General Comments. What do you think should be changed to improve future reports? (Indicate changes to organization, technical content, format, etc.) _____

CURRENT
ADDRESS

Organization

Name

E-mail Name

Street or P.O. Box No.

City, State, Zip Code

7. If indicating a Change of Address or Address Correction, please provide the Current or Correct address above and the Old or Incorrect address below.

OLD
ADDRESS

Organization

Name

Street or P.O. Box No.

City, State, Zip Code

(Remove this sheet, fold as indicated, tape closed, and mail.)
(DO NOT STAPLE)