A Performance Study of LS-DYNA on Vehicle Crash Simulation

Yih-Yih Lin
Hewlett-Packard Company
PMB 225
100 Powdermill Road
Acton, MA 01720
(800) 477-6111 Ext. 5278
ylin@rsn.hp.com

Cing-Dao Kan
FHWA/NHTSA National Crash Analysis Center
The George Washington University
20101 Academic Way
Ashburn, VA 20147
(703) 729-8365
cdkan@ncac.gwu.edu

Abbreviations:

ccNUMA	Cache Coherence/Non-Uniform Memory Access
CPU	Central Processing Unit
DOF	Degree of Freedom
DP	Double Precision
MPP	Massive Parallel Processing
OS	Operating System
SMP	Symmetrical Multiprocessing
SP	Single Precision
VPG	Virtual Proving Ground

Keywords:
Performance, Memory System, Precision, SMP and MPP, Cluster and ccNUMA

ABSTRACT

For computer companies to design cost-effective machines for LSTC to optimize the software and for the LS-DYNA user to make an intelligent choice of machines, various factors that affect LS-DYNA's performance must be investigated and understood. Therefore, we have embarked such an investigation with machines from both Hewlett-Packard and SGI. Many results are reported here. The factors considered include the following: memory systems, SMP vs. MPP LS-DYNA, single vs. double precision, cluster vs. ccNUMA configuration.

INTRODUCTION

It is well known that vehicle crash simulation is very computer-intensive [1]. For example, a crash model that has 2 million degrees of freedom (DOF) simulated for 150 milliseconds, a typical current production problem, would take more than 6 days on an HP N-Class 4000, the current leading RISC processor in the industry, running 1-CPU LS-DYNA. To effectively tackle this problem of long computing time requires efforts from computer companies, the software developer (LSTC), and the LS-DYNA user. First, computer companies should keep producing new machines that are cost-effective and perform better. Second, LSTC should optimize the code. Third, the LS-DYNA user should make a good choice of machines and use them intelligently. Clearly, for each of the three parties to achieve the stated goal, factors that affect LS-DYNA have to be investigated and understood. Therefore, we have embarked on such an investigation with machines from both Hewlett-Packard and SGI. The factors considered include the following: memory systems, SMP vs. MPP LS-DYNA, single vs. double precision, cluster vs. ccNUMA configuration.

APPROACH

Machines, Models and Codes

In this investigation, we have used an 8-CPU, 440 MHz HP N-Class 4000 (N4000); a 32-CPU, 440 MHz HP V-Class 2500 (V2500); and a 6-CPU, 250 MHz SGI ORIGIN 2000. In fact, the N4000 and the V2500 use the same CPU of HP PA-RISC 8500. The Neon model, in which a Neon car crashes into a barrier at 35 MPH, is used; see http://www.ncac.gwu.edu for details. The model has 1.7 million DOF, 325 components, and 6 contacts. The simulation time in this study is 30 milliseconds.

Furthermore, the versions of SMP and MPP LS-DYNA used in this study are both 940.2a. And it is noteworthy that the same SMP or MPP LS-DYNA executables have been used for the N4000 and the V2500.

Memory Systems

The N4000 and the V2500 have the same CPU at the same clock rate of 440 MHz and run on the same OS, HP-UX 11. The main difference is in their memory systems. Both the N4000 and the V2500 are shared memory and symmetrical multi-processing (SMP) machines. The N4000 has a simple memory bus. On the other hand, the V2500 has the elaborate memory scheme of an 8x8 crossbar, which is required to deliver high memory bandwidth. A simple memory bus costs less to build and has a smaller latency than a crossbar scheme. The only drawback of a simple memory bus is that there is not enough memory bandwidth to support an SMP system with a large number of CPUs. Hence the N4000 has a maximal number of CPUs of 8, while the V2500 has 32.

A single SMP system, such as the 8-CPU N4000 or the 32-CPU N4000, is called a *node*. It is possible to assemble a collection of several nodes with some interconnections to become a single SMP system under the control of one OS. Such an assembly is called an SMP multi-node system. Inevitably, the bandwidth of a simple bus is too small to sustain enough bandwidth for an SMP multi-node system. Instead, the current industry practice is

to use nodes with an elaborate memory scheme within a node level, such as the crossbar in the V2500. Interconnections for an SMP multi-node system is typically another layer of memory, so-called global shared memory. Presently, most SMP multi-node systems have a built-in scheme to ensure cache coherency (cc) and their memory access time is non-uniform (NUMA) depending on the location of a data word in memory [2]. Such a multi-node system is called a ccNUMA system. Since a single SMP node with elaborate memory scheme is essential for building a ccNUMA system, for the convenience of naming, we will also call such a single node as a ccNUMA system. By this token, both the V2500 and the ORIGIN2000 are ccNUMA systems, even though they have just a single node.

To study the effect of memory systems on the performance of LS-DYNA, we compare the performance of SMP LS-DYNA between the N4000 and the V2500. The result is shown in Table and Figure 1, where two conclusions can be drawn:

- At lower number of CPUs from 1 to 8, the N4000 outperforms the V2500 by a factor of about 1.7.

- However, as the number of CPUs increases, the speedup factor of the N4000 lags behind that of the V2500.

Machine	Execution Time in Seconds--Speedup / Number of Processors							
	1		2		4		8	
N4000/SMP	65,189	1.00	35,588	1.83	20,851	3.13	14,876	4.38
V2500/SMP	110,281	1.00	65,183	1.69	36,581	3.01	23,842	4.63
Ratio	1.69		1.83	0.92	1.75	0.96	1.60	1.06

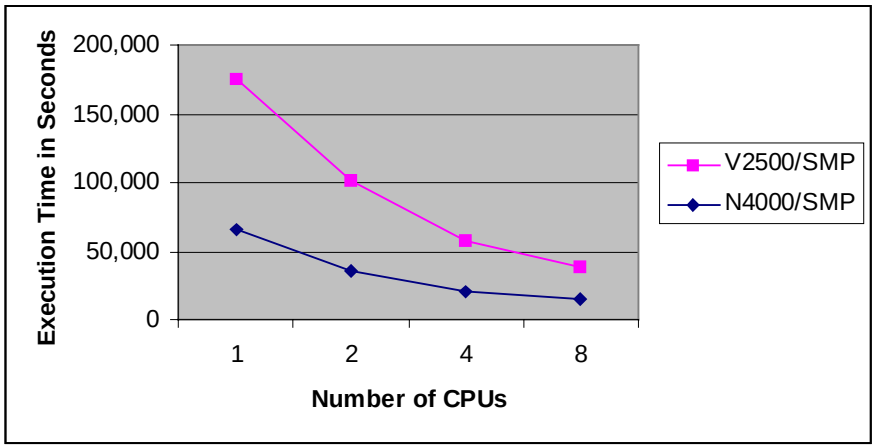

Table and Figure 1. Comparison of the performance of SMP LS-DYNA on two machines that have the same clock rate but different memory systems

These observations show the profound effect of memory systems on the performance of LS-DYNA. Considering a single SMP system, there are two opposing factors at work: memory latency and memory bandwidth. The memory latency favors an SMP system with a simple memory bus, but the memory bandwidth favors an SMP system with an elaborate scheme, such as crossbar, and allows it to have a better speedup, or in other words, scalability.

DISCUSSION OF RESULTS

The Crossover Point between MPP and SMP LS-DYNA

It is well known that at a larger number of CPUs with the MPP LS-DYNA, always outperforms SMP LS-DYNA. However, at a small number of CPUs, it is the opposite—The SMP outperforms the MPP LS-DYNA at a lower number of CPUs. Table and Figure 2 shows the performance of both SMP and MPP LS-DYNA at lower numbers of CPUs, from 1 to 8, on the three machines N4000, V2500, and ORIGIN2000. Several observations can be drawn:

- The execution time ratio between 1-CPU MPP and SMP LS-DYNA is about 1.3 for all three machines.

- The number of CPUs, at which the execution times of SMP and MPP LS-DYNA are equal, is 6 for both the V2500 and the ORIGIN2000, both of which are ccNUMA systems; and it is 8 for the N4000, which has a simple memory bus.

Machine/Code	Execution Time in Seconds / Number of Processors				
	1	2	4	6	8
HP N4000/SMP	65,189	35,588	20,851		14,876
HP N4000/MPP	88,568	55,203	26,960		14,874
HP V2500/SMP	110,281	65,183	36,581		23,842
HP V2500/MPP	140,133	82,608	40,555		21,269
SGI ORIGIN/SMP	134,569	69,251	45,620	38,972	
SGI ORIGIN/MPP	182,254	108,956	49,624	38,092	

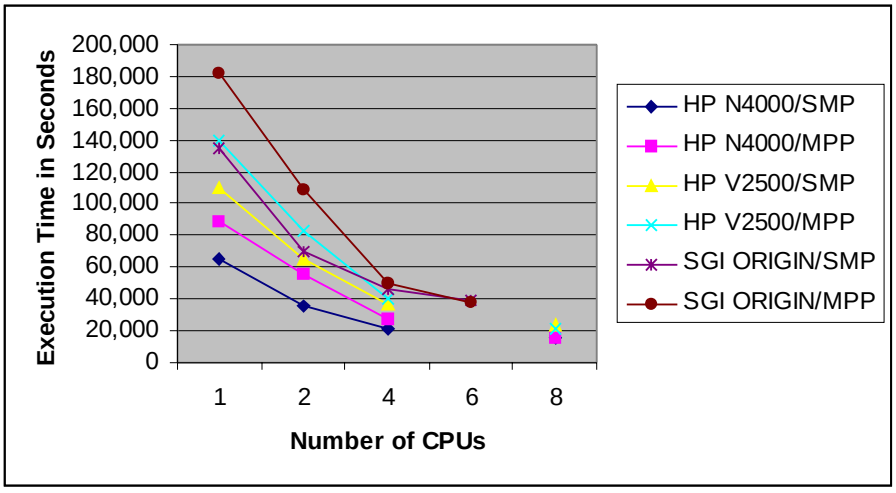

Table and Figure 2. The Performance of SMP and MPP LS-DYNA at a lower number of CPUs

We speculate that the main reason for the performance difference between SMP and MPP LS-DYNA is because they have used different algorithms in the computation of contacts. Clearly, since it is also true that presently SMP LS-DYNA has more features than MPP LS-DYNA does, the following advice to the LS-DYNA user could be made: If you have a system with no more than 8 CPUs, it is more advantageous to use SMP LS-DYNA than MPP LS-DYNA.

Precision

The numerical computation of LS-DYNA is an iterative process, which inevitably causes more accumulation of round-off errors as the simulation time increases. Although the commonly-used 32-bit, single precision (SP) version of LS-DYNA suffices most crash simulation problems, the 64-bit, double precision (DP) version is a must for the simulation of slowly moving vehicles that may last up to 5 seconds, as in ETA's Virtual Proving Ground (VPG) [3]. Modern CPUs, like HP PA-RISC 8500, will perform single and double precision arithmetic at the same speed. However, the double precision version not only requires the access of memory twice that of the single precision, but also renders the data cache only half effective as the single precision, because the fixed-sized cache can only hold half as many double elements as single. Consequently, the double precision version will be slower than the single precision version. Table and Figure 3 shows the performance comparison of single and double precision versions of SMP LS-DYNA on the N4000. It can be concluded that in general the double precision version is about 1.5 time slower than the single precision.

It will be also of great interests to see how the double precision version of MPP LS-DYNA compares with the single precision version. Since MPP LS-DYNA uses the domain decomposition paradigm, each processor in its n-CPU job will require memory about 1/n-th of its 1-CPU job. Additionally, because each processor of MPP LS-DYNA works only on one sub-domain, the locality of the data will be better preserve in MPP LS-DYNA than in SMP LS-DYNA. Therefore, we predict that the slowdown from single to double precision versions of MPP LS-DYNA will be much less than that of SMP LS-DYNA.

Codes	**Execution Time in Seconds / Number of CPUs**			
	1	**2**	**4**	**8**
SMP SP	65,189	35,588	20,851	14,876
SMP DP	94,602	51,959	30,641	23,126
Ratio	1.45	1.46	1.47	1.55

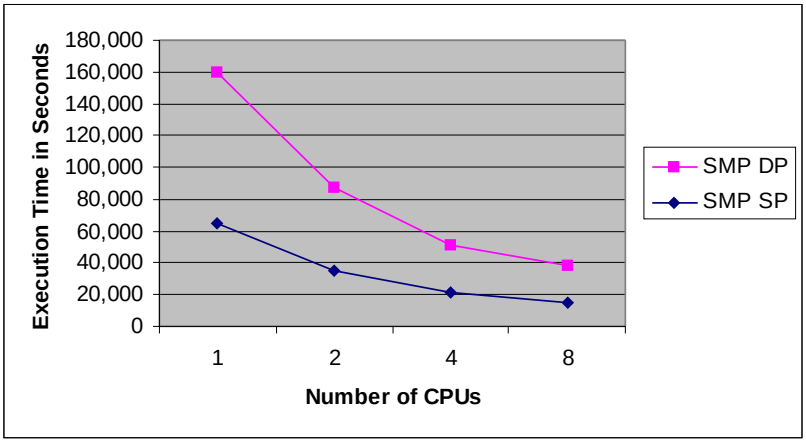

Table and Figure 3. Comparison of the performances of double and single precision SMP LS-DYNA

Cluster or ccNUMA for MPP LS-DYNA

As mentioned previously, both the N4000 and the V2500 have the same PA-RISC 8500 CPU at 440 MHz. Their main difference is in their memory architectures. While the N4000 has a simple memory bus, the V2500 has the elaborate memory scheme of

crossbar, which is essential to assemble several of them to become a ccNUMA system. On one hand, as shown previously, the crossbar has incurred performance penalty so that the N4000 consistently outperforms the V2500 up to 8 CPUs. On the other hand, because of not having the crossbar to sustain high bandwidth, the N4000 cannot be built as a single system with more than 8 CPUs. Does this mean that the N4000 cannot be used to harness the MPP LS-DYNA at a great number of CPUs? Fortunately, the answer is no.

Because it is based on MPI, a message-passing interface, MPP LS-DYNA can run on any collection of computers that are interconnected not by global shared memories, as in a ccNUMA machines, but by network switches. Such a collection of computers is called a cluster. In MPP LS-DYNA, the execution time of a job can be divided into two parts: compute time and communication time. In the case of cluster, communication time is basically the time required for messages passing through the interconnection. It is known that if the interconnection is the well-known Ethernet, then the communication time will dominate the compute time so that MPP LS-DYNA will scale badly. However, in this paper we will show that if the HP's very high-speed Hyperfabric network switch is used as the interconnect, the cluster configuration is a suitable technology for MPP LS-DYNA to gain scalability.

Machine	**Execution Time in Seconds -- Speedup / Number of Processors**											
	1		**2**		**4**		**8**		**16**		**32**	
HP N4000	88,568	1.00	55,203	1.60	26,960	3.29	14,874	5.95	9,193	9.63	6,636	13.35
HP V2500	140,133	1.00	82,608	1.70	40,555	3.46	21,269	6.59	12,646	11.08	7,068	19.83

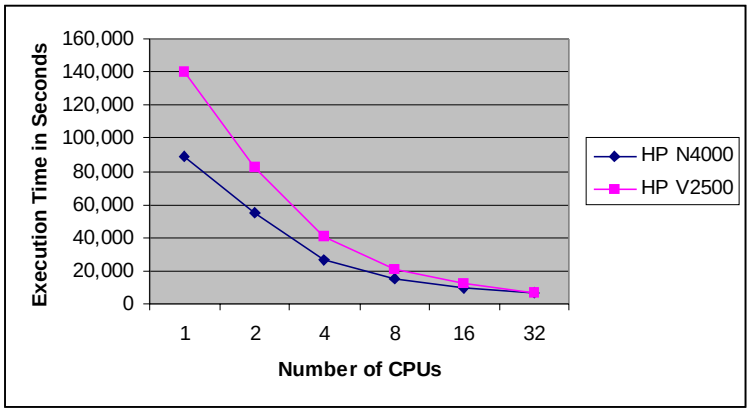

Table and Figure 4. The Performance of MPP LS-DYNA on a cluster of 4 N4000s and a V2500

Table and Figure 4 shows the performance of the 4-way cluster of N4000s and the ccNUMA system V2500, both of which have 32 PA-RISC 8500 CPUs, as mentioned before. Several observations can be drawn:

- Although the 4-way cluster of N4000s does not scale as well as the ccNUMA system V2500, the former still performs better than the latter from 1 up to 32 CPUs. The main reason is that at a lower number of CPUs—less than 8, the maximal number a single N4000 system can have—the N4000 is about 1.5 times faster than the V2500.

- In fact, at 16 CPUs, because the scalability of the cluster is close to that of the V2500 (9.6 vs. 11), the cluster outperforms the V2500 by a factor of 1.4. Furthermore, by extrapolating, the cluster, at 16 CPUs, can perform the full vehicle simulation of 150 milliseconds in about half a day.

- That at 32 CPUs the cluster does not scale as well can be explained by the communication cost. The cost of 4-way communication, as occurring in the 32-CPU run, is much higher than that of 2-way communication, as occurring in the 16-CPU run. This means that in order to improve the scalability of the cluster, running MPP LS-DYNA, beyond 16 CPUs, the major improvement is in the area of the communication speed.

Furthermore, it is noteworthy that a cluster configuration has several advantages over a ccNUMA system:

- A cluster costs less to build than a ccNUMA system of the same number of the same CPUs.

- When a ccNUMA system breaks down, the whole system is unusable. On the other hand, when one node of a cluster breaks down, all other good nodes can still be configured to form a new cluster to run MPP LS-DYNA. Thus, a cluster is less susceptible to downtime than a ccNUMA system.

- By its nature, when two MPP LS-DYNA jobs run at the same time on two disjoint sets of nodes of a cluster, they are truly independent from each other, unlike on a ccNUMA system, where some global interconnect is shared. Consequently, a cluster is a more ideal configuration than a ccNUMA for throughput.

CONCLUSIONS

First, we have shown that the memory system of an SMP machine plays an important role in the performance of LS-DYNA. Two opposing factors of memory systems are at work: memory latency and memory bandwidth. Memory latency favors a system with a simple memory bus, but memory bandwidth favors a system with an elaborate scheme such as crossbar. Second, by measuring the performance of SMP and MPP LS-DYNA on 3 computer systems from HP and SGI, we have found that it is more advantageous to use SMP than MPP LS-DYNA at a lower number of CPUs, from 1 to 8. Third, we have shown experimentally that the double precision version of SMP LS-DYNA is about 1.5 times slower than the single precision. This slowdown is predicted to be less with the double precision version of MPP LS-DYNA because it uses domain decomposition paradigm where memory requirement is less and locality is preserved. Finally, by measuring the performance of MPP LS-DYNA on a 4-way cluster and a ccNUMA system, consisted of the same type of CPUs, we have found that cluster is a viable technology ready to harness the power of MPP LS-DYNA.

REFERENCES

[1] LIN, Y. and KAN, C.D. [1998] Crash Simulation on Parallel Multiprocessors. 5th International LS-DYNA Users Conference.

[2] PATTERSON, D. and HENNESSY, J. [1996] Computer Architecture: A Quantitative Approach. Second Edition. Morgan Kaufmann Publishers, Inc, pp. 635-760.

[3] BENNETT, C. and PALMER, T. [1998] Computer System Requirements for Performing Virtual Proving Ground Analyses Using LS-DYNA. 5th International LS-DYNA Users Conference.