Grid-Based LS-DYNA Solutions

Authors:

Henry H. Fong, HPTC Marketing, Sun Microsystems, Inc., Menlo Park CA
Jonas Edberg, PhD. Client Services, Sun Microsystems, Inc., Stockholm, Sweden

Correspondence:

Henry H. Fong

(650) 352-5828 tel
(650) 786-6455 fax
henry.fong@sun.com

Keywords:

Sun's Opteron Servers/Workstations
LS-DYNA Performance on Sun Opteron & SPARC Clusters
Cold Rolling Simulation
Sun N1 Grid Engine
Grid Computing in Manufacturing

ABSTRACT

In addition to offering SPARC/Solaris systems for MCAE users, Sun Microsystems also markets Opteron-based servers and workstations. These Opteron solutions offer excellent price/performance for LS-DYNA simulations, and are particularly suitable for clustering. Users have a choice of three operating systems: Solaris 10 x64; Linux (SuSE or Red Hat); or Microsoft Windows. This paper describes briefly these computing systems, and show some typical LS-DYNA performances on standard benchmarks such the 3-car-crash problem. Also cited is a Mefos cold rolling benchmark, run using both a SPARC/Solaris Sun Fire V480 server as well as a cluster of Sun LX50 Xeon-based processors. Sun's cluster systems typically utilize Sun's N1GE grid engine load management software to schedule, manage, and prioritize compute jobs.

GRID-BASED VIRTUAL PROTOTYPING SIMULATION IN MCAE

Numerical simulation of real-world problems has become standard practice among manufacturers worldwide, allowing them to reduce costs, meet short time-to-market windows while improving product quality, optimizing materials usage, and addressing increasingly stringent governmental safety and environment regulations. This type of simulation is called MCAE (mechanical computer-aided engineering), or sometimes Virtual Product Development (VPD), and typically use finite element analysis (FEA) methods. With Sun's Grid Computing technologies, the latest AMD Opteron processor-based servers and workstations, Sun deliver outstanding performance and scalability for MCAE/VPD applications, allowing manufactures (automotive, aerospace, consumer goods, biomedical, etc.) to produce better and safer products at affordable costs.

THE MEFOS COLD ROLLING BENCHMARK

The model used in this Mefos benchmark is used for simulation of the cold rolling process. Due to symmetry, only half the upper roll and a quarter of the strip are analyzed. The roll is rotating with a peripheral velocity of 6.0 m/s at the outer radius. The strip in the model is 2.4 m long, 3.0 mm thick, and 1.0 m wide before the rolling. This model length is probably large enough to reach the steady-state condition which prevails in the central part of the real strip, which is much longer in reality. The outer diameter of the work roll is 62.35 mm, and the gap between the work rolls is 2.6 mm before rolling. The work roll is modeled as a rigid surface. There is a back and front tension applied to the strip of 235 and 300 MPa, respectively. The density of the materials is 25 times larger than the real density for steel, in order to speed up the calculations. The finite element model has 1,662,202 nodes and 580,000 solid elements.

The rolling process takes about 0.4 seconds. It is found that there is a steady-state condition relative to the rolling section, prevailing from about 0.12 to 0.28 seconds. In the benchmark, the simulation time is 0.01 second. Therefore, the execution time of a full LS-DYNA run is approximately 30 times longer than the one in the benchmark.

The friction along the slideline between the rolls and the plate is modeled as Coulomb friction. The friction coefficient is assumed to be 0.07 with lubricated rolls. Young's modulus for the plate material is assumed to be 206 GPa, and Poisson's ratio is set to 0.3. The yield criterion used is von Mises and the associated flow rule is used. Variable isotropic hardening is assumed. The effective stress-effective plastic strain relation is given from linear interpolation between the values in the following table:

Strain	Effective Stress (MPa)
0.00000	9.1000+8
0.15415	1.0100+9
0.27820	1.1330+9
0.44183	1.1660+9

The material is not very strain rate sensitive.

This Mefos cold rolling benchmark has been run using three Sun systems that use SPARC/Solaris or Linux:

- Cluster of 16 Sun Fire LX50 – Xeon-based, with copper gigabit Ethernet
- Sun Fire 6800 servers
- Cluster of 32 Sun Fire V480's – Myrinet interconnect

All these systems showed impressive LS-DYNA scalability on this Mefos cold rolling benchmark.

The benchmark results are summarized below:

1. Sun LX50 Cluster (Xeon-based), 1 cpu per system

#CPU's	*Total Seconds*	*Efficiency*
1	63,130	1.0
2	32,540	0.97
4	16,286	0.97
8	9,597	0.82
12	5,925	0.89
16	5,004	0.79

2. Sun LX50 Cluster, 2 cpu's per system

4	21,560	0.67
8	12,655	0,62
12	8,482	0.62
16	6,872	0.57
24	4,816	0.55
32	3,760	0.53

3. Sun Fire 6800 SMP-server, 24 x 1,200 MHz

#CPU's	*Total Seconds*	*Efficiency*
4	19,771	0.89
8	10,557	0.83
12	6,684	0.87
24	3,259	0.89

4. Sun Fire V480 Cluster, Myrinet interconnect

32 (6x2)	2,400	0.91
128 (32x4)	747	0.73

Note that on the Sun Fire 6800 server, the benchmark was never run on 1 cpu. The efficiency figure for 4 cpu's was taken on a Sun Fire 6800 server, 24 x 900 MHz, and then assumed to be the same for 4 cpu's on the Sun Fire 6800, 24 x 1,200 MHz system. The rest of the efficiency numbers for this system are calculated using this assumption. The same is true for the cluster of 32 Sun Fire 480's that use the same CPU's as the Sun Fire 6800 24 x 1,200 MHz system. Here, we are also assuming that single CPU performance on the Sun Fire 6800, 24 x 1,200 MHz and the Sun Fire V480 is identical.

In this example, LS-DYNA scales very well, even up to 128 cpu's. Of course, this excellent scalability may be due to the simplicity of the model. But, the model comes from real-world cold rolling applications, and the benchmark was not designed with scalability in mind. There is may similarities with the model presented in [1] that has been used for simulation of the hot rolling process. It is also noteworthy that a number of different domain decomposition strategies have been tried. None of the results, however,gave any real significant improvement over the default one shown in the tables above. The results clearly showed that LS-DYNA can be used in a large cluster without suffering too much from parallel overhead, provided that the model itself does not introduce any load imbalances.

SUN'S OPTERON-BASED SERVERS AND WORKSTATIONS

Sun Microsystems markets a family of AMD Opteron-based low-end servers and workstations:

- Sun Fire V20z server, up to 2 Opteron 200 series processors, up to 16 GB memory
- Sun Fire V40z server, up to 4 Opteron 800 series processors, up to 32 GB memory
- Sun Java Workstation W1100z, 1 Opteron 100 series processor, up to 8 GB memory
- Sun Java Workstation W2100z, 2 Opteron 200 series processors, up to 16 GB memory

Using the latest AMD Opteron processors, these systems offer outstanding performance, at competitive prices, on all the leading MCAE and MCAD applications. They are especially attractive for Manufacturing/MCAE organizations, because benefits include:

- Large memory capacity for complex models and data sets
- High-speed 3D graphics for pre- and post-processing
- Choice of Solaris 10 operating system, Linux, or Microsoft Windows
- 64-bit systems designed for high performance, with excellent scalability
  and 32-bit compatibility.

Racked clusters may be assembled using these Sun Fire V20z and V40z servers, in configurations of 8-, 16-, and 32-cpu's. Various interconnect options are available: gigabit Ethernet; Myricom/Myrinet; or Infiniband (Topspin, Voltaire, InfiniCon, etc.). Sun's N1GE Grid Engine software is used for load management, to efficiently distribute crash, metal forming, or drop impact workloads to the available computing resources in an organization. Sun's CRS (Customer Ready Systems) organization is available to help customers assemble desired configurations.

LS-DYNA PERFORMANCE ON 3-CAR-CRASH BENCHMARK

Sun's Opteron-based servers offer good floating point performance and large memory support. Clusters of these systems will provide users excellent scalability on codes such as LS-DYNA. The following typical 3-car-crash benchmark (car models from FHWA/NHTSA National Crash Analysis Center at George Washington University – benchmark from www.topcrunch.org, developed by Professor David Benson of University of California at San Diego) shows the following results on Sun's cluster of V20z Opteron-based servers [2004]:

CPU's	*Elapsed Time (hrs:min)*	*Scalability*	*Efficiency (%)*
1	79:04	1.00	100.0
2	41:14	1.92	96.0
4	20:37	3.83	95.8
8	12:07	6.52	81.5
16	5:56	13.3	83.1
24	4:40	16.9	70.4

These results were obtained using a Sun Fire V20z Myrinet cluster (AMD Opteron 248, 2.2 GHz), SuSE Linux, in April 2004. The latest AMD Opteron processor is at 2.6 GHz, and we expect the LS-DYNA results to scale and improve accordingly. The 3-car-crash model cited uses 794,078 elements, and the simulation time was 150 msec. The model was developed by LSTC's Mike Burger and Professor David Benson of UCSD.

LS-DYNA scalability is excellent on this 3-car-crash benchmark up to 16 cpu's, and fairly good even at 24 cpu's. Comparing elapsed times, these Opteron cluster results are equal to or better than other cluster systems with competitive processors – at significantly lower cost.

GRID COMPUTING IN MANUFACTURING

With increasingly complex designs and globally dispersed operations, manufacturers in virtually all industries need effective computational tools to facilitate collaboration of multiple teams, partners, and suppliers – while getting the most from their computational resources. *Grid Computing* technology can help by allowing disparate systems to be pooled and managed as a common computing resource. In particular, MCAE users benefit from increased access (e.g., using a Sun N1GE-based portal) to computational resources, while other design teams and partners get just the information they need. Ultimately, Grid Computing will help improve innovation, providing the resources as needed to get high-quality products to market quicker, while reducing costs and providing a rapid return on investment. To this goal, Sun Microsystems recently announced a *Utility Computing* initiative, where several large data centers, each offering thousands to tens of thousands of cpu's, can be accessed by customers to satisfy their computing requirements, on a pay-as-you-go basis. These data centers are initially targeted at vertical markets such as Energy, Life Sciences, Financial Services, and Manufacturing/MCAE.

CONCLUSIONS

This paper presented two recent LS-DYNA benchmarks – a cold rolling process simulation (using Sun SPARC/Solaris servers and Xeon-based clusters) and a 3-car-crash benchmark (using a Sun Fire V20z Opteron cluster). LS-DYNA scalability is excellent on both benchmarks. Sun's new Opteron-based clusters (running Solaris 10, Linux, or Microsoft Windows) offer MCAE users very cost-effective computing solutions for LS-DYNA users. Grid Computing is discussed, in the context of another cost-effective computing paradigm for Manufacturing/MCAE users.

ACKNOWLEDGMENTS

The authors would like to acknowledge Dr. Youn-Seo Roh of Sun Microsystems, and his technical expertise on LS-DYNA benchmarking and optimization. We also wish to express our appreciation to Mefos for permission to cite their cold rolling benchmark results. All brands, products names ,and trademarks are owned by the respective companies mentioned.

References

1. **Annika Nilsson, Agne Bogdanoff, Suvi Niemi, “A study on the ski end direction in plate mills”, Steel Grips 2 - Suppl. Metal Forming 2004, p 81-84, 2004.**