

LS-DYNA[®] Scalability Analysis on Cray Supercomputers

Ting-Ting Zhu

Cray Inc.

Jason Wang

LSTC

Abstract

For the automotive industry, car crash analysis by finite elements is crucial to shortening the design cycle and reducing costs. To increase the accuracy of analysis, in addition to the improvement in finite element technique, smaller cells of finite element meshes are used to better represent the car geometry. The use of finer mesh coupled with the need for fast turnaround has put increased demand on scalability of the finite element analysis. In this paper, we will use the car2car model to measure LS-DYNA scalability on Cray[®] XC30[™] supercomputers, an Intel[®] Xeon[®] processor-based system using the Cray Aries network. The scalability of different functions in LS-DYNA at high core counts will be analyzed. The MPI communication pattern of LS-DYNA will also be studied. In addition to that, we will also explore the performance difference between using one thread per core and two threads per core. Finally, we will explore the performance impact of using large Linux Huge Pages.

1. Introduction

Since May 2013, CRAY Inc. and LSTC have been working jointly with a customer on some very large (25M -100 M elements) explicit LS-DYNA analysis problems to evaluate LS-DYNA scalability on the CRAY XC30 system. Currently this problem takes nearly a month to complete on a Linux cluster at the customer site due to the large problem size and limited scalability of MPP LS-DYNA. To shorten their design cycle, the customer needs better hardware and software performance.

In the last decade, the number of cores per processor has been doubled in every couple of years, but the CPU clock has not been getting much faster. So to reduce the simulation turnaround time, we need highly scalable supercomputers and improved LS-DYNA parallel performance. Currently, due to the nature of load imbalance in LS-DYNA contact algorithm, MPP LS-DYNA scalability has been limited to 1,000 to 2,000 cores depending on the size of the problem. In order to improve LS-DYNA scalability, we need to use HYBRID LS-DYNA which combines distributed memory parallelism using MPI with shared memory parallelism using OpenMP.

The advantage of HYBRID LS-DYNA at low core counts has been demonstrated by Kendo and Makino [1], and Meng and Wang [2]. In this paper, we will use the car2car model from topcrunch.org to understand the scalability of MPP and HYBRID LS-DYNA at high core counts on the CRAY XC30 system.

2. Scalability Analysis of Pure MPP LS-DYNA on the CRAY XC30 System

First, we use the car2car model (with 2.4M elements) to test pure MPP LS-DYNA parallel performance at high core counts of 256 to 4,096 on the CRAY XC30 system.

2.1 Software and Hardware Configurations

The following hardware and software configurations are used for this evaluation:

Table 1: Summary of hardware and software configurations

System type	Cray XC30
Processor architecture	2.6 GHZ Sandy Bridge
sockets/node	2 sockets per node
Number of cores per socket	8 cores per socket or 16 core per node
Network	Aries Interconnect and Dragonfly Topology
Memory size per node	32 GB/node
OS	CRAY CLE5.1
LS-DYNA version	LS-DYNA R61_84148
Compiler version	Intel/13.1.3.192
MPI version	CRAY-MPICH/6.0.2

2.2 Pure MPP LS-DYNA Scalability

Figure 1 shows the scalability of pure MPP LS-DYNA in car2car at core counts from 256 to 4,096 on an XC30 system. Pure MPP LS-DYNA scales reasonably well up to 2,048 cores, but beyond that the scaling levels off. In the next section we will look at LS-DYNA functionality profiles to understand the limiting factor in pure MPP LS-DYNA scaling.

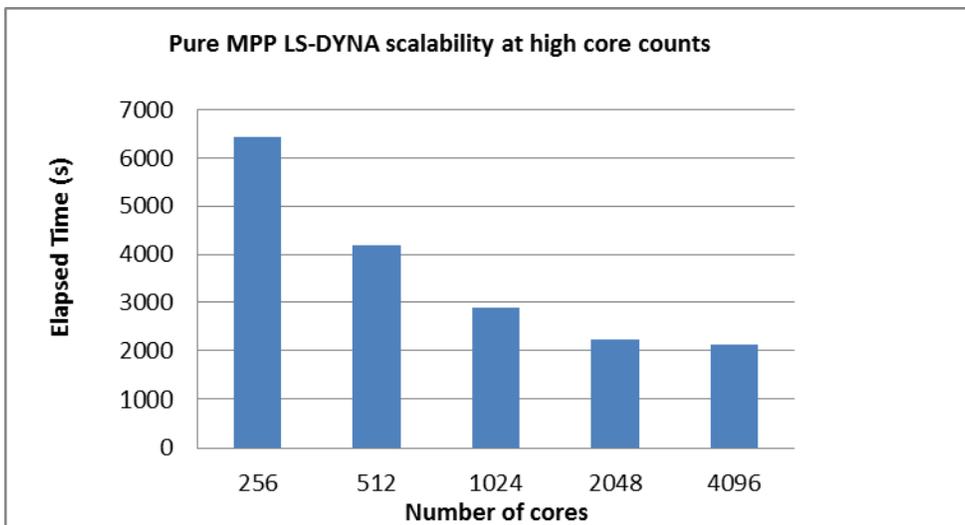


Figure 1: Pure MPP LS-DYNA scalability in car2car on a CRAY XC30 system

2.3 Scalability of Element Processing and Contact

LS-DYNA provides detail timing information of several major functions at the end of the log file. For car crash simulations, the functions of “element processing”, ”contact algorithm”, and “rigid bodies” consume more than 85% of total wall clock time at high core counts. Since there is no MPI barrier call between “contact algorithm” and “rigid bodies”, we will measure the combined time spent on both (referred to as “contact + rigid bodies” in this paper).

Now let’s take a look at how “element processing” and “contact + rigid bodies” of pure MPP LS-DYNA perform at high core counts. Figure 2 shows that “element processing” scales very well up to 4,096 cores; however, “contact + rigid bodies”, on the other hand, scale poorly for core counts of 256 to 2,048 and hardly scales after 2,048 cores. This result indicates that “contact + rigid bodies” scalability is the limiting factor for pure MPP LS-DYNA scaling beyond 2,048 cores.

Next, we will use CRAY MPI profiling tool called “profiler” to show the lack of scalability in “contact + rigid bodies”.

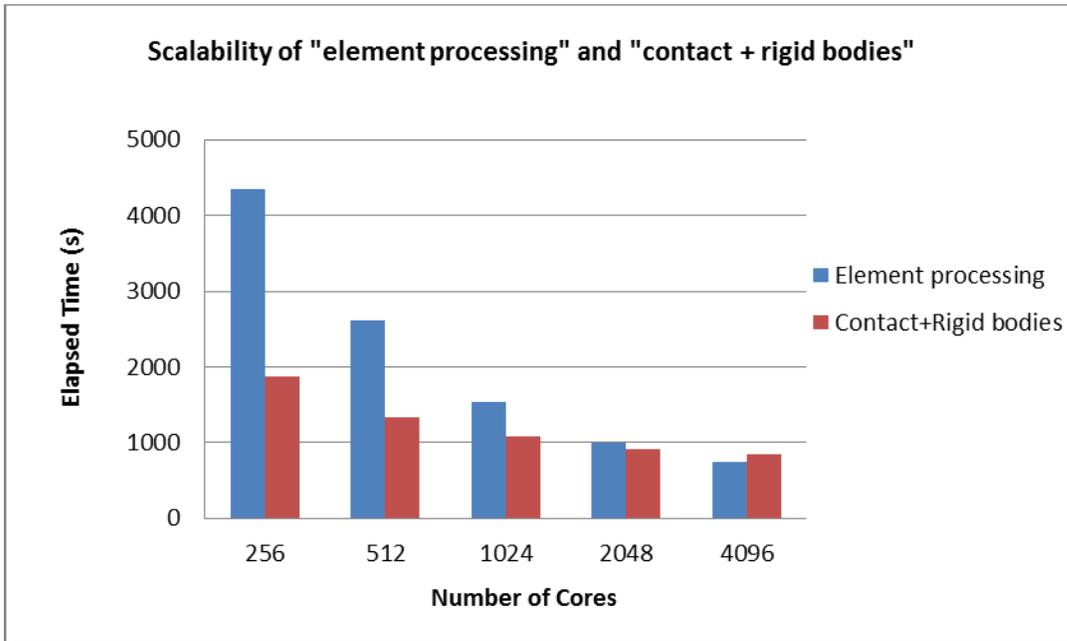


Figure 2: Scalability of “element processing” and “contact+rigid Bodies” in Car2car on a CRAY XC30 system

2.4 MPI Profile of Pure MPP LS-DYNA

CRAY MPI “profiler” reports the total elapsed time, computational time, and MPI time as well as MPI communication patterns. For collective calls, it further breaks down MPI time to MPI synchronization time (time that a process spends on synchronizing with other processes) and MPI communication time (time that a process spends on communicating with other processes). MPI synchronization time is proportional to the amount of the load imbalance in a program, and MPI communication time is disproportional to the speed of interconnect. For MPI time in point

to point communications, such as MPI_send and MPI_rcv, however, MPI synchronization time and communication time cannot be measured separately. Rather, measurement of MPI communication time in point to point MPI communication calls is the sum of MPI synchronization time and communication time.

Figure 3 shows the scalability of MPI time of pure MPP LS-DYNA in car2car on the Cray XC30 system. The total MPI time reduces from 2,309s at core count of 256 to 1,478s at core count of 2,048, but increases to 1,570s at core count of 4,096. The percentage of total MPI time to total elapsed time increases from 36% at 256 cores to 74% at 4,096 cores. Additionally, MPI synchronization time is about 65% of total MPI time. This result further indicates that load imbalance in contact algorithm of pure MPP LS-DYNA is the main cause of the poor scalability at high core counts.

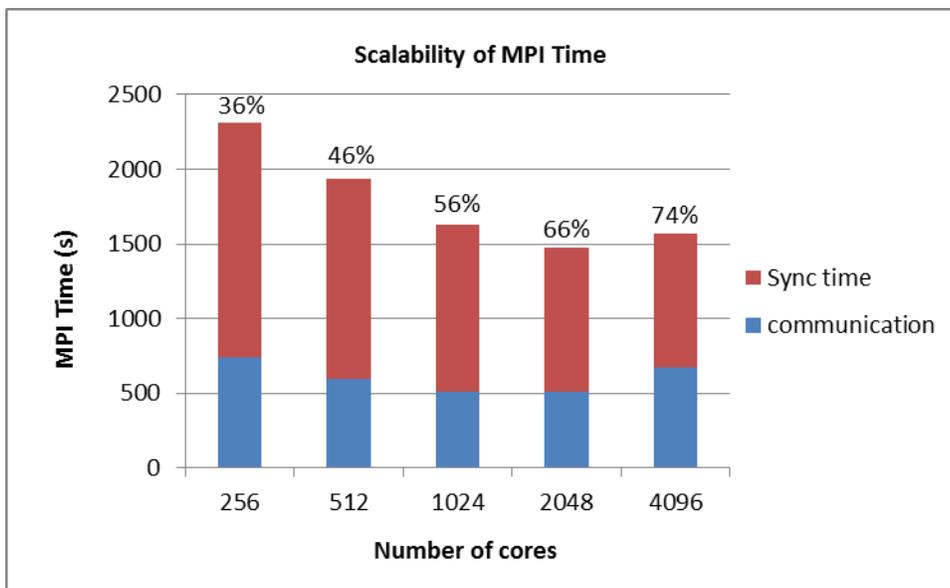


Figure 3: Scalability of MPI time in car2car on a CRAY XC30 system

From CRAY “profiler” reports, we find that 80% of MPI time in Pure MPP LS-DYNA is spent on MPI_bcast, MPI_allreduce, and MPI_rcv. Figure 4 depicts the MPI time, synchronization time and communication time on these three MPI calls. It reviews that MPI_bcast communication time is less than 2.5% of MPI_bcast synchronization time, and overall MPI_bcast time decreases as core count increases. MPI_allreduce communication time is less than 5% of MPI_allreduce synchronization time, and overall MPI_allreduce time increases with core count. MPI_rcv time, which is about 26% to 38% of total MPI time, reduces as core count increases from 256 to 1,024 cores. But MPI_rcv time increases as core counts increases from 1,024 to 4,096 cores.

As we have mentioned earlier, MPI_rcv synchronization and communication time cannot be measured separately. Based on the amount of synchronization time in MPI_bcast and MPI_allreduce, we believe that larger portion of measured MPI_rcv time is probably spent on

synchronization. This hypothesis implies that most of MPI time is spent on MPI synchronization meaning that load imbalance in LS-DYNA is what prevents it from scaling beyond 2,048 cores. This also means that without load imbalance issue, the speed of MPI communication in an XC30 system would be fast enough for pure MPP LS-DYNA to potentially scale to 4,096 and beyond.

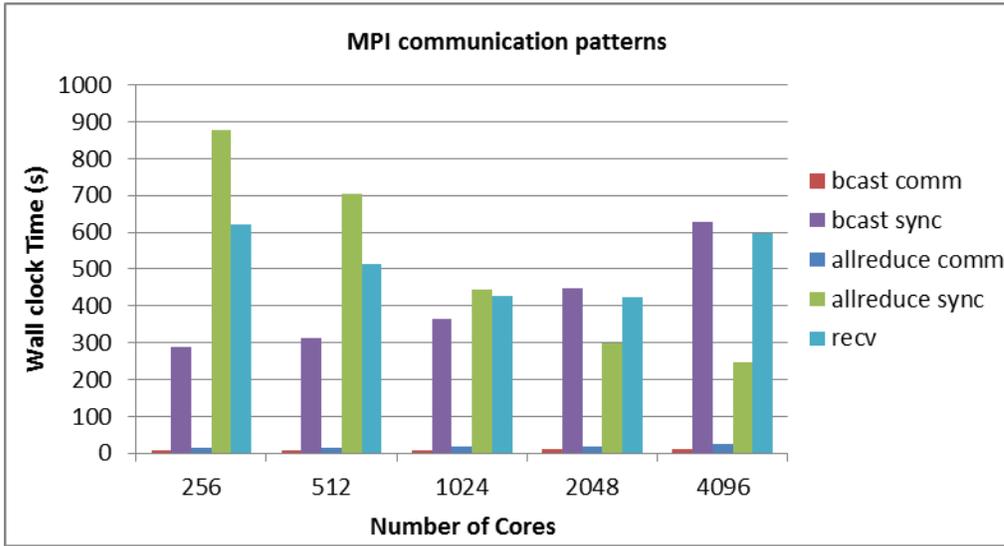


Figure 4: MPI communication patterns of Pure MPP LS-DYNA in car2car on a Cray XC30 system

One effective way to reduce the impact of load imbalance on LS-DYNA parallel performance is to use HYBRID LS-DYNA which combines MPI processes with OpenMP threads. With HYBRID LS-DYNA, we can push LS-DYNA scaling to beyond 2,048 cores. In the next section, we will discuss HYBRID LS-DYNA parallel performance in car2car, as well as its comparison with pure MPP LS-DYNA on an XC30 system.

3. Scalability Analysis of HYBRID LS-DYNA on a CRAY XC30 System

3.1 Advantage of HYBRID LS-DYNA over Pure MPP LS-DYNA

The shared memory parallelism of explicit LS-DYNA scales well to 8 threads. In the first evaluation, we run HYBRID LS-DYNA using 4 OpenMP threads with core counts from 256 to 4,096. Figure 5 shows that HYBRID LS-DYNA is 23% to 60% faster than pure MPP LS-DYNA for the same core counts.

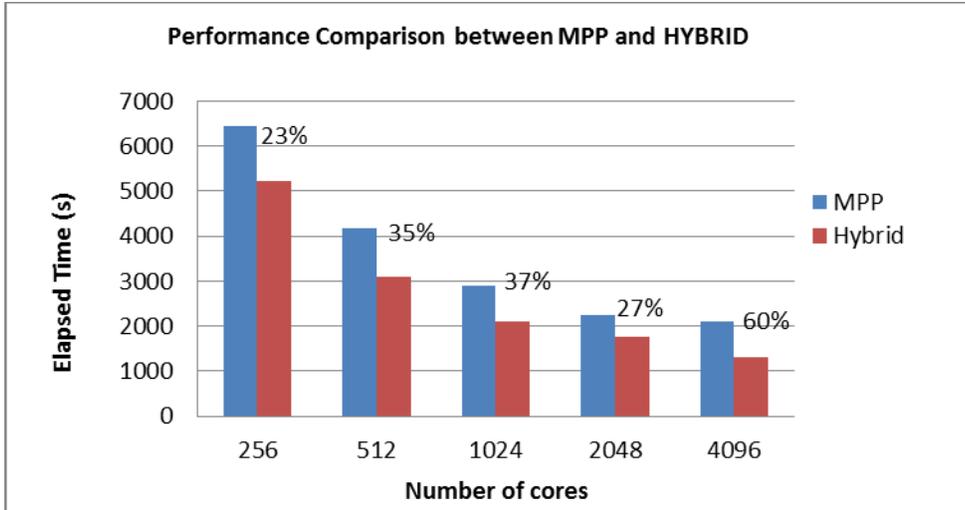


Figure 5: Performance comparison between Pure MPP and HYBRID LS-DYNA in car2car on Cray XC30

Figure 6 and Figure 7 show the performance comparisons of “element processing” and “contact+rigid bodies” between pure MPP and HYBRID LS-DYNA on an XC30 system. “Element processing” of HYBRID LS-DYNA is 35% to 58% faster than that of pure MPP LS-DYNA, while “contact+rigid bodies” of HYBRID LS-DYNA is 2% slower than that of MPP LS-DYNA at 256 cores and becomes 10% faster at 512 cores and reaches 43% faster at 4,096 cores. So for core counts of 256 to 512, HYBRID LS-DYNA gains its performance over MPP LS-DYNA mainly from “element processing”. For cores counts of 1,024 to 4,096, it gains its performance from both “element processing” and “contact + rigid bodies”.

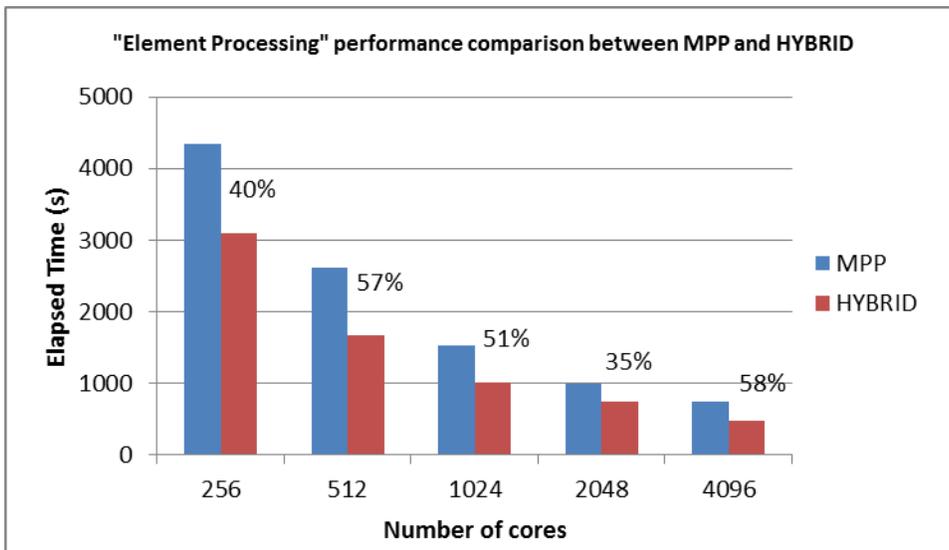


Figure 6: “Element processing” performance comparison between MPP and HYBRID LS-DYNA in car2car on an XC30 system

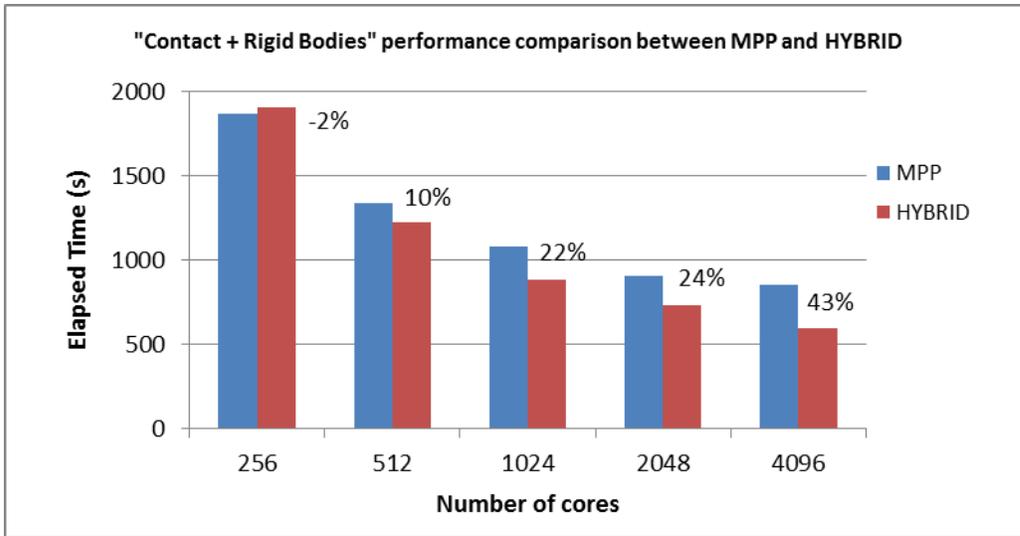


Figure 7: “Contact and rigid bodies” performance comparison between MPP and HYBRID in car2car on an XC30 system

In the next section, we will explore the effect of the number of OpenMP threads on HYBRID LS-DYNA performance at high core counts.

3.2 The Effect of the Number of OpenMP Threads on HYBRID LS-DYNA Performance

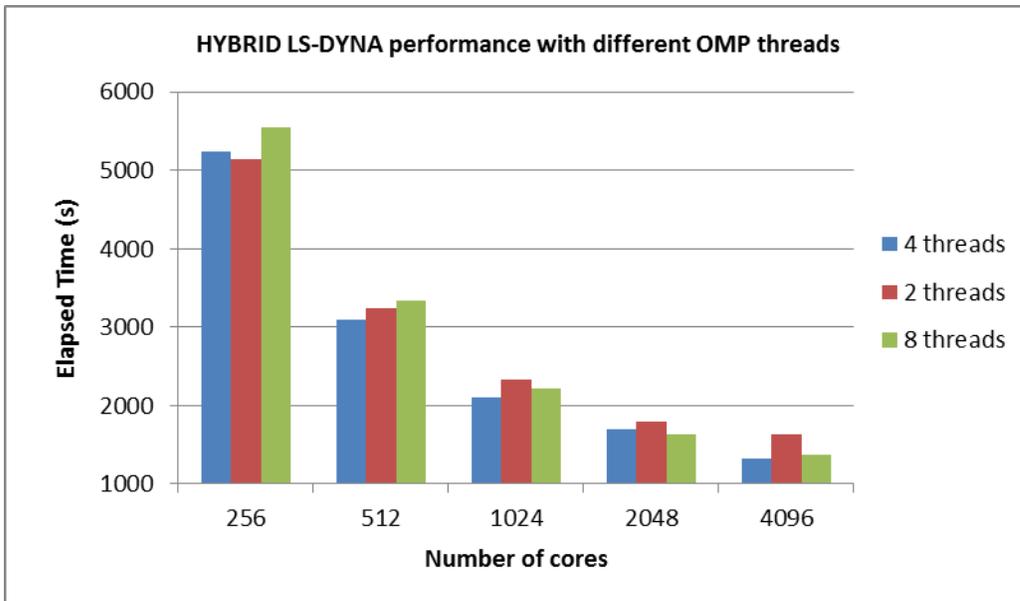


Figure 8: Hybrid LS-DYNA performance comparisons among different OpenMP threads in car2car on an XC30 system

In this study, we will keep the total number of cores, which is equal to MPI ranks \times OpenMP threads, the same while varying MPI ranks and OpenMPI threads. For an example, for 256 cores, we use three different combinations: 128 MPI ranks \times 2 OpenMP threads, 64 MPI ranks \times 4 OpenMP threads, and 32 MPI ranks \times 8 OpenMP threads. Figure 8 depicts the performance

difference among the three different OpenMP threads. At 256 cores, using 2 threads is 1.5% faster than using 4 threads and it is also 7% faster than using 8 threads. At 2,048 cores, using 8 threads is about 3% faster than using 4 threads, and it is also 9% faster than using 2 threads. At other core counts, using 4 threads is 5% to 23% faster than using 2 threads and it is also 5% to 8% faster than using 8 threads. So we recommend the use of 4 OpenMP threads for HYBRID LS-DYNA at high core counts.

4. The Effect of Intel Hyper-Threading Technology and Linux Huge Pages on LS-DYNA Performance

4.1 The Effect of Intel Hyper-Threading Technology on LS-DYNA Performance

Intel Hyper-Threading Technology is enabled on the CRAY XC30 system's BIOS. We run pure MPP LS-DYNA with 2 threads per core in car2car for core counts from 256 to 2,048. It is found that using 2 threads per core is about 5% to 29% slower than using only 1 thread per core for the same core counts. So it is not recommended to use 2 threads per core for pure MPP LS-DYNA at high core counts.

4.2 The Effect of Linux Huge Pages on LS-DYNA Performance

Cray XC30 software allows the use of different sizes of Linux Huge Pages. The default Linux Huge Page size is 4 kilobytes. We test pure MPP LS-DYNA in car2car using 2 megabyte Linux Huge Pages at core counts of 256 to 4,096, but no noticeable performance advantage is observed. This is because LS-DYNA is a cache friendly code.

Conclusion

Pure MPP LS-DYNA scaling stops at 2,048 cores on the CRAY XC30 supercomputer. The limiting factor is load imbalance in the LS-DYNA contact algorithm. HYBRID LS-DYNA extends LS-DYNA scaling to 4,096 cores and perhaps beyond on the XC30 system. Using 4 OpenMP threads in general yields close to the best or the best performance. Using 2 threads per core for LS-DYNA is not recommended on the CRAY XC30 system. Using large Linux Huge Pages has no effect on LS-DYNA performance.

Acknowledgement

The authors would like to thank CRAY Inc. for their support on this work.

References

- [1] Kenshiro Kondo and Mitsuhiro Makino, "High Performance of Car Crash Simulation by LS-DYNA Hybrid Parallel Version on Fujitsu FX1", 11th International LS-DYNA Users' Conference.
- [2] Nich Meng and Jason Wang, et al, "New Features in LS-DYNA HYBRID Version", 11th International LS-DYNA Users' Conference.