Faster Metal Forming Solution with Latest Intel® Hardware & Software Technology

Nick Meng [1], Jixian Sun [2], Paul J Besl [1]

[1] *Intel Corporation, Software and Services Group, Developer Relations Division*
[2] *Livermore Software Technology Company*

Abstract

Reducing part development time & cost and increasing quality & productivity have always been important goals in the metal forming industry. Quick response to frequent mould change is challenging metal forming simulation engineers in the modern high efficiency production environment. To achieve these goals, a requirement for nearly instant numerical simulation of metal forming processes is emerging. Unfortunately, poor scalability and slow I/O due to adaptive remeshing as well as long waiting times in corporate HPC job queues are all bottlenecks to instant numerical simulation of metal forming processes. Three optimized workstation solutions using new Intel® hardware and software technology are proposed.

The proposed solution addresses these issues through code optimization using the latest Intel® compiler technology, Intel® Xeon® processor E5-2600[1] product family processor, the 2nd generation Intel® Solid State Device (SSD) technology, and the new Advanced Vector Extension (AVX)[3] instruction set. The proposed solution allows metal forming simulations to be run on a local workstation with promising turnaround times. The performance of optimal configurations is discussed for real customer workloads in this paper.

Introduction

Numerical simulation of metal forming processes has become standard operating procedure in the automotive manufacturing industry. The technology allows sheet metal part designers to assess several alternative designs relatively quickly to optimize their part for (a) low cost manufacture, (b) fast time to market, and (c) high quality. Lean manufacturing is critical to the success of metal forming companies.

The traditional procedure of metal forming simulation in OEMs in the automotive industry involves metal forming engineers submitting their jobs to corporate HPC computing queues from their modeling workstation at the office via a batch job scheduler. Then metal forming engineers usually have to wait a few hours to finish the simulation, finally copying data back to their workstations for post processing & visualization. The disjointed-ness of the process tends to reduce metal forming simulation efficiency. Metal forming engineers hope to find an all-in-one (pre-processing, simulation, post-processing) solution that improves simulation efficiency and gives engineers more control over their own processes.

Increasingly, there is a shift in the stamping engineering function in the North American automotive industry from OEMs to tier 1 & tier 2 suppliers. This restructuring of the industry has made OEMs much leaner and efficient by outsourcing die/mould/part manufacturing to the suppliers. With these outsourcing activities, so go the engineering functions associated with these dies/moulds/parts. Furthermore, the state-of-the-art concept for total engineering in the stamping field now includes not only draw dies, but also line dies and stamping sub-assemblies (so called 'full cycle simulation'). The goal of this approach is to engineer all 'unknowns' in a release ***before*** a single die is cut. Stamping engineering releases must account for all unknowns, such as all dimensional stability issues, material properties, press line and assembly variations, etc. in every part of the tooling employed to produce the final product. These heightened engineering requirements have not only demanded that engineers perform at very high skill levels, but also have strained existing computing resources to the limit. This has also caused a surge in the requirement of more engineering computing resources mandating new approaches in software efficiency, and more importantly, in hardware innovations.

The work in this paper provides a practical and relatively inexpensive solution for OEMs and tier 1, 2 suppliers in meeting these new computing challenges.

Latest Intel® Hardware & Software Technology

The latest Intel® Xeon® processor E5-2600 product family, launched March 6, 2012, delivers more compute capacity and introduces an innovative i/o infrastructure that includes Intel® Integrated I/O[5] and Intel® Direct I/O[5] technologies. . This new processor succeeds the Intel® Xeon® X5500/X5600 series architecture offers up to a 1.7X performance improvement. Because of the major advances found in the Intel® Xeon® processor E5-2600 product family the opportunity for metal forming engineers to perform simulation efficiently in more powerful workstation in their own office is now possible.

- Eight cores per processor (2.4-3.1GHz)
- Up to 2x double-precision floating-point peak performance via AVX
- New decoded micro-operation instruction cache (decoded I-cache)
- Up to a 200%[2] improvement in effective Stream memory bandwidth
- Memory capacity increase up to 512GB@1600MHz DDR3
- Dual 8GT/s QPI channel links for excellent inter-processor communications.
- Up to a 300%[2] improvement in I/O bandwidth
- Integrated IO (40 PCIe lanes per processor socket)
- Highly efficient power management

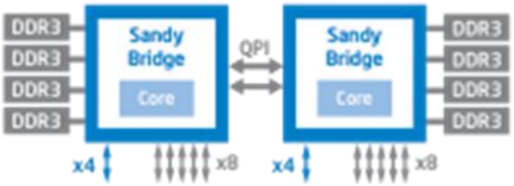

Figure 1: Intel® Xeon® Processor E5-2600 product Family

Secondly, Intel® launched a new professional compiler V12.1 in later 2011 in which the 2nd generation auto-vectorization tool is implemented. This version of the compiler has optimized the LS-DYNA® application from LSTC® more efficiently than past compilers. About 10% extra performance benefit is demonstrated on Sandy Bridge-EP system.

Finally, Intel® released 2nd generation of Solid State Disk in early 2012. The new Intel® SSD 520 series is outfitted with a SATA III interface with sequential reads and writes speeds in the 550MB/s to 520MB/s ranges, respectively. Meanwhile, the disk capacity is up to 480GB and the cost is lower than previous generations. As prices eventually drop, the Intel® 520 Series SSD will become an ideal investment for those who want to greatly improve their computer's overall performance. Along with integrated Direct I/O feature in Sandy Bridge-EP system, I/O performance can be improved with Solid State Disks in Intel® Sandy Bridge-EP workstation. Metal forming engineers will get much benefit from this for very large adaptive meshing models and the out-of-core implicit solver.

Intel SSD 520, aka Cherryville

Figure 2: Intel® 520 serial Solid State Disk

Overall, Sandy Bridge-EP Workstation is a huge step forward for metal forming market compared to the previously available platforms. These advantages are the fundamental basis of our faster metal forming solution.

Performance Analysis

Adaptive remeshing is a key feature in metal forming simulation. One of the features of metal forming simulation with adaptive remeshing is that the problem size is increasing constantly. Because of this, metal forming simulation takes much time during the initialization phase. The typical percentage of computing cost in initialization ranges is from 20-80%. The percentage is very large compared with the normal range for crash simulation, which is less than 1%. Here is a summary of an example of metal forming simulation:

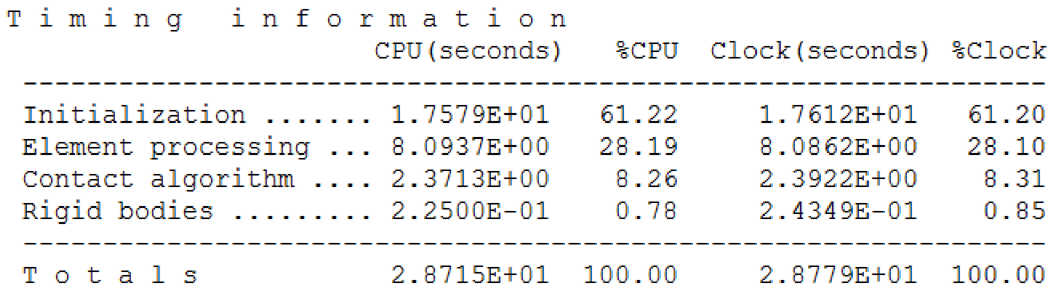

```
 T i m i n g   i n f o r m a t i o n
                        CPU(seconds)   %CPU  Clock(seconds) %Clock
  ----------------------------------------------------------------
  Initialization ....... 1.7579E+01   61.22     1.7612E+01   61.20
  Element processing ... 8.0937E+00   28.19     8.0862E+00   28.10
  Contact algorithm .... 2.3713E+00    8.26     2.3922E+00    8.31
  Rigid bodies ......... 2.2500E-01    0.78     2.4349E-01    0.85
  ----------------------------------------------------------------
  T o t a l s            2.8715E+01  100.00     2.8779E+01  100.00
```

The performance issue is suspected from slow I/O based on past experience. One medium size model is selected for analysis for confirmation. The initial problem size of the model is 236k elements, and the final problem size is 549k elements. The model does adaptive remeshing 58 times during simulation. A short test was done on a Sandy Bridge-EP workstation with two different file systems. One file system is a single SAS R10K disk; the other one is a RAID0 file system with 3 Solid State Disks (SSD's). However, the benchmark demonstrates that the problem is not from slow I/O, as shown below.

Mid-size model	1st run on SAS 10k rpm disk	2nd run on RAID0 file system
Initialization	560s	559s
Elapsed time	2173s	2170s

Mid-size model	1st run with original I/O subroutine	2nd run with tuned I/O subroutine
Initialization	533s	531s
Elapsed time	1977s	1957s

Table 1: Performance comparison of mid-size model for Initialization I/O analysis

It is clear that there is no significant difference from the file system comparison. There is only 1% improvement from optimized I/O subroutine. A large model (661K elements to 2.19M elements) with 265 adaptive remeshing WAS selected for further analysis, and the results are shown in Table 2 below:

large model	1st run with Single SAS 10k disk	2nd run with RAID0 & SSDs	ratio
Initialization	12893s	12785s	1%
Elapsed time	23364s	21591s	**8%**

Table 2: Performance comparison of large size model for Initialization I/O analysis

Obviously, it shows I/O wait time is at least one possible factor of performance issue for very large/huge metal forming model.

Based on results of initialization I/O comparison, the performance problem is not only from I/O, but also from decomposition work and MPI communication. To further analyze the behavior of the MPI communication, the large model was rerun with Intel® Trace Analyzer and Collector tool. The test was done on a Sandy Bridge workstation, with the results shown in the Figures below:

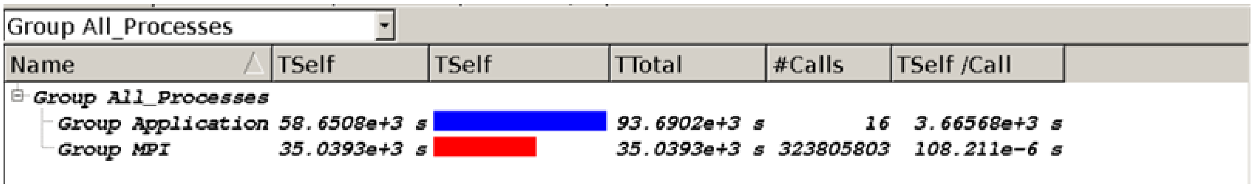

Group All_Processes

Name	TSelf	TSelf	TTotal	#Calls	TSelf /Call
Group All_Processes					
Group Application	58.6508e+3 s	🟦	93.6902e+3 s	16	3.66568e+3 s
Group MPI	35.0393e+3 s	🟥	35.0393e+3 s	323805803	108.211e-6 s

Figure 3: Profile of large metal forming model.

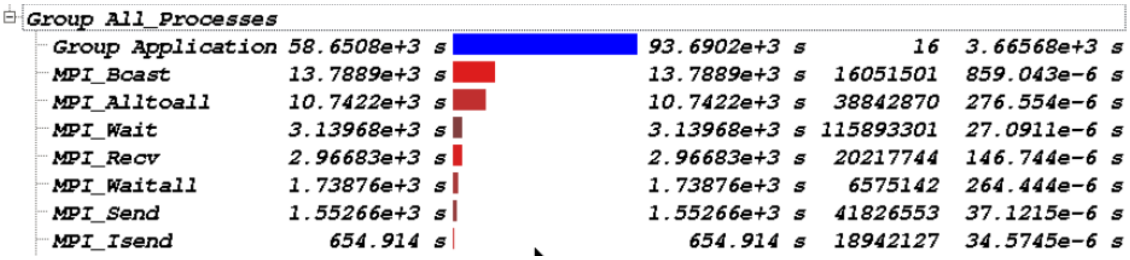

Group All_Processes					
Group Application	58.6508e+3 s	🟦	93.6902e+3 s	16	3.66568e+3 s
MPI_Bcast	13.7889e+3 s	🟥	13.7889e+3 s	16051501	859.043e-6 s
MPI_Alltoall	10.7422e+3 s	🟥	10.7422e+3 s	38842870	276.554e-6 s
MPI_Wait	3.13968e+3 s	🟥	3.13968e+3 s	115893301	27.0911e-6 s
MPI_Recv	2.96683e+3 s	🟥	2.96683e+3 s	20217744	146.744e-6 s
MPI_Waitall	1.73876e+3 s	🟥	1.73876e+3 s	6575142	264.444e-6 s
MPI_Send	1.55266e+3 s	🟥	1.55266e+3 s	41826553	37.1215e-6 s
MPI_Isend	654.914 s	🟥	654.914 s	18942127	34.5745e-6 s

Figure 4: Hot MPI functions in large metal forming model.

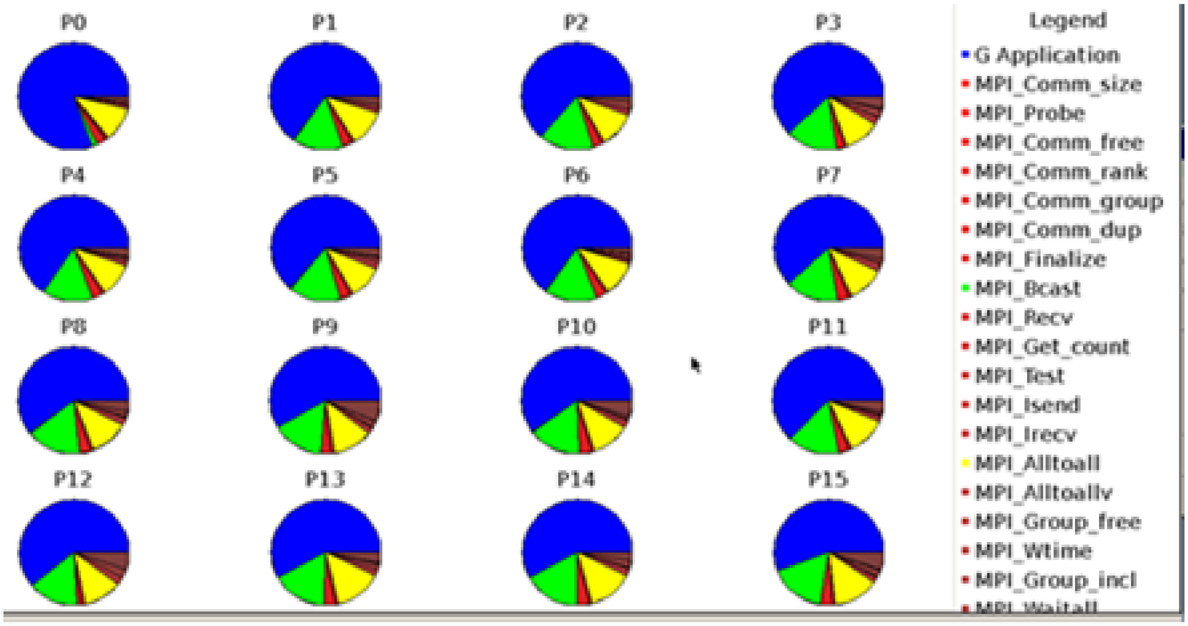

Figure 5: Hot MPI functions (pie mode) in large metal forming model.

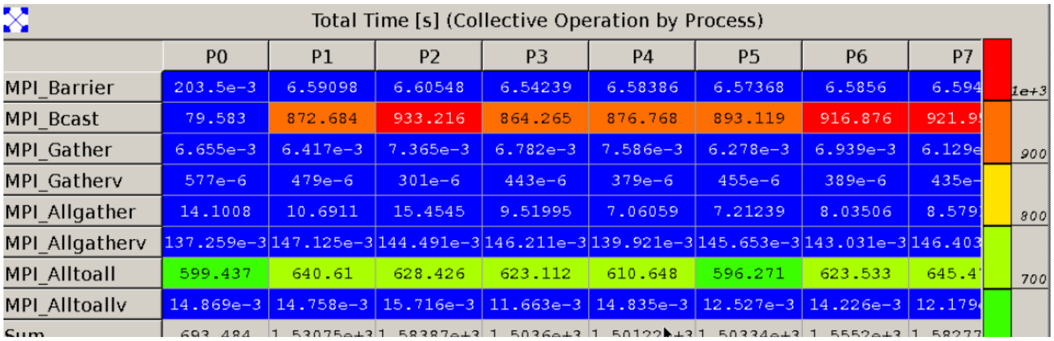

Total Time [s] (Collective Operation by Process)

	P0	P1	P2	P3	P4	P5	P6	P7
MPI_Barrier	203.5e-3	6.59098	6.60548	6.54239	6.58386	6.57368	6.5856	6.594
MPI_Bcast	79.583	872.684	933.216	864.265	876.768	893.119	916.876	921.9
MPI_Gather	6.655e-3	6.417e-3	7.365e-3	6.782e-3	7.586e-3	6.278e-3	6.939e-3	6.129e
MPI_Gatherv	577e-6	479e-6	301e-6	443e-6	379e-6	455e-6	389e-6	435e-
MPI_Allgather	14.1008	10.6911	15.4545	9.51995	7.06059	7.21239	8.03506	8.579
MPI_Allgatherv	137.259e-3	147.125e-3	144.491e-3	146.211e-3	139.921e-3	145.653e-3	143.031e-3	146.403
MPI_Alltoall	599.437	640.61	628.426	623.112	610.648	596.271	623.533	645.4
MPI_Alltoallv	14.869e-3	14.758e-3	15.716e-3	11.663e-3	14.835e-3	12.527e-3	14.226e-3	12.179
Sum	693.484	1.53075e+3	1.58387e+3	1.5036e+3	1.50122e+3	1.50334e+3	1.5552e+3	1.58277

Figure 6: Total Time of MPI Collective functions

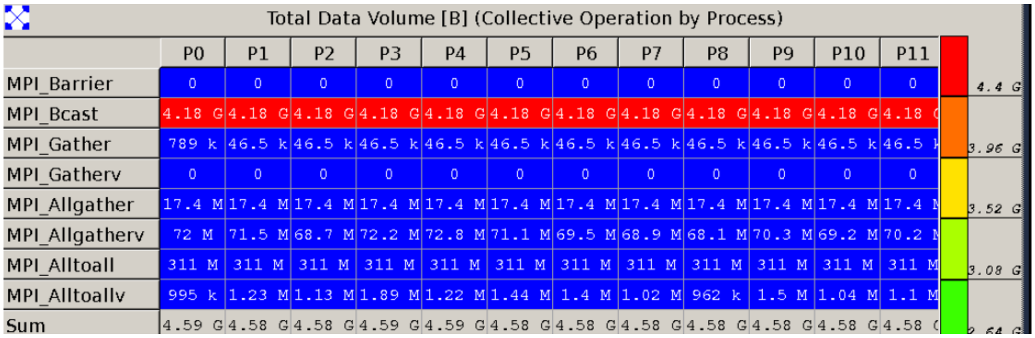

Total Data Volume [B] (Collective Operation by Process)

	P0	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11
MPI_Barrier	0	0	0	0	0	0	0	0	0	0	0	0
MPI_Bcast	4.18 G	4.18 G	4.18 G	4.18 G	4.18 G	4.18 G	4.18 G	4.18 G	4.18 G	4.18 G	4.18 G	4.18
MPI_Gather	789 k	46.5 k	46.5 k	46.5 k	46.5 k	46.5 k	46.5 k	46.5 k	46.5 k	46.5 k	46.5 k	46.5
MPI_Gatherv	0	0	0	0	0	0	0	0	0	0	0	0
MPI_Allgather	17.4 M	17.4 M	17.4 M	17.4 M	17.4 M	17.4 M	17.4 M	17.4 M	17.4 M	17.4 M	17.4 M	17.4
MPI_Allgatherv	72 M	71.5 M	68.7 M	72.2 M	72.8 M	71.1 M	69.5 M	68.9 M	68.1 M	70.3 M	69.2 M	70.2
MPI_Alltoall	311 M	311 M	311 M	311 M	311 M	311 M	311 M	311 M	311 M	311 M	311 M	311 M
MPI_Alltoallv	995 k	1.23 M	1.13 M	1.89 M	1.22 M	1.44 M	1.4 M	1.02 M	962 k	1.5 M	1.04 M	1.1 M
Sum	4.59 G	4.58 G	4.58 G	4.59 G	4.59 G	4.58 G	4.58 G	4.58 G	4.58 G	4.58 G	4.58 G	4.58

Figure 7: Total Data size of MPI Collective functions

Figures 3, 4 and 5 clearly show that the large model took about 27% computing time on MPI communication. Both MPI_BCAST and MPI_ALLTOALL are of particular concern. MPI_BCAST and MPI_ALLTOALL were called frequently in the domain decomposition phase. Figures 6 and 7 indicate that MPI_BCAST moved more than 4.4GB data in each MPI process totally. It is clear what happened in initialization (domain decomposition) phase. Even though MPI_ALLTOALL function only moves about 1311M data in each MPI process, the MPI_ALLTOALL function took a lot of time on waiting for completion (see number in green area in Figure 6).

Based on the analysis, the large model was rerun with optimal MPI settings for confirmation, and decent performance improvement (34%) is demonstrated. The benefit is showed in Table 3 below.

Large model 0.61M to 2.19M 265 adaptive remeshing 93.5ms	Elapsed Time With default MPI setting	Elapsed Time MPI_ADJUST_ALLREDUCE=5 MPI_ADJUST_ALLTOALL=4 MPI_ADJUST_BCAST=1	Speedup
Elapsed Time	8:02	5:59	**1.34**

Table 3: Performance comparison with default setting and optimal setting

The behavior of metal forming simulation demonstrated the root cause of performance bottleneck. Next step is to find out a faster/inexpensive metal forming solution.

Faster Metal Forming solution

Based on the analysis, I/O waiting time, MPI collective functions, and domain decomposition were identified as the root cause of performance issues, which will intensify as problem size increases. Furthermore, when problem size is big enough, computing cost will dominate the issues on performance. More cores will be needed for faster simulation. Meanwhile, memory requirement is exceeding the limit of high speed memory (e.g. 96GB@1333Mhz) of current Intel® Xeon® X5600 series based workstation for large/huge metal forming model. All of these issues above become bottlenecks to the improvement of the metal forming simulation procedure.

Along with launch of more powerful Intel® Sandy Bridge processor, Intel® professional compiler V12.1, and Intel® 2nd generation Solid State Disk, these bottlenecks are solved easily. Table 4 below shows the solution clearly.

Issues	Solution
MPI performance issue	Optimal MPI settings in Intel® MPI 4.0.3
I/O waiting time	Direct IO and SSDs
Computing cost	Up to 8 cores per processor 16/32 cores high end Workstation Intel® compiler V12.1
Memory limit	256+GB@1600Mhz DDR3

Table 4: Issues and Solution.

Three Sandy Bridge-EP workstations are setup for performance demonstration. Table 4 below is a configuration summary of these systems.

System 1	System 2	System 3
Single socket E5-1600 WKS	Dual socket E5-2600 WKS	Four sockets E5-4600 Server
6C@3.2GHz E5-1650 32GB@1600MHz DDR3 Single Large SSD RHEL 6.1	2 x 8C@3,1 GHz E5-2687W 128GB@1600MHz DDR3 RAID0 with 4 SSDs RHEL 6.1	4 x 8C@2.6Ghz E5-4650L 128GB@1600MHz DDR3 RAID 0 with 4 SSDs RHEL 6.1

Table 5: Sandy Bridge-EP system configuration

Three production metal forming models were collected from two customers for performance demonstration. The detail information of models is in the following table:

Model description	**Model 1**	**Model 2**	**Model 3**
Problem size in dof	0.24M to 0.54M	0.66M to 2.19M	0.36M to 3.77M
# of adaptive remeshings	58	265	83
Simulation time	45.8ms	65.37ms	93ms
MAXLVL	5	5	5
Elapsed time On older systems	**4396s** 6 cores X5670@2.93GHz workstation	**9hrs53min** 64 cores E5520@2.26GHz cluster	**23hrs 20min** 8 cores X5677@3.47GHz workstation

Table 6: Summary of benchmark models

These models were tested on different systems and decent performance improvement with the solution is demonstrated.

Solution for Model 1	**System**	**Elapsed Time**	**Speedup**
Baseline	6 cores X5670	**4396s**	**1.00**
New CPU / faster Memory Optimal MPI setting/ SSD	**6C System 1**	**3289s**	**1.34**
More CPUs	**16C System 2**	**1518s**	**2.89**

Table 7: Benchmark results of Model 1

Solution for Model 2	**System**	**Elapsed Time**	**Speedup**
Baseline	**64C** cores E5520 (8n)	**9:53**	**1.00**
New CPU / faster Memory Optimal MPI setting/ SSD	**16C System 2**	**9:31**	**1.03**
More CPUs	**32C System 3**	**5:59**	**1.65**

Table 8: Benchmark results of Model 2

Solution for Model 3	System	Elapsed Time	Speedup
Baseline	**8C** X5677	**23:20**	**1.00**
New CPU / faster Memory Optimal MPI setting/ SSD	**16C System 2**	**15:51**	**1.50**
More CPUs	**32C System 3**	**7:17**	**3.20**

Table 9: Benchmark results of Model 3

Summary

Three stamping models were evaluated on different configurations of Intel® Xeon®-based workstation. Three optimal configurations which were found below:

- Entry level configuration for small/middle metal forming model
  * 4C Intel® Xeon® Processor E3-1200 Product Family or 1 x 6C Intel® Xeon® Processor E5-1600 Product Family (1620, 1650, 1660)
  8 x 4 GB@1600MHz DDR3 memory
  1 x GFX card
  RAID0 file system with 2 x SAS/SSDs

- Middle level configuration for large/huge metal forming model and instant response for small/middle model
  2 x 8C Intel® Xeon® E5-2600 Product Family (E5-2643, E5-2667, E5-2670, E5-2680, E5-2690)
  16 x 8 GB@1600MHz DDR3 memory
  1 x GFX card
  RAID0 file system with 4 x SAS/SSDs

- High-end 4 socket server configuration for instant response for large/huge model
  4 x 8C Intel® Xeon® Processor E5-4600 Product Family
  16 x 16 GB@1600MHz DDR3 memory
  RAID0 file system with 4 x SAS/SSDs
  Plus Graphics Workstation[4]:
  1 x 4C Intel ® Xeon® Processor E3-1200 Product Family
  4 x 8GB@1600MHz DDR3 memory
  1 x GFX card
  RAID0 file system with 2 x SAS/SSDs

In summary, Model 2 took less than 6 hours on a single system for each run. This is good news for metal forming engineers. Intel and LSTC continue to work together for optimizing performance of LS-DYNA on Intel® Xeon-Based workstations in the future. The objective will be to complete two large models simulation in an 8-hour working day on single workstation, less than 4 hours per job.

Acknowledgment

Dr. Li Zhang of LSTC is acknowledged for providing valuable suggestion.

References

1. Intel® MPI Library for Linux OS Reference Manual, 2011
2. Intel® Trace Analyzer and Collector Reference Manual, 2012
3. LS-DYNA® 971 Keyword Users' Manual
4. http://ark.intel.com : Intel® product web site.

Footnotes

1 Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

2 (Generational Performance) Source: Performance comparison using SPECfp*_rate_ base 2006 benchmark. Based on Intel internal measured estimates using an Intel® Rose City platform with two Intel® Xeon® processor E5-2690, Turbo and EIST Enabled, with Hyper-Threading, 64 GB RAM, Red Hat* Enterprise Linux Server 6.1 beta for x86_6, Intel® Compiler 12.1. For more details, see: http://www.spec.org/cpu2006/results/res2011q1/cpu2006-20110131-14172.html

3 (AVX Performance) Source: Performance comparison using Linpack benchmark. Based on Intel internal measurements as of 4 August 2011 using a Supermicro* X8DTN+ system with two Intel® Xeon® processor X5690, Turbo Enabled or Disabled, EIST Enabled, Hyper-Threading Enabled, 48 GB RAM, Red Hat* Enterprise Linux Server 6.1 beta for x86_6. New score of Y 350.3 based on Intel internal measurements using an Intel® Rose City platform with two Intel® Xeon® processor E5-2690, Turbo Enabled or Disabled, EIST Enabled, Hyper-Threading Enabled, 64 GB RAM, Red Hat* Enterprise Linux Server 6.1 beta for x86

4. (HD Graphics) Intel® HD P4000 introduces 4 additional execution units going from 8 in the Intel® HD P3000 to 12 in the Intel HD P4000. Optimized Intel® HD Graphics P4000 only available on select models of the Intel® Xeon processor E3-1200 v2 product family. For more information, www.intel.com/go/workstation.

5 (Integrated I/O) Intel measurements of average time for an I/O device read to local system memory under idle conditions. Improvement compares Xeon processor E5-2600 product family (230 ns) vs. Xeon processor 5500 series (340 ns). Baseline Configuration: Green City system with two Intel® Xeon processor E5520 (2.26GHz, 4C), 12GB memory @ 1333, C-States Disabled, Turbo Disabled, SMT Disabled, Rubicon* PCIe* 2.0 x8. New Configuration: Meridian system with two Intel® Xeon processor E5-2665 (C0 stepping, 2.4GHz, 8C), 32GB memory @1600 MHz, C-States Enabled, Turbo Enabled. The measurements were taken with a LeCroy* PCIe* protocol analyzer using Intel internal Rubicon (PCIe* 2.0) and Florin (PCIe* 3.0) test cards running under Windows* 2008 R2 w/SP1.