Performance of LS-DYNA with Double Precision on Linux and Windows CCS

Author:

Yih-Yih Lin
Hewlett-Packard Company

Correspondence:

3000 Waterview Parkway
Richardson, TX 75080
Phone 1-978-841-7650
Fax 1-978-841-7650
Email yih-yih.lin@hp.com

ABSTRACT:

Although the majority of LS-DYNA jobs are simulated with single-precision arithmetic, the more accurate and robust double-precision arithmetic has to be used for certain situations. Currently, most public LS-DYNA benchmarks are with single precision and few with double precision. As a result, users often try to extrapolate the performance of LS-DYNA with double precision from that with single precision. This paper shows that such extrapolations are often misleading. Furthermore, a comparative performance study on LS-DYNA with double precision, using all of the industry standard processors—Intel Xeon, Intel Itanium, and AMD Opteron—and the two major operating systems—Linux and Windows CCS—is presented to provide users with information for choosing the right configuration to run LS-DYNA with double precision.

Keywords:

Double-precision Simulation, Performance

INTRODUCTION

For a given model, because there is a substantial performance difference between LS-DYNA with single-precision arithmetic and double-precision arithmetic, the user often chooses to use the former version over the latter. However, since double precision always offers more accuracy and robustness for any model than single precision does, several types of LS-DYNA models, as described in the next section, are necessary to be run with double precision. In fact, the trend to use the double-precision LS-DYNA is increasing.

Because of the single-precision preference, almost all published LS-DYNA benchmark results, including those in the well-known website http://www.topcrunch.org, are with single precision. As a result, users often try to evaluate the performance of the double-precision LS-DYNA performance on different platforms by extrapolating that of the single-precision LS-DYNA. This paper shows that such extrapolations are often misleading. Furthermore, a comparative study is presented to provide LS-DYNA users a more accurate picture on performance of the double-precision LS-DYNA.

WHY AND WHEN TO USE DOUBLE PRECISION

An LS-DYNA simulation, like any other numerical simulations, is necessarily performed on a computer system whose numbers are represented with a finite fixed number of digits—7 digits for single precision, 16 digits for double precision. The finite precision arithmetic causes two kinds of errors in LS-DYNA simulations: round-off error and truncation error.

Round-off error arises from rounding results of floating-point operations during a simulation. The simplest example for round-off error is the product of two 7-digit numbers in a single-precision simulation: The exact product (13 or 14 digits) cannot be used in subsequent calculations; and the more the floating-point operations are, the more the errors are accumulated. For LS-DYNA, the more the number of time steps is, the bigger the round-off error is. This is why any LS-DYNA model that requires a long-duration simulation time, such as rollover, requires double precision, for the round-off error eventually becomes so large that the solution diverges if single precision is used. The round-off error also explains why the implicit LS-DYNA has been released with double precision only, for the round-off error would make the inversion of matrices in single precision fail.

Truncation error, also called input error, arises because inputs to the simulation, such as geometry, velocity, and etc., cannot be represented accurately in finite-digit arithmetic. For a given precision, the truncation error is a function of the mesh size or, equivalently, the number of elements: The smaller the mesh size or the larger the number of elements is, the larger the truncation error is. Thus the maximal number of elements that LS-DYNA can accommodate is limited by the number of significant digits in precision. For example, the area of a 1 mm by 1 mm shell element is 10^{-6} m^2, the area of a car's exterior is approximately 5 m^2 (1.7 m wide and 3 m long), and their ratio is $2x10^{-7}$ (or 5 million inversely), which is close to the last significant digit of single precision. Consequently, for a car-crash model with an element size larger than 5 million, the double-precision LS-DYNA should be used to salvage the diverging effect of truncation error.

Finally, because the double-precision LS-DYNA simply has smaller round-off and truncation errors and produces more accurate results, users may prefer to use it over the single-precision LS-DYNA.

Performance of the Double-Precision LS-DYNA on Various Platforms

Today's industry standard processors are Intel Xeon, Intel Itanium, and AMD Opteron. HP manufactures and supports systems based on all the three processors, which currently are all multi-core. Clusters that comprise nodes of two dual-core processors from the three, interconnected with InfiniBand, are used in this study. Two of today's major operating systems for LS-DYNA are Linux and Windows CCS. For MPP LS-DYNA, HP supports both Linux and Windows CCS with the universal HP-MPI [1].

Shown in Figure 1 is the single-precision performance on the well-known three-vehicle-collision model (ref. the aforementioned website http://www.topcrunch.org), and shown in Figure 2 is the double-precision performance on a customer model. Every run in both figures was done with the same LS-DYNA version, 971.7600.398.

Two conclusions can be drawn from Figure 1 for the single-precision LS-DYNA performance: (1) The performances of Intel Xeon and AMD Opteron are comparable, regardless of operating systems; (2) the performance of Intel Itanium is somewhat lower. However, in Figure 2 the double-precision LS-DYNA performance of Intel Itanium is shown to be highest, contrary to the second conclusion. This clearly shows that to extrapolate the performance of the double-precision LS-DYNA from that of the single-precision LS-DYNA is misleading. The main reason that Intel Itanium performs well in double precision is that its floating-point unit is much more efficient than that of the X86-64 architecture, which includes both Intel Xeon and AMD Opteron. Also, it is interesting to note that HP BL460c with Windows CCS has a better double-precision performance than its counterpart in Linux; we have not had an explanation for this observation yet.

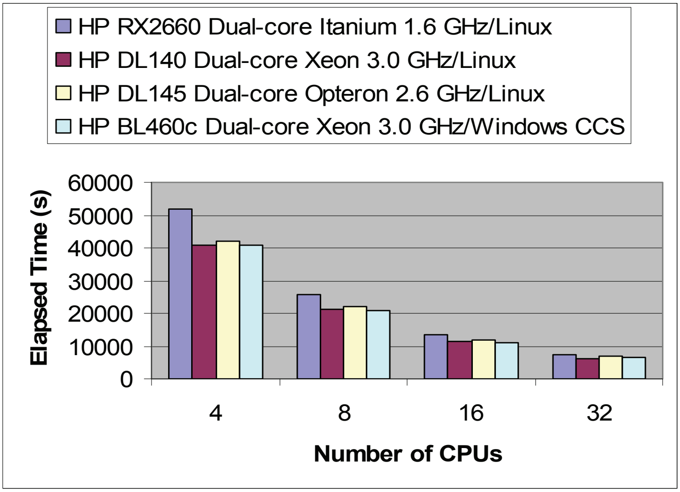

Figure 1: Performance of the single-precision LS-DYNA on the three-vehicle-collision model

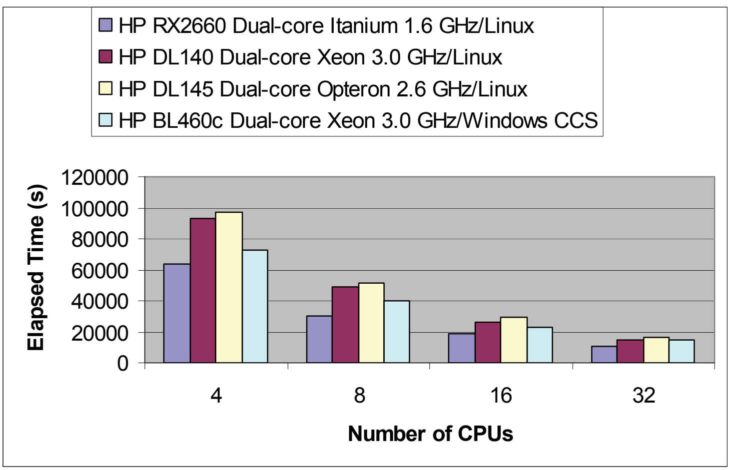

Figure 2: Performance of the double-precision LS-DYNA on a customer model

CONCLUSIONS

We have explained in this paper why and when to use the double-precision LS-DYNA. Furthermore, it is demonstrated that extrapolating the single-precision LS-DYNA performance to that of double precision is misleading. The Intel Itanium processor has a performance advantage in double precision over the X86-64 processors even though it is not the case for single precision, and therefore it should be seriously considered when using the double-precision LS-DYNA.

References

1. Y. Lin, "The Applicability of the Universal HP-MPI to MPP LS-DYNA on Linux Platforms," 5th European LS-DYNA Users Conference, Birmingham, 2005.