Cray® XC™ Supercomputer Accelerates LS-DYNA® Simulations for Aerospace

Partner
Livermore Software Technology Corporation

Scientific Field
Aerospace Manufacturing

Application
LS-DYNA® finite element program capable of solving complex, real-world problems

Key Results

- **1.9 to 2.7 times faster** performance on 24-million element model
- **34 times faster** performance on 80.6-million element model

"LSTC was privileged to work on this project with Cray. With LS-DYNA optimized for the Cray XC [system], our users can not only reduce runtimes but also gain insights into potential issues like load imbalance based on MPI communications. The ability to run LS-DYNA on large models in under a day can significantly change how users approach their simulation projects."

— John Hallquist
Livermore Software Technology Corporation
Founder

About Cray
Cray provides systems and solutions that help you solve your most difficult computing, storage and data analytics challenges. The company's comprehensive portfolio includes optimized cluster systems, powerful supercomputers, advanced storage systems and high performance data analytics and discovery platforms. Founded in 1972, Cray has focused exclusively on developing, building and supporting supercomputing technologies for over 40 years.

Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164
Tel: 206.701.2000
Fax: 206.701.2500
www.cray.com

Background

The failure of fan or compressor blades in an airplane jet engine is a major hazard. They release high-energy fragments that can perforate the engine case, damage fuel tanks and cause crashes. As a result, the Federal Aviation Administration requires that engine cases be capable of containing blade fragments. This situation makes "fan blade off containment" a critical research topic for the aerospace industry.

In an effort to improve fan blade off containment simulation, a team from Cray and Livermore Software Technology Corporation (LSTC) conducted a study using the finite element application LS-DYNA® and a Cray® XC™ supercomputer.

Challenge

Fan blade off containment simulation is technically challenging and computationally intensive. For example, a large 80-million-element simulation using LS-DYNA version R7.1.2 takes more than a month to complete. But to be useful, time-to-solution on this type of simulation ideally should be less than a day.

In order to enhance performance of the simulation, the team needed to make improvements to the most computationally expensive part of the simulation — surface-to-surface erosion contact.[1] To optimize the surface-to-surface process, they used the CrayPAT™ analysis tool to identify the most time-consuming subroutines. The resulting improvements included a 30 percent reduction in memory required for storing erosion contact surfaces, removal of redundant erosion calculations, and faster exterior surface calculation.

Then, they needed to test the improvements. To do this, they first compared the performance of the simulation between LS-DYNA R8.0.0 and the earlier LS-DYNA R7.1.2 on a medium-size model of 26.5 million nodal points and 24 million solid elements. Then, they analyzed the MPI communication patterns and load balance among the MPI processes. And finally, they carried out the 80-million-element simulation using the enhanced LS-DYNA version R8.0.0 on the XC system.

Results

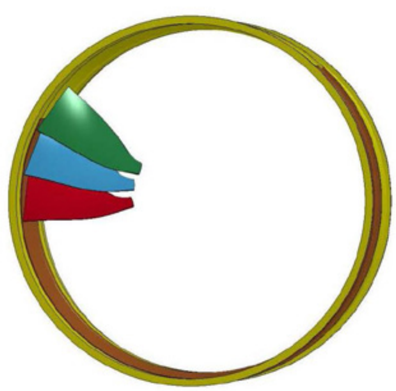

MEDIUM-SIZE MODEL: 24-million-element fan blade off simulation model with three blades.

LS-DYNA R7.1.2 versus R8.0.0

For the comparison test, the team focused on simulation performance during the collision phase (where the released blade collides with the engine case and the other two rotating blades), modeling these processes using surface-to-surface erosion contact. The team found that at 256 cores, the total elapsed time on LS-DYNA R8.0.0 was only about 1.9 times faster than LS-DYNA R7.1.2. But as core count increased, the speedup increased. At 2,048 cores, version R8.0.0 was about 2.7 times faster. The team then analyzed the LS-DYNA wall time based on its functionality to determine where the speedup came from. Their analysis showed the speedup came from the "contact" and "miscellaneous" functions which were direct outcomes of the code optimizations.

1 For this study, the team used LS-DYNA R8.0.0, the then-most recent version.

MPI Communication Patterns

Next, the team used Profiler, the Cray MPI profiling tool, to determine how much time was spent on MPI communication and which MPI calls dominated MPI time. Profiler revealed that total MPI time decreased as core count increased at lower core counts (256 to 1,024). The trend reversed at higher core counts of 1,024 to 2,048. This result revealed that MPI time is dominated by MPI synchronization time in the simulation and indicates load imbalance. Because of load imbalance, the parallel scaling of fan blade off simulation of the medium-size model is limited to 1,024 cores. However, for a large-size model, **fan blade off simulation can scale to 16,000 cores**.

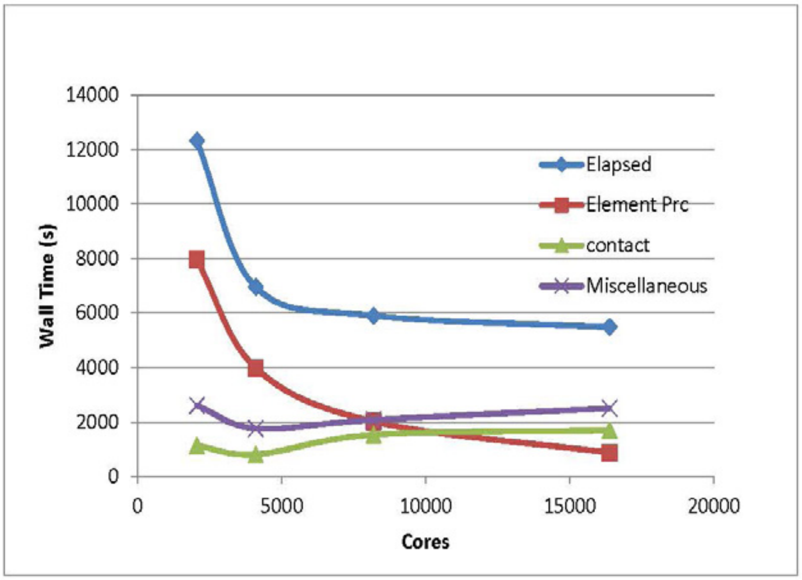

LARGE MODEL, HUGE SPEEDUP: Parallel scaling of an 80-million-element fan blade off simulation model using 16,384 cores on a Cray XC supercomputer. Lines represent elapsed time, element processing, contact and miscellaneous functions.

80-Million Element Model

For their final test, the team ran a large model of the fan blade off simulation using 82 million nodal points and 80.6 million solid elements. With the earlier LS-DYNA version, the large model took more than a month using 16,384 cores on the Cray XC system. Using the optimized LS-DYNA R8.0.0 the simulation took only 21 hours — or **34 times faster** than R7.1.2.

Read the full white paper at www.cray.com.