

Researchers run largest known transparent checkpointing process

November 23 2016

A team of researchers led by Jiajun Cao, a PhD candidate in the College of Computer and Information Science (CCIS) at Northeastern University, recently completed what appears to be the largest known instance of transparent checkpointing.

Transparent checkpointing allows [computer](#) scientists and engineers working on large projects to save and reopen programs without modifying any code. This assures researchers working across hundreds or thousands of computers that their work will be safe in case of a computer failure. Their programs are run on CPU cores, and computers can contain [multiple cores](#), allowing them to run one program simultaneously across multiple cores. Transparent checkpointing could simplify the work of computer scientists handling large amounts of data and using supercomputers to process that data. For example, with transparent checkpointing software, meteorologists can process and analyze billions of pieces of weather data without the fear that a computer crash could erase that work.

"The idea of checkpointing is that one can take a running computation, automatically stop it in the middle and save the state of everything to a file on disk," Gene Cooperman, a professor at CCIS and Cao's advisor, explains. "Then you can copy that file to another computer or keep it on the same one. When you restart, the program continues running from where it left off." Cooperman's work with Distributed Multi-Threaded CheckPointing (DMTCP) software, which is responsible for checkpointing, is now in its second decade.

What makes this example of transparent checkpointing significant is the massive amount of data that was run and saved in a short period of time. The MVAPICH software supporting the Message Passing Interface (MPI) was used to run the High Performance Conjugate Gradients (HPCG) program for linear algebra in parallel over 32,768 CPU cores on 2,048 computers. It used a total memory of 38 terabytes, and was checkpointed in 10 minutes and 53 seconds. A second program, Nanoscale Molecular Dynamics (NAMD), was run in parallel over 16,368 CPU cores on 1,024 computers, using a total memory of 10 terabytes. It was checkpointed in two minutes and 38 seconds. Checkpointing these amounts of data in 11 minutes or less is a breakthrough for scientists usually restricted by having to run their programs before modifying and saving them within 24-hour time slots.

These processes were carried out on the Stampede supercomputer at the Texas Advanced Computer Center (TACC). Stampede is one of the world's largest supercomputers. The research was supported by a grant from the National Science Foundation awarded to Cooperman's DMTCP project, under which Cao's checkpointing research falls.

"These results show how the Extended Collaborative Support Services from the National Science Foundation-supported Extreme Science and Engineering Discovery Environment can help scientists and developers improve the scalability and efficiency of their code on high performance computing clusters," says Jérôme Vienne, a research associate at TACC.

Dhabaleswar K. Panda, who leads the MVAPICH team at Ohio State, explains that "the results of this collaborative work push the existing capabilities of the MVAPICH2 library further in terms of fault-tolerance and check-pointing."

Cao's collaborators include Kapil Arya of Mesosphere, Inc.; Rohan Garg and Gene Cooperman of Northeastern University; Shawn Matott of the

Center for Computational Research at the State University of New York at Buffalo; Dhabaleswar K. Panda and Hari Subramoni of Ohio State University; and Jérôme Vienne of the Texas Advanced Computing Center at the University of Texas at Austin.

The paper, titled "System-level Scalable Checkpoint-Restart for Petascale Computing," is available to read online. This work will be published at the 22nd Institute of Electrical and Electronics Engineers International Conference on Parallel and Distributed Systems (ICPADS) in December 2016.

Provided by Northeastern University

Citation: Researchers run largest known transparent checkpointing process (2016, November 23)
retrieved 25 April 2024 from

<https://phys.org/news/2016-11-largest-transparent-checkpointing.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.