

# New programming approach seeks to make large-scale computation more reliable

October 8 2015, by Steve Koppes

---



Andrew Chien uses Midway, the Research Computing Center's supercomputing cluster, as an experimental vehicle to test how the project Global View Resilience handles errors. Credit: Research Computing Center

Moore's Law, the observation that integrated circuits halve in size every

two years, has been good to us. Prices for computers have dropped precipitously over the last few decades, even as their power has skyrocketed.

But as we approach the 50th anniversary of Moore's Law, that whole paradigm might be coming to an end: Today's circuitry is so small that it's brushing up against the limits of quantum mechanics. Future computers will need a new paradigm, argues Andrew Chien, the William Eckhardt Distinguished Service Professor of Computer Science and senior fellow in the Computation Institute, who is involved in several projects to pave the way for one. One such project is already bearing fruit, a concept called Global View Resilience—not designed so much to prevent errors as to allow a program to recover from them.

The traditional assumption among hardware and software experts in large-scale scientific computation was that they could depend on their computer hardware to be reliable, Chien explained. But the more circuitry brushes up against the quantum limit, and the more complex supercomputers—and the programs they run—get, the greater the odds that somewhere along the line something will go wrong. It could be a single bit error, corrupted data or a failure in flash memory—anything that interferes with getting the right data to the right place at the right time.

In the early days of computing, if your hardware failed you, you had no choice but to run the program again. More recently, researchers have been using a technique called checkpoint restart, which periodically saves the data at a given point mid-calculation. This is effectively the same method you use when you save a word document while working on it, but that only gives you a way to go back and restart the program—you have no way of knowing if the calculation has gone wrong until it's already finished.

But now, Chien said, computer scientists are looking at the possibility of such high rates of error that checkpoint restart is no longer viable. "You might have multiple different errors on your machine happening at the same time, or happening every few hours or few minutes or few seconds," he said. "You need to find a way of saving things as well as correcting things on the fly if you want your computation to succeed."

That's where GVR comes in. GVR enables applications to not only save the work underway, it also enables flexible error checking and allows the program to fix itself while still in operation. Applications can even specify which parts of a computation are more important than others and which need more care.

The GVR group, which includes postdoctoral scholars Nan Dun and Hajime Fujita and graduate student Aiman Fang, is using the Research Computing Center's supercomputing cluster Midway, located on the Hyde Park campus, as an experimental test vehicle. They run programs with different numbers of nodes or patterns of clusters, introducing errors along the way and seeing how well GVR allows the programs to recover. Virtually all of the errors in the test programs are injected by the researchers. "Our experience with Midway is that it's pretty reliable," Chien said.

GVR is already in use in some supercomputing centers in national labs, but in the long term, Chien sees a role for the concept beyond academia and research. In the future, even small computing devices like cellphones might become more unreliable as consumers keep them longer, since older devices are more error-prone, or want to run them using less energy, which correlates with more errors.

"We have the dream that these kind of techniques we're exploring in GVR will eventually have an impact not only in supercomputing and Facebook, Google or Amazon servers, but eventually even in the small

mobile devices that you and I use every day."

Directly experimenting on Midway, rather than using it as a tool to analyze other data, is a unique use for the cluster. Chien thinks it's unfortunate just how unusual that is.

"Computer scientists, who are the root of many of these computer systems innovations, don't often test them at scale of tens of thousands of nodes. The chemists or physicists tend to dominate use of supercomputers. And we, computer scientists, should be large-scale users of supercomputers for systems experiments at scale."

Provided by University of Chicago

Citation: New programming approach seeks to make large-scale computation more reliable (2015, October 8) retrieved 20 March 2024 from <https://phys.org/news/2015-10-approach-large-scale-reliable.html>

<p>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.</p>
--