

Bold idea for 'big data': Researchers take aim at data glut with customized optical network

September 20 2013



Rice University computer networking researchers are taking on "big data" with a customized, energy-efficient optical network called BOLD that will be capable of feeding rivers of data to Rice's supercomputers. Credit: Rice University/Photos.com

Computer networking researchers at Rice University have a new idea for

how to handle the mountains of data piling up in the labs of their fellow scientists around campus: Create a customized, energy-efficient optical network that can feed rivers of data to Rice's supercomputers.

The new [network](#) is called BOLD—short for "Big data and Optical Lightpaths-Driven Networked Systems Research Infrastructure"—and it's about to become a reality, thanks to a new grant from the National Science Foundation.

"Advances in computing and sensing technologies have led to a similar problem across many disciplines in science and engineering today," said BOLD principal investigator T.S. Eugene Ng, associate professor of computer science and of electrical and computer engineering at Rice. "Experiments produce mountains of data, and there is often no efficient way to process that data to make discoveries and solve problems.

"From a computing infrastructure perspective, the challenge goes beyond just moving data," Ng said. "We also need to develop transformative ideas in the network control software, operating systems and applications so that they can keep up with a faster network. Above all, for this network design to be appealing to industry, it has to be energy-efficient, scalable and nonintrusive to the end user."

BOLD will take advantage of [optical data](#)-networking switches, which have much higher capacity than typical electronic switches that are used mostly in Internet data centers. Optical switches are nothing new, but because of subtle differences in the way electronic and optical switches operate, the two technologies are not interchangeable.

"There's a trade-off," Ng said. "Optical networking devices consume very little power and can support enormous data rates, but they must first be configured, for example, by moving microelectromechanical mirrors into position, to establish a circuit. Electronic switches don't have

moving parts, so they don't have that pesky delay."

BOLD will be a hybrid network that combines both electronic and optical switches. It will also contain something new: a type of optical switch without the moving parts—and the delays—of traditional switches. These new silicon-photonics switches will be built in the laboratory of co-principal investigator (co-PI) Qianfan Xu, assistant professor of electrical and computer engineering at Rice, who specializes in creating ultracompact optical devices on chips.

"To make use of these three types of technology, we need an intelligent layer that can analyze data flow and demand, all the way up to the application layer, and dynamically allocate network resources in the most efficient way," Ng said.

The task of optimizing network design and performance will fall to Ng and co-PIs Alan Cox and Christopher Jermaine, both associate professors of computer science at Rice. Computational mathematician Bill Symes, also a co-PI, will help with both algorithm design and with testing how much BOLD can improve performance on "big data" problems.

Symes, the Noah Harding Professor of Computational and Applied Mathematics and professor of Earth science, directs the Rice Inversion Project (TRIP), an industry-funded consortium that solves complex seismic data processing challenges. For example, one type of operation called "adjoint state computation," which is used in 3-D seismic analyses, requires comparing two time-dependent simulations—one running forward in time and the other running backward. This type of computation, which is also used in aircraft design and meteorological research, routinely generates tens to hundreds of terabytes of intermediate data that must be loaded, cached, recalled, modified and saved many times over. For a sense of scale, 10 terabytes of data is about

the size of the entire print collection of the Library of Congress.

Ng said adjoint state computations are just one example of the extremely demanding data-intensive computations that BOLD can help streamline. The NSF grant runs for three years, but Ng said he hopes BOLD will improve the performance of computationally intensive research at Rice for years to come.

Rice's Ken Kennedy Institute for Information Technology helped facilitate the BOLD collaboration as part of its efforts to address ongoing challenges in computational science. Rice Information Technology's Networking, Telecommunications and Data Center group and the Rice IT Research Computing Support Group will help develop and support the BOLD network.

More information:

www.nsf.gov/awardsearch/showAward?AWD_ID=1305379

Provided by Rice University

Citation: Bold idea for 'big data': Researchers take aim at data glut with customized optical network (2013, September 20) retrieved 20 September 2024 from <https://phys.org/news/2013-09-bold-idea-big-aim-glut.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.