A Winning Combination

Wu-chun Feng

Team Leader and Technical Staff Member, Los Alamos National Laboratory, N.M.

www.lanl.gov

BS and MS in Computer Engineering, Penn State Univ., PhD. in Computer Science, Univ. of Illinois at Urbana-Champaign

HPCwire's Top People & Organizations To Watch List, Three R&D 100 Awards, and Innovative Supercomputer Architecture Award

He has written over 100 articles for textbooks, journals, and magazines. Over 100 have been written about him. What is it about Wu? Other than being intelligent and ahead of his time, he is also one really nice guy.

First off, let's start with being ahead of his time: Wu, a team leader at Los Alamos National Laboratory (LANL), N.M., conducts research in high-performance networking and computing (HPNC), with a focus on building efficient systems for HPNC, from hardware architecture to systems and application software.

One of these, mpiBLAST: A High-Speed Software Catalyst for Genetic Research, employs a process known as in-memory database segmentation, in which a database is segmented so that each compute node only has to search a distinct portion of the entire database. Due to its unexpected success—it has been downloaded over 10,000 times in the past two years—Wu and his team are now moving along two parallel tracks.

"First, we hope to continue to incrementally develop, maintain, and support mpiBLAST for the greater computational biology community," he says. "Second, based on our own critical evaluation of mpiBLAST, we are planning a major re-design, one that we expect will deliver a few more orders of magnitude of improvement. We are also looking into collaborations where mpiBLAST is being used as part of a longer pipe-line of programs."

On the high-performance networking front, Wu has been working on multiple aspects (hardware-software interfaces, end-host software algorithms, traffic characterization, and performance evaluation) in a number of environments (local-area, system-area, and wide-area network). His work with Intel on a 10-Gb Ethernet adapter—which can transfer information from one computer to another up to 23,000 times faster than a DSL connection—fits in all of the above contexts.

"My team and I are currently working with 10-Gb Ethernet vendors to figure out how to get information into and out of a compute node more efficiently, without significantly involving the main processor on the end-host computer, and while abiding by traditional Ethernet constraints, such as the maximum transmission unit size."

Ultimately, his goal is to better enable distributed computational grids that provide consistent and pervasive access to resources to enable sharing of computational resources, utility and autonomic computing, collaboration among virtual organizations, and distributed data processing. In short, a worldwide computer. "Part of enabling these goals requires developing tools and facilities that will troubleshoot and monitor such grids, which is the focus of our MAGNET research project. MAGNET stands for Monitoring Apparatus for General kerNel-Event Tracing. This tool is used for both troubleshooting and enabling self-aware adaptive systems."

Despite the grand accomplishments of all of his projects, the main challenge has been that they required "too much effort" to accomplish. "The holy grail is for systems to be self-aware, self-adapting, and self-healing so that results can be automatically achieved," he says.

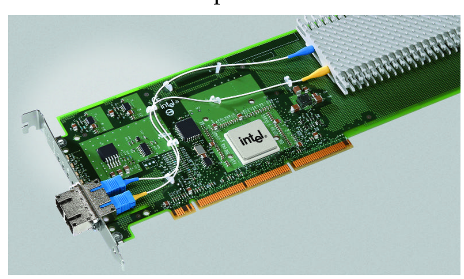

Wu and his team at LANL optimized Intel's 10-Gb Ethernet adapter and its associated subsystems, enhancing its performance nearly 300%.

Wu's laundry list of future technological advances he would like to see include eliminating the "digital divide" due to socioeconomics and age by improving the ease-of-use and high cost of the technologies. He'd also like to see the seamless integration of technology and its environment, with tools like tablet PCs and networks for a virtual classroom, instead of blackboards. A desire to see quantum computers, 100 Mb/sec speeds to the home, and a worldwide digital library tops his list as well.

—Lorraine Joyce