New software design technique allows programs to run faster

April 5 2010

(PhysOrg.com) -- Researchers at North Carolina State University have developed a new approach to software development that will allow common computer programs to run up to 20 percent faster and possibly incorporate new security measures.

The researchers have found a way to run different parts of some programs - including, for the first time, such widely used programs as word processors and Web browsers - at the same time, which makes the programs operate more efficiently.

In order to understand how they did it, you have to know a little bit about computers. The brain of a computer chip is its central processing unit, or "core." Computing technology has advanced to the point where it is now common to have between four and eight cores on each chip. But for a program to utilize these cores, it has to be broken down into separate "threads" - so that each core can execute a different part of the program simultaneously. The process of breaking down a program into threads is called parallelization, and allows computers to run programs very quickly.

However, some programs are difficult to parallelize, including word processors and Web browsers. These programs operate much like a flow chart - with certain program elements dependent on the outcome of others. These programs can only utilize one core at a time, minimizing the benefit of multi-core chips.

But NC State researchers have developed a technique that allows hard-to-parallelize applications to run in parallel, by using nontraditional approaches to break programs into threads.

Every computer program consists of multiple steps. The program will perform a computation, then perform a memory-management function - which prepares memory storage to contain data or frees up memory storage which is currently in use. It repeats these steps over and over again, in a cycle. And, for difficult-to-parallelize programs, both of these steps have traditionally been performed in a single core.

"We've removed the memory-management step from the process, running it as a separate thread," says Dr. Yan Solihin, an associate professor of electrical and computer engineering at NC State, director of this research project, and co-author of a paper describing the research. Under this approach, the computation thread and memory-management thread are executing simultaneously, allowing the computer program to operate more efficiently.

"By running the memory-management functions on a separate thread, these hard-to-parallelize programs can operate approximately 20 percent faster," Solihin says. "This also opens the door to development of new memory-management functions that could identify anomalies in program behavior, or perform additional security checks. Previously, these functions would have been unduly time-consuming, slowing down the speed of the overall program."

Using the new technique, when a memory-management function needs to be performed, "the computational thread notifies the memory-management thread - effectively telling it to allocate data storage and to notify the computational thread of where the storage space is located," says Devesh Tiwari, a Ph.D. student at NC State and lead author of the paper. "By the same token, when the computational thread no longer

needs certain data, it informs the memory-management thread that the relevant storage space can be freed."

More information: The paper, "MMT: Exploiting Fine-Grained Parallelism in Dynamic Memory Management," will be presented April 21 at the IEEE International Parallel and Distributed Processing Symposium in Atlanta.

Provided by North Carolina State University