

New system greatly speeds common parallelcomputing algorithms

July 3 2017, by Larry Hardesty



A new system dubbed Fractal achieves 88-fold speedups through a parallelism strategy known as speculative execution. Credit: MIT News

The chips in most modern desktop computers have four "cores," or processing units, which can run different computational tasks in parallel.



But the chips of the future could have dozens or even hundreds of cores, and taking advantage of all that parallelism is a stiff challenge.

Researchers from MIT's Computer Science and Artificial Intelligence Laboratory have developed a new system that not only makes parallel programs run much more efficiently but also makes them easier to code.

In tests on a set of benchmark algorithms that are standard in the field, the researchers' new system frequently enabled more than 10-fold speedups over existing systems that adopt the same parallelism strategy, with a maximum of 88-fold.

For instance, algorithms for solving an important problem called max flow have proven very difficult to parallelize. After decades of research, the best parallel implementation of one common max-flow algorithm achieves only an eightfold speedup when it's run on 256 parallel processors. With the researchers' new system, the improvement is 322-fold—and the program required only one-third as much code.

The new system, dubbed Fractal, achieves those speedups through a parallelism strategy known as speculative execution.

"In a conventional parallel program, you need to divide your work into tasks," says Daniel Sanchez, an assistant professor of electrical engineering and computer science at MIT and senior author on the new paper. "But because these tasks are operating on shared data, you need to introduce some synchronization to ensure that the data dependencies that these tasks have are respected. From the mid-90s to the late 2000s, there were multiple waves of research in what we call speculative architectures. What these systems do is execute these different chunks in parallel, and if they detect a conflict, they abort and roll back one of them."



Constantly aborting computations before they complete would not be a very efficient parallelization strategy. But for many applications, aborted computations are rare enough that they end up squandering less time than the complicated checks and updates required to synchronize tasks in more conventional parallel schemes. Last year, Sanchez's group reported a system, called Swarm, that extended speculative parallelism to an important class of computational problems that involve searching data structures known as graphs.

Irreducible atoms

Research on speculative architectures, however, has often run aground on the problem of "atomicity." Like all parallel architectures, speculative architectures require the programmer to divide programs into tasks that can run simultaneously. But with speculative architectures, each such task is "atomic," meaning that it should seem to execute as a single whole. Typically, each atomic task is assigned to a separate processing unit, where it effectively runs in isolation.

Atomic tasks are often fairly substantial. The task of booking an airline flight online, for instance, consists of many separate operations, but they have to be treated as an atomic unit. It wouldn't do, for instance, for the program to offer a plane seat to one customer and then offer it to another because the first customer hasn't finished paying yet.

With speculative execution, large atomic tasks introduce two inefficiencies. The first is that, if the task has to abort, it might do so only after chewing up a lot of computational cycles. Aborting smaller tasks wastes less time.

The other is that a large atomic task may have internal subroutines that could be parallelized efficiently. But because the task is isolated on its own processing unit, those subroutines have to be executed serially,



squandering opportunities for performance improvements.

Fractal—which Sanchez developed together with MIT graduate students Suvinay Subramanian, Mark Jeffrey, Maleen Abeydeera, Hyun Ryong Lee, and Victor A. Ying, and with Joel Emer, a professor of the practice and senior distinguished research scientist at the chip manufacturer NVidia—solves both of these problems. The researchers, who are all with MIT's Department of Electrical Engineering and Computer Science, describe the system in a paper they presented this week at the International Symposium on Computer Architecture.

With Fractal, a programmer adds a line of code to each subroutine within an atomic task that can be executed in parallel. This will typically increase the length of the serial version of a program by a few percent, whereas an implementation that explicitly synchronizes parallel tasks will often increase it by 300 or 400 percent. Circuits hardwired into the Fractal chip then handle the parallelization.

Time chains

The key to the system is a slight modification of a circuit already found in Swarm, the researchers' earlier speculative-execution system. Swarm was designed to enforce some notion of sequential order in <u>parallel</u> <u>programs</u>. Every task executed in Swarm receives a time stamp, and if two tasks attempt to access the same memory location, the one with the later time stamp is aborted and re-executed.

Fractal, too assigns each atomic task its own time stamp. But if an atomic task has a parallelizable subroutine, the subroutine's time stamp includes that of the <u>task</u> that spawned it. And if the subroutine, in turn, has a parallelizable subroutine, the second subroutine's time stamp includes that of the first, and so on. In this way, the ordering of the subroutines preserves the ordering of the atomic tasks.



As tasks spawn subroutines that spawn subroutines and so on, the concatenated time stamps can become too long for the specialized circuits that store them. In those cases, however, Fractal simply moves the front of the <u>time</u>-stamp train into storage. This means that Fractal is always working only on the lowest-level, finest-grained tasks it has yet identified, avoiding the problem of aborting large, high-level atomic tasks.

More information: Fractal: An Execution Model for Fine-Grain Nested Speculative Parallelism. ISCA '17, June 24-28, 2017, <u>DOI:</u> 10.1145/3079856.3080218

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: New system greatly speeds common parallel-computing algorithms (2017, July 3) retrieved 2 May 2024 from https://phys.org/news/2017-07-greatly-common-parallel-computing-algorithms.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.