


# A new parallel strategy for tackling turbulence on Summit

November 14 2019



An illustration of intricate flow structures in turbulence from a large simulation performed using 1,024 nodes on Summit. The lower right frame shows a zoom-in view of a high-activity region. Credit: Dave Pugmire and Mike Matheson, Oak Ridge National Laboratory

Turbulence, the state of disorderly fluid motion, is a scientific puzzle of great complexity. Turbulence permeates many applications in science and engineering, including combustion, pollutant transport, weather forecasting, astrophysics, and more. One of the challenges facing scientists who simulate turbulence lies in the wide range of scales they must capture to accurately understand the phenomenon. These scales can span several orders of magnitude and can be difficult to capture within the constraints of the available computing resources.

High-performance computing can stand up to this challenge when paired with the right scientific code; but simulating turbulent flows at problem sizes beyond the current state of the art requires new thinking in concert with top-of-the-line heterogeneous platforms.

A team led by P. K. Yeung, professor of aerospace engineering and mechanical engineering at the Georgia Institute of Technology, performs direct numerical simulations (DNS) of turbulence using his team's new code, GPUs for Extreme-Scale Turbulence Simulations (GESTS). DNS can accurately capture the details that arise from a wide range of scales. Earlier this year, the team developed a new algorithm optimized for the [IBM AC922 Summit](#) supercomputer at the Oak Ridge Leadership Computing Facility (OLCF). With the new algorithm, the team reached a performance of less than 15 seconds of wall-clock time per time step for more than 6 trillion grid points in space—a new world record surpassing the prior state of the art in the field for the size of the problem.

The simulations the team conducts on Summit are expected to clarify important issues regarding rapidly churning turbulent fluid flows, which will have a direct impact on the modeling of reacting flows in engines and other types of propulsion systems.

GESTS is a computational fluid dynamics code in the Center for

Accelerated Application Readiness at the OLCF, a US Department of Energy (DOE) Office of Science User Facility at DOE's Oak Ridge National Laboratory. At the heart of GESTS is a basic math algorithm that computes large-scale, distributed fast Fourier transforms (FFTs) in three spatial directions.

An FFT is a math algorithm that computes the conversion of a signal (or a field) from its original time or space domain to a representation in the frequency (or wave number) space—and vice versa for the inverse transform. Yeung extensively applies a huge number of FFTs in accurately solving the fundamental partial differential equation of fluid dynamics, the Navier-Stokes equation, using an approach known in mathematics and scientific computing as "pseudospectral methods."

Most simulations using massive CPU-based parallelism will partition a 3-D solution domain, or the volume of space where a fluid flow is computed, along two directions into many long "data boxes," or "pencils." However, when Yeung's team met at an OLCF GPU Hackathon in late 2017 with mentor David Appelhans, a research staff member at IBM, the group conceived of an innovative idea. They would combine two different approaches to tackle the problem. They would first partition the 3-D domain in one direction, forming a number of data "slabs" on Summit's large-memory CPUs, then further parallelize within each slab using Summit's GPUs.

The team identified the most time-intensive parts of a base CPU code and set out to design a [new algorithm](#) that would reduce the cost of these operations, push the limits of the largest problem size possible, and take advantage of the unique data-centric characteristics of Summit, the world's most powerful and smartest supercomputer for open science.

"We designed this algorithm to be one of hierarchical parallelism to ensure that it would work well on a hierarchical system," Appelhans said.

"We put up to two slabs on a node, but because each node has 6 GPUs, we broke each slab up and put those individual pieces on different GPUs."

In the past, pencils may have been distributed among many nodes, but the team's method makes use of Summit's on-node communication and its large amount of CPU memory to fit entire data slabs on single nodes.

"We were originally planning on running the code with the memory residing on the GPU, which would have limited us to smaller problem sizes," Yeung said. "However, at the OLCF GPU Hackathon, we realized that the NVLink connection between the CPU and the GPU is so fast that we could actually maximize the use of the 512 gigabytes of CPU memory per node."

The realization drove the team to adapt some of the main pieces of the code (kernels) for GPU data movement and asynchronous processing, which allows computation and data movement to occur simultaneously. The innovative kernels transformed the code and allowed the team to solve problems much larger than ever before at a much faster rate than ever before.

The team's success proved that even large, communication-dominated applications can benefit greatly from the world's most powerful supercomputer when code developers integrate the heterogenous architecture into the algorithm design.

## **Coalescing into success**

One of the key ingredients to the team's success was a perfect fit between the Georgia Tech team's long-held domain science expertise and Appelhans' innovative thinking and deep knowledge of the machine.

Also crucial to the achievement was the OLCF's early access Ascent and Summitdev systems and a million-node-hour allocation on Summit provided by the Innovative Novel and Computational Impact on Theory and Experiment (INCITE) program, jointly managed by the Argonne and Oak Ridge Leadership Computing Facilities, and the Summit Early Science Program in 2019.

Oscar Hernandez, tools developer at the OLCF, helped the team navigate challenges throughout the project. One such challenge was figuring out how to run each single parallel process (that obeys the message passing interface [MPI] standard) on the CPU in conjunction with multiple GPUs. Typically, one or more MPI processes are tied to a single GPU, but the team found that using multiple GPUs per MPI process allows the MPI processes to send and receive a smaller number of larger messages than the team originally planned. Using the OpenMP programming model, Hernandez helped the team reduce the number of MPI tasks, improving the code's communication performance and thereby leading to further speedups.

Kiran Ravikumar, a Georgia Tech doctoral student on the project, will present details of the algorithm within the technical program of the 2019 Supercomputing Conference, SC19.

The team plans to use the code to make further inroads into the mysteries of turbulence; they will also introduce other physical phenomena such as oceanic mixing and electromagnetic fields into the code in the future.

"This code, and its future versions, will provide exciting opportunities for major advances in the science of turbulence, with insights of generality bearing upon turbulent mixing in many natural and engineered environments," Yeung said.

**More information:** K. Ravikumar, D. Appelhans, and P. K. Yeung, "GPU Acceleration of Extreme Scale Pseudo-Spectral Simulations of Turbulence using Asynchronism." Paper to be presented at the 2019 International Conference for High Performance Computing, Networking and Storage Analysis (SC19), Denver, CO, November 17–22, 2019.

Provided by Oak Ridge National Laboratory

Citation: A new parallel strategy for tackling turbulence on Summit (2019, November 14) retrieved 23 April 2024 from

<https://techxplore.com/news/2019-11-parallel-strategy-tackling-turbulence-summit.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.