

## Turbulent Times

### TACC's Ranger supercomputer helps P.K. Yeung explore the physics of flow

Turbulence, which manifests itself in bumpy airplane rides, industrial smoke discharges, and scenic waterfalls, has inspired and frustrated generations of physicists. When asked what he would like to know from god, Werner Heisenberg, theorist of the uncertainty principle, said, "I am going to ask him two questions. Why relativity? And why turbulence? I really believe he will have an answer for the first."


Characterized by complex, disorderly motions over a wide range of scales in time and space, turbulence is a grand challenge question that cuts across numerous disciplinary boundaries, from the science of atmospheric phenomena, to the physics of combustion as an energy source for automobiles and jet engines.

"We actually experience turbulence in many aspects of daily life," noted Professor P.K. Yeung of the Georgia Institute of Technology's School of Aerospace Engineering and Division of Computational Science and Engineering, and a leading scholar in the field. "If we want to demonstrate turbulence, all we have to do is stir up a cup of coffee or boil water in a kettle and observe the convoluted paths that the steam takes before it dissipates."

Yet, because of its wide range of scales and highly non-linear nature, understanding and modeling turbulence is a daunting challenge. The difficulty is especially great in flow conditions corresponding to large Reynolds numbers, which is a non-dimensional parameter closely related to the range of scales present in the turbulent fluctuations. Fortunately, this is where rapidly advancing computing power helps.

Together with his research partners, Yeung recently began using Ranger, the world's largest supercomputer for open-science research, at the Texas Advanced Computing Center (TACC) to resolve a turbulence

grid of 68 billion (40963) grid points. The simulations will use up to 16,384 processor cores per simulation (or 131 trillion floating point operations per second), and this number will increase to 32,768 cores in the future (262 teraflops). When completed, this study is expected to be a truly unique resource for the international research community, and will play a key role in helping re-establish US leadership in large-scale turbulence studies, which in the last few years has been challenged by the Earth Simulator in Japan.



A very important effect of turbulence is efficient small-scale mixing, which is significant in flows with chemical reactions, where the local scalar dissipation rate of concentration fluctuations is a key parameter. This image is a 3D color rendering from a 20483 simulation (performed at San Diego Supercomputer Center, SDSC), which shows the complex, fine-scale detail of high-activity regions in the flow. Image reproduced by courtesy of A. Chourasia, SDSC and D.A. Donzis, Georgia Tech (now at University of Maryland).

To compute turbulence on a high-performance computing (HPC) system like Ranger, fluid dynamicists use the technique of Direct Numerical Simulation, where exact physical laws (representing conservation of mass, momentum, and energy) are applied to a

## Turbulent Times

large number of grid points to compute flow evolution over a period of time.

“We can calculate almost anything we like, including many quantities — like details of fine-scale motion in space, or the relative motion between material fluid elements in chaotic motion — which are difficult to measure precisely in the laboratory,” Yeung said. However, the simulations take incredible amounts of computational power and generate huge volumes of data that need processing to ensure the greatest scientific benefit.

To effectively use Ranger’s more than 60,000 cores, Yeung and his team partition their numerical solution into a very large number of sub-units, structured to allow efficient interprocessor communication of large messages. Yeung’s code is based on a two-dimensional domain decomposition implemented by Diego Donzis, of the University of Maryland, and TeraGrid strategic consultant, Dmitry Pekurovsky, from the San Diego Supercomputer Center. Typically, as simulations scale to larger core counts, the percentage of work per core lowers because of the latency of communication between cores. But extensive benchmarking efforts have allowed the researchers to identify configuration parameters that have led to demonstrated scalability as high as 90 percent in many cases. A high scalability percentage leads the researchers to believe that their simulations can be scaled up to even larger core counts in the future, with accurate results.

“The simulations recently begun at TACC are the first phase of a larger effort towards the full exploitation of Petascale

architectures for turbulence research,” Yeung said. Future work will continue the drive towards high Reynolds numbers while incorporating additional physics such as buoyancy, chemical reaction, and solid boundaries.


According to Yeung, untangling the turbulence riddle would allow experts to design aerodynamic devices with low drag, predict pollutant dispersion in an industrial accident or terrorist attack, and develop devices that are cleaner, burn fuel more efficiently and generate fewer byproducts. But without powerful enough computers, none of this would be possible.

“We are embarking on a brave new journey and we cannot do it with the other existing TeraGrid machines. It requires resources at the level of Ranger,” Yeung said. “TACC is making a difference by motivating scientists to think about new possibilities and helping them to perform ambitious simulations that could not be done before, or previously were thought to be ‘too expensive.’”

With twice the processing power of the entire NSF TeraGrid, Ranger catapults turbulence research forward, helping Yeung and his colleagues address fundamental physics with increasing fidelity, and hence contribute to accurate flow modeling in many problems of importance to society.

###

*The Ranger supercomputer is funded through the National Science Foundation (NSF) Office*



P.K. Yeung, Georgia Tech  
 Diego Donzis, U. Maryland  
 Dmitry Pekurovsky, San Diego Supercomputer Center



## **Turbulent Times**

*of Cyberinfrastructure "Path to Petascale" program. The system is a collaboration among the Texas Advanced Computing Center, The University of Texas at Austin's Institute for Computational Engineering and Science, Sun Microsystems, Advanced Micro Devices, Arizona State University, and Cornell University.*

*Ranger is a key resource of the NSF TeraGrid ([www.teragrid.org](http://www.teragrid.org)), a nationwide network of people, resources and services, also sponsored by the NSF Office of Cyberinfrastructure, which enables discovery in U.S. science and engineering. The TeraGrid provides scientists and researchers expertise in and access to large-scale computing power, networking, data-analysis, and visualization systems.*

**Aaron Dubrow**  
**Texas Advanced Computing Center**  
**Science and Technology Writer**  
**March 07, 2008**