


# BURSTS *of* STELLAR TURBULENCE

With PSC's Cray XT3, Paul Woodward and David Porter figured out how to do the hardest thing, run small to medium-sized problems really fast

TURBULENCE AND CONVECTION // BURSTS OF STELLAR TURBULENCE



David Porter and Paul Woodward



Nathan Stone and Raghurama Reddy, PSC developed PDIO, software that routes data from the XT3 in real-time to remote users.

**What if** as a computational scientist you had the ability to ask "what if" questions and get answers not in days or weeks but minutes? And what if you could see the results from your questions in real time and, depending on what you see, shift your inquiry on the fly and get answers almost as quickly as you're able to process what you see. It would be like having a powerful computer wired into your inquiring mind, and your ability to ask new questions as they occur to you, based on rapid results, would transform your ability to do science.

Using PSC's Cray XT3 and with help from PSC-developed software, Paul Woodward and David Porter have been doing this, running interactive simulations of turbulence and "steering" them to explore the features that most interest them. "The XT3 has the fastest processor interconnect in a machine with thousands of CPUs," says Woodward, "and this feature is all important to enable interactive steering of flow simulations."

For many years, Woodward and Porter, astrophysicists at the University of Minnesota's Laboratory for Computational Science and Engineering (LCSE), have studied turbulent astrophysical flows. "The consistent thread in our research," says Woodward, "has been using large-scale supercomputer simulations to understand and model turbulent convection in stars."

Their ambitious long-range goal is to accurately simulate in detail the turbulent dynamics of an entire giant star, stars similar to the sun. As a step toward this, their recent work has

focused on small-scale turbulence, studies that can identify parameters from which to build an accurate model of turbulence on large scales. "This need to represent the large-scale effects of small-scale turbulence," says Woodward, "arises in many areas of science and engineering, from flow in pipes and combustion engines to meteorology and ultimately to the sort of 'stellar weather' we have been simulating for many years."

Their focus on small-scale turbulence led the researchers to a significant breakthrough in massively parallel processing. With the Cray XT3, they've solved the problem of "strong scaling" — which, essentially, is getting a large number of processors to work together efficiently on a moderate-sized problem. "Getting a small problem done fast is the hardest thing," says Woodward.

In January 2007, Woodward and Porter used 4,096 XT3 processors (plus eight input/output nodes), the whole system, to simulate turbulent shear between two fluids. They used a computational grid of  $576^3$  cells, fine enough to resolve the small-scale turbulence they want to understand — a run that would take weeks or months on an average cluster. With performance of 2.32 gigaflops (billions of calculations per second) on each XT3 processor, 9.5 teraflops overall, the run of 6,000 time steps completed in 7.7 minutes.

The drop in per-processor efficiency in using the whole machine was less than 5-percent — a credit to the XT3's interprocessor communication. "That's pretty damn amazing," says Woodward. "It's a testimonial to the effectiveness of your Cray XT3 interconnect."

Although impressive, running a problem of this size in eight minutes, Woodward acknowledges, is overkill. The real significance lies in the potential it demonstrates. "You can take that and say 'Suppose we do a bigger problem not by assigning more work to each processor but by having more processors.' That leads to petascale computing that runs really well on reasonably sized grids. The principle challenge of petascale computing is strong scaling."

## Shear Between Two Gases

From simulation of turbulent mixing between initially isolated layers of two gases of different density traveling across each other at Mach 1/2, these graphics show the "shear layer" that forms between them. The lower fluid is 2.5 times more dense. Only mixtures are visible, with regions of pure denser or lighter fluid transparent. Color goes through yellow to red where the denser fluid predominates and aqua to blue where the lighter fluid predominates.



### CONVERGING ON TURBULENCE IN GIANT STARS

As an ultimate goal, Woodward and Porter want to simulate an entire giant star, which means modeling the swirling, pulsating bands of gas that surround a star's thermonuclear core. These volatile turbulent phenomena underlie and govern the birth and death of stars, which provide not only the light energy that sustains life, but also the elements that are its constituents. The primary means scientists have to learn about these processes is computational simulation.

As a necessary step toward simulating an entire giant star, their recent work with PSC's XT3 has focused on "subgrid-scale" turbulence. To radically simplify, subgrid-scale means turbulence that happens in very small spatial dimensions compared to the large domain of the overall problem, but which nevertheless affects the large scale and must be accounted for in a large-scale model.

With subgrid-scale turbulence, their aim is to run simulations at fine enough resolution (large quantity of grid cells) that the results correspond to what happens in nature. Their test is "convergence" — getting results that no longer significantly change (as measured statistically) as the grid becomes progressively more fine.

One kind of turbulent mixing, known as Rayleigh-Taylor instability, occurs due to gravity when a heavy gas is on top of a lighter one. This sequence shows four time intervals from an XT3 simulation (on a  $768^2 \times 1536$  grid), with regions of predominantly heavier fluid (red), lighter fluid (blue), more even mixture (yellow), and 50-50 (white). By the time of the fourth image, the mixing, though confused, shows descending "mushroom caps" that produce organized very fine-scale mixing.



This thin slice [above left] through the mixing domain at a late time interval from an R-T instability simulation on a fine grid ( $768^2 \times 1536$  cells) shows striking agreement in larger features with the same slice (below left) from the same simulation on a coarser grid ( $384^2 \times 768$ ). The qualitative differences are minor, which suggests convergence — that all reasonable simulation codes would arrive at similar conclusions. "If so," says Woodward, "we may be confident in the simulated results and use them to validate subgrid-scale models of turbulent multifluid mixing."

### TO ASK "WHAT IF?" QUESTIONS ON THE FLY IS A NEW CAPABILITY IN TURBULENCE RESEARCH.

several ways that they demonstrated first at conferences, iGrid2005 in San Diego and again at SC05 in Seattle.

To do simulations in real time as exploratory runs — to ask "what if?" questions on the fly — was a new capability in turbulence research. To make it possible, Woodward and Porter worked closely with PSC scientists Nathan Stone and Raghurama Reddy, who developed specialized software — called Portals Direct I/O (PDIO) — to route simulation data from the XT3 in real time to remote users.

As numbers crunch in Pittsburgh, PDIO assembles the resulting data streams and routes them for Woodward and Porter to volume-render and display as real-time images. "The ability to have instant

### PETASCALE & STRONG SCALING

To appreciate the leap to strong scaling, it helps to know about its opposite, weak scaling. Over the next few years, supercomputing infrastructure will evolve to "petascale" — the ability to do a quadrillion ( $10^{15}$ ) calculations per second. This amount of computing — which will enable major advances in many areas of science — will come from systems comprising tens or hundreds of thousands of processors linked with each other so that scientists can use them in teams to address very large problems. In computational lingo, this is "scaling" — as you "scale up" a computational problem, you apply more processors.

Weak scaling is the easy approach, at least relatively speaking. It means, basically, that as you scale up you make the problem larger, but each processor does the same amount of computing. If you double the size of a 3D grid in each dimension, for instance, you need eight times more processors.

Strong scaling, on the other hand, as Woodward says, is not easy. It means that for a given problem as you scale up, say, from 100 to 800 processors you apply this greater number of processors to the same grid, so that each processor is now doing 1/8 as much work. You would like the job to run eight times faster, and to do that requires restructuring how the program divides the work among processors and increases the communication between them.

### THE XT3 & RESTRUCTURED CODE

Woodward and Porter appreciated the XT3 — fast processors with a fast interconnect — from its inception as a TeraGrid resource in 2005, and worked with it from early on. Using their turbulence code, PPM (Piecewise Parabolic Method), they developed the ability to do interactive, steerable runs remotely with real-time visualization — a breakthrough in

response from supercomputer simulations is very useful," says Woodward. "You can, for instance, change the Mach number, and almost immediately see what that does. This is a whole pipeline of utility programs that we tied together in an automated fashion. The support from PSC is outstanding."

With real-time interactive capability in place, over the past year Woodward and Porter extended their effort to rapidly solve small-scale problems and made a dramatic leap when they found a way to restructure their PPM code. The happy discovery came as result of working with the Cell processor, a multi-core processor — multiple processors on one "chip" — with much potential for supercomputing applications.

The approach they arrived at, says Porter, is "maximum cache reuse" — not easy to explain, but it has to do with "vector length," strings of identical mathematical operations. From early days of supercomputing, code performance was optimal with vector-lengths as long as possible. With the Cell processor, however, Woodward and Porter found — to their surprise — that shortening the vector-length to conserve cache also speeded up performance.

"Over the past year," says Woodward, "we have completely redesigned our PPM simulation codes, and improved performance — about one gigaflop per processor — by dramatically reducing the data granule on which the processors operate." The performance improvement depends on the XT3's fast interconnect, because maximum cache reuse makes increased demands of inter-processor communication. The combination of restructured code and the XT3, as Woodward and Porter demonstrated with 9.5 teraflops sustained on 4,000 processors, is a big step in the direction of petascale.

### MORE INFORMATION:

[www.psc.edu/science/2007/turbulence.html](http://www.psc.edu/science/2007/turbulence.html)