Columbia's architecture also lends itself well to large-scale molecular dynamics simulations involving millions to billions of atoms. These simulations can be important for designing new materials for thermal protection systems or for building future spacecraft. Preliminary tests on 1,920 processors of the 2048 system have achieved incredible results including a 1.4 million- and 18.9 billion-atom simulation using a space-time multi-resolution molecular dynamics algorithm.

63 Teraflops and Beyond

To remain a viable resource capable of meeting the agency's and the country's growing high-end computing requirements for advanced modeling and simulation, NASA is planning to work with other vendors to explore emerging architectures and cluster solutions.

While Columbia has already increased NASA's compute capability ten-fold and has attained science results not previously possible, the agency's demand for high-end computing continues to grow rapidly. To keep up with this demand for high-end computing capabilities, NASA will continue working with the National Coordination Office and sister agency programs like the U.S. Defense Advanced Research Projects Agency's (DARPA) High Productivity Computing Systems (HPCS) program to bring sustained petaflop-scale computing to the U.S. by the end of the decade.

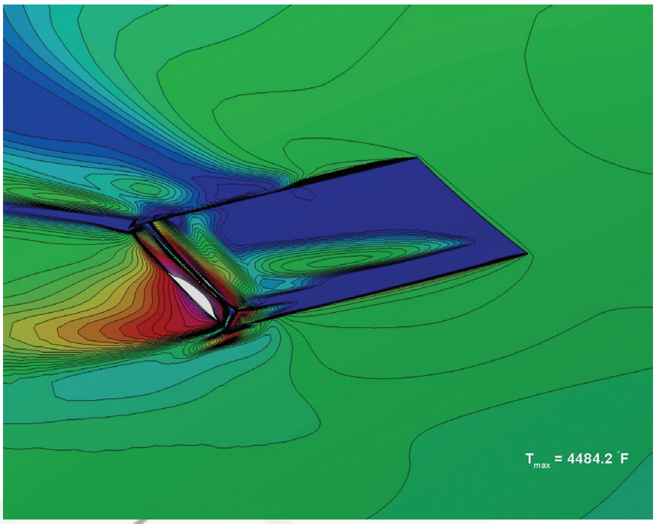

Figure 3: A new, rapid aerothermal computational fluid dynamics (CFD) analysis capability using Columbia permits near-real-time analysis of observed Orbiter damage during flight. Shown here: temperature contours of a shuttle tile cavity case traveling at Mach 22.9. *Image courtesy of the NASA Debris Transport Team.*

Newton as a Single System

Kevin Roy
Senior HPC Consultant, Manchester Computing, University of Manchester

CfS (the consortium that provides the CSAR service) has continually pushed Newton to expand its potential and make it a more productive, easier to use system for those who use the system. As a result we have implemented our final planned hardware alteration; coupling both halves of the system into a single system with full NUMAlink interconnect and running a single operating system.

Before Newton became a 512 processor machine it had really only operated as two 256 processor machines with batch jobs running on one of the two hosts and never across both of them. The effect of this was that the maximum job size was restricted to the size that would fit on each machine.

The reason for this limitation was that although each of the two 256 processor machines were connected using SGI's Numalink interconnect, which gives 3.2Gb/s between pairs of processors, the machines were connected to each other with a single gigabit ethernet connection, which needed to be shared amongst all the processors.

The two 256 processor machines have now been reconfigured architecturally as a single 512 processor machine allowing batch jobs to span the whole system; this allows jobs of up to 496 processors to run. Newton is a 512 processor machine but some processors are reserved for interactive (4) and system processes (12), leaving the remaining 496 available for batch. This has also enabled the scheduling to improve, providing better turnaround for medium and small jobs.

We now have a single point of entry into the machine for all processors, allowing networked applications (such as grid projects) to connect directly to the internet.

The simplicity of a single system is evident in the file systems too. Operating and system software is now identical across the whole machine and all local temporary disk can now be viewed from the interactive area of the machine.