Center for Applied
Scientific Computing

Achieving High Sustained Performance in an Unstructured Mesh CFD Application

Gordon Bell Special Award

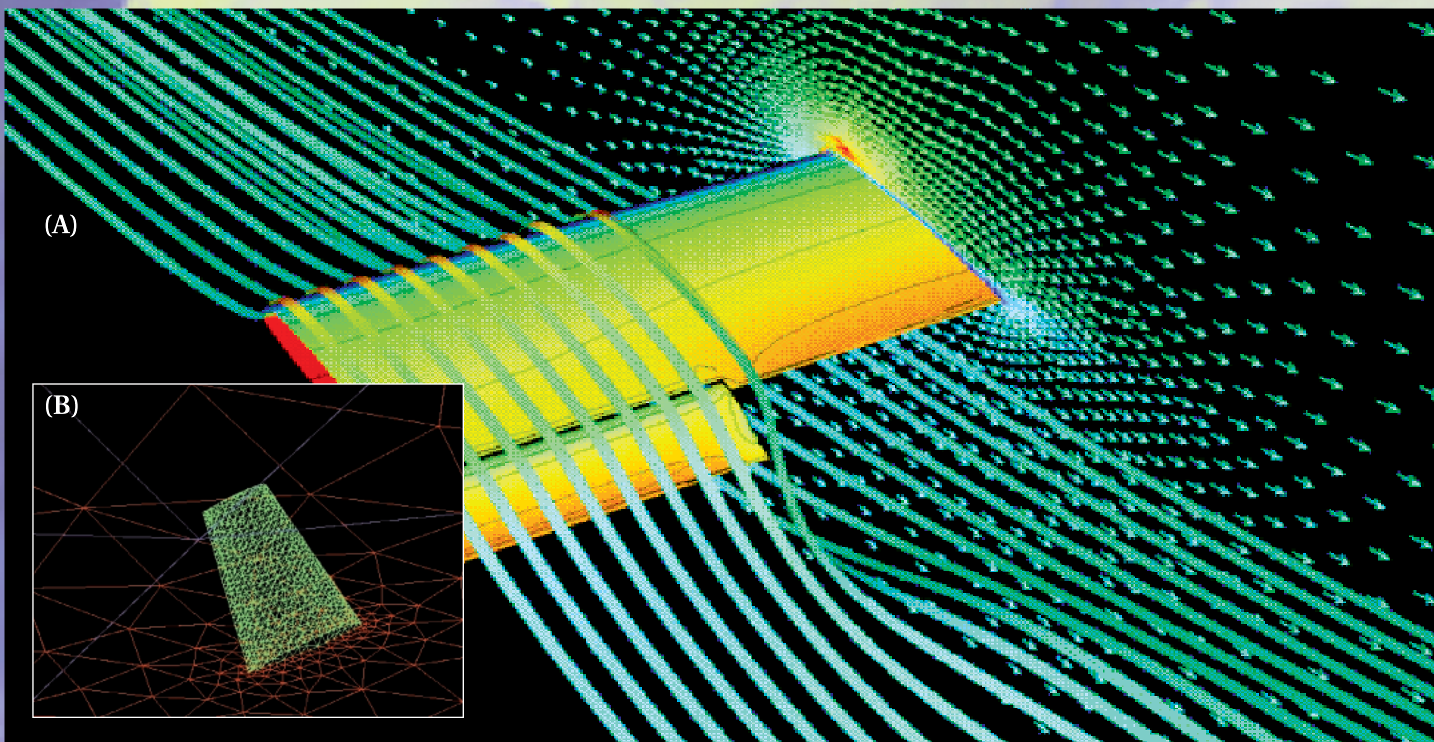

Figure 1A (large image). The research team that created Achieving High Sustained Performance in an Unstructured Mesh CFD Application *adapted a legacy F77 computational fluid dynamics code to the distributed memory parallel environment and applied it to problems in computational aerodynamics. This implicit unstructured grid simulation won a Gordon Bell Special Award at SC99.*

Figure 1B (inset image). The simulation solved for the airflow over an ONERA M6 wing surface outlined in green triangles, as shown, with tetrahedral volume elements. The largest computational domain comprised 2.8 million vertices. For clarity, a grid with about 9 thousand vertices is shown here.

David E. Keyes
Lawrence Livermore National Laboratory

Dinesh Kaushik
Old Dominion University

Kyle Anderson
NASA

William Gropp and Barry Smith
Argonne National Laboratory

A collaborative research team created *Achieving High Sustained Performance in an Unstructured Mesh CFD Application*, the simulation project that won a Gordon Bell Special Award at SC99. Various Gordon Bell Awards are presented annually at the international "SC" supercomputing meeting, and SC99 was held in November 1999 at Portland, Oregon. A Gordon Bell Special Award recognizes exceptionally well executed practical supercomputing simulations that achieve new levels of performance.

This presentation highlighted a three-year project by an interdisciplinary team that began with a legacy F77 computational fluid dynamics code. The research demonstrated that implicit unstructured grid simulations can execute at rates not far from those of explicit structured grid codes, provided attention is paid to data motion complexity and the reuse of data positioned at the levels of the memory hierarchy closest to the processor, in addition to traditional operation count complexity.

UCRL-TB-139276
http://www.llnl.gov/CASC/news/gordonbell.html

Achieving High Sustained Performance in an Unstructured Mesh CFD Application

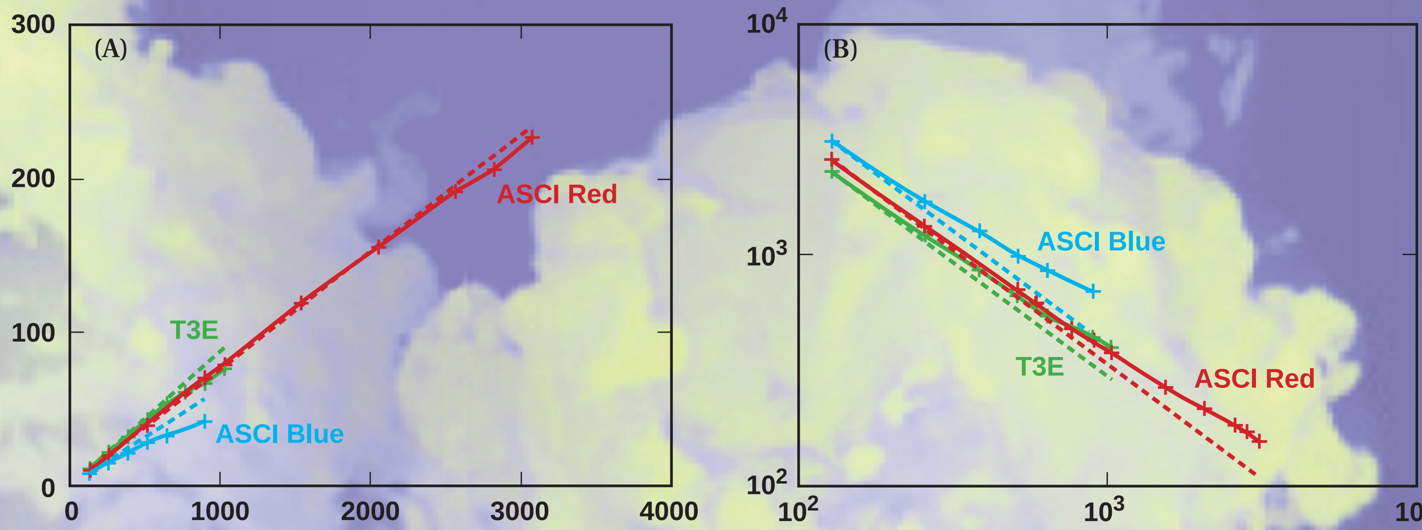

Figures 2A and 2B. A plot showing aggregate Gigaflop/s versus number of processors and a log-log plot showing execution time versus number of processors for the largest case on the three most capable machines used in this simulation are shown above in 2A and 2B, respectively. The dashed lines show the ideal behaviors. Note that although the ASCI Red flop/s rate scales nearly linearly, a higher fraction of the work is redundant at higher parallel granularities, so the execution time does not drop in exact proportion to the increase in flop/s.

Methodology

The winning entry, which simulated the flow of air over an airplane wing, obtained an unprecedented level of performance for an unstructured grid application. The highest performance run was 0.227 Tflop/s (trillion floating operations per second) on an unclassified partition of 3,072 nodes of the Department of Energy's (DOE) ASCI Red supercomputer at Sandia National Laboratory. The research team also ran the same simulations on other high-end computers, including LLNL's ASCI Blue–Pacific. ASCI is the Accelerated Strategic Computing Initiative, a national effort to develop dependable terascale computing for scientific applications.

Demonstration code originated with NASA (FUN3D), while the enabling parallel hardware and (freely available) parallel solver toolkit (PETSc) came from the Department of Energy. Despite these different sources, the research methodology should be broadly applicable, and the hardware limitations exposed should allow programmers and vendors of parallel platforms to focus with greater encouragement on sparse codes with indirect addressing.

A snapshot of ongoing work showed a performance of 15 microseconds per degree of freedom to steady-state convergence of Euler flow on a mesh with 2.8 million vertices using 3,072 dual-processor nodes of ASCI Red, corresponding to a sustained floating-point rate of 0.227 Tflop/s.

High sustained scalable performance has been demonstrated on simulations that use implicit algorithms of choice for unstructured PDEs. In the history of the peak performance Bell Award competition, PDE-based computations led (or were part of leading entries containing multiple applications) in 1988, 1989, 1990, and 1996. The results for all of these leading entries were obtained on vector or SIMD architectures, and all were based on structured meshes. The last (1996) and most impressive of these PDE-based entries was executed on 160 vector nodes of the Japanese Numerical Wind Tunnel (NWT), and ran at 111 Gflop/s. The 227 Gflop/s sustained performance of this unstructured application on a hierarchical distributed memory multiprocessor in the SPMD programming style exceeds that of the 1996 entry by a factor of two.

The achieved flop/s rate is less important to computational engineers than are solutions-per-minute of discrete systems that are general enough to be employed in production design — as PETSc–FUN3D is now employed. In addition, PETSc–FUN3D is a portable message-passing application that runs on a variety of platforms with good efficiency, thus lowering the total cost of achieving high performance over the lifetime of the application.

A multi-year investment by DOE in parallel scientific software enabled the investigators to implement algorithmic ideas in software that can be reused on related applications. Software employed in the *Unstructured Mesh CFD Application* can be reused successfully on related applications such as combustion, radiation transport, atmospheric and ocean modeling, petroleum reservoir modeling, or semiconductor device simulation.

Collaborating toward Success

The research team included investigators representing Old Dominion University (a DOE subcontractor and ASCI academic alliance partner), NASA, Argonne National Laboratory in Illinois, and LLNL. David Keyes, acting director of LLNL's Institute for Scientific Computing Research and chairman of Old Dominion University's Mathematics and Statistics Department, represented LLNL on the research team. Other members included Dinesh Kaushik of Old Dominion, and Kyle Anderson of NASA, as well as William Gropp and Barry Smith from ANL.

Keyes notes that, "Our simulation was not the fastest overall submitted in the SC99 competition, but the three-judge panel created a special category to recognize such strong scalability in the implementation of an implicit unstructured grid solver." In addition, the researchers demonstrated that such high performance could be obtained from a general-purpose library of parallel code modules, on a variety of architectures.

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.