Sailing Past a **Billion**

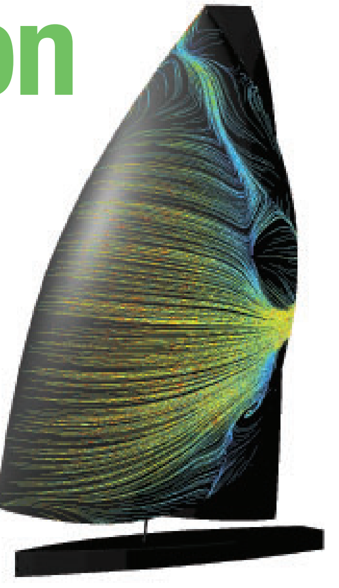

Oil-flow pathlines just above the yacht model surface. The observed tracks, colored by velocity, are painted by the wind and simulate a classic wind tunnel experiment. Converging lines show separating or re-attachment regions.

Racing yacht design researchers push flow simulation past a meshing milestone.

By Ignazio Maria Viola, Yacht Research Unit, University of Auckland, New Zealand,
Raffaele Ponzini, High-Performance Computing Group, CILEA Consortium, Milan, Italy, and
Giuseppe Passoni, Maritime Hydrodynamics Department, Politecnico di Milano, Italy

Over last few decades, the development of techniques in computational fluid dynamics (CFD) together with the increasing performance of hardware and software have helped engineers understand the role of geometrical and mechanical factors on external aerodynamics in ways that were nearly intractable in the past. In recent years, several leading America's Cup sailing teams have become top-shelf users of flow simulation software by pushing the envelope of existing meshing and solver technology. Just a decade ago, experiments on physical models — using wind tunnels and towing tanks — were the main tools for the top teams in their external aerodynamic and hydrodynamic investigations. The option of simulating a number of boat designs in a virtual environment has been shown to have several advantages, including full control of all the parameters involved, repeatability of the measurements, and the ability to simulate non-standard sailing condition scenarios.

In the 2003 America's Cup, in New Zealand, only a few racing syndicates had adopted fluid flow simulation as an effective design tool, though by the 2007 Cup, in Spain, almost all of the 12 competing teams had recognized the value of investing resources in both experimental tests and computational research. Nevertheless, for several technological reasons, there is still a reliability gap between experimental- and simulation-based results. One of these is the extremely complex flow around a racing yacht, particularly in downwind conditions.

To design the sail plan for an International America's Cup Class yacht, a model-scale boat is commonly tested in a wind tunnel. To perform the same test in a virtual environment, all of the turbulent scales of the wind need to be simulated — from the largest, which draw energy from the mean flow, to the smallest, which are associated with the viscous dissipation that extracts energy as heat. It is possible to estimate the overall number of cells required to simulate all of the turbulent scales. This theoretical cell count is directly related to the Reynolds number, which is the ratio of inertial forces to viscous forces, and it is of the order of 10 billion. If such a number

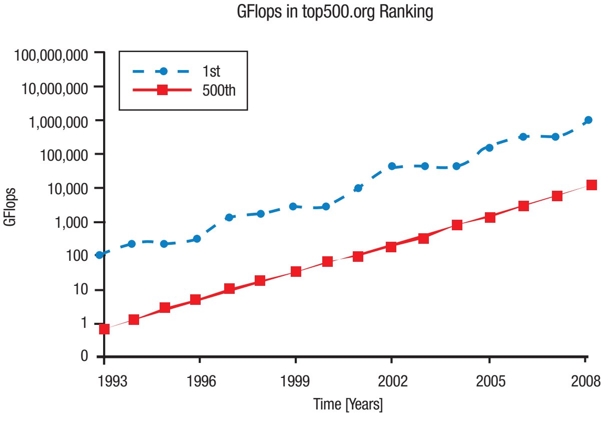

Increasing computational capabilities of the last 15 years, expressed in Gflops (billions of floating point operations per second) and published by the official worldwide ranking top500.org

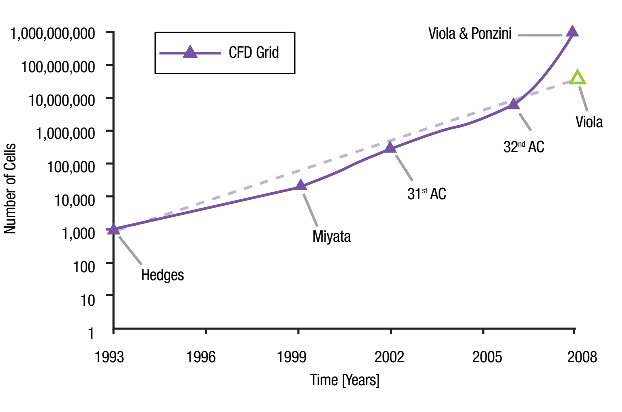

The increasing trend of the number of cells adopted in downwind CFD simulations. Very similar behavior is shown compared to increasing computational capabilities, with the exception of the groundbreaking billion-cell computation.

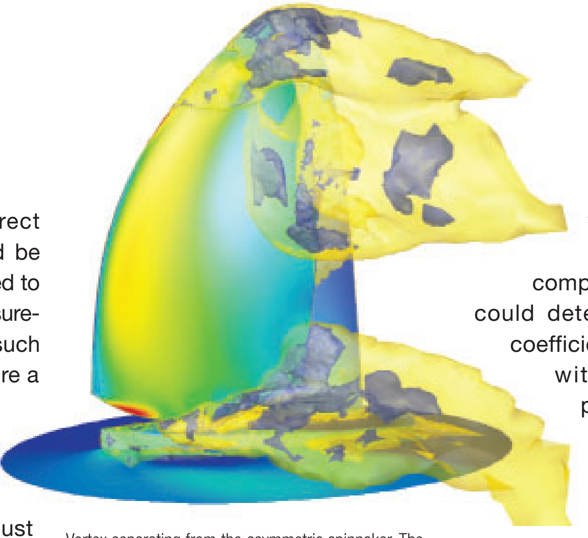

Vortex separating from the asymmetric spinnaker. The yellow region is evidence for separated flow, and the grey regions show the low-speed zones.

of cells were achievable, a direct numerical simulation (DNS) could be achieved, which is widely considered to be as accurate as a full-scale measurement. Unfortunately, generating such a high number of cells would require a huge amount of memory to create the mesh and to then perform the fluid dynamics computation. Prior to the 2007 Cup races, a state-of-the-art mesh had just 10 million cells, meaning that it was necessary to use turbulence models to account for the effects of the smallest turbulent scales on the mean flow.

In 2008, a researcher involved in design for a leading contender in the 2007 Cup [1] collaborated with members of the CILEA inter-university consortium and the Maritime Hydrodynamics Department from Politecnico di Milano to achieve the most realistic simulation of a racing yacht thus attempted [2]. The resulting 1 billion-cell CFD model was therefore two orders of magnitude greater than the previous state-of-the-art mesh size in wind engineering. To achieve such an enormous cell count, the researchers reconstructed the sail shapes along with a simplified model of the hull and rig using GAMBIT and TGrid pre-preprocessors from ANSYS, which resulted in an initial grid of 16 million tetrahedral cells. This grid was then imported into ANSYS FLUENT flow simulation software and partitioned into 512 parallel processes so that it could be run on CILEA's powerful supercomputer, known as Lagrange. Using a hanging-node algorithm, each tetrahedral cell was subdivided into eight smaller cells, which grew the mesh to 128 million cells. By repeating this procedure a second time, the team arrived at a final mesh of just over 1 billion cells.

Running on 512 CPUs, the job occupied 2 terabytes (TB) of RAM for just over a week (170 hours) to complete the calculation of flow velocities and pressures. Performing such a large calculation in parallel was imperative, as the time necessary for a serial process to complete the same computation would be more than 10 years. By calculating the pressure at each cell of the computational mesh, the team could determine the aerodynamic coefficients and compare them with experimental tests performed in the Politecnico di Milano's twisted flow wind tunnel.

Running an ANSYS FLUENT simulation with a billion cells — the first commercial simulation of its kind focused on a single computational model — shows the possibility of performing very accurate CFD modeling in the aerodynamics of downwind sails using leading-edge hardware and software resources. Though this simulation had 100 times more mesh density than other recent CFD computations performed in America's Cup boat design, the mesh was still 10 times coarser than what would be needed to resolve all turbulence scales using DNS and, hence, numerical models, with their inherent assumptions, were still used to simulate the flow. As high-performance computing becomes cheaper and more accessible, however, the DNS goal is in sight. Another important goal in the near future of racing yacht modeling is to couple the CFD computations with a fully 3-D shape optimization procedure. This would overcome the sail designer's requirement to perform a physical wind tunnel test to determine a finite number of trims for different sails beforehand. Until that time arrives, though, the benefits of complementing the global accuracy of wind tunnel testing with the local insight of the flow simulation continue to be affirmed by top America's Cup design teams from around the world. ■

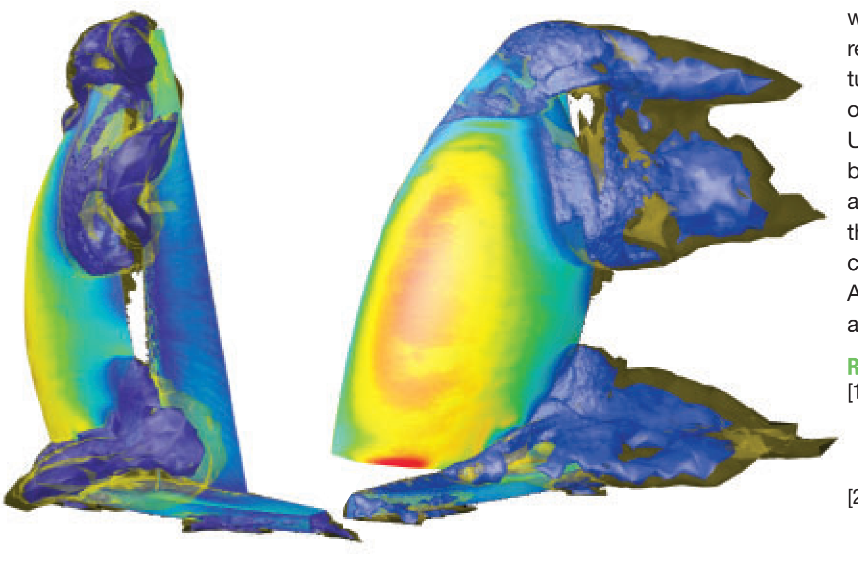

Two leeward (downwind) views, from behind (left) and front (right) showing the air flow velocity over the yacht. The low-velocity regions in blue show vortices from the asymmetric spinnaker and the hull.

References

[1] Viola, I.M., Downwind Sail Aerodynamics: A CFD Investigation with High Grid Resolution, *Journal of Ocean Engineering,* to be published in 2009.

[2] Viola, I.M., Ponzini, R., and Passoni, G., Downwind Sail Aerodynamics: Large Scale Computing vs. Large Scale Wind Tunnel Test, *Journal of Wind Engineering and Industrial Aerodynamics*, submitted in 2009.