
BSC
Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

CRUNCHING THE NUMBERS

In the quest to improve aerodynamic efficiency, SEAT's CFD engineers have turned to one of the most powerful computing facilities in the world

WORDS BY **CHRIS PICKERING**

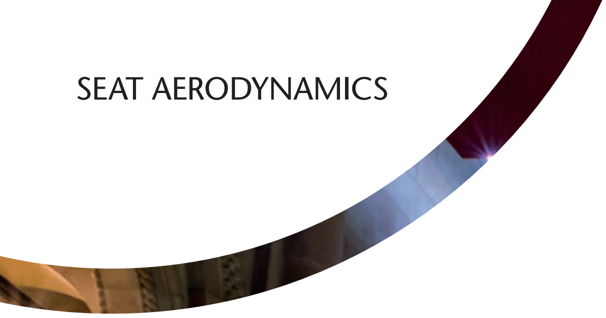

With its atmospheric lighting, ornate columns and Romanesque arches, the interior of the Barcelona Supercomputing Centre (BSC) looks more like an art gallery than a research facility. Yet hidden inside this 19th century former chapel lies one of Europe's most powerful supercomputers – MareNostrum 4. This immense number-crunching machine consists of 165,888 individual processors, which enable it to carry out up to 13.7×10^{15} floating point operations per second (flops) – equivalent to the power of around 40,000 desktop PCs.

There is a logical reason for BSC's choice of location. The old stone building helps to maintain the consistent levels of temperature and humidity required for this vast machine, which sits behind 19 metric tons of glass. We suspect, however, that the visual impact wasn't entirely lost on whoever chose to house the supercomputer here.

Since going online in 2017, MareNostrum 4 has been used to work on everything from genome research to climate-change predictions. It's also used by engineers from SEAT's aerodynamics department to carry out the most processor-intensive modeling on the Spanish firm's latest models.

"External aerodynamics is one of the most computationally expensive problems, because it needs to calculate the whole air volume of a wind tunnel, which means solving fluid dynamic equations for many millions of elements," says SEAT aerodynamics engineer María García Navas.

One particular case that SEAT's aerodynamics department is currently working on involves an investigation into air turbulence in car wheels. The process involves working with a sliding mesh, where the computational elements that are used to simulate the volume of air inside the wheels have to rotate relative to the rest of the model. That means the computer needs to calculate the solution thousands of times for every turn of the wheel and communicate the results to the rest of the mesh each time.

"The tools we have at present only allow us to simulate the rotational speed of the wheels, so they do not allow us to see what happens when the spokes push the air and create turbulence," says García Navas. However, things are rapidly changing. Computational fluid dynamics (CFD) codes are becoming more capable and increasingly processor-efficient. Meanwhile, the computers themselves are

Facts and figures

The supercomputer	The facility
3,456 nodes	180m²
6,912 chips	24°C ambient temperature
165,888 processors	36% relative humidity
13.7 petaflops	19 metric tons of glass
78,000kg in weight	26 metric tons of steel

ABOVE: The supercomputer used by SEAT is the most powerful computer in Spain, and the seventh in Europe

BELOW: Housed inside a deconsecrated chapel, with the power of 40,000 PCs, MareNostrum 4 must be kept at just the right temperature

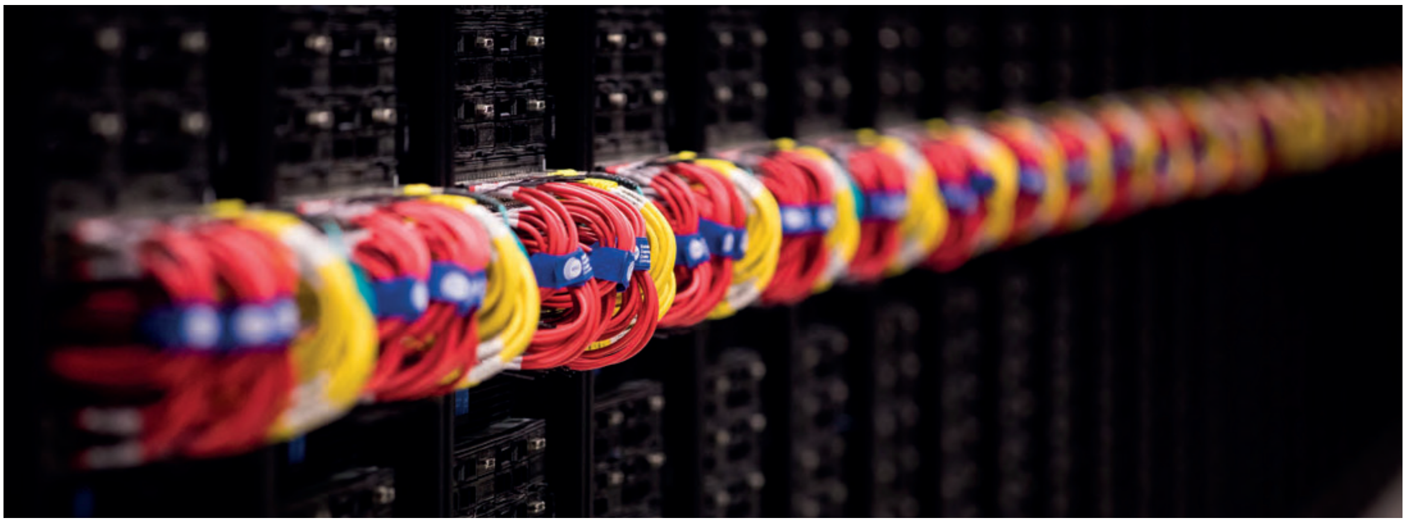

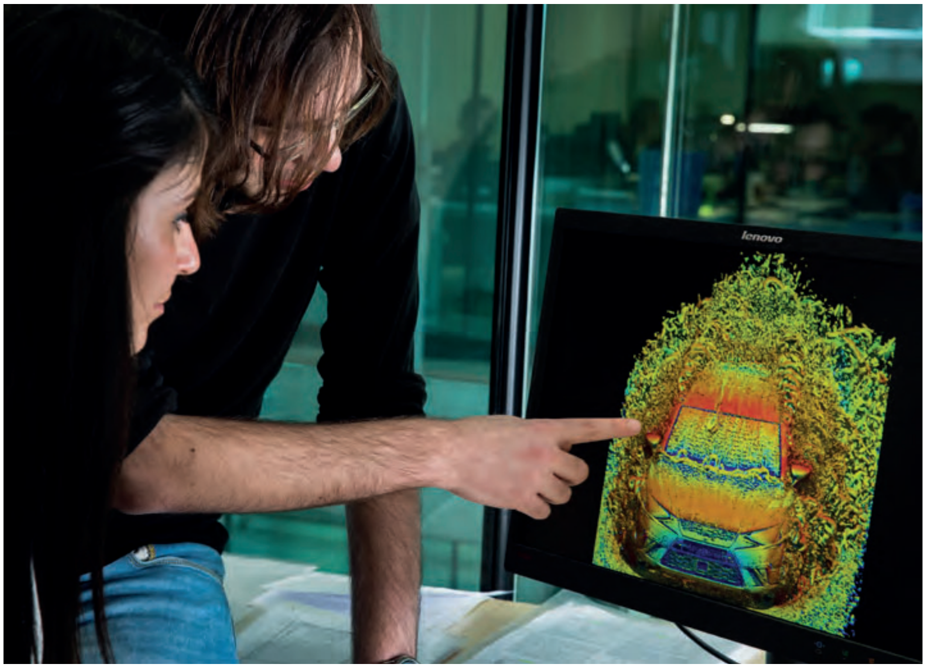

LEFT: Simulation using the supercomputer opens new avenues of research in vehicle aerodynamics

evolving even faster. Work is already underway on MareNostrum 5, which will have around 20 times the capacity of the current system.

"In the future we'd like to be able to simulate everything at once: the airflow around and inside the car, the reactive combustion inside the cylinder, the aftertreatment and even the occupants sitting inside," says BSC researcher Dr Oriol Lehmkuhl. "At the moment that's not possible so we look at each of the problems in isolation and link those simulations using boundary conditions. In the future, with the rise of exascale machines – systems capable of a billion billion calculations per second – we would be able to simulate everything together."

DOING THE MATH

He's not talking about the far future, either. Lehmkuhl says it's possible that we could have computers 1,000 times more powerful in as little as 15 years. By this point, facilities such as BSC are likely to be well into the exascale, where the computer's floating point capability is measured in quintillion (10^{18}) operations per second rather than the petascale (10^{15}), where most currently reside. Already, though, at least one machine is edging into this area, with the US Department of Energy's Oak Ridge National Laboratory's Summit OLCF-4 supercomputer reportedly achieving a 1.8x10^{18} flop calculation while analyzing genomic information in 2018.

This not only means that it will become practical to solve larger, more complex models, but also that different techniques will be available. At present the majority of automotive CFD simulations use a Reynolds-averaged Navier Stokes (RANS) turbulence model. This time-averaged method is dramatically more processor efficient than fully resolving the Navier Stokes equations with direct numerical simulation (DNS). The downside is that it remains an approximation. As processing power has increased, however, there has been growing interest in the use of large eddy simulation (LES). This provides a middle ground between the two – fully resolving the Navier Stokes Equations at larger scales to improve accuracy, while using averaged methods at smaller scales to reduce the computational demands compared with DNS.

"LES resolves most of the Navier Stokes equations but is a transient simulation, so still requires a lot more computing power than RANS," says García

"The tools we have at present only allow us to simulate the rotational speed of the wheels, so they do not allow us to see what happens when the spokes push the air and create turbulence"

María García Navas, aerodynamics engineer, SEAT

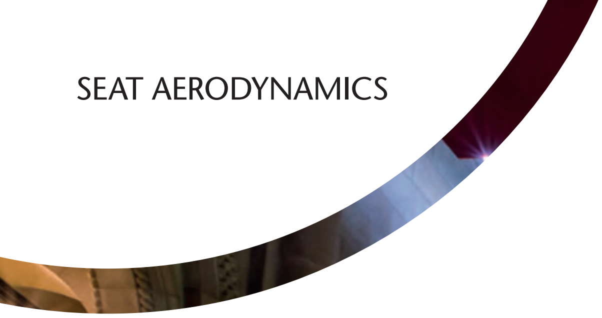

Navas. “It may have to resolve the equations half a million times, compared with RANS, which is steady state and might only have to solve them 10,000 times.”

BETTER THAN THE REAL THING

The practical upshot of this increasing CFD capability is that engineers will be able to generate more accurate results at a much faster rate. A decade ago, SEAT had its own computer cluster with fewer than 100 cores. At that time, a whole-car simulation (consisting of approximately 20 million elements) could take six or seven days to solve using 16 of those cores. Now, the simulations can have more than 100 million elements, run in 960 cores each and take less than a day. What’s more, access to 20,000 cores in the BSC supercomputer allows the aerodynamics engineers to run 5 or 10 simulations in parallel.

These improvements are rapidly shifting the balance between CFD and physical testing. “As CFD simulation becomes more reliable, it will play a larger role, because it’s much less logistically complex [than wind tunnel testing] and takes less time,” says García Navas. “However, the physical tests will not disappear as they return a real result that we can use to correlate the simulation. The idea is to test multiple design possibilities virtually, which we can later corroborate in a real tunnel.”

This comes at a time when aerodynamic development is arguably more important than ever. The move toward electric and hybrid powertrains may eventually provide a solution to tailpipe emissions, but reducing energy consumption will still be paramount in improving battery range and reducing lifetime CO_2 emissions.

“Climate change and CO_2 reduction have given aerodynamics the vital importance it has today,” says García Navas. “Our challenge is to improve the development processes so that the result is as competitive as possible.”

Growing processing power is helping to accelerate this development and all major car manufacturers now have access to high-performance computing. It’s doubtful, however, that any of the others look quite as charismatic as the BSC facility that SEAT uses, hidden away inside a deconsecrated chapel. ❮

Physical testing complements virtual analysis. As computing power increases, the gap between simulation and real-world measurements continues to decrease

Rise of the supercomputer

Given that the average office now has dozens of computers – each millions of times faster than those used to guide the Apollo missions to the moon – you might assume that the number of dedicated supercomputers would be dwindling. In fact, the demand for hugely powerful computing facilities is on the up.

The Argonne National Laboratory, operated by the US Department of Energy, is soon to be home to one of the largest. Aurora, as the new machine is known, will push well into the exascale, with 5 to 10 times the power of the most powerful supercomputers that exist today. It will be open to scientists and engineers all over the world once it goes online in 2021.

“With Aurora’s increased computational power and capabilities, researchers will be able to address complex challenges that are not feasible on even the most powerful supercomputers today,” says Dr Katherine Riley, director of science for the Argonne Leadership Computing Facility.

Automotive applications range from in-cylinder CFD models to nanoscale materials research for new battery chemistries. “Researchers will have access to a system that supports simulation, data analysis and learning capabilities in one environment,” says Riley. “This means they’ll be able to use the massive data sets produced by the simulations and experiments to train machine learning and deep learning models, enabling new paths for transportation-related research. New approaches like this will help scientists overcome the limitations of today’s computational methods, to carry out multiscale, multiphysics simulations that accurately account for complex processes like turbulence and combustion taking place in realistic engine geometry, or to gain a better understanding of materials’ relationships in targeted applications.”