Sun, Spain and supercomputing

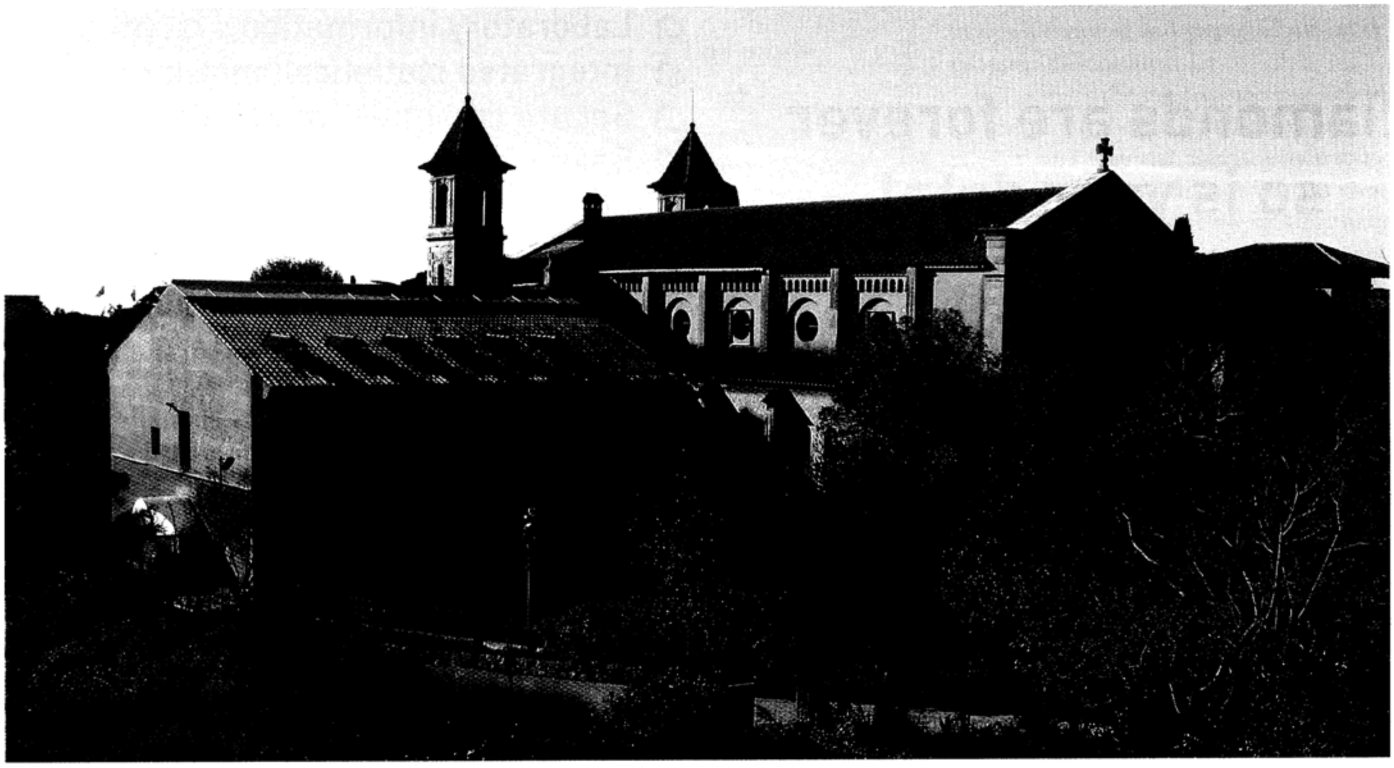

This 19th century chapel houses MareNostrum, Europe's largest supercomputer.

There's more to this 19th century chapel than meets the eye, as David Robson discovers

Built in the mid 19th century, and rebuilt after the Spanish civil war, the Torre Girona is a stunning neo-romantic chapel that was once the focal point for the spiritual life of the Girona-Agrafel family of Barcelona. In as late as 1960, the chapel still served the Roman Catholic Church, acting as the place of worship for the Collegi de Monges de l'Ascenció school of monks.

It may be a sign of our changing priorities that, since 2004, the chapel has served a very different – although equally inspirational – purpose. Torre Girona is now home to MareNostrum, Europe's largest supercomputer and the one-time fifth most powerful computer in the world, where scientists flock (in person and virtually) in a scientific pilgrimage to gain answers to some of the most pressing questions of our age, including the meaning of our genetic code (see panel opposite) and predictions of our future in the face of global warming.

The Barcelona Supercomputing Centre (BSC), which installed, maintains and oversees the use of MareNostrum, is the product of years of collaboration between the Technical University of Catalonia (Universitat Politècnica de Catalunya or UPC), the Spanish Government and two of the biggest names in the computing industry. While the majority of this work has gone into developing powerful resources for scientists and engineers, the latest of the agreements, with Microsoft, has recently given birth to a new research centre that could bring high-performance computing to our desktop PCs and mobile phones.

The first of these collaborations in the early nineties, between the UPC and two R&D support agencies – the Spanish CICYT (Comisión Interministerial de Ciencia y Tecnología) and the Catalan CIRIT (Consell Interdepartamental de Recerca i Innovació Tecnològica) – resulted in the European Centre for Parallelism of Barcelona (CEPBA), a precursor to the BSC that established the centre's strong roots in computing research. The following years saw a flurry of research at the new centre that tested the ground of supercomputing in areas such as signal theory and communications, nuclear physics and structural engineering, in further collaborations with regional research institutes.

Years later, at the start of the new millennium, the centre embarked on a series of collaborations with IBM, which finally resulted in the construction of MareNostrum in 2005 using some of IBM's extensive experience in computing hardware. After a gestation period of 14 years, the Barcelona Supercomputing Centre was finally born as a supercomputing facility with national and international standing. Formally established as the Spanish National Supercomputing Centre, MareNostrum also gained an initial ranking in the top five of the global Top500 list of the world's fastest supercomputers.

Right from the start, the BSC chose to use innovative approaches to supercomputing architecture that would become global trends throughout the HPC community years later. In contrast to many centres from that time, MareNostrum is a commodity cluster of commercially available components from different vendors, allowing high power processing at a reduced cost. It also allowed a more compact architecture – MareNostrum occupied less than half the footprint of the Yokohama Earth Simulator in Japan, which offered roughly the same power at that time.

It is debatable whether the centre could have achieved its global reputation without these well-placed, strategic collaborations. As Mateo Valero, director of the facility, explains: 'It's very important to collaborate with the best. There are two benefits: they provide funding once you have established a formal collaboration, and you can work with researchers within the company to talk about topics that are difficult to solve. It's very convenient.'

However, the BSC has two key strengths, and MareNostrum is just one of these. The oldest string to the BSC's bow is its reputation for conducting cutting-edge research into supercomputing hardware and software, to help scientists make the most of the facilities that are available to them and to direct the installation of future supercomputing centres, and it is possibly this capacity that has provided the big attraction for both IBM and Microsoft.

Andrew Herbert, of Microsoft Research, in Cambridge, explains: 'The BSC has a long-established research facility in multicore processors and hardware. We are connecting more for their research expertise than for MareNostrum itself.'

Some of the BSC's current research activities concern finding ways to exploit the different ways of constructing a cluster from various types of computer processors. In contrast to previous architectures, which would connect many chips with the same processor to spread the load of a computation, researchers are now investigating the possibility of integrating these with 'hardware accelerator' devices to further accelerate the processing. Likely candidates for hardware accelerators include field programmable gate arrays (FPGAs), which allow users to reconfigure the processing hardware to suit a particular problem, and graphics processing units (GPUs).

By harnessing the advantages of each technology, this 'heterogeneous' (as opposed to homogeneous) approach presents many new opportunities for computer engineers, but it does come at a cost: each component is best suited to a particular type of job, and the programming of these computers must be altered accordingly to account for this. For example, reconfiguring an FPGA processor is laborious but it is very quick at processing the same job again and again, so it is more suited to repetitive calculations. 'This approach can speed up algorithms in a better way than multi-purpose processors. It should need less power and less space, and it should be cheaper than alternative methods,' says BSC's Valero. 'However, it needs an additional effort with the software. You need to split the computation into hundreds of thousands of pieces, and you need to distinguish which part of the problem is performed by which processor.'

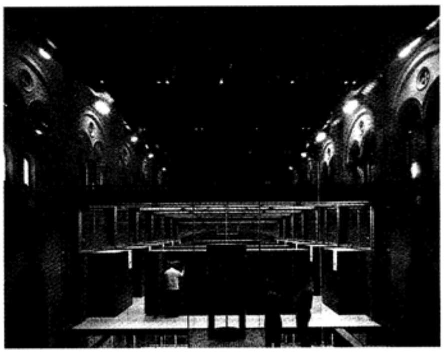

MareNostrum from inside the Torre Girona chapel.

It is inevitable that this research will also feed into the BSC's own supercomputing facilities. The BSC plans to use this approach in the next upgrade to MareNostrum, with a particular interest in the use of Sony's Cell processor, found in the PlayStation 3. The computer engineering research also benefits the other scientific projects that are run on MareNostrum, in areas such as the earth- and life-sciences, because the research groups can share computer scientists with knowledge of how to exploit parallel processing.

Cracking the genetic code

Much has been made of 'cracking' the genetic code, as if all the secrets of the human body are somehow locked within a three billion-long sequence of letters that can be translated once we have found a suitable key. Many scientists hope to find the key through a careful analysis of common patterns that run throughout the sequence of bases found in the human genome to find a 'signature' that could be used to identify further examples.

This has provided some success, but often the sequence of bases varies too much from region to region for this to be of much use. Gene promoters, which determine when and how the code of a gene is expressed into a real biological effect, are one of the most prominent examples. Being able to pinpoint these regions would allow researchers to understand the process and to learn how to control it artificially, but so far the problem has been too complicated to be tackled with much success.

Now, researchers from a collaboration between the Barcelona Supercomputing Centre and the Institute for Research in Biomedicine think they have found a way around this. Using the huge processing powers of MareNostrum, the team faced the Herculean task of modelling the physical properties of the DNA molecule throughout the genome. The team found that DNA molecules all exhibited similar flexibility characteristics at promoter regions, allowing them to identify previously unknown sections for further analysis.

'It's the first time researchers have used the structure to analyse DNA,' says David Torrents, principal researcher of the IRB Barcelona – BSC Joint Programme. 'This is very important – it's the starting point for many projects.'

➤ The new BSC-Microsoft Research Centre, officially founded this January, is likely to extend these computer science research activities in new directions, with more dedicated researchers and a new innovative approach to the way computers are built. Rather than developing the hardware first, and then developing the software to fit the hardware, the centre is looking to develop both simultaneously, 'to get the best interface between the two,' according to Microsoft Research Cambridge's Herbert.

This will include work on transactional memory, which involves finding a technique to establish a shared memory resource without the different processors competing to access the same part of a database at the same time. Rather than facing as a problem that solely concerns the operating system, the BSC-Microsoft Research Centre will investigate how the hardware can be built so the different processors do not interfere with one another.

'It's a fairly new area of research, which we've been talking about from a theoretical perspective for a couple of years,' says Herbert.

'By harnessing the advantages of each technology, this heterogeneous approach presents many opportunities'

As more and more commercial appliances, including laptops and even mobile phones, make use of multicore processing, it is hoped that the research conducted at the new centre will have applications that reach much further than the scientific computing community.

'The supercomputing model is applicable at every level,' says Herbert. The results, he says, could be magnificent: 'When desktop computers have tens or hundreds of processors, they will have a much richer user interface.' He cites the possibility of advanced speech recognition and handwriting recognition as being the normal way of communicating with computers in the future.

However, the extent of the BSC's influence does not end in Catalonia, with the BSC acting as a key hub for both Spanish and European supercomputing. The centre is the administrative headquarters for the Spanish National Supercomputing Network, with facilities in Barcelona, Madrid (currently at number 89 in the Top500), Cantabria, Málaga, Valencia, Zaragoza and the Canary Islands.

It is also part of the DEISA consortium of supercomputing centres that allows users from across Europe to access its facilities, often from the comfort of their own hometown. The centre will also form a key part of PRACE, the Partnership for Advanced Computing in Europe, which plans to build on DEISA with a tighter infrastructure and greater funding for future upgrades, starting from early 2010.

It is this spirit of collaboration that has allowed the BSC to develop into Europe's leading supercomputing facility. The name MareNostrum ('our sea' in Latin – the Roman name for the Mediterranean) reflects this perspective, and with its key placement within PRACE it looks like this attitude is set to continue for some time yet.