

update

more online at www.spectrum.ieee.org



The Church of Microsoft

The software maker tries to get ahead of the move to hundreds of processor cores per chip in a deal with the MareNostrum supercomputer

THE IBM MareNostrum supercomputer sits in a Gothic-style chapel on the outskirts of Barcelona, Spain. It may not be the world's fastest—although it is in the top 20—but it is certainly the world's most beautiful computing machine [see “Solving the Oil Equation,” January]. And if all goes according

to plan, this is where future generations of Microsoft's Windows operating system will be born.

For Microsoft, MareNostrum's more than 10 000 IBM microprocessors and 20 terabytes of memory are the ideal testing ground for the software that will run the kind of multicore and many-core microprocessors that will

hit our desktops in the next few years. Those CPUs are expected to be made up of hundreds of processor cores, so it takes a supercomputer with thousands of processors to simulate them for software development. Which is why Microsoft and the Barcelona Supercomputing Center, which runs the MareNostrum, struck a deal in late January to form a joint research center dedicated to solving the vast array of problems associated with programming for multicore processors.

To make ever more powerful processors, the chip industry once relied on simply shrinking

MODERN GOTHIC:

MareNostrum's 10 000 processors will be key to figuring out how to program processors with hundreds of cores.

PHOTO: ROLAND HALBE/
BARCELONA SUPER-
COMPUTING CENTER

update

a single processor core and ramping up its clock speed. But a few years into the new century, it became clear that this was a dead end: performance was not improving fast enough, while power consumption was accelerating out of control. The solution was to put more than one processor on the same chip and run them both at moderate speeds.

Two- and four-core processors are common now. “We know how to use these,” says Andrew Herbert, managing director of the Microsoft Research Laboratory in Cambridge, England. The question is how to make the best use of the hundreds of cores that will appear on chips in the next 10 years. Microsoft hopes to find out by simulating the problem and various solutions on the MareNostrum.

For decades, computer languages have been conceived and designed with the expectation that a sequence of instructions will be executed essentially one after another. This approach makes sense when a calculation is carried out on a single microprocessor. But when there are 100 processors, how should this sequence be divided up? Answering that question is at the heart of the joint research center’s mission. “There are lots of good ideas out there which we want to explore,” says Herbert.

In some cases, it’s easy to see how the work can be divided, says Tim Harris, a computer scientist

at Microsoft Research who is involved in the MareNostrum collaboration. For example, when rendering a scene from a computer game, the instructions can be easily divided among cores by giving each a portion of the scene to be rendered.

With other tasks, things aren’t so straightforward. One problem is how to give parallel computations access to shared data without them all trying to access the same chunk of information at the same time.

The conventional solution is to lock the memory so that only one computational thread has access to it at a time. But lock-based programming is notoriously hard to do in practice and can cause bottlenecks.

To cope with this problem, one of the ideas Microsoft is testing in Barcelona is transactional memory, which allows free-for-all access to shared memory in the hope that each thread will want different pieces of data. If a conflict arises, the transactions involved are halted and started again. “This is one of the hot topics in parallel computing,” says Harris.

Transactional memory can be built into the hardware. Indeed, at February’s IEEE International Solid State Circuits Conference in San Francisco, Sun Microsystems reported the first server processor utilizing a type of hardware-enabled transactional memory. Sun’s move is the kind of thing Microsoft may want to see more



DUST-DEVIL DYNAMO

RESEARCHERS AT the University of Michigan in Ann Arbor showed that natural electric fields as strong as 160 kilovolts per meter play a role in determining the amount of dust that makes it into the atmosphere, where it can influence weather and climate. Electric charge reduces the amount of wind needed to blow sand around and can directly lift dust off the ground. Scientists first noticed electricity’s influence when examining dust devils in the desert but had to develop a new kind of electric field analyzer to quantify it.

More at <http://www.spectrum.ieee.org/mar08/dynamo>

of. One of the goals of the MareNostrum project is to “explore a top-down approach in which the software requirements determine the hardware architecture rather than the other way round,” says Herbert.

Such an approach could lead to some radical departures in design, says David Patterson, an IEEE Fellow and expert on parallel computing at the University of California, Berkeley (who is not involved in the

collaboration). He suggests using cores with different architectures on the same chip. “It may be that one type of architecture is best for speech recognition and another for image processing,” he says.

At this point, almost any idea can be entertained, and Microsoft will surely try many of them. The next few years are “a rare opportunity to reinvent computing entirely,” says Patterson.

—JUSTIN MULLINS