Who wants to buy a computon?

The Economist | 12 March 2005

Grid computing: Electricity is sold by the kilowatt-hour. Now a researcher has proposed that computing power should be sold by the computon

LAST month, Sun Microsystems, a big computer-maker, announced the details of its plans to rent computers over the internet. Instead of crunching numbers on their own in-house machines, customers of the “Sun Grid” service can pay $1 for every hour that they use a processor on one of Sun's computers, and $1 per month for every gigabyte of storage. This is the sort of thing people have in mind when they talk about grand, but often vague, visions of “utility computing” or “grid computing”, in which computer power is supplied when needed, like electricity, over a network by a central provider. Unveiling the service on February 1st, Jonathan Schwartz, Sun's number two, submitted a computing job to the Sun Grid. In a flash, the answer was ready, at a cost of $12: in a few seconds, 12 processor-hours of work had been carried out by hundreds of processors.

It is all very clever. But offerings such as Sun Grid, while novel, do not solve the ultimate problem: the efficient allocation of networked computing resources. People do not think of their computing needs in terms of, say, 50 processor-hours; instead, they have specific tasks of varying importance and urgency, and want to get those tasks done economically, using whatever resources are available.

The issue therefore comes down to economics as much as technology. As long as the number of computers and users is small—as in a “cluster”, rather than a genuine grid—resource allocation can be done socially, or by an “omniscient” administrator who simply decides who will be allowed to do what. But as soon as the grid becomes big, any such arrangement will fail for the same reason that the Soviet Union's economy broke down. Valuing and prioritising millions, perhaps billions, of different transactions is too complex for any central planner. Only a market mechanism of some sort can maintain order.

Building just such a market mechanism is the nut that Bernardo Huberman, a researcher at Hewlett-Packard (HP), and his team have been trying to crack. The key, Dr Huberman realised, was to have a system that can allow users to assign different priorities to tasks, to reflect their importance. This rules out any system that would simply give each user a priority without differentiating that user's many tasks. It also rules out a reservation-style system of the sort that airlines use, since a lot of processor cycles (like aeroplane seats) would end up unused, and the system would not be able to accommodate new tasks as they arose, even if they were extremely urgent. In a grid, it must be assumed, demand is changing constantly and unpredictably, and so is supply (since individual host computers on the grid come and go).

Mr Huberman's answer is Tycoon, a piece of software for computing grids that turns them into a sort of “stockmarket or clearing house”, he says. Users start by opening a bank account and getting credits. They then open a screen that shows all the available processors, their current workloads, and a price list. Users place bids for various processors, using a sliding price dial that looks like a volume control. Allocation is proportional, so that if one user bids $2 and the other $1, the first gets two-thirds of the resource and the second one-third. If the deadline of one task suddenly moves forward, the user can up his bid and immediately get more processor cycles for that task. As users consume cycles, the software deducts credits from their account.

The HP team has so far tried out Tycoon on a cluster of 22 Linux servers distributed between HP's headquarters in Palo Alto, California, and its offices in Bristol, England. Tycoon did well in these tests, and several amusing animated films were rendered using its system. HP has now given Tycoon to CERN, the world's largest particle physics laboratory and a hotbed of grid-computing research, for more testing.

This is only the beginning, of course. Mr Huberman reckons that Tycoon, in its current form, could run clusters of 500 host computers with perhaps 24 simultaneous users. But the ultimate vision of grid computing is for one gigantic network spanning the globe and accommodating unlimited numbers of users. So a lot still needs to happen.

For a start, the metaphor for computing grids as “utilities”, similar to water or electricity supplies, is misleading, since there is no equivalent of litres or kilowatt-hours. Processor cycles are just one component of computing resources, alongside memory, disk storage and bandwidth. Mr Huberman would like to combine all of these factors into one handy unit, which he wants to call a “computon” (a cross between “computation” and “photon”, the name for a packet of electromagnetic energy). Tycoon's descendants would then help to allocate computons across the grid's global market. Of course, Mr Huberman adds, that will happen in a different decade.