Science, Deep computing and the RS/6000
For an IBM fall launch campaign
Ghosted by Robert Fripp, 2001

I want to use this paper to map out what scientists, researchers and engineers can expect from IBM in the next few years. The news is good and getting better.

Briefly, IBM is rapidly overtaking Silicon Graphics/Cray, vendors who long dominated the science, research and development market. Two factors are pulling us ahead: major investments in technology; and major investments in the people needed to build that technology into premium solutions for each and every customer.

Before I explore these two major investment areas I want to touch on a recent trend. Computer science is seeing real convergence: between computing for the business of science, and computing for the science of business. That convergence has led IBM to coin the term “Deep Computing” to describe challenges which employ four interlocking components: powerful, super-fast hardware; high-performance software; expert domain knowledge; sophisticated algorithms.

Business and science both benefit from this convergence. Business Intelligence (BI) computing gives us a technique called “horizontal” or “total relationship marketing.” This lets a bank, for example, respond to a specific customer's total relationship with it rather than having individual departments--accounts, loans, mortgages--marketing to that customer independently. At IBM we have adopted a similar horizontal approach to a customer's requirements, letting us pinpoint people-resources right across the RS/6000 Division and IBM Research to handle any given problem.

We can summarize this trend at IBM as introducing refined business management skills to the allocation of compute-resources in science, research and engineering.

This leads us in three directions, each with distinct investment needs: the research segment covers government laboratories, educational institutions and pharmaceutical companies. Many of these customers are doing what we continue to call “grand challenges,” numeric-intensive research at the frontiers of science in all sorts of fields, from integrating oceanic and atmospheric circulation models to simulating the “docking” or repulsion of virtual antibodies by virtual viruses; and, of course, the design of new prescription drugs.

The petroleum sector includes oil and gas reservoir mapping based on seismic traces, extraction modeling and resource processing. And hard rock mining, too. The cost of blasting and excavating hard rock in nickel and gold mines is such that computer-assisted optimal resource utilization means the difference between profit and loss.

Our third sector, manufacturing, comprises everything from metal forming to directing numerical cutting machines; crash and safety analysis; automotive and aerospace design and manufacture. Just one example: Boeing built its 777 without blueprints, giving it a place in aviation history as the first "paper-less" civilian airplane. 238 design-build teams created the 777. Many of those designers used CATIA software on almost 4,000 RS/6000 workstations. Boeing estimates savings in error and rework alone fell between 60% and 90%. And that was five years ago! We have advanced a long way since then.

In computer-assisted design and manufacture (CAD/CAM), IBM may already be the high-end market leader. Also in electronic data analysis (EDA), which is used to design computer chips and test electronic circuit simulations.

Right across the board we are finding that science and engineering sectors benefit enormously from the convergence and coupling of computing skills with business-oriented techniques, such as data warehousing and mining. The product of this marriage, the "Deep Computing" I mentioned earlier, is penetrating many markets, testing new applications as yet unknown. Which brings us back to my starting point, those two critical investment sectors: in people; and in technology.

In May, IBM opened an Institute of Deep Computing within IBM Research. Our Director of Mathematical Sciences, William Pulleyblank, has brought together a team of over a hundred top-flight people who understand the implications of science computing: they are helping customers migrate their projects to the RS/6000 SP. Our researchers are also pressing forward in new directions, such as bioinformatics, the new numeric-intensive field in the life sciences. Genome research comes into this heading, gene-mapping, in humans and many other species.

Since a primary focus of genome research is to map the individual characteristics of genes--traits, differences and similarities--you can readily see where such a science can benefit from an infusion of BI techniques; as, for example, discovering subtle differences and similarities of customer preferences in numeric-intensive business functions such as data mining. Both focus on discovering miniature needles in mountain-sized haystacks.

So Bill has people in deep computing working on large-scale business functions. When you bring together BI algorithms with scientific algorithms, really exciting things start to happen. Those people are charting new paths to brand new applications for very large supercomputers, in business and in science. Steering new directions along new paths depends on striving for technological breakthroughs allowing the hardware to keep up with human intellect and imagination. Well, we have those, too. From the first, what has made the RS/6000 platform so excellent for science and engineering is the fact that the system is fully scalable from bottom to top. You can sketch out a project on a laptop and eventually migrate it to the hundreds or thousands of nodes on an SP, if that is the direction the research needs to go.

Our latest offering here is the Nighthawk node. This is going to be IBM's flagship product in the near future. Nighthawk is an 8-way symmetric multiprocessor (SMP) high node for the SP; a node with a very high performance memory switch which lets it scale up to eight processors in its first incarnation. We have plans to take it further. This technology has evolved very rapidly. Just over a year ago we released the PowerPC 604e (Silver) node, unleashing five times the processing power of Deep Blue, the computer which won the chess series against World Champion Garry Kasparov. (We are considering the possibility of a rematch.) Then came the Winterhawk node, and now the Nighthawk. You can see that IBM developers are in an aggressive frame of mind!

In fact we are so aggressive about prospects for the Nighthawk node that we have “cloned” it. Its architecture is brilliant, so good that we have repackaged the Nighthawk into an 8x2-way SMP called the T70, which fits right into our product line as the perfect departmental technical server. You can draft your project on a RISC laptop; develop it on an RS/6000 workstation; do a small-scale run by pulling in extra cpu from the T70 down the hall; then scale up massively to a Nighthawk-based SP. I can already see the manufacturing sector making heavy use of Nighthawk-based machines and the T70, especially in automotive and aerospace where, typically, you have many designers creating components at workstations--just as Boeing designed the 777. Furthermore, these sectors typically run independent software vendor (ISV) applications, all of which are ported to parallel processors, so these are in excellent shape to take strong advantage of big SMP machines.

Just as deep computing is the product of converging antecedents, the Nighthawk node also has a solid pedigree rooted in converging architecture. A wealth of new applications for the manufacturing sector is drawing a lot of excitement because of the Power 3 processor on which the Nighthawk node is based. The Power 3 itself merges the traditional Power 2 processor with IBM's Power PC. So the coming together of those two architectures will accelerate everything--and simplify things, too!

It is important for customers to know that we have a strong, well-directed road map. Much work in the R&D sector is done on customer-developed applications. Customers take comfort in knowing that IBM RS/6000 architecture scales as far as you want to go; and they take more comfort knowing that IBM has a strong product road map. The Nighthawk node and the T70 departmental server are now key mileposts along that road. Our current road map began with the ASCI Project, the recent instal of a 3.9 teraflop computer at Lawrence Livermore National Laboratory based on PowerPC 604e (Silver) node technology. After that system was up we installed our Power 3 processor. From there the road map jumps by multiple steps of Power 3, to be followed in 2001 or 2002 by Power 4, running at over 1 gigahertz.

One final item: IBM's Deep Computing Institute is making major breakthroughs in algorithmic scaling, not just getting more out of each processor, but scaling those applications to higher and higher levels of performance than ever before. The focus here is not to do the same things faster--though that's good--but to evolve the tools to do completely new science and research.

This raises the prospect of creating whole new fields of endeavor, which in turn depends on all of us thinking way past where we have been, truly understanding the potentials which computer simulation offers. When we truly come to grips with whatever the future state of the art holds for us, we will have to boost the tools available to bring it about. For example, we will have to give more attention to data management and storage. And it is more than likely that we will borrow at least some of those tools from BI. But that lies in the future. My immediate mandate is to make quite sure that the research segment of IBM's RS/6000 Division is positioned to make the best of today--as well as tomorrow. I can assure you of this: I firmly believe that we are making the right investments in the right places to place IBM as the world leader in scientific and engineering computing, tomorrow, as well as today.

-30-