David Gelernter's Romance With Linda

John Markoff

January 19, 1992

Last fall, computer scientists at the Lawrence Livermore National Laboratory in California unplugged a Cray supercomputer and replaced it with 14 I.B.M. RS/6000 work stations, each the size of an orange crate, and wired them together.

When the scientists turned on the new network, which cost about $1 million, they discovered that it was just as powerful as the old Cray X/MP, which had cost $20 million several years before. And the scientists believe that for solving some types of scientific problems, the network can easily be made as powerful as the latest-model Cray, the Y/MP C-90, which costs $20 million to $30 million.

Linking together a platoon of small computers to replace a supercomputer, an unthinkable act just a couple of years ago, represents a rapid melding of forces that are transforming the computer industry and the way that businesses get work done with computers.

Each time a network like the one at Lawrence Livermore is turned on, it is sweet vindication for David Gelernter, the 36-year-old computer scientist at Yale University who as a graduate student 13 years ago wrote a landmark computer programming language called Linda that made it possible to link computers together to work on a single problem. He has now emerged as one of the seminal thinkers in the field known as parallel, or distributed, computing.

In parallel computing, a task is divided into many pieces and several computers are put to work simultaneously, each to accomplish one of the mini-tasks and to bring the finished pieces of work together at the end. A host of new start-up companies have emerged to offer a variety of software products to make parallel computing possible.

Parallel computing is gathering momentum as more and more businesses are linking their desktop computers into networks within the office. And there is a growing realization by computer scientists — with Mr. Gelernter leading the way — that each of these networks contains vast pools of computational power that sit largely unused.

The Bid for Power

Most office computer users, whether printing a mailing list or performing calculations on a spreadsheet, are using only a tiny fraction of their desktop computer's total computational powers in the course of a day. A midsized insurance agency, for example, with two dozen work stations hooked together in a local area network, or LAN, would have as much power as a supercomputer if each desktop computer were working in tandem with the others during idle time.

So much idle power suggests the possibility of new companies that contract with networks to broker the idle power, selling it to the highest bidder. Some computer scientists have devised software that would allow one computer on a network to bid for power from others. "The world at large hasn't grasped the possibilities of this yet," Mr. Gelernter said.

"Parallel computing is still in its infancy," said Kenneth Kennedy, a computer scientist at Rice University in Houston and an expert on parallel computing. "But Linda has had a great deal of influence. You can take a handful of work stations in an office and solve real problems."

Some fleet-footed entrepreneurs in the computer industry — as well as some of the computer giants — are now hurrying to establish markets for products based on Mr. Gelernter's ideas and using software that is intellectually, if not directly, descended from Linda.

The International Business Machines Corporation, for one, is now competing head to head with the Cray Supercomputer Corporation for oil industry exploration business. Supercomputers have been used to digest vast amounts of seismic data about underground formations in the search for new pools of oil. I.B.M. is now trying to take away Cray's customers with networks of its RS/6000 work stations, just as happened at Lawrence Livermore.

Weidlinger Associates, a New York engineering firm that often works for oil companies, has used Linda software on a variety of scientific and engineering work stations to speed up the analysis of oil field seismic data.

"Suddenly the question of using free work-station computing cycles is becoming important," said Victor Pereyra, a computer scientist who is a consultant for Weidlinger. "We are showing a perfect speedup." In one set of seismic calculations, Mr. Pereyra said, he reduced the time needed from four days on a single machine to four minutes on a cluster of four work stations linked together.

Corporate research laboratories, including those at Boeing Computer Services, American Telephone & Telegraph and United Technologies, are also experimenting with computer networks and Linda software.

United Technologies researchers used a network of eight work stations to run a complex program that simulates air turbulence created by rotating helicopter rotor blades. They found that the blade program, named Freewake, ran seven times faster on the network than if it were run on a single work station. Secrecy on Wall Street

Wall Street brokerage houses have also begun to use networks of work stations and programs like Linda, according to computer scientists with expertise in distributed computing who have worked as consultants on Wall Street. The work is being carried out with great secrecy by the firms, which seek a competitive advantage in stock and bond trading.

A computer expert at one large Wall Street firm, who described his company's work in the field only on the condition that he or his company not be identified, said the firm is using a network of work stations to run its daily trading activity. But Linda is also being used to harness idle time on the work stations throughout the day for work on large computational projects involving mathematical models of financial market behavior. After the markets close and the traders go home, Linda takes over all the power of the work stations to perform more of the market simulation calculations through the night, doing work that otherwise might require a supercomputer. "We call it MIPS pooling," the computer expert said. MIPS stands for millions of instructions per second, a standard measure of computer power.

The Wall Street computer expert said that his firm had gained the power of a supercomputer while also giving traders a powerful tool to use at their desks when the

markets are open. And it had all been accomplished at a substantial savings over the cost of buying a supercomputer.

Until now, Linda software has been available for high-end computers like work stations. But a wave of Linda-based programs is being introduced that will link less expensive desktop computers together and combine their power so they can jointly tackle problems.

At an exhibition in San Francisco last week for products that run on Apple Computer's Macintosh, Ray Dream Inc. of Mountain View, Calif., and Specular International of Amherst, Mass., introduced software programs that chain Macintoshes together in networks to more quickly render high-resolution synthetic photographs. Computer graphics programs can generate scenes so realistic they look like a photograph. But doing so requires vast amounts of computational power.

Ray Dream executives said that Apple plans soon to test its program on a network of 300 Macintoshes at its research laboratory in Cupertino.

Specular's program, called Backburner, allows the user to draw a realistic image in three dimensions. It links an unlimited number of Mac intoshes together and the speed at which an image can be generated increases almost directly in proportion to the number of Macintoshes in use, Specular executives said.

The foundation for much of this activity has been a handful of companies that have used Linda, which is now in the public domain, to create parallel computing applications. Linda's Variations

The largest is Scientific Computing Associates in New Haven, which was formed in 1980 by several Yale computer scientists. (Mr. Gelernter sometimes works as a consultant to S.C.A., but has no equity in the firm.) S.C.A. has taken the lead in adapting the original Linda programming language and selling it to private corporations and Government agencies for use in parallel processing. The company is privately owned and does not disclose its sales figures. S.C.A. owns the Linda trademark and competitors must pay a licensing fee if they want to use the name.

Three other small, privately held companies are now selling their own versions of the Linda language. They are Torque Systems Inc. of Palo Alto, Calif.; Cogent Inc. in Beaverton, Ore., and VXM in Boston. Torque calls its language Tuplex, although it is based on Linda, and VXM markets its language under the name Balans.

Two hardware makers, Hewlett-Packard of Palo Alto and Intel Scientific of Beaverton, a subsidiary of the Intel Corporation, sell versions of Linda that will run on their computers. In the future, NCUBE, another Beaverton company, and Meiko Scientific, a British computer maker, plan to distribute Linda with their machines.

Though no one has a good estimate of how much it is spending, the Federal Government is also investing research money in distributed processing network software.

At the National Center for Supercomputing Applications in Urbana, Ill., scientific calculations are automatically distributed over a network of work stations so that idle time can be used efficiently for the center's work. The National Science Foundation is also helping finance a Southern California fiber optic link between supercomputers

located at the California Institute of Technology in Pasadena, the San Diego Supercomputer Center and the University of California at Los Angeles.

Networks of work stations set up by national laboratories or Wall Street brokerage houses are but the tip of the iceberg, Mr. Gelernter believes. Much larger computer networks already exist and, like any network, at any particular time they have vast amounts of computational power sitting idle. A Global Network

There are computer networks for airline reservations, automated bank tellers, credit card verification and corporate electronic mail. There are telephone networks, news and information networks, air traffic control networks, supercomputer networks and private networks operated by large corporations. Indeed, what Mr. Gelernter calls a "computational membrane" is quickly spreading around the globe, linking together a world in which every type of information will exist in digital form and can move at high speeds.

Several years ago, Lawrence Livermore scientists studied their work stations and found that on average only 2 to 3 percent of their power is actually used. Mr. Gelernter and his colleagues see this as a great waste of vast amounts of computer power, somewhat like all the potential automotive power wasted by commuter cars driven a short distance to and from work each day and left in the office parking lot at work or in the garage at home the rest of the time.

Mr. Gelernter visualizes all these computer networks linked together — along with all the desktop computers that are not now linked to anything. When that happens, his piece de resistance will go to work: a software program that constantly goes from computer to computer and from network to network seeking out idle computer power and putting it to work.

"The general issue is that huge clouds of excess idle computer power hover over even small office networks," Mr. Gelernter said. "The real challenge is to find out what to do with it." He envisions computers at the center of "mirror worlds," a concept he discussed in a book by the same name, published last fall by Oxford University Press. Mr. Gelernter imagines that virtually all information would be captured in digital form and that it would be available instantly.

Like others, he sees a new direction for the computer industry in this. "The amount of wasted power is enormous," he said. "It's that waste which is parallel computing's foot in the door to the commercial establishment. We are able to make dollars-and-cents arguments to the people who pay for the equipment."

The early steps to create such a program — the successor to Linda — are now underway at Yale. Programmers have named it "Piranha Linda," from a remark by Gordon Moore, vice chairman of Intel, who once described parallel computing as "piranhas chewing on a carcass." The program is designed to automatically move a task around a computer network to wherever there is free computer time.

Piranha Linda has been running in the Yale computer labs and is expected to be available commercially from S.C.A. sometime this year. S.C.A. pays Yale a small licensing fee.

Personal computer makers have also been moving toward Mr. Gelernter's ideas. Last spring, several computer and software makers, including Compaq Computer, Digital Equipment and Microsoft said they would work jointly on the design of a new type of personal computer.

Instead of sitting by itself on the top of a desk, the new computer, the companies said, would be connected into an electronic web with other computers and have far more muscular software that would provide the individual user with power similar to that available in the past only to computer users connected to mainframe computers or supercomputers.

Last summer, I.B.M. and Apple announced that they would work together on a similar project to build a personal computer capable of distributed computing. At Issue: Free Choice

The computer makers' plans and Mr. Gelernter's vision of a world wired together into one giant computer is considered a potential nightmare by people who worry about computer privacy.

What is to keep Piranha Linda or its descendants from being subverted by someone who wants to tamper with another computer or steal information? And what if an individual doesn't want to share a computer? Indeed, a generation of computer users embraced desktop computers in the 1980's precisely because they were suddenly freed from sharing a single mainframe computer with hundreds or thousands of others.

Mr. Gelernter and his graduate students have already encountered these concerns in their own lab. They have had trouble persuading personal computer and work station users to share their idle machine time.

David Kaminsky, a Yale graduate student working on Piranha Linda, hands out chocolates to fellow students as an inducement to free their work stations when he needs their computational power.

Said Nick Carriero, a computer scientist at Yale: "People say, 'I bought this computer. Who says you can have my cycles?' "

Privacy experts say the issue is a broader one: being able to choose whether to participate at all.

"The critical test for any technology is whether it leaves you the ability to retreat into a private sphere," said Marc Rotenberg, Washington director for the Computer Professionals for Social Responsibility. "If you can't turn the system off, you're trapped."

But trends already taking hold in the computer industry are likely to sweep aside or simply overrun such concerns. The growth of networks is expected to continue as more and more corporate data processing executives turn to Mr. Gelernter's ideas about parallel computing.

In the scramble between American and Japanese manufacturers to dominate supercomputers, designers are turning increasingly to multiple processors, in some cases thousands of processors, hooked together inside one machine. A crucial goal: Creation of a machine that can perform a trillion calculations a second, a measure known as a teraflop.

"Teraflop machines aren't years off into the future," said Leigh D. Cagan, a vice president at S.C.A. "They're already here, spread out across departments and sometimes whole buildings, waiting to be discovered and begging to be used." THE PARALLEL LOVELACES

As generations of college professors can attest, sometimes there is no accounting for graduate student humor. And that explains how Linda, the parallel processing language, got its name.

When David Gelernter began his doctoral studies at the State University of New York in Stony Brook in 1978, the major parallel processing programming language was named Ada. It had been developed by the Defense Department, and Mr. Gelernter considered it inelegant and bulky.

He developed his new language around a set of six simple additions that could be made to existing computer languages — Fortran, for example — so that software authors would not have to start from scratch learning a new language if they wanted to use parallel processing.

But Mr. Gelernter's new language might have been little more than a graduate student thesis if he had not met Nick Carriero, who arrived at Stony Brook to begin his graduate studies just as Mr. Gelernter was finishing up. Both men became captivated with the potential of parallel computing and Mr. Carriero followed his mentor to Yale University and became Mr. Gelernter's first graduate student.

When it came time to name the language, Mr. Gelernter said he noted that Ada was named after Ada Augusta Lovelace, the daughter of Lord Byron, the English poet. Miss Lovelace is regarded as the first computer programmer because she worked for the computer pioneer Charles Babbage.

Another woman named Lovelace was in the news when Mr. Gelernter was casting about for a name — Linda Lovelace, a star of pornographic films. So he named the language Linda and it stuck. Asked about it now, Mr. Gelernter grins and shrugs. "I was a graduate student at the time," he said.

John Markoff
David Gelernter's Romance With Linda
January 19, 1992

The New York Times, Jan. 19, 1992, Section 3, Page 1 of the National edition.
<www.nytimes.com>