

High Performance Computing

—And Communications Week—

KING COMMUNICATIONS GROUP, INC. • 627 National Press Building, Washington, D.C. 20045 • (202) 638-4260 • Fax: (202) 662-9744

Thursday, July 8, 1993

Volume 2, Number 27

Gordon Bell Makes His Case: Get The Feds Out Of Computer Architecture


BY RICHARD McCORMACK

It's an issue that has been simmering, then smoldering and occasionally flaring up: will the big massively parallel machines costing tens of millions of dollars prove themselves worthy of their promise? Or will these machines, developed with millions of dollars from the taxpayer, be an embarrassing bust?

It's a debate that occurs daily—even with spouses in bed at night—but not much outside of the high-performance computing industry's small borders. Literally thousands of people are engaged in trying to make massive parallelism a viable technology. But there are still few objective observers, very little data, and not enough experience with the big machines to prove—or disprove—their true worth.

Interestingly, though, one of the biggest names in computing has made up his mind. The MPPs are awful, and the companies selling them, notably Intel and Thinking Machines, but others as well, are bound to fail, says Gordon Bell, whose name is attached to the most prestigious award in parallel computing.

These big, massively parallel computers "are crap," he says. Worse, "the scientific community is basically



Bell: "You've got huge forces telling you who's mainline and how you build computers."

afraid to talk about [the situation] because they know they've conned the world and they have to keep lying to support" their assertions that the technology needs government support, says the ever-quotable Bell. "It's really bad when it turns the scientists into a bunch of liars."

Bell, 58, who sponsors two, \$1,000 awards each year for achievement in parallel computing, is a recipient of

(Continued on page 3)

IEEE Picks Gordon Bell Prize Finalists

The finalists for the 1993 Gordon Bell Prize for significant achievement in parallel processing were named last week, and by all indications, the award is growing in popularity. There were 23 entries in this year's competition, nearly double that of any previous year.

Moreover, the performance reported by the top entry was more than eight times that of the previous best, reports the IEEE Computer Society, which administers the award. Two entries exceeded 60 billion floating-point operations per second. "More significantly," says the IEEE, "12 entries achieved sustained rates exceeding 25 percent of the theoretical peak performance of the machines they ran on. Price/performance improved by almost 600 per-

cent, to 7.5 GF/s/million dollars."

"I want to congratulate them on reaching 60 gigaflops," says Gordon Bell on Tuesday afternoon. "That's about 50 percent of the peak on the [CM-5], so that's great...The users are finally getting something out of the computer."

Those making the award say that the number and quality of the submissions, as well as the increased performance of the machines, are "indicative of the rapid progress in the application of parallel processing to scientific and engineering problems."

The 1993 finalists are:

- P.S. Lomdahl, P. Tamayo, N. Gronbech-Jensen and D.M. Beazley of Los Alamos National

Laboratory for "45 Gigaflops Molecular Dynamics on the Connection Machine 5;"

- L. Long and M. Kamon of Pennsylvania State University and D. Dahl, M. Bromley, R. Lordi, J. Myczkowski and R. Shapiro of Thinking Machines Corp. for "A Deterministic Parallel Algorithm to Solve a Model Boltzmann Equation;"

- Robert Means, Bret Wallach and David Busby of HNC Inc. and Robert Lengel of Tracor Applied Sciences for "Bispectrum Signal Processing on HNC's SIMD Numerical Array Processor;"

- Gary Sabot, Skef Wholey, Jonas Berlin and

Paul Oppenheimer of Thinking Machines Corp. for "Parallel Execution of a Fortran 77 Weather Prediction Model;"

- And Cray Research for the numerous entries that utilized the company's C90 supercomputer. "These entries were too close in overall quality for the judges to pick the most outstanding one, so the fifth finalist is Cray Research," says the IEEE. It is the first time that a company has been honored in the Gordon Bell Prize competition.

For more information, contact the IEEE's Marilyn Potes at 714-821-8380.

**Start Receiving Your Own Copy
Of HPCC Week
Call Dawn Miles at 202-662-9724**

Gordon Bell Makes His Case... (From front)

the 1991 National Medal of Technology, awarded to him by President Bush. "Bell has displayed more vision than any other computer designer in the field, in the area of understanding how a particular architecture, as yet not created, could address an important set of requirements of a broad community," says the U.S. Department of Commerce in its Medal of Technology citation.

Bell sat down with *HPCC Week* editor Richard McCormack at his home recently in Los Altos, Calif., and spoke about where he thinks the industry is headed. While he is not very bullish on the prospects of those companies making message-passing multi-computers, he does think there will be an explosion in clustered workstations and shared memory computers. He is particularly fond of Kendall Square Research, a company for which he works occasionally as a consultant.

No doubt, there will be plenty of people who may wince when they read Bell's remarks on the inside pages of this publication (including Bell). But he says in a recent paper he has written on government support of computer companies that "my attack on State Computers is unfortunately taken as an attack on people and their organizations. It shouldn't be, because the designers are all bright enough and given enough time and money will eventually learn to design


good computers." He goes on to say, however, that it is "not fair to society to have to fund their learning...as a taxpayer with a growing tax burden, I don't want to continue this funding."

It is people like Bell who keep the computer industry fresh and alive. Larry Smarr, director of the National Center for Supercomputing Applications, describes Bell, his former boss at the National Science Foundation as "incredibly conflicted." But the computer industry itself is "incredibly conflicted." And thank goodness; without conflict or chaos there is no creativity or progress. It is what keeps the industry so healthy and vibrant.

HPCC Week last week shared Gordon Bell's thoughts with the executives of Thinking Machines and Intel and asked for their responses, which are published in full. We also interviewed the users of the big new CM-5 at Los Alamos National Laboratory and the Paragon at Oak Ridge National Laboratory, and will present their stories separately in next week's issue. It should be noted now, however, that both users insist their machines are experimental, and that it is still too early to pass judgment on them.

We invite any of our readers to respond to Bell's comments. Call Richard McCormack at 202-662-9721. Bell can be reached at 415-949-2735.

Question: You've become an outspoken critic of DARPA's computing program. Yet there are a number of outstanding DARPA-funded success stories: SUN, MIPS and Silicon Graphics, for instance.



Bell: Those companies are the product of the way to do it. But it isn't 10 years ago, because SUN is 11 years old now. There are some AI companies started in the last decade around Carnegie [Mellon University]. But if I were DARPA I would worry about getting together some good recent success stories.

Q: Thinking Machines is one.

Bell: That's a massive failure because here you've got \$350 to \$500 million that has been poured into that all together, from DARPA, investors and from government purchases. If you look at the amount of work output and look at whether any taxes have been returned other than payroll taxes—hey I want these systems to pay back financially—then Thinking Machines is easy. I just look at the financial return of that and I'd say it was an absolute disaster.

It's a little over 10 years old and they haven't paid back. My guess is that Intel is way in the hole on its venture. And the tragedy about Intel is they will never be successful. Never.

I like to say there are two campuses of DARPA U.; one is in Portland, Oregon, of a bunch of bright circuits guys who had never programmed a computer in their life, and the guys in Cambridge, [Mass.], who came out of [MIT] who never made a computer that anybody had ever used.

Fine. These are all bright people and given enough money, bright people can be educated or produce something useful. These people are kindergarten architects, for which we pay a large amount to fund the bastards.

I'm not against specialized computers. But when you're putting a lot of money into it and a lot of people are programming and reprogramming it, then you have to think if any of these have any generality to it.

Q: So what happens to them?

Bell: What happens to basically all the multicompiler guys is they're all going to get wiped out by ATM and plain old computers because if you strip the machines down to their bare nothing, sitting there is basically a simple computer and a high-speed switch. Guess what? ATM is a high-speed switch.

The only thing that is going to save them is if ATM doesn't come in very fast or isn't as good. But over time, plain old switches and plain old computers connected together just eliminate the special machines.

Q: Even with the problems of latency and bandwidth?

Bell: There are ATM switches on the market that have the bandwidth and latency of the switches that people are using now. They have trouble scaling them to 1,000. But if there were reason economically to scale them to 1,000, why then you could do it. But my thesis is that none of these [MPPs] can operate effectively at the 1,000-level anyway, so 30 is as much as you can get.

(Continued on next page)

Bell Interview... (From three)

The thing about me is that I'm massively optimistic about everything. People think I'm Dr. Gloom, but parallelism is really, really hard. And getting effective use at these large numbers with these long latencies of these switches has turned out to be very hard. My feeling is the machines are not generation scalable, so all the software has to be retuned.

Q: Three of the five finalists last year of the Gordon Bell prize, which is named after you, were working on what was the biggest parallel machine in existence at the time, the Touchstone Delta. How do you explain?

Bell: In a sense it's tragic that either the C90 or the NEC didn't enter because they would have won. There are about three prizes there. The original one was given for parallelism and I think we accomplished in what it set out to do which is to stimulate people in parallelism. We now understand that if you make the problem big enough, and it will still fit in memory, then anything will work. Even a bunch of workstations can be run in parallel if you can put enough memory on every one. Now you may not have enough time to ever get the problem solved because as the problem gets bigger the time goes up cubically.

Q: But the people who run the Touchstone Delta say that their winning your prize is an indication of the success of the machine.

Bell: The one regret I have with the prize is that it may have focused too much on winning the prize. I don't think there are a lot of people spending a lot of time trying to win it—science ought to be their main

Bell's Bio

Gordon Bell, now an independent consultant based in Northern California, received his B.S. in electrical engineering in 1956 from MIT and his M.S. in electrical engineering from the same institution one year later.

From 1960 through 1966, he was the manager in charge of computer design at Digital Equipment Corp., and later became vice president of engineering at DEC from 1972 through 1983. He worked as a professor of electrical engineering and computer science at Carnegie-Mellon University from 1966 through 1972.

After leaving DEC in 1983, Bell became the chief technical officer at Encore Computer Corp., a position he held until 1986, when he became the assistant director of the Directorate for Computer and Information Science and Engineering at the National Science Foundation. In 1987 he became vice president of research and development at Ardent Computer Corp. in Sunnyvale, Calif.

He is best known for his role in creating the minicomputer industry, having lead the engineering teams that designed the PDP 4-8, which formed the basis of DEC's growth from \$25 million in 1966 to more than \$1 billion in 1977. He was also instrumental in the development of Digital's widely successful VAX line of minicomputers.

goal. But I've got a rule for machines: people are still smarter than machines, so given any collection of megaflops, somebody can write a program that will get the megaflops out of it; that's all it takes.

I think you have to look at the log at Caltech, and I looked at the log, and it didn't look very good. I claimed that four processors on a C90 will get vastly more megaflops per day done and handle a lot more users and not give the users severe pain in reprogramming.

Q: Why then are so many people buying the parallel machines over the C90 for instance?

Bell: I won't say a lot of people are buying them. We don't have a huge flourishing market.

Q: But they're taking business away from Cray.

Bell: First off, I think Cray has incredibly bad marketing and that they have not been aggressive about doing benchmarking and looking at real work. I think there have been a lot of cases of scientific user community saying, "Oh, we've got to get into this. We know it's painful and we may not get any output for a long time." They're getting paid to go work on these things rather than being paid to do science.

Q: So what happens to the idea of massively parallelism?

Bell: It think it will be a continually slow, slug-it-out kind of race. The only solution is to put a lot of relatively small parallel machines—30 to 100 processors—into the university community and not in the computer science community. That is the wrong place for these machines because computer scientists don't really care about parallelism other than parallelism translates to research proposals or research dollars. The more money we have to put into computer science, the worse things are going to get because that is generating programs for little tiny computers. Rather you give it to users who have problems that can be solved in a unique way.

The other thing that's hard about parallelism is that too many good alternatives get in the way. I built multiprocessors for 30 years and every one of them has been a disappointment because as an engineer you say I want to get a simple system and replicate that like crazy and that's how I'm going to get power. But every time you build one, it takes a little longer to build so that it works that way and the guy with just the single processor comes up and kills you.

Well I thought this time we had them. There was absolutely no way you can do it with microprocessors. But the same thing is continuing to happen. The enemy of parallelism is fast serial processing. And people who are off telling you to reprogram are just dreaming. There is no way to retrain a paradigm of programming, or there is no systematic attempt to reprogram everybody in how to program in explicit parallelism, a la, message passing.

So you really have to see some problem and then have the time and do the work to do it that way. It's so much easier to go sequential with processors getting so fast.

If MIPs Silicon Graphics produce that 350-megaflop single processor chip, then that's going to have a tremendous impact on technical computing and it will go

(Continued on next page)

Bell Interview... *(From four)*

right to the heart of the parallel guys. Why bust your ass on this parallelism stuff when you can just get your job done?

I just see a slow steady progress and I don't see any magic at all in it, and the only thing people building computers can do is to continue to fight to keep the latency down so they don't lose ground in parallelism.

I think there is a chance the next generation of parallel machines will lose ground. They may go backwards because the processors will have gotten faster. But unless they've got the networks going faster they're going to be right back having to retune everything and finding out they're not going to work well and they'll drop back.

Q: So you don't see the industry doubling or tripling in size?

Bell: It's not clear that the massively parallel guys are going to take that much market share from the supercomputer guys. And right now I would certainly conjecture that if you've got a large fixed amount of money—\$10 to \$30 million—to spend that you're going to get about the same amount of work with a supercomputer that is eight times slower peak as with the massively parallel machine.

If you compare the PAP—peak announced performance—to the actual average power delivered, I'd say you'd be lucky to get anywhere near as many flops per month out of a machine that is almost an order of magnitude faster at least on peak. It just is really coming out that way.

The problem is this: if it had been a natural market, a lot of different things would have happened: a lot of people wouldn't be building crazy compilers; we wouldn't be off trying to generate a new dialect of Fortran; the companies wouldn't have been oversized relative to what they had to do; there would probably be fewer companies and therefore more mass in the companies that are there; and the companies would be better managed and under better control.

So I think the government has significantly impeded the progress in parallel processing: just clear as hell.

Q: The people funding these programs would very much argue with you on that point, that they impeded the market.

Bell: I'm sure they would. But it's an argument. You can't ever go back and run the experiment. I'll say that the computing industry grew totally by the other method and that it is generally a healthy industry and I'm sure that if any of these other environments had been under government control, they wouldn't have grown that way.

Q: But a lot of people have this vision that parallelism is the way to go and there's a large ship pointed in that direction. Has the ship changed course?

Bell: That ship is launched. Parallelism is here and in whatever form it's going to continue. It's a long, slow journey, it's not an end point and you have to fight for it in everything you do.

Q: So where should the emphasis be?

Bell: I'm for providing a lot of machines to the people who have the problems and can justify spending

the effort to get the power out of those.

Q: And the problems may only require an SGI Indigo, for instance?

Bell: Exactly, and hey, if it dies, it dies. But the important thing is to go back and refocus on the problems themselves and then let the guys who want to go build the machines compete feverishly for those bucks. But don't force people into these unnatural solutions by saying you have to use this computer to do it.

Q: Stephen Squires at ARPA argues that the agency has never put anybody in that position.

Bell: They've just made money very attractive to do things and so when a lab is looking for a mission and you're going to throw out a few hundred million dollars a year, then they're going to be attracted to it. They'll work on anything—gallium arsenide or HiPPI or fiber optics or whatever—so long as the money is there.

I know some guys who are working in a parallel environment with plain, ordinary workstations and the neat thing is anybody can be doing that kind of work.

I just see the large laboratories as the wrong way to go. Los Alamos doesn't have any problems. All they've got is buildings, buildings and a lot of people. They have a miserable network up there. That's why I would put it down basically into the NSF-style program, and not at the [supercomputer] centers either. The centers have reached the plateau; the four centers seemed to have plateaued at 5,000 users a year. I wouldn't perturb the centers program, but I wouldn't put big machines there. I'd far rather invest in eight 128-node machines than in one 1,000-node machine. You've got eight times the parallelism—in the sense of graduate students, undergraduates and people solving problems.

One of the things I heard that is going on at the NSF Centers is that Larry Smarr is giving up his Cray and is going to run his workload on Thinking Machines. My feeling is if he gives up his Cray he's not going to be able to handle much of a workload. That would be the sure way to kill the centers because your ability to do work drops very rapidly.

Q: But only the large 1,000-node systems are going to solve the grand challenge problems that require that amount of power.

Bell: But if you look at how often they're using those in a 1,000-node configuration, I'd argue that they're probably used very little. The 500-node machine at Caltech wasn't used that way. Occasionally they had somebody using it that way, but mostly it was partitioned into four or five different subsystems

(Continued on page 6)

Bell Interview... (From five)

running at a maximum of 100 processors. So I'd say these machines are already being partitioned and are being run in these smaller partitions.

Q: ARPA's program managers say that virtually all of the machines they've funded are small processor machines.

Bell: That's great.

Q: Squires says the program changed a long time ago; that he's not funding Intel and Thinking Machines exclusively now; and that he is funding the researchers who want to use the machines for science. They just awarded Maspar a \$20 million contract.

Bell. They didn't give them a contract; they opened a purchase rack. That's great because it speaks to a crazy government where people can't buy what they want and you've got to have some purchasing agency to allow them to buy. That's nonsense.

Q: It's a way to get around the procurement rules.

Bell. The tragedy was the perception that you had to buy it through an ARPA-kind of thing or that anybody was favored over anybody else. That Maspar contract is a testimony to idiocy. If they're going to be the purchasing agency—the GSA of parallel computing—then put everybody on the list.

Q: If you were to go out and buy any of the big computers, which would you buy?

Bell: I am a multiprocessor guy, that's my bias and I wouldn't fart around with anything else, period. When Danny Hillis [founder of Thinking Machines] told me about his machines, he said he uses SPARC chips. And I said, "Danny, tell me one thing, you didn't take the heads off the workstations did you?" He said, "Yes, you'd never do that." And I said, "Yes you would because guess what happens? When you can't parallelize, you've got yourself a workstation." I told IBM that for God's sakes, don't take the heads off your workstations, leave them there and you've got some nifty stuff.

(Continued on page 7)

Thinking Machines' Response

In response to the Gordon Bell interview: I am less insulted by being called a "bastard" and a "kindergarten architect" than I am by his argument that Thinking Machines' success has been at the expense of the taxpayers. Bell's argument seems to go as follows: I, Gordon Bell, believe that massively parallel machines could not possibly work, yet people are spending hundreds of millions of dollars on massively parallel machines and reporting good results. Therefore, somebody must be cheating.

In a perverse twist of logic, Bell uses the fact that people have bought hundreds of millions of dollars of Thinking Machines products as evidence of his theory. Contrary to Bell's implication, almost all of these dollars have come in the form of commercial sales, not government grants. Thinking Machines' government contract research over the past 10 years has averaged about \$2.3 million a year, for which the government has received good value.

Our government sales are primarily to DOE, NSF and to classified applications where they have already made significant contributions to the nation's defense. They are selected in hard-won competitions against other vendors and other technologies. As anyone who has ever tried to get time on the Los Alamos or NCSA Connection Machines knows, these machines are fully utilized and there is a substantial backlog of users waiting to use them.

One of the best indicators that the customers are getting good results is that most of them have followed up their initial purchases with additional orders for Connection Machines. The fact is, contrary to Bell's continuing prediction, massive parallelism works. Connection Machines and other massively parallel computers routinely achieve significantly higher performance than vector computers at a



Thinking Machines Founder Danny Hillis

significantly lower cost. That's why people are buying hundreds of millions of dollars of our products. The customer successes with massively parallel computing are the reason why companies like Cray, Fujitsu and IBM have announced plans to offer massively parallel machines of their own.

Bell states that he will do whatever he can to "kill" us in the marketplace, and I do not doubt that this is the case. No doubt he will succeed in specific cases in delaying the transition to parallel computers, but I think that it is sad that one of the leading architects of the minicomputer revolution has decided to devote the latter part of his career to trying to diminish the successes of large-scale parallel machines. It is even sadder that he seems more willing to believe that customers are lying about performance and competitors are "playing foul" than he is willing to admit the simple truth that he just turned out to be wrong.

—W. Danny Hillis
Founding Scientist
Thinking Machines

Bell Interview... *(From six)*

I know some guys who are working in a parallel environment with plain, ordinary workstations, and the neat thing is anybody can be doing that kind of work.

Q: So you wouldn't buy one of the new massive-parallel machines?

Bell: The thing that I know about technology is that enough people have to go in one direction to beat down all the problems. Multiprocessor computers are the mainline of computer development. They all work independently and communicate for various reasons—like file servers or print servers—and some day we'll be able to use them in parallel a little bit.

The other mainline is if I ever want to work a single job the best way is to keep it in one memory and keep putting processors in there to absorb memory bandwidth. The breakthrough to me has been the fact that KSR has built a scalable version of that.

When you're trying to solve one set of problems, we have enough problems already—keeping all of the memory together, keeping the resources together, having one operating system—without having to have this enormously complex operating system for all of these independent computers.

So I'm a shared memory multiprocessor guy as the main line and there is no evidence that that's not what's going to happen. You've got huge forces telling you who's mainline and how you build computers.

Then there is the experimental scalable multicomputers with non uniform memory access and they're all going to disappear. That's Fujitsu, Intel, Meiko, Ncube and TMC. These guys all disappear because they're not mainline. The market tells you who's mainline.

The enemies are incremental parallelism from the traditional vendors: DEC if it ever decides to wake up; and IBM's workstation farms will have a tremendous impact because we haven't heard from them yet, except a little bit at Cornell. IBM is vicious when they decide to compete. They know how to give stuff away and get in.

Q: But Intel is busy selling its Paragon.

Bell: I don't think Intel has a single acceptance of it's current machine.

Q: They say they do—60 percent of the 25 or 26 machines delivered have been accepted.

Bell: You can look at what the acceptance criteria

(Continued on page 8)

Intel's Response

In his colorful, eccentric style, Bell has made quite a few comments that simply aren't true. But even more important is setting the record straight on what role Gordon Bell plays in this industry. It is no longer appropriate to position him as an impartial industry spokesperson. Unfortunately, many people don't realize he has a vested interest in Kendall Square Research—to the point of maintaining an office at KSR's headquarters. It should be widely recognized today that Bell has taken on the role of a very opinionated lobbyist for KSR and his comments have to be considered in that context.

In addition, Gordon Bell has no direct experience in large-scale parallel systems and therefore is hardly qualified to be a critic. Having an award named after him does not make him an expert. As I recall, Bell's last two companies were a minicomputer maker (Encore) and a workstation maker (Ardent). Encore's attempt to build a scalable, shared memory system (Gigamax) was not successful.

Back around 1986, Bell claimed that he could build a shared memory machine with 1,000 processors. Nearly seven years later, no such machines have been built while a number of 1,000 processor distributed memory machines have been built with more coming. Those of us who have been around the industry for some time certainly respect the enormous contributions Gordon Bell made as a minicomputer architect, but perhaps it's time for Bell to admit that his predictions and views on massive parallelism are wrong.

Bell further states that Intel's supercomputer operation "will never be successful." His comments are simply unfounded. Many prominent and credible analysts can confirm that our supercomputer business is already successful. The facts are:

- Intel has an installed base approaching 400 systems—with many repeat buyers who have come back for second and third systems from Intel;
- Of the 30 Paragon systems shipped, 21 have already been accepted and the others are going through acceptance tests with out customers;
- Demand for our systems in 1992 was twice that of 1991. We entered 1993 with the largest backlog in our history;

Next, Bell questions the acceptance of the installed Paragon supercomputers. Again his concerns are not based on any facts. As stated earlier, 70 percent of the 30 Paragon supercomputers shipped have been accepted and revenue. Of those not yet revenue, customers are waiting for particular software or hardware features which are still to be released for production.

Bell also implies that the only market for Intel's systems has been the U.S. government. He fails to recognize commercial accounts such as Prudential Securities, a user of Intel's systems since 1989, and our many installations in Europe and Japan.

There is one last point that needs to be addressed. While I don't feel I need to defend Intel's track record with the Touchstone Delta System, I must defend the many scientists who have performed outstanding work on the Delta and have applied such efforts competing for the popular Gordon Bell prizes, only to have Bell discredit their efforts with his comments. We continue to support the efforts of these scientists to advance the state of parallel computational science—with or without such prizes.

— **Michael E. Bernhardt**
Director, Marketing Relations
Intel Supercomputer Systems

Gordon Bell Vents Steam... (From page seven)

are. But to me, acceptance is that it is at least eight hours MTBF [mean time between failure]. And I haven't talked to anybody that believes eight hours with more than one user. I don't think they're doing that now.

Q: Intel says they've sold 40 machines, so there has to be a good market for them.

Bell: Well, how much is the market totally government controlled? How much of it is an internal government thing where you're stimulating a bunch of people to try to use these beasts? I'd love to know something about the character of that market and what's really going on there and how much real work is coming out of it.

I was hoping the GAO would say [in its recent report on ARPA's computing program] that we want the National Academy to sponsor a study that really is going to look at the work output of the machines.

As far as I can see absolutely everybody cannot talk about it. It's like you have this crazy aunt in the closet that you have to do something with because it's everybody's livelihood; the bureaucrats have to maintain the myth so they can maintain the funding; the scientists have to maintain the myth so they can keep doing what they're doing. The systems people love it because there are years to try to make the things work.

I just see that it's a program that has fundamentally gotten out of control. The problem is the scientific community is basically afraid to talk about it because they know they've conned the world and they have to keep lying to support this crap. It's really bad when it turns the scientists into a bunch of liars. So I think it's been fundamentally bad all over.

Q: We've learned a lot about parallelism in the process though.

Bell: One of the things I can see coming out of this is that we've gotten a much better mastery of parallel-

ism across a wide range of people, programs and programmers and training people for the next big computer structure of fast centralized switches. This next computer structure is in fact going to be all pervasive, everybody will have a lot of fast machines.

Q: So what do you see for a decade from now?

Bell: Basically a much more simple environment than we have today of basically going back to a simple switch and very powerful computers everywhere. In a computer room—just simply for security and for maintenance—there will be bays and bays and bays of independent simple computers and bays and bays of disks.

All of the things that are on the desktop will look an awful lot like a telephone.

Q: Given your views, what kind of reception do you get?

Bell: Right now, I suspect I'm persona non grata in Washington, which is okay by me because I call that a no-fly zone anyway. I don't need to go to Washington. I spent enough of my life in Washington. They're not going to change their policies. So what am I going to do? It's pretty easy. I can interact with the guys out here. I'm going to go help Henry [Burkhardt of KSRI], I'm going to go help anybody I can, and I'm going to kill them in the marketplace. Because they don't have all the money in the world, and good ideas kill dumb ideas all the time.

But it's the first time I've ever had to compete with people who are basically playing foul. I always used the marketplace. I always dealt with the market. It's just totally foreign to me. Sure I used to bitch about IBM. Sure IBM would threaten their users and stuff like that. But that's nothing compared with today.

This parallelism is crap. It's crap. Hey, it's good science, it's hard stuff, but for Christ sakes don't spend money designing weird-ass computers. I love to build computers, but it's a waste of money.

HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS WEEK IS A PUBLICATION OF KING COMMUNICATIONS GROUP, INC.

Llewellyn King, Publisher.

PUBLISHED WEEKLY. COPYRIGHT 1993 (ISSN: 1069-2096).

SUBSCRIPTION PRICE: \$597 PER YEAR. FAX PRICE: \$897 PER YEAR.

EDITORIAL STAFF:

Richard A. McCormack (Executive Editor) (Tel. 202-662-9721), Mark Crawford (202-662-9730), Ken Jacobson (202-662-9746), Jim Schubiner (202-662-9746).

CALIFORNIA: Sheridan Tatsuno, (408-685-8188). LONDON: David Fishlock.

SPECIAL CORRESPONDENT: Paul Kemezis, (615-695-5485).

BUSINESS STAFF:

Jane Connolly (Comptroller), Rita Robinson, Rosemarie Hargett, John Sodergreen, Mary Lou Probka, Karen Dahl.

CIRCULATION And CUSTOMER SERVICE:

Dawn Miles (202-662-9724), Sharee Wharton, (202-662-9711)

PRODUCTION:

Tony Campana, Clyde Pelzer, Anthony Russell, David Holloway, Keith Carter.

REPRODUCTION OF THIS NEWSPAPER BY ANY MEANS IS STRICTLY PROHIBITED, EXCEPT THROUGH COPY CLEARANCE CENTER. Call 202/662-9724 FOR INFORMATION.