

Faster Chips Are Leaving Programmers in Their Dust

By [JOHN MARKOFF](#)

Published: December 17, 2007

REDMOND, Wash. — When he was chief executive of [Intel](#) in the 1990s, [Andrew S. Grove](#) would often talk about the “software spiral” — the interplay between ever-faster microprocessor chips and software that required ever more computing power.

[Skip to next paragraph](#)



Kevin P. Casey for The New York Times

At Microsoft, from top, Craig Mundie, Burton Smith and Dan Reed are working on the next generation of computing power.

The potential speed of chips is still climbing, but now the software they run is having trouble keeping up. Newer chips with multiple processors require dauntingly complex software that breaks up computing chores into chunks that can be processed at the same time.

The challenges have not dented the enthusiasm for the potential of the new parallel chips at [Microsoft](#), where executives are betting that the arrival of manycore chips — processors with more than eight cores, possible as soon as 2010 — will transform the world of personal computing.

The company is mounting a major effort to improve the parallel computing capabilities in its software.

“Microsoft is doing the right thing in trying to develop parallel software,” said Andrew Singer, a veteran software designer who is the co-founder of Rapport Inc., a parallel computing company based in Redwood City, Calif. “They could be roadkill if somebody else figures out how to do this first.”

Mr. Grove’s software spiral started to break down two years ago. Intel’s microprocessors were generating so much heat that they were melting, forcing Intel to change direction and try to add computing power by placing multiple smaller processors on a single chip.

Much like adding lanes on a freeway, the new strategy, now being widely adopted by the entire semiconductor industry, works only to the degree that more cars (or computing instructions) can be packed into each lane (or processor).

The stakes are high. The growth of the computer and consumer electronics industries is driven by a steady stream of advances in both hardware and software, creating new ways to handle audio, video, advanced graphics and the processing of huge amounts of data.

Engineers and computer scientists acknowledge that despite advances in recent decades, the computer industry is still lagging in its ability to write parallel programs.

Indeed, a leading computer scientist has warned that an easy solution to programming chips with dozens of processors has not yet been discovered.

“Industry has basically thrown a Hail Mary,” said David Patterson, a pioneering computer scientist at the University of California, Berkeley, referring to the hardware shift during a recent lecture. “The whole industry is betting on parallel computing. They’ve thrown it, but the big problem is catching it.”

The chip industry has known about the hurdles involved in moving to parallel computing for four decades. One problem is that not all computing tasks can be split among processors.

To accelerate its parallel computing efforts, Microsoft has hired some of the best minds in the field and has set up teams to explore approaches to rewriting the company’s software.

If it succeeds, the effort could begin to change consumer computing in roughly three years. The most aggressive of the Microsoft planners believe that the new software, designed to take advantage of microprocessors now being refined by companies like Intel and [Advanced Micro Devices](#), could bring as much as a hundredfold computing speed-up in solving some problems.

Microsoft executives argue that such an advance would herald the advent of a class of consumer and office-oriented programs that could end the keyboard-and-mouse computing era by allowing even hand-held devices to see, listen, speak and make

complex real-world decisions — in the process, transforming computers from tools into companions.

The chip industry will continue to be able to add more transistors to a silicon chip for the foreseeable future, but the problem lies in the amount of power they consume and thus the amount of heat generated. That will limit the rate at which processing speeds increase.

The need to get around what the industry is calling the “power wall” has touched off a frantic hunt for new computing languages, as well as new ways to automatically break up problems so they can be solved more quickly in parallel.

Although the Microsoft effort was started about five years ago by Craig Mundie, one of the company’s three chief technical officers, it picked up speed recently with the hiring of a number of experts from the supercomputing industry and academia.

Mr. Mundie himself is a veteran of previous efforts in the supercomputer industry during the 1980s and 1990s to make breakthroughs in parallel computing. “I’m happy that by hiring a bunch of old hands, who have been through these wars for 10 or 20 years, we at least have a nucleus of people who kind of know what’s possible and what isn’t,” he said.

The more recent arrivals at Microsoft include luminaries like Burton Smith, a supercomputer designer whose ideas on parallel computing have been widely adopted, and Dan Reed, an expert on parallel computing.

Dual-core microprocessors are already plentiful in consumer devices. For example, both Intel and A.M.D.’s standard desktop and portable chips now have two cores, and even the [iPhone](#) is reported to have three microprocessors.

Microsoft sees this as the company’s principal opportunity, and industry executives have said that the arrival of manycore microprocessors is likely to be timed to the arrival of “Windows 7.” That is the name the company has given to the follow-on operating system to Windows Vista.

The opportunity for the company is striking, Mr. Mundie said, because manycore chips will offer the kind of leap in processing power that makes it possible to take computing in fundamentally new directions.

He envisions modern chips that will increasingly resemble musical orchestras. Rather than having tiled arrays of identical processors, the microprocessor of the future will include many different computing cores, each built to solve a specific type of problem. A.M.D. has already announced its intent to blend both graphics and traditional processing units onto a single piece of silicon.

In the future, Mr. Mundie said, parallel software will take on tasks that make the computer increasingly act as an intelligent personal assistant.

“My machine overnight could process my in-box, analyze which ones were probably the most important, but it could go a step further,” he said. “It could interpret some of them, it could look at whether I’ve ever corresponded with these people, it could determine the semantic context, it could draft three possible replies. And when I came in in the morning, it would say, hey, I looked at these messages, these are the ones you probably care about, you probably want to do this for these guys, and just click yes and I’ll finish the appointment.”

There are those who argue that there will be no easy advances in the field — including some inside Microsoft.

“I’m skeptical until I see something that gives me some hope,” said Gordon Bell, one of the nation’s pioneering computer designers, who is now a fellow at Microsoft Research.

Mr. Bell said that during the 1980s, he tried to persuade the computer industry to take on the problem of parallel computing while he was a program director at the [National Science Foundation](#), but found little interest.

“They told me, ‘You can’t tell us what to do,’” he said. “Now the machines are here and we haven’t got it right.”