## Innovation for Innovators

*an occasional column exploring principles, models, and theories of innovation in business and management*

**Roger Smith**

# Supercomputing On Your Desktop

Ray Kurzweil has garnered a great deal of attention by exploring the future path of computer technologies. His books use past trends and current information to predict what the world will be like ten, twenty, and one hundred years from now. His analysis is so in-depth and his reasoning so persuasive that his ideas are treated as models of the future rather than optimistic speculation. One of his more popular graphs illustrates the date by which a personal computer will have reached a level where it has the capacity to think like a human. He believes that current computers are able to process as much information as a dog's brain and they will reach human levels by 2020 (Kurzweil, 2006).

Quite apart from Kurzweil's predictions, Intel and AMD have delivered computer chips that make major leaps forward in computation every year. Both companies have made it clear that their future processors will offer something different. Instead of doubling or tripling the linear speed of their chips, they will be producing chips that contain multiple processors inside each one (Held et al, 2006). Chips with two or four processing cores are available now, but within a few years they will deliver 32 or 64 processors to a standard consumer desktop. Both companies plan to double the number of processors, also known as "cores", inside of a chip every 18 to 24 months. So for example, if we have 4 core machines in 2008, then we will have 8 in 2009, 16 in 2011, 32 in 2013, 64 in 2015, and 128 in 2017. Within the very short period of nine years we will have gone from a few processing cores to over one hundred. The consumer desktop begins to sound like a small supercomputer. It is a little personal version of a NORAD command center, NSA code breaking machine, Wall Street financial cluster, or Google search farm – right on your desktop.

All of these processing cores present a serious challenge and an industry changing threat to the companies that make the software for these machines. From Microsoft to the one-man programming shop, everyone is going to have to learn to work with this new paradigm of computation. Almost all computer programs are designed to do computation linearly because they have been built to run everything through a single processor. Microsoft Windows and Office, the Firefox Bowser, TurboTax, and computer games are all predominantly a single thread of computation that runs through the computer similar to thread running through a sewing machine. Some programs spin off a few tasks as independent and parallel threads, but there has been limited advantage to this when a computer has only one processor. Eventually everything has to squeeze through that one processor in a linear fashion.

But the personal computers of the future will offer multiple processing cores that do work simultaneously and in parallel. Even though there are 100+ processors in a computer, a

traditional linear program will not speed up because it does not know how to spread its work across all of the cores. These multi-core processors require a new generation of software applications that know how to use this power and a new generation of computer programmers who know how to create multi-threaded or parallel computing code. This is a major concern all through the software industry. How do you parallelize the operating system, office productivity software, the web browser, or the next computer game? The companies that master this will be the first to harness the new computing power and deliver "mind blowing" performance to customers. They become the "must have" applications in the next decade.

Multi-processor machines and parallel computing are far from new, they are just not common at the consumer level. They have existed in government, military, and academic products for decades, as well as industrial research labs, but they have never been a staple of consumer-grade equipment nor been part of the average person's experience with a computer. Programs like SETI@Home have allowed anyone to become part of a global, multi-processor computing project that leverages millions of idle computers around the world. But the average consumer has never been able to harness this kind of computing for their own personal needs and entertainment. What does it mean to give this power to every consumer in every home or every employee in every company? How can it be useful?

The recent wave of "web 2.0" applications has given us some clues to what people would do with nearly unlimited access to free computing on a company server and the network bandwidth to access it. Millions of people used these resources to publish their own memoirs, news feeds, diaries, personalized billboards, and encyclopedias. But all of this computation and storage exists in a public space on someone else's computers. It is by nature a public square for public performance, hence the social applications that have emerged with this power. The 128 core desktop machine will be a personal space that you control and can entrust with more private information. Your own personal finances, gambling habits, stock investments, hobbies, family information, and personality can be represented and served.

With every improvement in computer hardware, computer software companies immediately begin to search for ways to exploit this in an application that will attract customers. With 128 processors on a consumer or employee desktop, the possibilities for what can be done with this power are extreme. But it is not clear what the next application with a market the size of MS Office or Internet Explorer will be. During the dot.com boom this was often dubbed the search for the "killer app" (Downes and Mui, 2000). Some killer apps that emerged in the past are shown in Table 1. Office applications like word processing, spreadsheets, and presentations have been at the top of this list since they were first invented. In fact spreadsheets like VisiCalc, introduced in 1979, and Lotus 1-2-3, introduced in 1983, are credited with launching the success of the Apple II and the IBM PC as serious business machines rather than niche hobbyist tinker toys. Killer apps are essential because they create a reason for the consumer to buy the machines that can run them. More recently, high-end PC sales have been driven by the demands of visually compelling computer games. Without such games there is little

reason to upgrade to a new machine every couple of years. Most applications will run just fine on a five year old machine. Therefore, the killer app has the potential to double, triple, or quadruple hardware sales for the segment of the market that adopts those applications.

What are the killer apps that will drive the average consumer to purchase a 128 core desktop supercomputers? Two broad categories are generally reliable business areas – entertainment and data analysis. Weather modeling is one example of a universally interesting area for data analysis. The Weather Channel and Weather.com are hugely successful because of the shared interest in the topic across the entire population. Imagine that all of the data that drive these weather sites is available to the average consumer in real-time and that they are offered a program that allows them to easily run their own models of future projections and to explore their own customizations of the data. Imagine that instead of seeing the weather predications for their county, the home user can run predictions for their own city block. What will be the local wind speeds, volume of rain, and accumulated ground water from the latest hurricane in the area? 128 processors on your desktop may allow you to predict this better than your local weather station does today. Consumers just need the tools so they do not have to work with the raw data and write the computer code themselves.

Ever since the 401(k) retirement plan was introduced in 1980, the entire population has become more and more invested in the stock market. These investors watch the market as keenly as they watch the weather. But analysis of stock data is something for the professionals who have access to the data, have created the computer models, and possess the computer hardware necessary to handle these complex models. But given 128 processors, that same data, and the right software, every consumer can run their own analysis of the markets that are as in-depth as those run by brokerage houses today. They can build a picture of their investments and risk levels, and explore potential futures for themselves.

In the entertainment space, fantasy sports competitions attract nearly 30 million people just in the United States. These people's hobby is creating their own professional sports teams and using the performance of real athletes in real games to estimate the outcomes of their own customized teams in fantasy games against each other. These teams and their computations were originally done on paper forms with hand calculators. They are now run on spreadsheets and web servers. The average player uses a spreadsheet, web service, or custom program to study all of the players, build his team, and plan for the future. What if these 30 million people each had access to 128 processors to make their predictions? Would they value a fantasy sports program that allowed them to be 128 times more productive or rigorous than their competitors?

Imagine turning photos and videos into 3D virtual worlds in which the trees in the images are represented realistically in a space that can be navigated. The people in the videos actually walk through the space in the same way that they walk around in the movie. Your son's soccer team becomes a 3D experience that you can watch from any seat in the stadium or down on the field. Military and industrial surveillance videos become a 3D

space that can be viewed from all angles. All of this puts the current wave of movie and photo processing software to shame. But it may offer such a compelling experience that everyone has to have the computer hardware and software that can make it happen.

Finally, for the computer gamers, we offer a "digital buddy". Combine the 3D worlds, realistic avatars, artificial intelligence, access to real-time data via the network, voice recognition, voice generation, and lip synching that is available in pieces now to create a computer generated person that is custom built to be your best friend. This is the person who greets you when you come home; asks what you think about last night's football game; tells you that there is a club meeting at the local steakhouse; discusses your views on current events; and listens to everything you want to say, while giving you the kind of feedback you have programmed it for. This sounds impossible now. But each feature required to make this work already exists today and most of them run on the current high-end gaming machines. With 128 processors they could be combined into a single digital buddy that is as real as any person, but customized to your needs. Perhaps this would attract far more customers than even the 10 million players of World of Warcraft that tops the charts right now. The digital buddy might supplant both the computer game and the television as the most popular form of entertainment. Specialized characters could be added to act as business coach, weather man, medical consultant, athletic trainer, or girlfriend. In industry these buddies may replace the training department with customized digital trainers who are knowledgeable on every subject.

Each of these ideas attempts to answer the question – what can I do with the power of a supercomputer on my desktop? It tries to go beyond the current four processor machines and offer applications that have the potential to become more valuable as the number of processors multiplies. For those who can afford the machines, there is a very bright and exciting future ahead as long as the software industry steps-up to create the programs that can leverage this power. Unfortunately, even at consumer PC prices, there are many in society who will not be able to afford these devices and will have no access to the advantages they offer. The "digital divide" in society will become bigger. The "haves" will benefit from more information, better understanding, and more options than the "have nots".

How much electric power will it take to run these machines? We have generally adopted computers and electronic entertainment with little regard to their power consumption. But as electricity becomes more expensive and devices draw more of it, there will be a personal and a social pressure not to expend this power unnecessarily. We may be able to afford the machines and the software, but not the monthly power bill to run them.

From the perspective of the hardware vendors and futurists like Ray Kurzweil, a future with hundreds of personal processors and some really amazing applications running on them lies ahead of us. But other pioneers like Google offer a future driven by online services in which the average consumer needs a less powerful personal computer, not a more powerful one. They suggest that all of the computational, storage, and networking power that you need will reside in "the cloud." If this is the future then computer power will cease to be a consumer product and will become a corporate product offered to

consumers as a service. These two opposing visions seem to pit Intel and AMD against Google and Yahoo! in direct competition over the future of computing. The shift to multi-core computing could be as transformative as the first adoption of personal computers in the 1980's. It is clear that the hardware will exist. But it is not clear who will own it or what type of software will be created for the average consumer or professional . The technology creators like Intel, AMD, Google, and Microsoft are wrestling with the very difficult problems associated with creating software for these advanced machines. But the consumer and industrial customers will be faced with the opportunity to do much more complex work on their desktops or online – work that was previously reserved for supercomputers.

Table 1. Past and Future Software Killer Apps

| Wave | Applications |
|---|---|
| Office | Lotus 1-2-3, Visicalc, WordStar, Word Perfect, MS Word, PowerPoint |
| Communication | E-Mail, FTP, Web Browser, News Groups |
| Entertainment | Digital Video, Digital Audio, MP3, Computer Games |
| Socialization | Social Networks, Blogging, Wikis, RSS Feeds, PodCasts |
| Analysis | Weather Modeling, Sports Data Analysis, Personal Rendering, Financial Analysis, Simulation, Digital Buddy |

**References**

1. Kurzweil, R. (2006). *The Singularity is near: When humans transcend biology*. New York: Penguin Books.
2. Downes, L and Mui, C. (2000). *Unleashing the killer app: Digital strategies for market dominance*. Cambridge, MA: Harvard Business School Press.
3. Held, J.; Bautista, J.; and Koehl, S., Editors. (2006). From a few cores to many: A Tera-scale computing research overview. Intel white paper at http://download.intel.com/research/platform/terascale/terascale_overview_paper.pdf

*Roger Smith is the chief scientist and chief technology officer for U.S. Army Simulation, Training, and Instrumentation, in Orlando, Florida. He has also served as a group-level CTO for Titan Corporation and as a vice president of technology for BTG Inc. A member of RTM's Board of Editors, Smith has led technology innovation for software and computer systems for military training and command systems. He holds a Ph.D. in computer science and master's degrees in technology management, business administration and statistics.*