

Blue-Collar Computing: HPC for the Rest of Us

STANLEY C. AHALT AND KATHRYN L. KELLEY

Commercial forces influence all technologies. The HPC market, which has gone through many changes in recent years, is no exception. Indeed, the first do-it-yourself white box clusters look quite a bit different than the current blade systems available today. And yet, by all accounts, cluster technology is still in its infancy. Understanding the challenges that lay ahead is critical if the full potential of the market can be realized. Fortunately, there have been similar evolutionary technologies in the past from which we may learn some valuable lessons.

For example, let's look at the onset of the commercially viable automobile. Early models were as varied as their creators' visions. In 1900, wealthy people bought cars for pleasure, comfort, and status. Rural Americans liked cars because they could cover long distances without depending on trains. One example of vehicular pulchritude, the Duesenberg, was considered one of the best, most expensive American cars made in the early 1900s. The "Duesy" sold for \$15,000 at a time when a Ford cost \$500 and the Auburn Automobile Company produced around 1,000 of these automobiles. Likewise, the 1921 Winton was a low-silhouette luxury car, costing more than \$4,000; only 325 were built that year.

The Ford Model T, made between 1908 and 1927, cost less than most models of the time but was sturdy and practical. The Model T looked like an expensive car but actually was very simply equipped. And more Model Ts were sold than any other type of car at the time — over 15 million. Farmers, factory workers, schoolteachers, and many other Americans changed from horses or trains to cars when they bought Model Ts.

So the early automobile market was heavily skewed to the low end of the market. Many, many inexpensive models were sold, but relatively few more expensive automobiles were sold, regardless of their capability or appeal.

Compare that early automobile market to today's market. While there are a number of bare-budget cars on the market, the biggest-selling car models are those that are available in the mid-range of capability, power, and options. And on the far end of the scale, relatively few models are available to satisfy the discerning automobile customer looking for finely tuned, high-end sports and luxury vehicles. Moreover, the number of high-end autos sold is minus-

cule compared to the large number of mid-range cars that are sold.

As the automobile market matured, factors developed that increased demand from a full spectrum of the buying public, creating a bell curve of price and performance — most automobiles that are sold today are mid-priced, and have mid-level performance characteristics. Later, we'll argue that similar demands may cause the high performance computing (HPC) market to mature in a similar fashion.

HPC Commercialization Today

What was considered high end computing yesterday has quickly become part of everyday life. Who could have guessed 50 years ago that computing would be used as ubiquitously as pencils? While we are all familiar with the use of computers in most business offices, other uses of computers were unimaginable at the time computers first became commodity items.

Examples abound — human genomics, climate modeling, and jet propulsion, to name a few, have changed the way that science impacts our lives, the environment, and our economies — and each is based on computation and computational models.

However, the benefits reaped from HPC research and its resulting applications have not transferred to some industries that desperately need an infusion of computation. On the contrary, computational technology has been viewed by some to be at least partially to blame for massive workforce reductions. News sources and economists differ on how many jobs have been displaced due to outsourcing and improved technologies — both of which allow companies to improve productivity by either employing fewer employees or by employing a less expensive workforce that, in most cases, computes and communicates in a different country.

One of the industrial sectors most profoundly affected by these trends has been manufacturing. Nationally, the United States has lost almost three million jobs in manufacturing; the states with the most losses include California, Texas, and Ohio. Rarely discussed is how

manufacturing might use HPC to improve workforce productivity and manufacturing technologies in order to produce radically improved products and processes. For instance, research involving the development of advanced metals that are only nanometers thick, yet stronger than their thicker counterparts, hold great promise in manufacturing (See mse-po12.eng.ohio-state.edu/~fraser/html/programs.html). According to analysts, the market is ripe for higher-end manufacturing and industrial engineering and they expect “...manufacturing will grow faster than the overall economy.” (www.mapi.net/html/prelease.cfm?release_id=1492).

Fortunately, a growing number of companies are beginning to consider the proposition that HPC may be a key tool in increasing competitiveness and improving business. A July 2004 white paper commissioned by the Council on Competitiveness (CoC) and conducted by the International Data Corporation (IDC) surveyed 33 chief technology and information officers from aerospace, automotive, life science, electronics, pharmaceutical, and software companies, to determine the HPC needs of U.S. industry. The survey results are compelling:

- Over 70 percent indicated their companies could not function without HPC
- Over 25 percent of companies could quantify HPC’s ROI to their business, whether it saved millions of dollars, shortened production development cycles, or provided faster product-to-market timing (over 50 percent of affirmative responses).

Thus, while we can already see an emerging industry for HPC applications and HPC software that supports industrial and engineering work, there are cautionary notes as well. There are a number of interesting barriers that must be addressed before HPC is widely viewed as an essential component of our economy.

Industrial Barriers to Entry

According to the CoC study, business demand for HPC is still a relatively underdeveloped market. Over 65 percent of the reporting companies have important, but currently unsolved computational problems; the rest (35 percent) need faster computers for their problems. The need for HPC is obvious. Let’s discuss some of the traditional and nontraditional barriers that prevent industry from fully utilizing HPC.

■ One of the traditional barriers to industry HPC use has been the lack of technical expertise from the exist-

Problems with Building the Industrial HPC Market

Traditional Barriers

- Technical Expertise/Education
- Pricing and Support
- Intellectual Property
- Security
- No Immediate ROI

Nontraditional Barriers

- Cultural Barriers
- Little HPC/Industry Collaboration
- Lack of Imagination
- Not Forging New Tools and Utilities
- Risk Aversion

ing and emerging workforce. According to the CoC study, company technology officers cite a lack of trained computational scientists to apply HPC to their company’s problems. Many computer science graduates do not take an extensive number of science and engineering courses; likewise, engineering and science graduates are not traditionally trained in computational methods — and they have only been marginally represented in HPC. Scientists and engineers learn domain applications, while CS students learn programming skills — however, both sets of skills are needed in many industries. Put simply, engineers need to know how to write software, and computer scientists need to know engineering fundamentals. The same arguments hold true across many disciplines — it can be argued that what the pharmacology industry really needs is biologists who know how to compute, or computational scientists who understand biology.

What is really needed are sophisticated computational science curricula that integrate and use concepts that have been developed in multiple science and engineering domains. For example, simulation and modeling will be applied across science and engineering domains, coupled to visualization, data-mining, and statistical analysis. HPC should be an integral part of upcoming computational science programs and/or integrated in existing computing curricula at the undergraduate levels.

In addition, most engineering and computer science college programs cater to what industry needs now, not what industry might need in the future. Engineering and computer science curricula in most U.S. colleges focus on mainstream industry requirements, and doctoral students in engineering and science are channeled into very specialized domains of discourse — they are

not encouraged to seek the breadth that is needed for computational scientists to be effective.

2 It can be argued that many typical programming jobs are now routinely outsourced overseas. Without adding to the controversy surrounding this recent phenomenon, it can be argued that something other than jobs may need to drive the need for more industry-relevant computer science graduates. As *Wired* reporter Chris Anderson states, computers have always produced creative jobs in the workforce (See www.wired.com/wired/archive/12.02). Even though there is some evidence that some IT work will be brought back to the U.S., an AMR Research Inc. study of 125 large and mid-size companies found that 41 percent use a mix of offshore and onshore outsourcing, while 17 percent send all of their IT work overseas (“IT Labor Boomerangs Back Home,” available at www.eweek.com/article2/0,1759,1642854,00.asp). So instead of needing workers trained to do relatively straightforward computer tasks, industry might now require a workforce that is trained to take advantage of HPC-enabled innovation.

3 Another barrier to the extensive industrial use of HPC is pricing and support for industry computational systems and software. Third-party vendors and their costs for licenses differ from those available for industry users. Academic software licenses are not transferable to commercial use — even for testing. This situation is a barrier for mid-level companies that wish to engage academia in research that uses HPC for innovative industrial research.

4 Intellectual property and the security needed for handling industrial data is yet another barrier. There are documented cases in which legal confidentiality issues have extended the process of providing HPC services to industrial partners by 18 months or more. Thus, industry HPC users have tremendous confidentiality issues to surmount before they can outsource or share the best-equipped HPC centers. Simply stated, most companies are not ready to address the legal and security issues that accrue when they consider moving their most challenging problems to HPC centers.

5 Return on investment (ROI) is calculated differently in educational and governmental HPC research labs and industry. Industry’s timing is driven by their need for immediate returns and deliverables. Given that most CEOs must answer to shareholders who expect a yearly, if not quarterly, ROI, there is reduced flexibility in embracing long-term solutions to design and production issues. In addition, some companies cannot work within an academic time-share approach for HPC resources.

6 Non-traditional barriers involve culture clashes that occur between industrial and HPC communities. The reason why industry has not tapped the expertise of national or regional HPC centers can be boiled down to a simple issue of tech-transfer. The industrial community is accustomed to solving problems on the desktop, so many industrial leaders — including those in management, design, production, and testing — have simply not been exposed to HPC. Little HPC/industry collaboration currently takes place, which is problematic considering how industry might tap HPC experts to help with some of their most challenging problems.

7 Another stumbling block may be a lack of imagination. Given the current reliance on desktop computing as the foundation of most industrial communication and computing, it is not surprising that most engineers and designers don’t consider what problems could be solved more efficiently, more economically, or more comprehensively by using far more potent computing than that which is found on their desktops.

Tools: The Biggest Barrier

Given the above-mentioned structural and culture obstacles, it should come as no surprise that HPC opportunities are also lost due to issues with the tools currently used by the HPC industry. While it is generally conceded that HPC hardware and software are hard to use, HPC companies have little reason to forge new tools and utilities, or even use those tools that are widely available and compatible with all of the current available platforms. However, if HPC systems and software were easier to use, the CoC study found that that industry could tackle more complex models in a much wider context — thus creating more business opportunities. One exacerbating factor is the cost of HPC tools vs. other business investments. And as mentioned earlier, avoiding risk is common in many business cultures. HPC is an uncharted territory for many in industry — so convincing companies to invest in the development of new HPC tools is inherently problematic.

It’s a vicious cycle. Barring the development of new tools and software, HPC systems will continue to be hard to use. But developing HPC tools is expensive, and the market is limited, so companies have little incentive to develop the needed tools. And if HPC systems are hard to use, only those who can currently justify the essential use of HPC are willing to live with relatively crude tools and thus reap the benefits of HPC.

Hard to use means hardly used — at least by the broader community.

Refocusing HPC

Science and engineering research has historically offered significant advances in both theory and computer hardware, resulting in greater expectations. We understand physical phenomena better, and we have more capable computers, so we should be able to combine theory and hardware and reap the benefits! But there are real challenges. Companies must now begin to understand that HPC may offer many benefits in many application domains. And while there are codes that will (and do) benefit from the speed-ups available from even a modest number of processors, some extremely large problems may need thousands of processors.

Unfortunately, the link between improved theory and powerful HPC hardware is the software, and here is where we have a profound problem. There is a surprisingly limited demand that the capability of the software match the power of the hardware. Interestingly enough, while the U.S. leads the world in hardware engineering, Europe and Japan are investing more strategically in computer science research focused on software. The European Union plans to sink \$63 million into universities and research labs to make grid computing work for industrial projects.

HPC software that utilizes the most powerful hardware in a user-friendly, domain-oriented way is needed — and this requires an entirely new programming paradigm. An anonymous writer in a recent HPCWire article argued that if programming doesn't change radically, "parallel computing will be essentially dead within ten years." The vast majority of software applications do not take advantage of parallel computing for environments, and because conversion of serial codes requires major effort, the entire spectrum of HPC applications could reap the benefits derived from improved parallel programming models. It's one thing to come up with programming algorithms and quite another to make it available to the common user on a parallel machine. Furthermore, once you have developed a parallel algorithm, installing and maintaining it on multiple HPC platforms can be difficult if not impossible.

At the high end of the spectrum — the really hard problems that are being computed on parallel machines — benefits from HPC come from "grand challenge" problems that cannot be otherwise tackled. Grand challenge problems have been the bread and butter of the national research labs for the last decade or more, bringing federal resources and funds to bear in order to solve high-end computational problems. The solutions to grand challenges usually represent several orders of magnitude improvement over previous capabilities. The fundamental scientific problems that are represented in the grand challenges currently being explored 1) generate increasingly complex data, 2) require more realistic simulations

of the processes under study, and 3) demand greater and more intricate visualizations of the results. So interestingly, a special barrier to these extremely high-end computation challenges is the inherent difficulty of "large-code" programming, which will be exacerbated by dramatic increases in the number of processors needed to solve the problems — perhaps hundreds of thousands of processors. National interests mandate heroic programming efforts, and the continued investment of significant long-term funding indicates this aspect of HPC will persist into the future. That is, heroic computing will remain a fundamental part of the HPC ecology.

HPC in an Industrial World

Given the current economic climate and industry demands, the U.S. has reached a critical juncture with regard to HPC. Room for dramatic growth exists between current industry HPC use and "heroic computing." We argue that there is an increasing need for a partnership between industry and "heroic" HPC. We must find a way to promote HPC as a full-spectrum industry — and one way to achieve this goal is by focusing on high productivity computing languages and education.

"We are significantly expanding capabilities in computational modeling and computer-aided engineering, so we can do an increasing percentage of product and process design through virtual simulation," said A.G. Lafley, President and CEO of Proctor & Gamble at a 2003 Wall Street analysts meeting. Many large, forward-thinking firms are already making significant investments in advanced computational approaches to design and knowledge discovery. Through virtual simulation, production and process design is cheaper, quicker, and results in better products. Tom Lange, Associate Director of Corporate Engineering Technologies at P&G, states that innovation is his company's lifeblood. P&G spend \$1.6 billion a year in research and development. "Explore digitally, confirm physically" is the mantra for the company that has benefited from coupling supercomputer systems with knowledge in computational fluid dynamics and biomechanics to make innovative, competitive products.

Given the success of companies such as P&G, General Motors, Morgan Stanley, Merck & Co., Boeing, and Lexis-Nexis in integrating HPC into their research and development cycle, it's easy to see how focusing national research labs on a full-spectrum HPC market will greatly improve our national competitiveness. High-end computing will be increasingly important in making industries competitive in the global marketplace; companies that have found a way to leverage this advantage already know this. Just as the use of HPC has strengthened national security through stimulation of the field by way of federal grants,

we should now focus our innovations, advances, and education on the entire application spectrum — not just those at the high end of the spectrum. The “small” jobs of today will become the large jobs of tomorrow. Indeed, the greater impact will be felt across the entire computing market; that is, if applications can be scaled up and scaled down depending on the problem that needs to be solved.

The federal government is moving toward HPC as a solution to the problems of outsourcing and struggling industries. The White House is instructing executive-branch heads to give priority to supercomputing and cyberinfrastructure research and development in their fiscal 2006 budgets. In a memo, Office of Science and Technology Policy director John Marburger III and OMB director Joshua Bolten requested that supercomputing R&D “should be given higher relative priority due to the potential of each in further progress across a broad range of scientific and technological applications.” Agency plans in supercomputing should be consistent with a recent report of the High-End Computing Revitalization Task Force that describes a coordinated R&D plan for high-end computing (www.informationweek.com/showArticle.jhtml?articleID=29106042). The memorandum from the President’s Office give priority to research that aims to create new technologies with broad societal impact, such as high-temperature and organic superconductors, molecular electronics, wide band-gap and photonic materials, and thin magnetic films. According to the President’s Council of Advisers on Science and Technology (PCAST) Subcommittee on Information Technology Manufacturing and Competitiveness, the country must maintain a strong base of university R&D, educating the workforce in advanced tools and techniques in order to be competitive. These steps can also pave the way to creating well-paid, interesting jobs.

**Blue Collar™
High Performance Computing**

Economic forces will continue to shape High Performance Computing (HPC). It is clear that the United States has reached a critical juncture with regard to HPC, and the central challenge will be to sustain sources of funding so that our leadership position in HPC is not diminished. However, another view of this challenge might be more illuminating: can HPC be realistically viewed as one of the critical economic drivers for our future? Is this view of HPC realistic or even realizable? And if HPC is one of a relatively small number of critical economic differentiators, what level of national investment in HPC is justified by its economic potential?

Blue Collar Computing is high performance computing for industries that do not currently have the expertise or the time to be an HPC incubator or research new HPC ap-

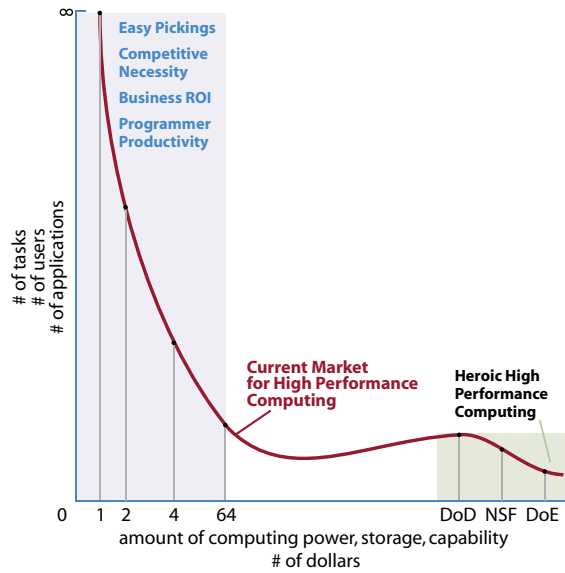


FIGURE ONE The current HPC market composed of easy pickings and heroic computing

plications. High performance programming languages, training, and collaborations are required to open up greater capability and competitiveness to business, science, and engineering users. Blue Collar Computing scales up the number of processors beyond the one or two CPU boxes that companies usually run. The focus should be on high-productivity languages, industry and supercomputer lab collaborations, and the training needed to provide the expertise needed to allow the greater capability and efficiency of HPC to be utilized by business users.

National labs that do large-scale “heroic computing” — the Departments of Energy and Defense (DOE and DoD) and the National Science Foundation (NSF) — will continue to concentrate on monumental, “grand challenge” problems. However, if new resources are focused on industry HPC requirements, everyone would benefit in the long run. Blue Collar Computing is the computing needed in order to sustain the nation’s aspirations, and the computing needed to ensure that the US will remain an economic and scientific leader.

Looking at *Figure One*, we notionally describe the current HPC market as follows. We currently have relatively few users that utilize the nation’s most powerful HPC installations. And on the lower end of the spectrum, use of more than one to two processors is relatively rare. Programmer productivity is maximized when only a few processors need to be harnessed, but a large number of applications are available for those few processors. The need for more processors has not been demonstrated, since the vast majority of average users are generally happy with using only one proces-

sor, since it offers immediate ROI at a minimal cost.

Looking at *Figure Two*, Blue Collar Computing demands from industry may allow us to push the mid-range of the market to a higher level — especially in the auto, petroleum, financial, and pharmaceutical industries. Collectively, there could be as much — or more — need for HPC in industry as now found in the DoD, NSF, and DOE. We have the heroes — DoD, NSF, DOE — still using HPC to take research and supercomputing to new heights. Yet industry could partake in the benefits of increased HPC investment, processing power, and compatible software.

The future of industrial competitiveness must move beyond the one- or two-processor model if we are to visualize what is shown in *Figure Three* (page 18). Industry will be working on bigger and better solutions using computing languages and employees educated to fit in this new economy. In an ideal market for industrial HPC, training and education is a major component — even chemists, engineers, and other scientists who use research labs and HPC centers need training in computer and computational science. This environment can be a true economic generator, moving away from traditional manufacturing to real “knowledge economy” applications.

Ultimately, the entire HPC market will grow by enabling industry to solve problems and develop better products more quickly. In the next decade, this is where we will see dramatic gains: increased productivity gains in industry and engineering and increased gains in scientific discovery from those hero HPC applications that will always push the envelope.

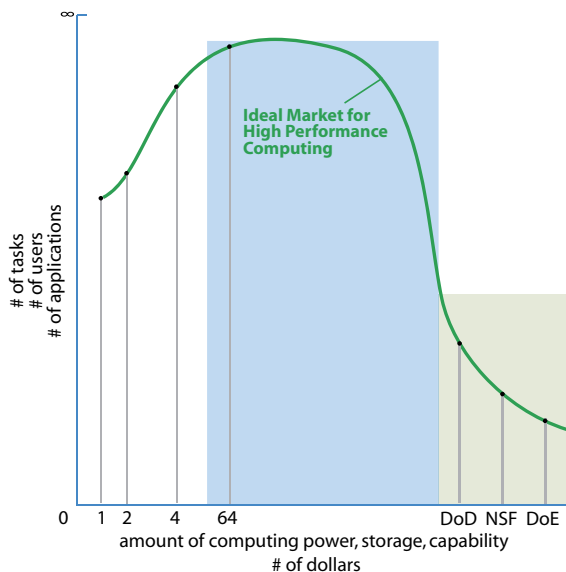


FIGURE TWO The ideal HPC market

Blue Collar HPC Applications

If we can expand the HPC market — change the shape of the market to reflect a greater use of Blue Collar Computing — we can guarantee that parallel processing, clusters, and all of the associated software is made available to industry. However, the computing market today is like the automobile market of the early 1900s. We produce a huge number of very reliable, very useful Model Ts and relatively few high-end Duesenbergs. But, if we can follow the lead of the automotive market and focused on increasing the functionality of HPC by providing standardized interfaces and highly reliable software, we may be able to transform the HPC market so that it feeds a constantly improving cycle of innovation. Once the needed tools are developed and software interfaces are made more accessible to today’s domain scientists and engineers, the move to significant industrial use of HPC may take hold.

How do we create the tools? First of all, we need a public/private partnership to work on sustaining interfaces and software tools. Second, individual corporate entities may have the capability to solve problems on one or two processors, but they may not be able to spend the funds required to look at more advanced applications that will utilize many more CPUs. However, companies might be convinced to invest funds in long-term HPC research if there is a concomitant matching investment by government.

In order to use these applications, we must remove the industrial barriers we have discussed above. The result will be an ability to create better products, help our innovators to “think faster,” and actuate the next long-lived productivity expansion. The heroes and the rest of the HPC community will also reap the benefits.

Summary

We have argued that a fundamental shift in the HPC market — a shift to Blue Collar computing — needs to take place in order to revitalize U.S innovation. We are at the beginning of a new phase in the evolution of computing, and we’ve only begun the journey. We have

Resources

Ohio Supercomputer Center Blue Collar Computing

- www.osc.edu/bluecollar

Council on Competitiveness HPC Users Advisory Group

- www.compete.org/hpc

International Data Corporation

- www.idc.com

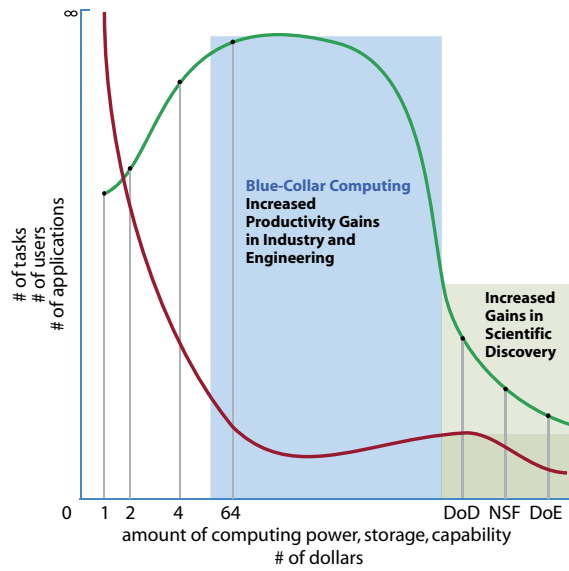


FIGURE THREE Blue Collar Computing

proposed some general solutions that involve high productivity languages and education; however, a focused attempt to solve specific industrial HPC problems and barriers is vitally needed.

We propose that the HPC community start working on a public/private partnership to develop the elements of Blue Collar Computing that are most pressing. Next steps involve focusing on possible implementation plans for high productivity languages and radically changing computer science education. There is a need for both undergraduate and graduate personnel with computational science expertise, which takes an extraordinary level of institutional cooperation. One possibility could be to allow companies to provide curricular material to stimulate teaching in the direction of solving current and future industrial “grand challenges.”

Ultimately, the first question that needs to be asked is if are we ready for a shift to Blue Collar Computing. Can a commitment be made to embrace it? A transition to Blue Collar Computing will naturally break down many of the current barriers to entry, and provide the nation with a vitally important economic edge. This shift will allow the United States to maintain its economic leadership in the global marketplace. But, it requires a paradigm change in our way of thinking, our way of teaching, and our way of approaching HPC.

Stanley C. Ahalt is Executive Director of the Ohio Supercomputer Center and can be reached at ahalt@osc.edu. Kathryn L. Kelley is Director of Government and Community Relations and can be contacted at kkelley@osc.edu. Figure credits: Ian MacConnell.