Department of Computer Science Technical Reports

Department of Computer Science

1995

Computational Science and the Future of Computing Research

John R. Rice
Purdue University, jrr@cs.purdue.edu

Report Number:
95-045

Rice, John R., "Computational Science and the Future of Computing Research" (1995). *Department of Computer Science Technical Reports.* Paper 1220.
https://docs.lib.purdue.edu/cstech/1220

Computational Science and the Future of Computing Research

John R. Rice
Computer Sciences Department
Purdue University
West Lafayette, IN 47907

CSD-TR-95-045
July 1995
Revised September 1995

Computational Science and the Future of Computing Research

John R. Rice
Department of Computer Sciences
Purdue University
West Lafayette, IN 47907, USA

1. The Environment for Research Funding

It is evident that the environment for funding scientific research is changing. There is widespread concern that levels of industrial and government funding will decrease and that there may be shifts in emphasis for research funding. The latter comes from the emerging focus on funding *strategic research* which means research that has the potential for impact on the economic, health, environmental, educational, etc., well being of the nation. Strategic research is still fuzzily defined, but it refers to areas of research and not to types; it is orthogonal to concepts like basic, applied, long-term, or short-term. The nature of strategic research is illustrated by the six fundamental and over-reaching goals for all federal science and technology investments (from the National Science and Technology Council [3] which reports to the President):

- *A healthy, educated citizenry,*
- *Enhanced national security,*
- *Harnessing information technology to support all of the other societal goals,*
- *Improved environmental quality,*
- *Job creation and economic growth,*
- *World leadership in science, engineering and mathematics.*

The leaders of the scientific community argue strongly [1] for across the board support of research in science. They probably view the strategic research concept as leading to making choices among the fields and sub-fields of science, a prospect they oppose and a process they are unlikely to participate in. They purpose that the reason for the decline in levels of research funding comes from three phenomena: ignorance, the end of the cold war, and misguided expectations about science. It is certainly true that the general public and political leaders are, by and large, poorly informed about science and research. The cold war has ended and there has been a significant decrease in defense expenditures. There certainly are some unrealistic expectations about science due to a lack of understanding of where the frontiers of science are located. For example, we see complaints about the lack of progress in combating the AIDS virus by those who feel that all viruses should be equally easy to control.

These phenomena all exist but it is not likely that they are the reasons for reducing the level of support for scientific research. In fact, science research has fared very well in federal budgets so far. Its strong support has kept its funding from declining as much as for most discretionary items. The end of the cold war has primarily affected the Department of Energy through the reduced need for nuclear weapons. Even there the ultimate effect on research is unclear. Thus, science research should be viewed as doing very well in the current budgetary environment. But there is a shift of emphasis embodied in the focus on strategic research. The emphasis is moving towards funding research that has plausible prospects for reasonable return on the taxpayer's investment.

Thus we conclude that the pressure to decrease science research budgets is due almost entirely to the pressure to decrease the federal deficit. There is, however, another intriguing but nebulous theory, namely, that the science establishment is reaching its mature size measured as a proportion of human activity. This establishment has grown exponentially for about four centuries and exponential growth cannot continue forever. Perhaps we are reaching the time when the scientific community stabilizes in size. It appears to some observers that the problem is not so much the lack of money but the ever increasing flow of bright, young researchers competing for research support.

The thesis of this paper is that all subareas of computing research can prosper in an era of focusing government funding on strategic research. We argue for this thesis by examining science and engineering applications but we believe there are equally compelling arguments based on other application areas. However, prosperity will not come automatically. The computing research community has been described as inward looking [4] and many, perhaps the bulk, of its researchers have avoided applications entirely. This can be justified by the fact that this young field needed time to establish its own foundations firmly. Now is the time to become more outward looking and to appreciate that computational science (and other) applications will essentially involve and greatly challenge all subareas of computing research.

2. The Context for the Future

The growth in computing power continues to be astounding and shows no signs of abating. That this growth is unprecedented in recorded history is illustrated in Table 1 where quantitative changes in computing are compared to changes in public transportation, explosive power, energy production, construction and education. The growth of computing power of the next two decades, coming on top of five decades of explosive growth, will be more than the growth in speed of transportation from the times when everyone walked to the supersonic jets projected for the early 2000's.

The nature of this growth is illustrated by a simple application: compute where the cooling water pipes should go in an automobile engine block (see Figure 1). This is a real world problem that has been "solved" by experimental and analog methods for many decades. It involves one of the best understood physical phenomena, heat flow. One just has to solve the Poisson problem for a complicated three dimensional object. Methods and machines were available in 1940 that could, in principle, solve such a problem. I estimate that this computation (for just one engine block), in 1940 would have cost the entire wealth of the United States. When I first encountered this problem in 1963 there had been enormous progress in both computing hardware and algorithms since 1940. Nevertheless, the computation was not yet economically feasible. Today, the cost of

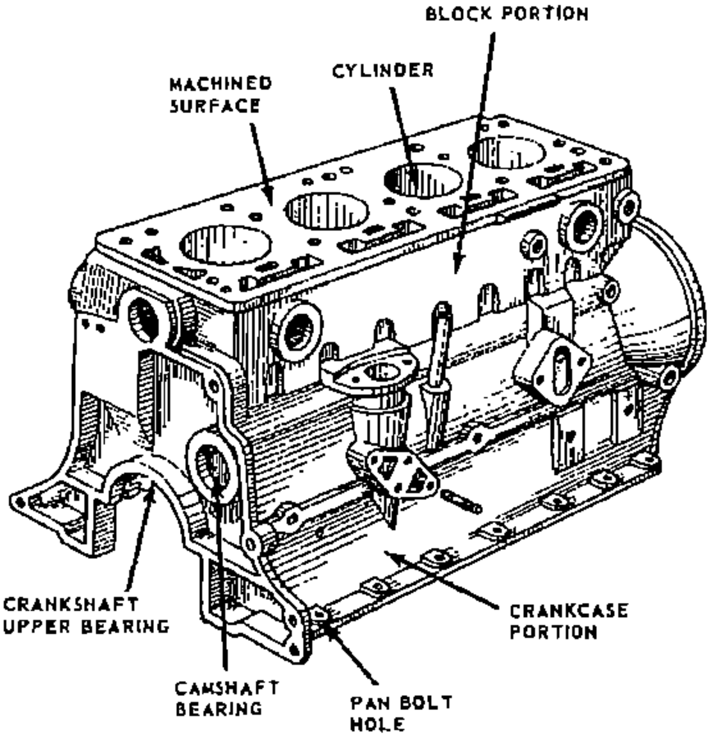

Figure 1: A typical automotive engine block.

Table 1: Comparison of technology changes from ancient times to the near future.

Area	Ancient Times	1890	1950	1970	1990	2010
Transportation (Miles per day)	40	200	6,000	35,000	35,000	150,000
Computation (Multiples per second)	0.005	0.04	40	10 Million	5 Billion	20 Trillion
Explosive Power (Tons of TNT)	0.0003	0.5	1 Million	100 Million	100 Million	100 Million
Energy (horsepower/day)	0.15	0.5	3	5	7	9
Construction	Great Wall	Suez Canal	Fort Peck	Aswan Dam	US Highways	?
Education (Years, U.S.A.)	None	1	8	10	12	?

this computation is a few tens of dollars. It is very significant that algorithmic progress has been a larger factor in decreasing the cost than the progress in computing hardware speed.

One can quibble about whether the progress in computing speed for the next 20 years will be a factor of 1,000 or 5,000. There is little doubt that the power of 10 megaflops with 10 megabytes of memory will cost the order of $5. One should visualize that every computer one is using now will have 999 others beside it in 2015 to provide better service.

3. Blast Furnaces

Four example applications are considered to illustrate the nature of future computational science applications. The first is the control of blast furnaces, a traditional numerical application of the not far distance future. Figure 2 shows a schematic of a typical furnace. There is a lot of complicated machinery outside the reaction chamber where the steel is extracted. The problem is to determine what is going on inside in order to adjust the input (oxygen, coke, ore, etc.) and controls so as to optimize throughput and quality. The idea is to monitor the outside and simulate the inside to determine the effects of actions and thereby provide more precise control.

The current technology is to rely on the skill and experience of operators who apply heuristics accumulated from years of observing blast furnaces. Simulations of the inside have started but they are not yet accurate enough or fast enough for actual control. This simulation involves complex 3D geomety, moving fluids and solids, combustion and other chemical reactions, melting and mixing of materials, and real time constraints. It is a complex but standard partial differential equations application. This application involves four subareas of computing: simulation of physics (many levels), high performance computing, control theory, and AI/expert systems. The last is involved in managing computational resources and for incorporating as much of the operator heuristics as possible because the simulations will be less than perfect.

It is plausible that within a decade every blast furnace will be computer controlled. A rule of thumb is that accurate control can improve productivity by 30% or so. A blast furnace costs millions and a 5–10 gigaflops machine will cost less than $100,000 within a decade. Computer control will become an economic imperative for blast furnaces.

4. The Design of Physical Objects and Mechanisms

The first application presented few challenges for most subareas of computing; the next is closely related in many ways and yet it involves essentially all of computer science. The problem is to design reasonably complex physical objects and mechanisms using electronic prototyping. Figure 3 shows a collage of images from a current system [2] of this nature. The key problem is illustrated along the bottom where the shape of the end of a piston rod is to be optimized. The size of the end is to be reduced while maintaining adquate strength.

At the top level, this application involves the following subareas of computing:

- *Simulation of physics.* There are many kinds of phenomena and they must be modeled more accurately than for a successful blast furnace control system. Further, everything must be

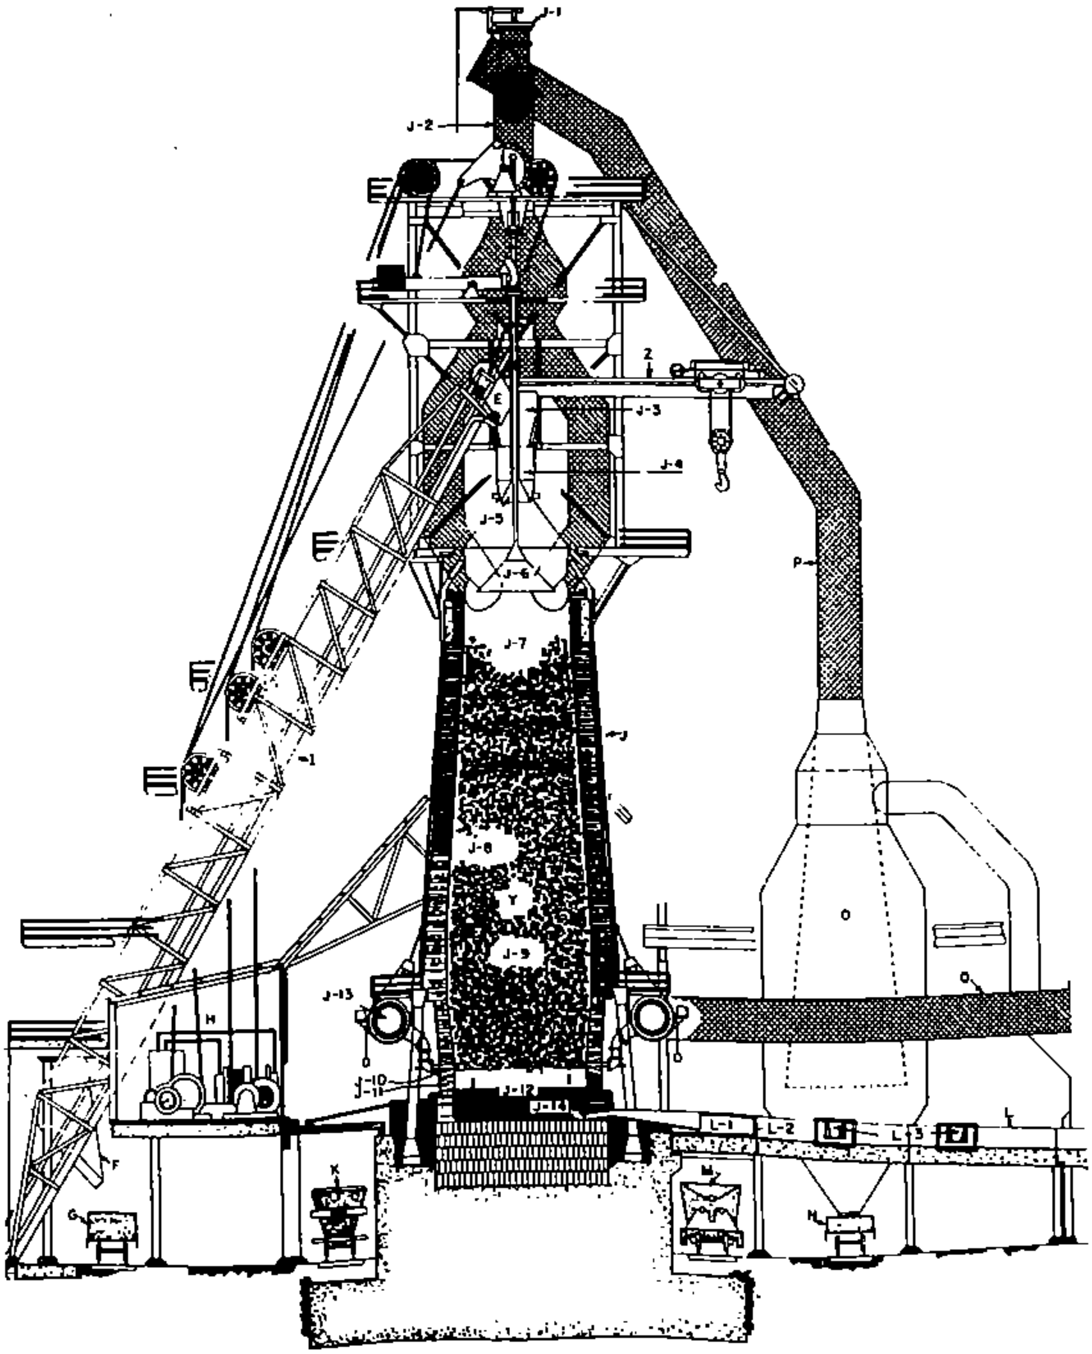

Figure 2: Idealized cross-section of a typical modern blast furnace.

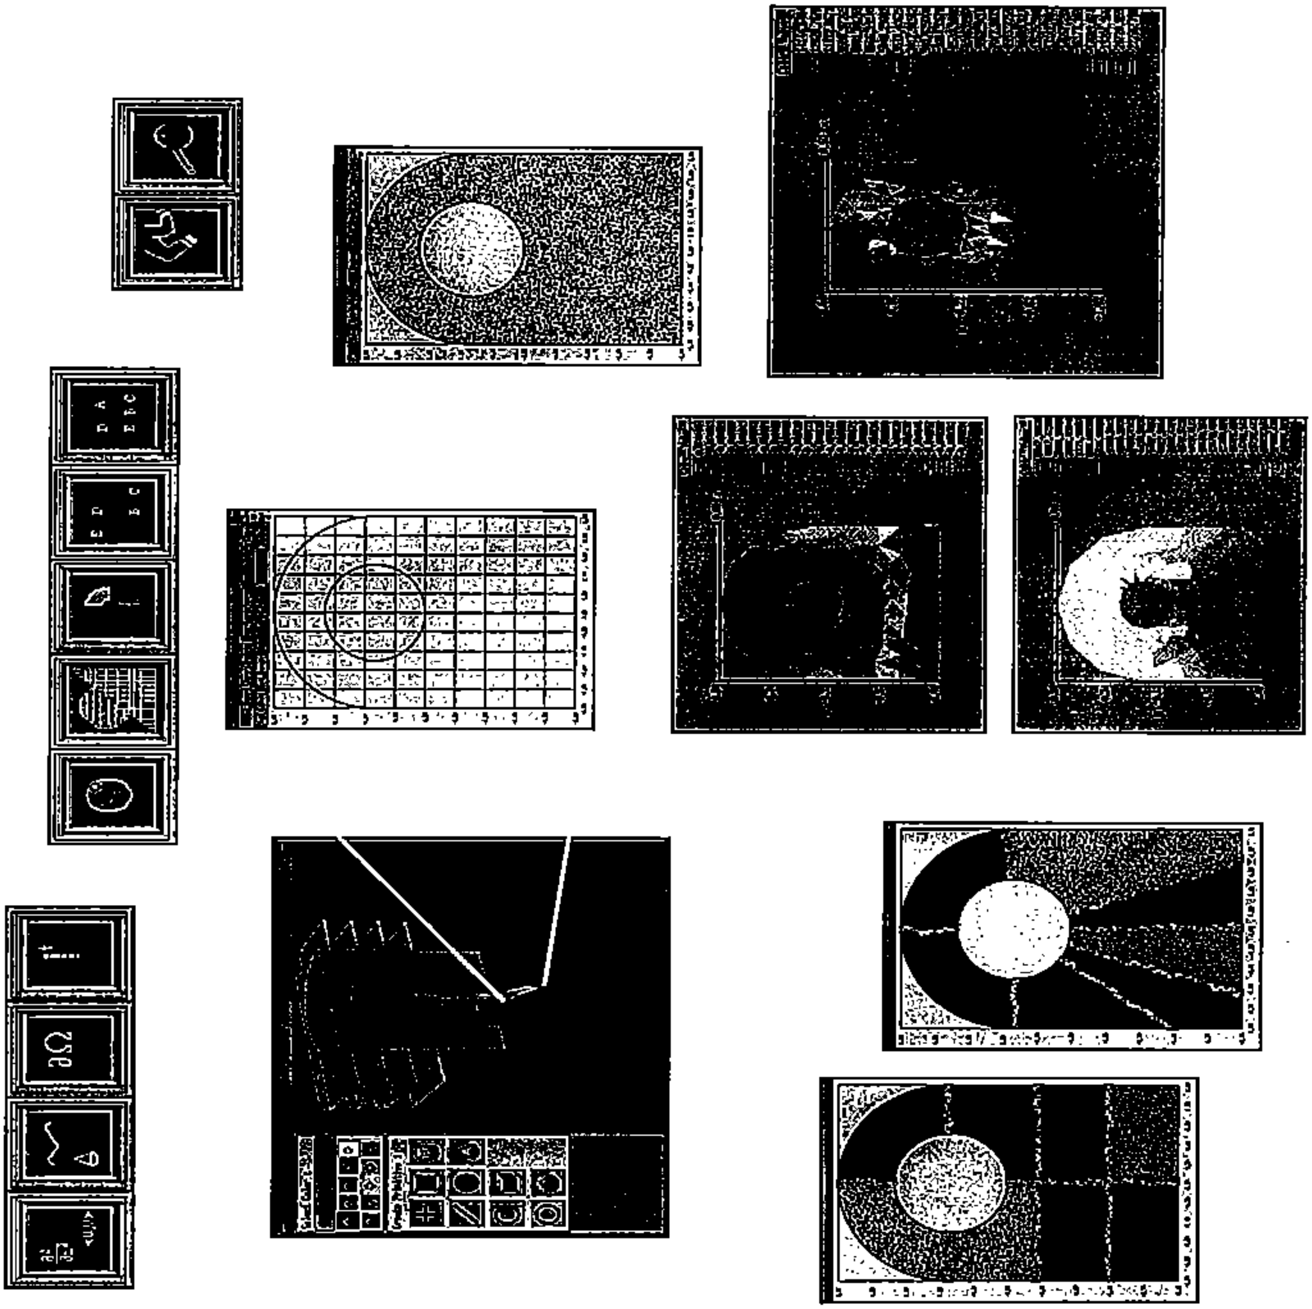

Figure 3: A collage of screen images from the PDELab system showing steps in the optimization of the shape of a part of a simple engine.

simulated, not just one or two key parts of the mechanism.

- *High performance computing.* A useful rule of thumb is that it takes from 100 to 1,000 times as much computing power to optimize a design as just to simulate one instance. Thus computers that deliver 10s, 100s, and 1,000s of gigaflops are necessary.

- *Artificial intelligence and expert systems.* Expertise is needed throughout such an application. It is clear that many components of good design are not yet codified in a way that they can be used routinely. Thus a large array of heuristics about design must be incorporated. Furthermore, mechanisms must be manufacturered and this places a further large number of constraints on a design. These constraints are even less well understood than those of design and they must be represented primarily by heuristics. Shape optimization involves many variables interrelated in complex, nonlinear ways. Such optimization can consume enormous computing resources. It is plausible that the emerging idea of application specific optimization algorithms will be needed. Here one applies optimization algorithms to a set of similar problems (e.g., optimizing the shape of a piston rod or a crank handle) and "learns" those tactics that are effective for this particular set of problems. This is a promising, simply stated idea that is a stiff challenge to carry through. Finally, even the management of the computing resources for such a compute intensive application requires sophisticated optimization and heuristics.

- *Geometry and graphics.* It is obvious that physical design requires extensive geometry and graphics facilities. It is less well recognized that "geometric computing" is grossly underdeveloped compared to numerical, symbolic and logical computing. The manipulation of even simple shapes cannot be done in hardware – or at hardware-like speeds. Yet this capability is sorely needed for design applications.

- *Database.* The items relevant to mechanical systems are not like simple records of banking systems. Further, an enormous number of such items must be available as part of the knowledge about how things are designed and how things are manufactured.

- *Human interfaces.* A physical design system will involve millions of lines of code, thousands of software modules and subsystems. Yet a designer should have easy, natural and responsive control of the design process. This surely presents a great challenge to the designers of human interfaces.

This application has a substructure which involves a much broader range of subareas of computing. These are mentioned along with a brief indication of how they are involved in this design application.

- *Algorithms and data structures.* There are hundreds of these needed for all the specialized representations and manipulations performed on complex objects.

- *Parallel algorithms.* Parallel computing is the only hope to provide the computing power needed. It will not be easy to supply the power needed in a responsive way.

- *Knowledge bases, smart algorithms, adaptivity and learning.* The substructure of the "smart systems" involved in this application will include most techniques and methodologies of artificial intellegence applications. These methodologies have been under development for several decades with mixed results. It is now the time to produce effective systems for adaption and learning; it is plausible that the increased computing power available can make this happen.
- *Symbolic systems.* Mathematics is the basis of most modeling of physical systems and is heavily used in some approaches to geometry computations.
- *Big, complex data structures.* The objects involved have many attributes and mountains of associated data. The detailed nature of these structures is not known in advance so the definition must be dynamic and the large sets of them must be self-organizing so that access paths to the structures are dynamically determined.
- *Languages.* Language of many forms is involved: jargons from application areas, cues, natural language, and, above all, visualization.

This application is harder yet because it involves:

- *Distributed design.* Even modestly complex mechanisms are designed by teams so that networks of collaborating people and computers must be supported. Each application will have its own operating system, one with more demands and as much complexity as the generic operating systems of the 1980s. There will be more heterogeneous resources to manage, more deadlines, and more synchronization constraints.
- *Performance analysis.* The performance of such systems will be awful at first, there are hundreds of places to lose.
- *Security.* There is no information more valuable to a company than the designs and specifications of the products they plan to introduce.

Finally, this application involves the central unsolved problem of computer science: *how to engineer software effectively.* This application involves millions of lines of code from multitudes of sources and the system is to be reliable and efficient. We still have a long way to go in order to understand how to build such systems efficiently.

5. Reality, Simulation and Virtual Reality

Reality is the starting point of these two applications and simulation is supposed to compute what would happen in reality. The blast furnace application is mostly reality; most of the furnace is available to be measured as needed and only a part (but the most crucial part) is simulated. The design application is mostly simulation, the contact with reality comes only when the object is manufactured and tested.

Virtual reality is a form of simulation that is a direct extension of the design system. Whatever is involved in the virtual reality environment is simulated well enough that the human sensory inputs are (nearly) the same as for reality. As this methodology advances, it will involve all the human senses, not just vision. And, as more senses are involved, the simulation must be more complete. For example, walls in a building now are simulated by idealized planes with color and texture superimposed. When sound and touch are included then the walls must simulate much of the physical structure of real walls. Virtual environments are similar to virtual reality in that everything must seem "real" to the humans in the environment. However, they may combine many actual objects with simulation. For example, in pilot training simulators the cockpit is real but the motion, the rest of the airplane, and the views out the windows are simulated. A virtual environment may also be completely or partly artificial. For example, one could be that of a boat navigating through the blood stream of a person or through the molten materials inside a blast furnance. A completely artificial virtual environment could be based on a pseudo-physical representation of the flow of money (and other financial instruments) in the economy of a city. Then a person in the environment could directly "observe" these flows as the economy changes.

Within two decades virtual reality and environments will provide very high levels of realism using accurate and complete simulations of the physical (or pseudo-physical) environment. This will involve all the subareas of computing that physical object design involved and with more demanding performance for most of them.

6. Robots

The final application considered is robots, which will appear within 20 years with reasonable speech and vision capabilities. Their movement and touch capabilities will be useful. Their capabilities to access information and do computations will be enormous. None of these capabilities will be anywhere close to human capabilities, but that is not necessary for them to be very useful. Recall that the vision system of a frog is only black and white and that a frog can see things only when they are moving. In spite of this primitive vision system, frogs get along quite well. Robots will be able to do likewise with their limited capabilities.

The computational problems for robots are much more difficult than for virtual reality. Compare the requirements for walking down a hall by a robot and within a virtual reality system. The principal activities are

Control. In virtual reality, a person uses natural (existing human) control mechanisms for balance, path determination, moving, etc., while walking. A robot must compute its path after recognizing the environment and then control its motion in order to follow this path.

Vision. In virtual reality, a person sees a scene created from a known data struture which is designed to make scene display efficient. A robot must observe an arbitrary scene and identify the major components (walls, doors, obstacles, stairs, etc.). Such scene analysis computations have proven to be one of the more difficult challenges for computing research.

Sound. In virtual reality, a person hears sounds created from a known data structure or simply recorded previously. A robot must analyze the sounds and extract the important components (speech, footsteps, objects colliding, direction, etc.).

Touch. In virtual reality, a person senses objects from forces created from a known data structure. To compute and deliver these forces accurately is a major computational and mechanical challenge. A robot must have tactile sensory devices and must be able to interpret their input in terms of its environment and objectives. This is an even greater challenge.

The four areas of capability for a robot are currently in different states of development. The mechanical control and motion problems have been studied for a long time and much progress has been made. While there are still unsolved problems, this does not appear to be a major hurdle. These vision problems also have been studied for a long time but less progress has been made. It is plausible that vision requires more computation power than previously expected and that more rapid progress can be made with the continuing increases in this power. The auditory problems have also been studied for some time, primarily related to speech recognition. This effort has been much less than for vision but it appears that, at least, speech recognition is now feasible. It is no doubt a major challenge to extend this technology to general sound analysis. The problem of touch seems to be much less studied; one can, however, hope that useful robots can be made with primitive touch capabilities.

We have focused on the high level computational problems of creating robots but, just like in the previous applications, there is a large, complex substructure based on the subareas of computing. A robot will be controlled by a network of powerful processors with a specialized operating system, databases, semi-autonomous processes, etc.

7. Conclusion

It is likely that the focus on strategic research is not a passing fad, even though the final definition of this term is still unclear. A computing researcher who cannot connect to strategic research in a direct way probably does not deserve research funding. If one does not have enough creativity to make this connection, then one probably does not have enough creativity to do significant research either. Many computing researchers might have to put some effort into establishing the connection to the larger world but this task is just part of living in an ever changing, dynamic world.

References

[1] Committee on Science, Engineering and Public Policy, *Science, Technology, and the Federal Government: National Goals for a New Era*, National Academy of Sciences, Washington, D.C., 1993.

[2] C.M. Hoffmann, E.H. Houstis, J.R. Rice, A.C. Catlin, M. Gaitatzes, S. Weerawarana, N-H.L. Wang, C.G. Takoudis and D.G. Taylor, Softlab – A virtual laboratory for computational science, *Math. Comp. Simulation*, **36**, (1994), 479–491.

[3] Committee on Science, Engineering, and Public Policy. National Academies of Science and Engineering, 1993.

[4] John R. Rice, Is computing research an isolated science? *Computing Research News*, **2**, No. 2 (1990), 1.