

The Future of Computing: *Cyberspace*

Jens Pohl

CAD Research Center, Cal Poly, San Luis Obispo, California, USA

Abstract

This paper traces trends in the technological advances of computer and communication systems and examines the promises of the *Information Society*: global information sharing; and, intelligent decision-support. The technological developments that will lead to the realization of *Cyberspace*, an information rich environment in which virtual reality capabilities couple directly to the human senses, is explored in terms of five essential components: information processing requirements; communication networks; computing devices (i.e., platforms); hardware and software user-interfaces; and, the meaningful representation of information.

Attention is drawn to the critical role played by information representation in a *Cyberspace* environment. The author argues that the communication infrastructure must become more than a message passing facility. It must have some *understanding* of the information it is transmitting. If this fundamental requirement is met then *Cyberspace* will present human society with an unprecedented potential for leveraging the capabilities of the individual members of society for their own benefit and the collective benefit of mankind.

Keywords

agents, body nets, communication, computers, Cyberspace, data, future, information, Information Society, knowledge, LAN, local loops, networks, projections, representation, requirements, WAN, wireless networks

Introduction: The *Information Society*

Over the past 12 years the Collaborative Agent Design Research Center at Cal Poly, San Luis Obispo, California has been intensely focused on the design and implementation of collaborative decision-support systems. In these systems human decision makers and computer-based agents opportunistically assist each other in the exploration, analysis and solution of problem situations in which there are many variables with complex relationships and dynamic information changes. Such problems are of increasing interest in an *Information Society* that is rapidly building a global infrastructure of communication networks and the facilities to access and transmit data from one node to any other node. Through this infrastructure an individual person who is in possession of a surprisingly inexpensive access device has the potential to extend his or her sphere of activity beyond any geographical boundaries.

While these communication capabilities present unprecedented opportunities for the smallest business to propel itself into world-wide markets, it also forces that same business to come to terms with complex planning and execution situations. The increased complexity is related to the larger number and kinds of interrelationships that exist both within the problem situation and tie major components of the problem space to external factors. For example, the initiation and coordination of a business venture that involves the export of multiple products that are being

designed in one country, manufactured in another country, and marketed in several countries, is a complex undertaking that can certainly be carried out by a single person with access to the Internet and other global communication facilities. It will require a strategic planning and execution process that must consider not only the interplay of the relevant technical variables, but also many economic and cultural factors that may have a profound impact on each other as well as the technical variables. Less than 20 years ago such undertakings were the sole province of large corporations and government agencies that could afford to bring substantial human and material resources and time to bear on their solution. Today, very small groups or even single persons have the ability to contrive more and more convoluted and complicated ventures, and at the same time hold fast to the expectation that they will be highly successful in a relatively short period of time.

It can be argued that the *Information Society* is still in its earliest formative stages. While it has provided global person to person and data access, it has provided only the most primitive tools for converting these data into information and knowledge. Contrary to human cognitive processes which are built on powerful symbolic reasoning capabilities, our global computer networks are still largely limited to the transmission of text streams and similar data elements whose meaning is understood only by the human sender and receiver. In other words, current communication networks are designed for the more or less optimum transmission of data between and among nodes, on the assumption that these nodes are attended by intelligent human beings on whom rests the sole responsibility for understanding the information that is being conveyed by the data streams and recognizing the intricate relationship patterns that are embedded in the data. If we add to this the time-critical nature of many decision situations, whether life threatening (e.g., crisis management, warfighting, medicine) or subject to financial imperatives (e.g., economic survival and profit in business planning and management), then it is not surprising that human decision makers are increasingly complaining about an 'information overload'.

If the *Information Society* is to deliver on its dual promises of meaningful global information sharing and intelligent decision-support (Fig. 1), then its infrastructure has to become much more than a message passing facility. What is needed is the ability to pass information and knowledge in a form that allows computer-based components within the communication environment to apply at least elementary reasoning processes. In other words, the communication infrastructure must have some *understanding* of the information it is transmitting. Of course it can be argued that computers, and therefore communication systems because neither is possible without the other, cannot have an *understanding* in the human sense of the word (Penrose 1989, Dreyfus 1997). Similar arguments are often applied with equal conviction to the words *learning* and *intelligence*.

This author would humbly submit that while these arguments may be perfectly valid in a theoretical sense, they are not intended to imply that computers **cannot** *learn* anything, *understand* anything, or have any semblance of *intelligence*. This is obviously a matter of degree and definition. If we mean by *learning* the apparently unlimited human abilities to abstract knowledge and to apply vast intuitive and creative processes to problem solving then it would indeed be difficult to believe that current computer technology directions could ever reach that level. If, on the other hand, we accept that there are levels of *learning* and that the ability to recall and effectively apply a previous solution to a current problem situation, and to detect

problem and solution patterns, are a kind of *learning* then it is certainly conceivable that computers could have some level of *understanding* of the information that they are processing.

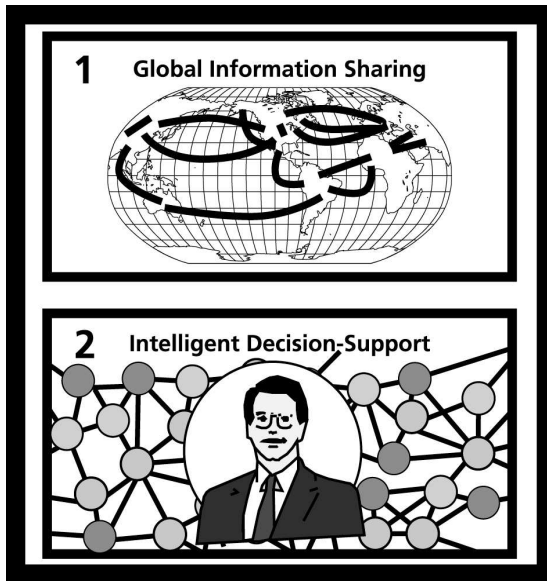


Fig. 1: Promises of the Information Society




Fig.2: What computers cannot do easily!

The early expectation of the Artificial Intelligence (AI) field was that computer systems would be able to deal with *commonsense*. It was postulated that everyday knowledge could be represented symbolically as a set of formal rules that can be manipulated without intuition and understanding. This expectation turned out to be overly optimistic. What human beings often refer to as ‘simple commonsense’ is in fact not at all simple, nor is it necessarily logical. The problem surfaced in the early 1970s in attempts to develop computer programs that could understand children’s stories (Goldstein and Papert 1975). Difficulties arose when it became necessary to relate a particular pronoun (e.g., ‘it’) to one of several previous nouns (e.g., ‘dog’ or ‘ball’) in a sequence of sentences. For example (Fig.2): *When Billy looked over the fence he saw one of Mr. Jones’ puppy dogs playing with a red ball. He wanted it so much.* To determine whether Billy wanted the puppy dog or the red ball requires a deeper knowledge of Billy’s beliefs that cannot be represented in a manageable set of rules nor with a less than exhaustive number of facts.

In our work in the Collaborative Agent Design Research Center over the past decade or so, we have found that many useful functions can be performed by computer-based decision-support systems that do not pretend to deal with the *commonsense* problem. Such functions are not limited to the monitoring of events and the evaluation of the consequences of courses of action, but can extend to the development and implementation of action plans within a collaborative multi-agent environment. We have found that lower levels of *understanding*, *learning*, and *intelligence* can be achieved in computer systems through the use of internal symbolic representation and the interaction of human users. In such collaborative systems computer-based reasoning facilities continuously monitor, evaluate, propose, and communicate events as they assist each other and the human users, while the latter assist the computer-based facilities to

interpret and understand whenever such interpretation and understanding cannot be obtained through reasoning alone.

The design and implementation of such collaborative systems has directed our attention to the representation of information and knowledge in a form that will support the necessary interaction between the human users and the computer-based assistance components. We have found that a high level internal representation of the real world objects that are pertinent to a particular application environment is an absolute pre-requisite for achieving a useful level of *intelligent* assistance in decision-support systems. Over the years, as we have moved from exploratory feasibility studies, to proof-of-concept models, to prototypes, to large-scale in-service production applications, we have gained confidence in our conviction that a discernible level of internal *understanding* can be achieved in such systems. This has increasingly focused our interest on gaining some understanding of the manner in which an *intelligent* communications infrastructure can and will evolve over the next few decades into what is popularly referred to as *Cyberspace*.

What is *Cyberspace*?

Cyberspace is a loosely coined term that describes an emerging environment in which much, if not most, information about physical objects such as manufactured products, buildings, processes, organizations, artifacts, human beings, and dialog between human beings, is accessible on-line through computer-based communication systems. Furthermore, it is expected that in *Cyberspace* both the user and the environment itself will be able to analyze, synthesize and evaluate information in a virtual reality environment that couples to the human senses through, at the very least, sound, speech synthesis, three-dimensional space, and animation.

A *Cyberspace* environment should allow its human users to immerse themselves into a virtual world that both emulates and simulates the real world to the degree necessary for these users to participate in certain activities without the need for physical presence. This does not necessarily require all five senses to be simulated. Sight and hearing are of critical importance, while taste, smell and touch are desirable but not essential. In *Cyberspace* we would expect the virtual environment to couple to the users with body nets that directly access the human senses. It is conceivable that such body nets would be able to detect and influence emotional changes, automatically store and recall what its wearer hears and sees, and continuously assist in the analysis and synthesis of this information as well as the formulation and evaluation of courses of action. In this respect body nets serve not only as receivers and recorders of information that reaches the human users through their *Cyberspace* access nodes, but also as intelligent personal assistants that continuously analyze, synthesize and relate the incoming information to the interests of its wearer.

In *Cyberspace*, whether through body nets or other user-interface facilities, human users should be able to have their interests represented by software agents that can be tasked to navigate the virtual environment with either general or specific objectives. Such objectives might include: searches for specific information; establishing contacts with other agents or users; offering services both for-profit and not-for-profit; requesting help and assistance; detecting patterns of activity; monitoring the behavior of users and other agents; intelligence gathering; and so on.

For example, it is certainly in the realm of possibilities for a *Cyberspace* user to own one or more expert agents that are capable of providing services to other users or agents for monetary profit. These agents could be launched into *Cyberspace* and advertise their services in much the same

way that human service providers have done in the real world for hundreds of years. They would collect their fees and credit these to the bank account of their human owner without the need for the latter to be actively involved in any of their interactions and transactions. This rather obvious and simple example demonstrates the enormous potential for leveraging the capabilities of a single human user within a *Cyberspace* environment. It also immediately identifies the high value that such an environment will place on the acquisition and exploitation of intellectual skills and knowledge. Limited only by the bounds of human intelligence and motivation, it would allow its users to create virtual worlds of business, industry, education, and government, with both virtual and real products (Fig.3). Clearly, the progressive realization of these prospects will precipitate adjustments and the imposition of controls that will fundamentally change the nature of human society.

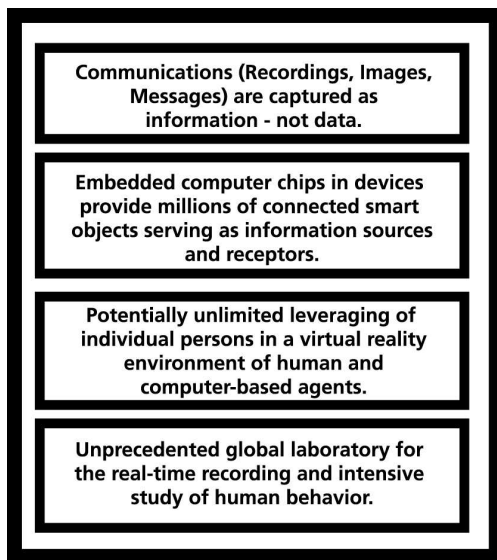


Fig.3: Essential features of *Cyberspace*

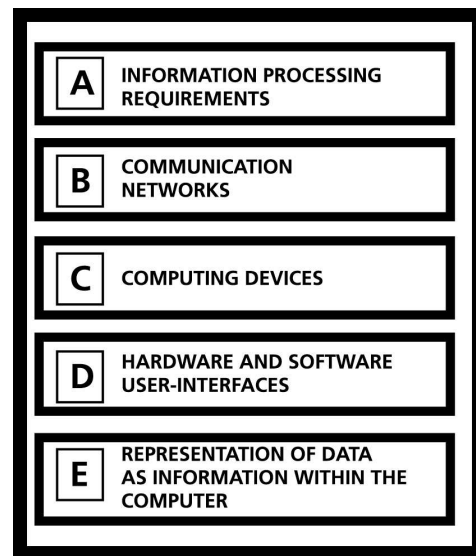


Fig.4: Principal components of *Cyberspace*

For example, the notion of *presence* will be decisively redefined in a *Cyberspace* society. In recent years we have seen the gradual acceptance of a new concept of *presence* that does not require the physical relocation of persons. Major sporting events and entertainment shows are more conveniently viewed on television from the home. Typically, in the case of sporting events, the quality of the televised presentation of the competition is greatly improved by technical enhancements such as instant replays and close-up shots of particularly skillful maneuvers, explanations and analyses by informed commentators, and short profile films of the best competitors.

Electronic mail, Internet chat groups, telephone and video conferencing facilities, and facsimile (FAX) transmissions, have reduced the need for face-to-face meetings. Commercial companies are gradually being forced to reassess the need for a centralized workplace. Why pay the considerable overhead costs associated with maintaining office space for employees, if the employees could equally well perform their work at home? Computer-based messaging services and global connectivity have already reached a level of reliability and convenience that is more than adequate for business communications.

Nevertheless, the passenger transportation industry has experienced unprecedented growth over the past two decades. This is not a paradox, but rather a symptom of two concurrent forces. Firstly, the increased expectations of individuals for an improved quality of life together with greater control over working hours has generated a burgeoning recreational services market. Secondly, it will probably take at least another decade or two of technological development in hardware and software user-interfaces before *Cyberspace* meetings will be able to effectively compete with physical *presence*.

High Presence	Low Presence
High cost. Fixed location. Fixed performance schedule. Dramatic social event. Strong sense of occasion. Direct active contact in context.	Much lower cost. Flexible location. Flexible performance schedule. Convenient private event. Weaker sense of occasion. Decontextualized passive contact.

Table 1: Economy of *presence* (adapted from: Mitchell and Strimpel 1997)

A similar situation also still exists in the late 1990s in the performing arts industry. Here the physical *presence* during the live performance of a play or concert can add a quality of drama and excitement that cannot currently be preserved in electronic transmissions to remote sites. While it is likely that there will always be individual persons willing to sacrifice time and convenience for the emotional experience of attending a live performance, the vast majority of persons will be increasingly satisfied with the *Cyberspace* equivalent of this experience. In this context Mitchell and Strimpel (1997) stress the economical forces that define the distinctions that drive the choice between “high presence” and “low presence” forms of theatrical performance (Table 1).

To explore the feasibility of *Cyberspace* and estimate when such a vision might become reality we need to examine at least the following five essential components or building blocks of a *Cyberspace* environment (Fig.4): information processing requirements; communication networks; computing devices (i.e., platforms); hardware and software user-interfaces; and, the meaningful representation of information within the *Cyberspace* environment.

1. Information Processing Requirements

Before considering the current status and projected advances in computing and communications technology that are expected to support a *Cyberspace* environment it might be appropriate to estimate the data rates and storage requirements that might be generated in such an environment. Taking the data production values proposed by Bell and Gray (1997) for a single human being as a convenient unit of measurement and assuming an active day of 16 hours for seven days per week over a lifetime of 75 years, Table 2 shows the approximate storage requirements for various data types per hour, per day, per week, per year, and per lifetime.

Since the extrapolated values (Table 2) assume a continuous generation of each data type they need to be adjusted for the kind of activity pattern that we might expect during a typical day for a

person living in a global *Cyberspace* community. As expected, the requirements for video far exceed the combined requirements for speech and text. However, while the figure of five peta bytes constitutes almost 95% of the sum of all of the human generated data, if it were to be generated concurrently, it cannot necessarily be taken as a satisfactory estimate of the total data storage requirements (Figs.5 and 6).

Data Type	Data Rate		Approximate Storage Requirements			
	(per sec)	(per hour)	(per day)	(per week)	(per year)	(per lifetime)
speech text (@ 120 wpm)	12 B/sec	43 KB	700 KB	5 MB	2 GB	128 GB
read text (few pictures)	50 B/sec	180 KB	3 MB	20 MB	7 GB	537 GB
video (compressed)	500,000 B/sec	1800 MB	29 GB	202 GB	72 TB	5 PB

Table 2: Approximate data rates and storage requirements per person.

The question that has to be addressed is whether the *Cyberspace* community will be intent on recording the daily activities of its members on film (i.e., video). If yes, then the estimate for video storage in Table 2 will be in addition to the storage of all other forms of data and could well be closer to seven or eight peta bytes per person. It is not difficult to find arguments in support of the automatic monitoring of spaces and persons in a technically advanced society. There are many situations in which the personal security, safety and well being of individuals would be greatly enhanced by continuous visual contact. Examples that are already apparent in today's world include the ability of parents to watch their children in day-nurseries through Internet 'web-browsers' that transmit live images from video cameras mounted in the day-nursery to the parent's computer screen. Retail shops routinely use video cameras to detect thieves. Likewise many public spaces are continuously filmed to maintain security and capture evidence in case a crime is committed.

In a society where the costs associated with the recording, transmission and storage of information are trivial there is likely to be an emphasis on the collection and analysis of data. *Cyberspace* will seduce society to study itself more intensely than in any previous time period, and to archive its day-to-day behavior for analysis and as a testament of its endeavors and accomplishments. Increasing expectations of quality and efficiency will drive the quest for greater knowledge and understanding of human behavior and the interaction of society with the natural and built environment. In this regard, *Cyberspace* could easily become a real life experimental laboratory of human psychology on an unprecedented global scale. A laboratory in which the collection and analysis of data is undertaken on a continuous and real-time basis for purposes of improved productivity, energy conservation, political sampling, optimized marketing, and better understanding of human behavior.

From a practical point of view the information storage requirements of the *Information Society* are unlikely to explode to such unprecedented levels overnight. Rather, they will increase gradually as the *Cyberspace* access capabilities spread from the technologically more advanced nations to Third World countries. While the number of Internet users has increased dramatically to about 60 million (1997) over the past decade it is still only about 1% of the world population of just over six billion (1997) or just over 16% (31 million) of the adult (18 years and older) population in the US . If the telephone network with its current 650 to 700 million world-wide

terminations (Cerf 1997) can serve as a near term guide for Internet growth, then we could expect perhaps 10% of the world's population (i.e., 800 million) to be active participants in the evolving *Cyberspace* environment within 10 to 15 years. Relatively few of these projected users are likely to require anywhere near the five to eight peta bytes of estimated lifetime storage (Table 2) within the next 20 years.

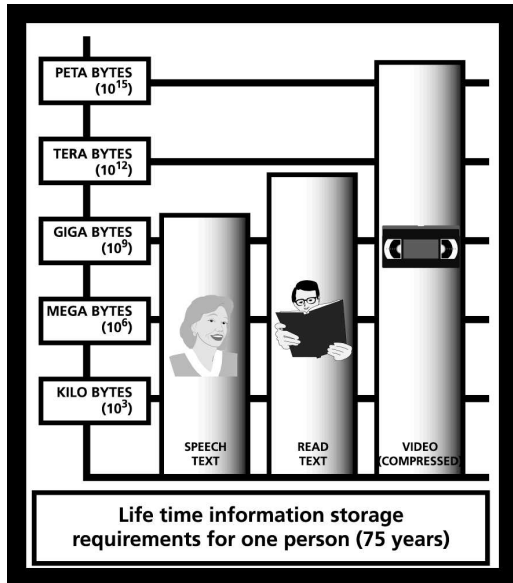


Fig.5: Lifetime information storage requirements for one person (75 years)

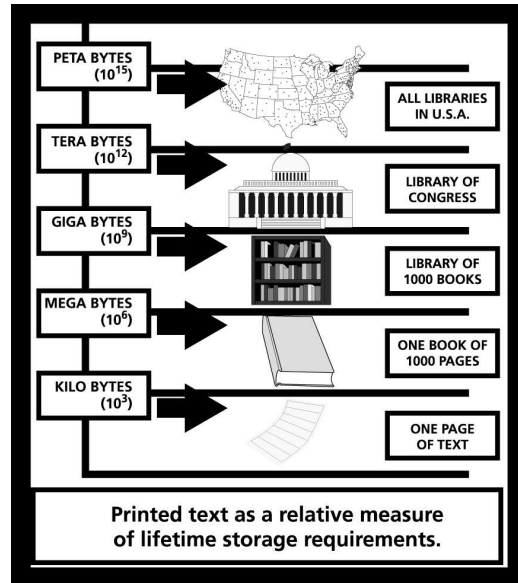


Fig.6: Printed text as a relative measure of lifetime storage requirements

2. Communication Networks

The backbone of Cyberspace is networking technology that allows computers to communicate with one another. At least four interconnecting network types (Figs.7 and 8) will be required in support of a Cyberspace environment: wireless networks that connect mobile and portable access devices; local area networks that connect systems within limited site boundaries; local loops that connect central distribution points to users (e.g., telephone and television); and, wide area networks that connect thousands of central distribution points. These network types do not necessarily constitute hierarchical levels of network connectivity. Although wide area networks are generally long-haul networks that provide gateway connections to thousands of more localized and customized networks, the nature of a network is more accurately defined by its functional purpose and its throughput capacity.

Wireless networks are needed to provide Cyberspace access to transient users that are either mobile or remote from fixed local area network or local loop access points. Cellular telephones are a rapidly expanding example of wireless networks, while remote television and VCR control units are examples of convenient point-to-point wireless connections. Wireless networks rely on radio and infrared technologies, and have the potential for completely changing the communications infrastructure. Not only could they support truly portable and mobile computing, but in combination with global positioning systems (GPS) they would be able to automatically and accurately determine the geographical location of each access point and user.

The realization of this potential would probably require a reallocation of the radio spectrum to reserve the limited bandwidth for communication applications that would benefit most from wireless access. In other words, fixed access devices such as television sets would be largely excluded from wireless communication. Theoretically, existing radio frequency bands are capable of providing capacities from 500 Mbps to 50,000 Mbps (i.e., 50 Gbps). In practice, the actual capacity depends on the geographical cell size that allows space-sharing (Bell and Gray 1997).

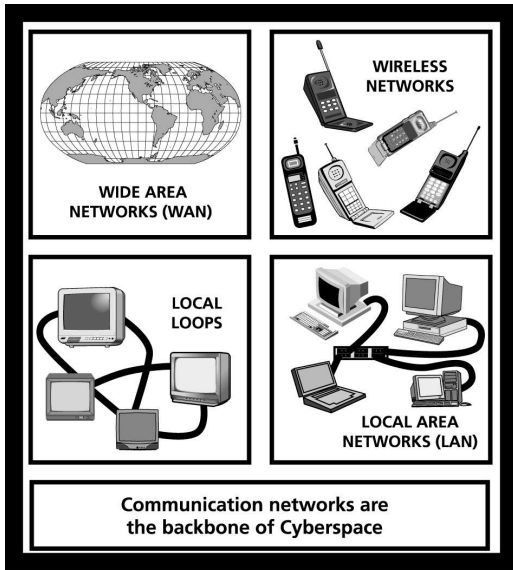


Fig.7: Communication network types

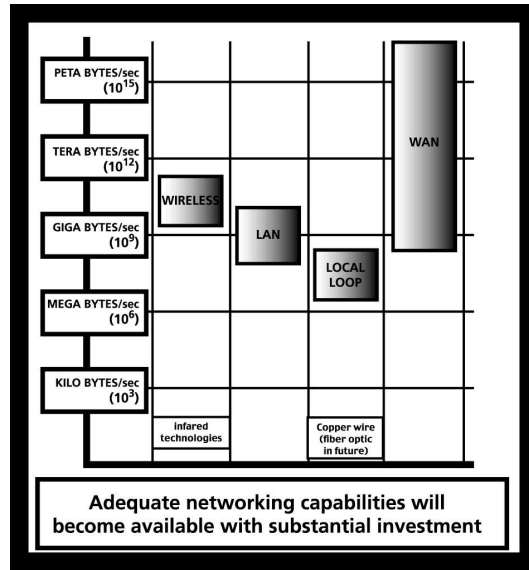


Fig.8: Bandwidth projections

Local area networks (LAN) are typically used to connect computer-based platforms (i.e., computers and appliances with embedded computers) within a local site such as a home or an office. Local area network bandwidth has approximately doubled every three years during the past 20 years. When Ethernet was first introduced in the early 1980s it operated at 10 Mbps (mega (or million) bits per second). It increased to 100 Mbps in 1994 and 1000 Mbps (i.e., 1 Gbps) in 1997.

Local loops are widely used to connect central distribution offices such as telephone exchanges and cable television stations to local subscribers. Currently, most of these networks utilize copper wires and constitute a bottleneck that extends to approximately a four-mile radius. Local loop networks are greatly hampered by monopoly and regulation. Over the next two decades the transmission rate of copper wires might reach 20 Mbps, which could support high-resolution video. More likely, with increasing deregulation much of the local loop infrastructure will be replaced with fiber-optic and wireless transmission, as we move toward the goal of a 'single dial tone' communication system that supports voice, video and data in the same network. The realization of this goal, however, is unlikely to be achieved before the year 2010 (Bell and Gemmell 1996).

Wide area networks (WAN) connect localized networks to each other, to ultimately form a global network of interconnectivity. The transmission capacity of wide area networks is expected to increase to around 1000 Gbps (i.e., 1 tera bits per second) by the year 2010. The deregulation of

wide and local area networks has greatly accelerated their development in recent years, in contrast to local loops that are still largely monopolized and regulated.

While the expectations of achieving at least adequate networking capabilities in support of a *Cyberspace* environment appear to be quite promising, the realization of these expectations will necessitate a substantial investment by industry and government. However, this investment is well justified by the potential profits that are likely to accrue to the investors.

3. Computing Devices

Computers can be described as universal machines that can be built on top of each other in hierarchical layers to produce more and more complex computing devices (Fig.9). At the lowest level of these virtual processor building blocks is the physical hardware, which is configured in terms of gates that translate electrical impulses into the most primitive components of logical sequences (e.g., AND, OR). These micro-code sequences represent the Instruction Set Architecture (ISA) of the computer and constitute the transition from hardware to software. Above this layer resides a software Operating System (OS) that serves as a control interface for all virtual processors and devices. Higher levels provide programming languages and similar software development tools. At the highest level are the applications that allow users to utilize the computer system for the practical purposes of commerce, industry, education, and personal entertainment.

Early computers took advantage of the ability of one computer to simulate several virtual computers by timesharing a single computer to many users. This was an effective strategy for sharing the considerable cost of even one computer among many users. Today, with the relatively low cost of computers, users rarely share the same computer unless it has special capabilities or facilities that are not readily available on other machines.

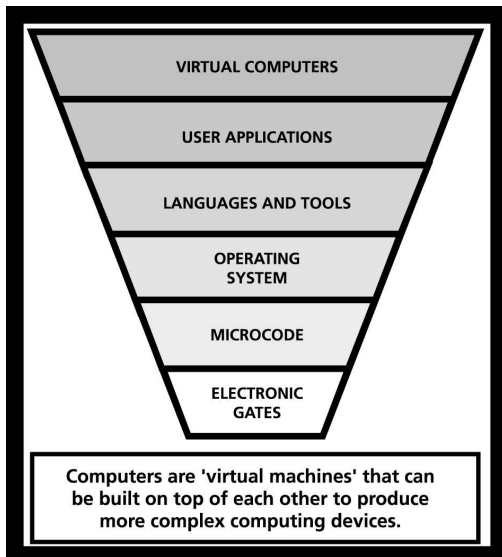


Fig.9: Computing devices

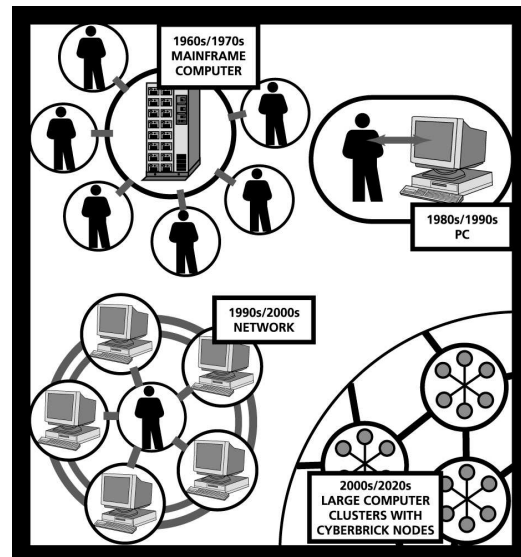


Fig. 10: Computer clusters

With the availability of many low cost computers the converse of shared machines is a much better strategy. An almost unlimited number of computers can be seamlessly linked together as

an integrated system with far more computing power than any single computer. Distributed operating systems are available to connect many powerful independent computers into a cluster, with high performance System Area Networks (SAN). These clusters (Fig. 10), using the spare processing capacity and storage space of its nodes to provide a degree of redundancy, become the server nodes of intranets which in turn connect to form a global internet. It has been predicted (Semetech 1994) that by the year 2010 the nodes of clusters will serve as 15 Gips (giga instructions per second) *cyberbricks* with 30 GB (gigabytes) of memory (Gray 1996). Already today, in 1998, clusters of hundreds of computers have replaced single supercomputers for processing very large commercial databases and for computing intensive scientific applications.

As shown in Fig. 11 computer processing speed has been increasing by a fairly constant factor of about 10 every five years. It is expected that this rate will continue, or increase slightly, over the next 50 years (Bell and Gray 1997). The most important technological gains in computing have occurred in semiconductor circuit density and magnetic information storage density. These gains have allowed computers to operate faster at lower cost, because the smaller the distances the faster the instructions will travel within the computer, and the smaller the electronic components the greater the production volume.

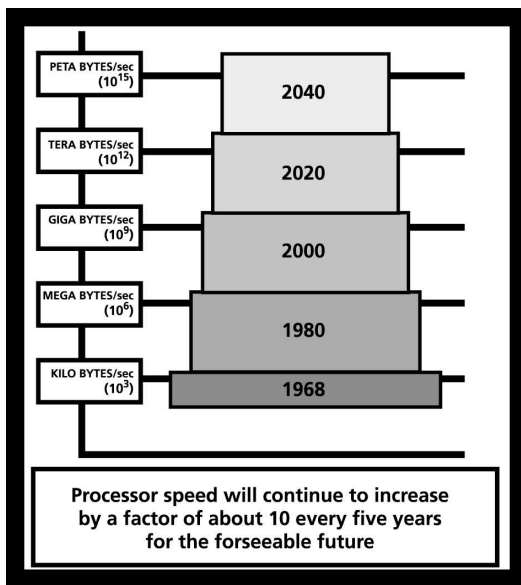


Fig. 11: Processor speed projections

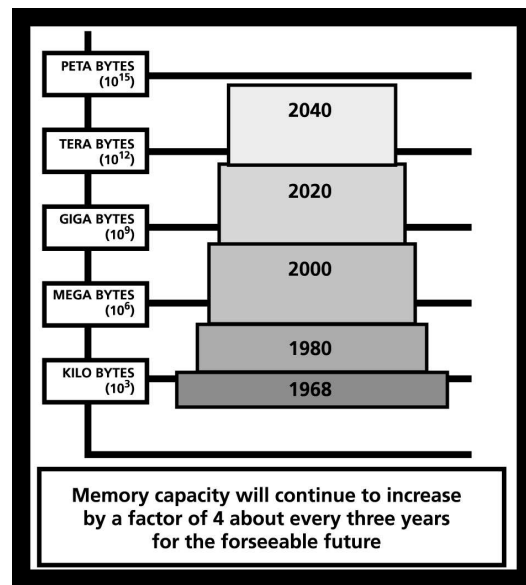


Fig.12: Storage density projections

Since 1960, semiconductor memory capacity (Fig.12) has increased nearly fourfold every three years (Moore 1996). Within a computer, Dynamic Random Access Memory (DRAM) holds current data being processed by the Central Processing Unit (CPU) and typically operates at about the same speed as a processor. Static Random Access Memory (SRAM) serves as a cache and holds recently used data that may have come from slower secondary memory, such as a disk drive. Magnetic disks with access times in milliseconds are often used for the storage of larger files and databases. Electro-optical storage devices (e.g., tapes and some disks) with access times in seconds are used for archiving data.

Mass storage capacity per unit area for disk drives has approximately doubled every two years since the 1960s when IBM released its RAMAC-350 disk system with a storage density of just

over 100 bytes per square inch. In 1996 the same company released a disk drive with a storage density of 100,000,000 bytes (100 MB) per square inch. Electro-optical disk technology currently provides one giga byte (1 GB) of storage at the cost of a compact disk (CD). While CDs can currently store about 400,000 pages of text or 10,000 frames of video images, this storage density is expected to increase non-linearly to around 20 tera bytes (20 TB) in 2050 (Bell and Gray 1997).

If we relate these secondary memory densities to the potential per person information storage needs projected in an earlier section of this paper, then we might expect *Cyberspace* participants in the year 2050 to archive a video record of their daily activities over one entire year (i.e., about 70 TB) on no more than four CDs. Alternatively, if the equivalent amount of information were required to be stored dynamically with more rapid access speeds then this would likely require less than one square inch of the effective storage area of a disk drive. Clearly, under these conditions the information storage requirements of a *Cyberspace* society are well within the realm of even the most pessimistic technological predictions.

4. Hardware and Software User-Interfaces

According to Bell (1975, 1997) new computer classes or families have appeared every 10 to 15 years. Table 3 includes among past, present and expected future computer families: mainframes; minicomputers; microprocessors; Web computers; body nets; and, robots.

As mentioned earlier in this paper much remains to be accomplished in the development of hardware and software interfaces that will allow users to participate effectively and conveniently in a *Cyberspace* environment. Most promising in this regard is the concept of embedding computer chips in all kinds of devices and in body nets. In this way appliances, books, pictures, toys, and similar commodities become *smart objects* with the ability to interact with each other and users. Body nets, on the other hand, provide a direct interface to the human senses for purposes of personal health care, control, assistance (e.g., prosthetic devices), security, communication, and the general enhancement of human functions.

Weiser and Brown (1997) point to the emergence of *ubiquitous computing* as a third phase in the evolution of computer systems, following the mainframe and the PC-based distributed computing eras. We are already witnessing the early formative stages of *ubiquitous computing* with the increasing number of small microprocessor chips that are controlling the smallest appliances and the largest conveyances. Today, these embedded computer chips are still largely stand-alone processors. Once they are linked together into integrated, local monitoring and controlling *systems* and these in turn are connected via local and wide area networks, then we will have truly entered the realm of *ubiquitous computing*. The implications of *ubiquitous computing* are enormous, representing the availability of millions of connected interfaces that can serve equally well as information sources and information delivery nodes. Regardless of whether these devices are components of body nets or the integrated, intelligent control systems of homes they will serve to multiply human capabilities to a degree comparable with the invention of the printing press.

How will the development of software evolve to keep pace with these *ubiquitous computing* trends? Since the emergence of the first electronic computing devices, some 50 years ago, the programming of computers has evolved from 'plug boards' to low level Assembler language, to higher level languages such as Fortran, Cobol, C and C++, and increasingly to a later generation

of tools that allow the description of a result rather than the step by step generation of a solution. The emphasis has been gradually shifting from providing the programmer with tools to develop programs, to providing the user with tools to interact directly with the computer to solve problems. In this respect, we are essentially expanding the notion of ‘programmer’ to potentially include anyone using the computer. Rather than specifying the steps for solving a problem, we are providing a set of tools and a functional description of how these tools can be applied to gain a better understanding of the dynamically changing problem situation and progressively implement the most promising solution strategy.

Year or Period	Computing Devices	Significant New Technologies	User-Interfaces
1950s - 1960s	mainframe computer	transistor, core, drum magnetic tape	punched cards, tape
1960s - 1970s	minicomputer	integrated circuits, disk drive	teletype, video display monitor
1070s - 1980s	microprocessor (PC) workstations	diskette drive	WIMP interfaces
Mid-1990s	WWW (World Wide Web) access via PC workstation	server	browser
Late-1990s	Web computer, SNAP (Scaleable Network and Platforms), SAN (System Area Network) for clusters	client software from server using JAVA; Active X, etc. server provisioning	Video-phone TV access to WWW, speech synthesis
2000s	PDA (Personal Digital Assistant), computer control by speech	hand-held devices	speech as primary data type
2010s	one information dial tone	video-capable devices of all types	video as primary data type
2020s	embedded speech and vision functions		books/pictures that identify themselves
2030s	body nets	implanted sensors	vision, speech and gesture control
2040s	robots for home and workplace	appliances with robotic capabilities	mobility, vision and speech

Table 3: Past, present and future computer families (adapted from: Bell and Gray 1997)

Clearly, we are moving from programming to problem solving. One way of characterizing problem solving is as the incremental process of gaining an understanding of the interactions among the components of a problem. Whereas in simple problems these interactions are

progressively clarified as a more or less automatic by-product of our solution of the components, this is typically not true for complex problems. In complex problems the complexity is a direct result of the interactions (i.e., relationships) among the problem components, and these relationships are often indirect, usually dynamic, and always multi-faceted. A reasonable course of action in such problem situations is to try and simplify the components of the problem as a means of gaining a better understanding of the relationships among them.

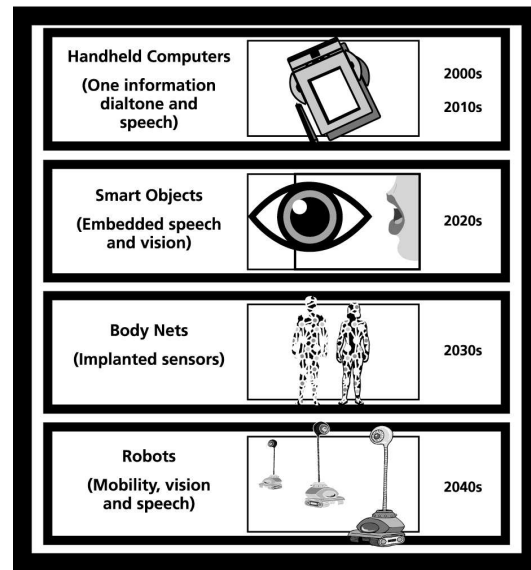



Fig. 13: Computer families and user-interfaces (1950s to 1990s)

Fig. 14: Computer families and user-interfaces (2000s to 2040s)

The traditional methods for simplifying a problem would have us predefine both the problem conditions and the internal relationships. In this way we are able to exercise a high degree of control on, not only the nature of the interactions, but also on the conditions under which such interactions might occur. Unfortunately, such methods often lead to a gross oversimplification of the problem and an outcome that has few if any of the ingredients of a solution. A better approach is to treat the problem components as contributors to an evolving problem solution that continuously adapts to changes in the problem environment. The objective is to simplify the components, but not the problem. Rather than predefine the relationships, the approach focuses on predefining the capabilities and expertise of the components. In this respect the components can be characterized as simple tools that make their contributions opportunistically based on the current state of the problem. This presupposes that each tool has some understanding of the problem conditions under which it can and should contribute; namely, the ability to interpret current problem conditions and compare these with a set of objectives (i.e., trigger conditions).

The cooperating components constitute, in effect, a small society in which their interactions are governed by rules and some common set of objectives. Conceptually, this is not unlike a human society, although there are two major differences within the constraints of current technological capabilities. Firstly, there are many limitations due to practical implementation obstacles, such as *representation* (i.e., How to represent information within the system so that the system

components can reason about the information?), *knowledge* (i.e., How to provide each component with a rich knowledge base to draw on for its reasoning activities?), and *goals* (i.e., How to convey to each component the goals and objectives (i.e., the intent) of the problem situation?). Secondly, there is an absence of free will, motivation, and emotion, both in respect to the individual components and the collective sense of the system. While, on the positive side, this makes for a more predictable and controllable problem solving environment, on the negative side, it greatly reduces the ability of the system components to make much headway when there is little information available.

In support of this notion of many cooperating parts we are beginning to see a major shift in the nature of computing systems, from the predetermined design of a software system to the addition of components to a loosely defined network of nodes. The components are capable of communicating, as well as requesting and contributing information. In this regard, they behave as agents with capabilities that are matched to their function and purpose. Again, in general terms, this kind of behavior is not dissimilar to the behavior of human beings. The principal differences between computer-based agents and human agents lie in the degree of useful autonomy that they can apply to the situation at hand. Within the bounds of current and near-term technological capabilities, computer-based agents are limited in their useful autonomy to tasks that are either logical derivations of existing knowledge or largely predetermined solutions (or solution approaches) to novel situations. In either case, the actions of the agent are based on logical inferences that are both predictable and reproducible. The useful autonomy of human agents, on the other hand, extends well beyond the limits of logical reasoning into areas of intuition and creativity that are neither predictable nor reproducible. Whether the available information is sufficient or insufficient for a logical solution approach, the decisions of the human agent will always be at least partially governed by an inert feeling about the problem situation that is based on an emotional disposition rather than a factual analysis (Fig. 15).

5. Meaningful Information Representation

Although technological advances in computer hardware and communication systems have been truly astounding over the past 20 years, the utilization of these advances has been less than remarkable. The fact is that we are still using computers largely as data processing devices that perform only the most menial and least intelligent data transmission and manipulation tasks. While computers are performing these tasks with great speed and accuracy, and while they are able to provide connectivity among a virtually unlimited number of access points, the higher level and much more rewarding tasks of analyzing, interpreting and abstracting data as information and knowledge is almost entirely left to the human users (Fig. 16).

This serious deficiency has become increasingly apparent as technological advances have increased computing power, data storage capacities, and data transmission speeds by orders of magnitude in such a short period of time. Convenient global access to users and data has increased the need for information filtering, so that individuals might take advantage of the opportunities for material and personal profit that this connectivity and processing power present to the user. Needless to say, the capabilities of a computer to assist in the intelligent assessment of information are basically non-existent if the computer processes this information as bitmaps and alphanumeric text strings. Surely, a truly Cyberspace environment carries with it the expectation that information is held within the system environment in a representational form

that is, if not equivalent to, at least compatible with human cognition.

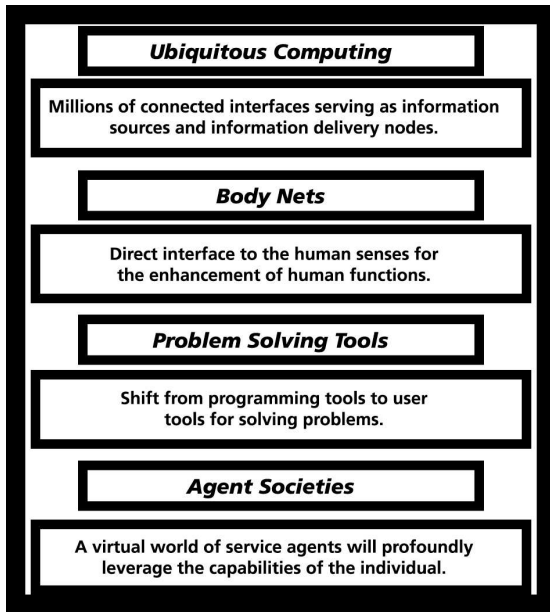


Fig.15: Evolving computer-human partnership

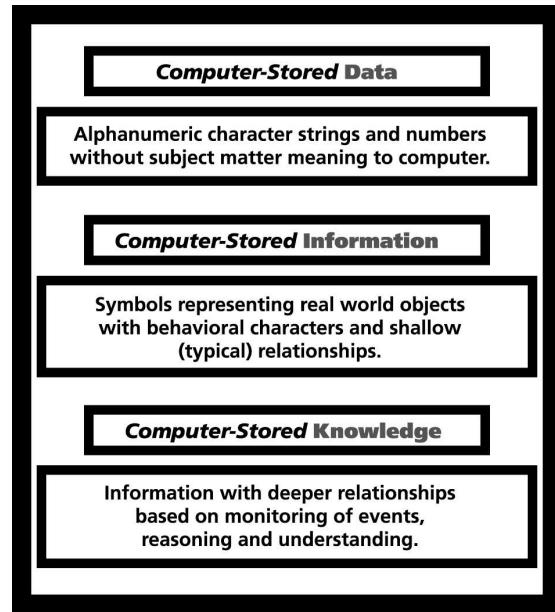


Fig. 16: Data-information-knowledge

The current approach for achieving this objective is to represent information in the computer as objects with behavioral characteristics and associations to other objects (Myers et al. 1993). While this approach is hardly sophisticated it does allow real world objects (e.g., car, building, chair) to be represented symbolically so that computer software modules can reason about them. Abstract concepts such as privacy, beauty and power, are less amenable to this approach since their meaning and role in our day-to-day activities is less easily defined. For example, the characteristics of *privacy* are neither static nor can they be accurately described in relational terms. They depend on a wide range of factors that relate to both environmental and personal circumstances and dispositions.

Nevertheless, even this relatively weak form of representation of real world objects can provide the basis of usable problem solving support and decision making assistance. Improvements are possible with the addition of knowledge bases and user interaction. In the latter case the user becomes as much a helper to the system as the system serves as an assistant to the user. However, this occurs in quite different ways. The system uses its computing and logical reasoning capabilities to monitor, analyze and evaluate the actions, requests and interests of the user in an opportunistic manner. The user, on the other hand, helps the system to understand the nature of the objects and relationships that it is processing in a more deliberate manner (Pohl 1995).

Conclusion

How long will it take for a *Cyberspace* environment with substantial global information sharing, intelligent decision-support, and virtual reality capabilities to evolve? Certainly, we are far removed from that state today, in 1998. The most serious current deficiencies exist in two critical areas: information representation; and, user-interfaces. We are still only in the transition stage

between the First Wave and the Second Wave of software (Pohl 1996). The evolution of Second Wave software will place increasing emphasis on the high level internal representation of information, in terms of real world objects with behavioral characteristics and dynamically generated relationships. The need for computer-based systems to have some *understanding* of the information that they are processing is being recognized only gradually, mostly in response to the overwhelming increase in information availability, rather than as a result of an intellectual realization that representation is at the core of the *Cyberspace* vision.

While computer user-interfaces have come a long way from punched card beginnings in the 1960s to windows, icons, mouse, and pull-down menus (WIMP) in the middle 1990s, they are still relatively primitive and awkward. The type of virtual reality facilities that are expected to be available in *Cyberspace*, coupled to the human senses through body nets, still require significant technical advances. Speech synthesis is more mature, following two decades of development, and is starting to be used for various commercial applications such as airline reservation systems.

How long these necessary developments will take is difficult to predict. The field of computing is about 50 years old. During those 50 years, most of that which is of interest to us today has been developed over the past decade. Clearly, technical developments such as computing power, network capacities, and storage capabilities have been predicted with considerable accuracy for the past 20 years (Moore 1996, Semetech 1994). However, predictions relating to the application of computers in areas requiring cognitive capabilities have been very poor. For example, in 1950 Turing predicted that by the end of the Twentieth Century it would be possible to program a computer that could not be distinguished from a human being through its responses to questions (Turing 1950). In the 1960s it was believed that *thinking* machines would be a reality within two to three decades. These predictions have been far off the mark.

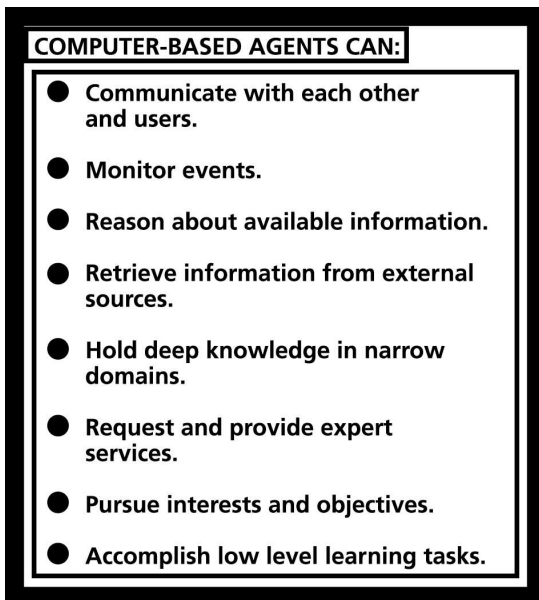


Fig.17: What computers can do!

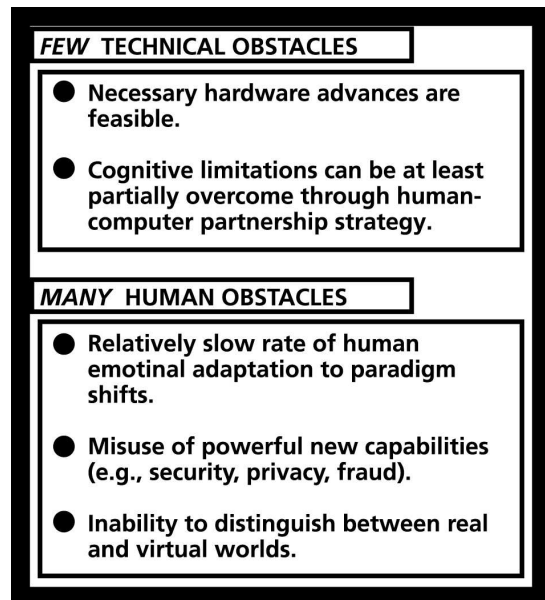


Fig.18: How long will it take?

It is the author's belief that the evolution of *Cyberspace* is likely to be governed at least as much by the ability of society to adapt to the human emotional responses evoked by an *ubiquitous*

computing environment, as by the need for technical advances that will make such an environment possible (Figs. 17 and 18). The degree to which the technical progress will surpass the necessary human emotional adjustment will be a significant determinant of the level of social upheaval that we are bound to experience over the next 50 years. *Let us hope for a smooth ride!*

References

Bell G. and J. Gemmell (1996); 'On-Ramp Prospects for the Information Superhighway Dream'; Communications of the ACM 39, No.7, (55-61).

Bell G. and J. Gray (1997); 'The Revolution Yet to Happen'; in Denning and Metcalf (eds.) 'Beyond Calculation: the Next 50 Years of Computing'; Springer Verlag, New York, (5-32).

Cerf V.G. (1997); 'When They're Everywhere'; in Denning and Metcalf (eds.) 'Beyond Calculation: the Next 50 Years of Computing'; Springer Verlag, New York, (33-42).

Denning P. and R. Metcalf (1997); 'Beyond Calculation: the Next 50 Years of Computing'; Springer Verlag, New York.

Dreyfus H.L.(1997); 'What Computers Still Can't Do: A Critique of Artificial Reason'; MIT Press, Cambridge, Massachusetts.

Frankston R. (1997); 'Beyond Limits'; in Denning and Metcalf (eds.) 'Beyond Calculation: the Next 50 Years of Computing'; Springer Verlag, New York, (43-57).

Gelernter D. (1997); 'The Logic of Dreams'; in Denning and Metcalf (eds.) 'Beyond Calculation: the Next 50 Years of Computing'; Springer Verlag, New York, (117-126).

Moore G. (1996); 'Nanometrics and Gigabucks - Moore on Moore's Law'; University Video Corporation Distinguished Lecture (<http://www.uvc.com>).

Myers L., J. Pohl, J. Cotton, J. Snyder, K. Pohl, S. Chien, S. Aly and T. Rodriguez (1993); 'Object Representation and the ICADS-Kemel Design'; CAD Research Center, Technical Report (CADRU-08-93), Cal Poly, San Luis Obispo, California.

Penrose R. (1989); 'The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics'; Penguin Books, New York.

Pohl J.G. (1995); 'The Representation Problem in CAD Systems: Solution Approaches'; in Pohl J. (ed.) Advances in Cooperative Computer-Assisted Environmental Design Systems, focus symposium: 8th International Conference on Systems Research, Informatics and Cybernetics, Baden-Baden, Germany, August 16-20.

Pohl J.G. (1996); 'Agents and their Role in Computer-Based Decision Support Systems'; 9th International Conference on Systems Research, Informatics, and Cybernetics, Baden-Baden, Germany, August 14-18.

InterSymp-98: Advances in Collaborative Decision-Support Systems for Design, Planning and Execution;
Baden-Baden, Germany, August 17-21, 1998 (pp. 9-28) *RESU63*

Semtech (1994); 'The National Semiconductor Roadmap'; Semtech and Semiconductor Industry Association.

Turing A.M. (1950); 'Computing Machinery and Intelligence'; *Mind*, 59(236); reprinted in Hofstadter D. and D. Dennett (eds.), *The Mind's I*, Basic Books, Harmondsworth Middlesex, England.

Weiser M. and J. Brown (1997); 'The Coming Age of Calm Technology'; in Denning and Metcalf (eds.) 'Beyond Calculation: the Next 50 Years of Computing'; Springer Verlag, New York, (75-91).

**Collaborative Agent Design Research Center (CADRC), California
Polytechnic State University (Cal Poly), San Luis Obispo, CA 93407
Telephone: (805) 756-2841/1310 FAX: (805)756-7568
E-mail: jpohl@calpoly.edu; Web-Site: <http://www.cadrc.calpoly.edu>**