INFORMATION TECHNOLOGY PLATFORM FOR THE 1990's

Stuart E. Madnick

August 1989 WP#3061-89-MS

INFORMATION TECHNOLOGY PLATFORM FOR THE 1990's

Stuart E. Madnick
Professor of Management Science / Information Technologies
5 August 1989

Theme

In surveying the trends that are most strongly influencing the evolution of information technology for the 1990's, a common theme has emerged from our research efforts, research sponsors, and related literature. This theme can be articulated as: "*Advances in information technology provide opportunities for dramatically increased connectivity, enabling new forms of interorganizational relationships and enhanced group productivity.*"

This theme emerges as the end result of two forces at work: 1990's business forces and 1990's information technology opportunities as depicted in Figure 1.

1990's Business Forces:

Although articulated many different ways, four business forces are exerting increasing influence and will, for many organizations, shape their destiny in the 1990's.

The first of these forces is the rapidly increasing growth in globalization, whereby the scope of organizations are expanding beyond their traditional geographic boundaries. In the case of current multinational firms, the entire corporation and all of its subsidiaries are finding needs to be increasingly coordinated to provide maximum impact, such as in manufacturing and supply activities, as well as in marketing and distribution. As a by-product of globalization, its inverse effect in the form of worldwide competition is also on the rise. This means that the number of competitors one must face in each marketplace and geographic region has increased

"ADVANCES IN INFORMATION TECHNOLOGY PROVIDE OPPORTUNITIES FOR DRAMATICALLY INCREASED CONNECTIVITY ENABLING NEW FORMS OF INTER-ORGANIZATIONAL RELATIONSHIPS AND ENHANCED GROUP PRODUCTIVITY."

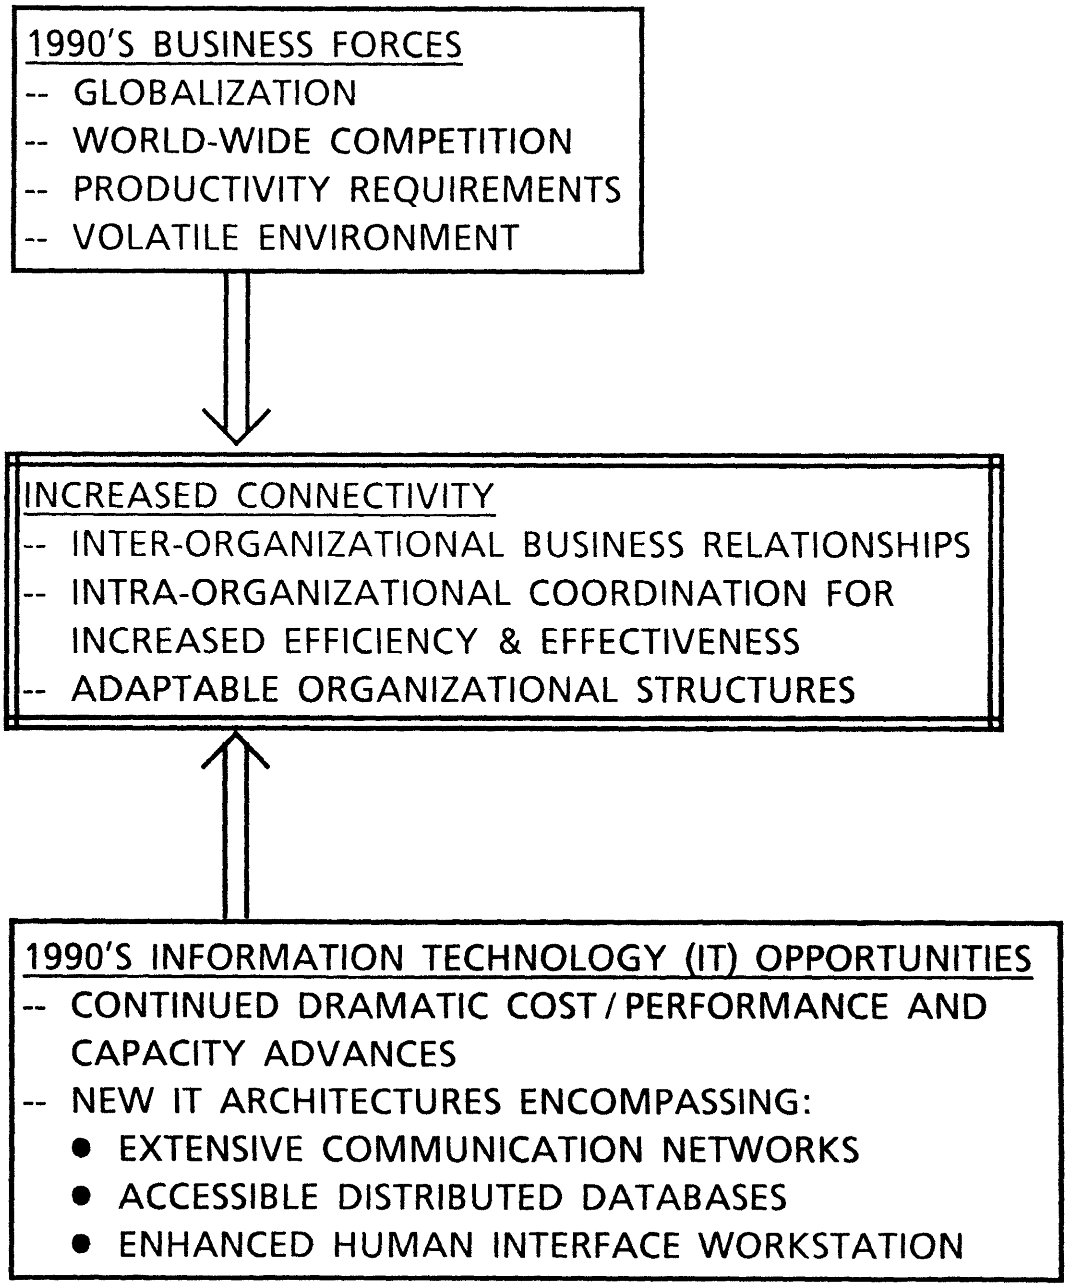

Figure 1

by the entry of corporations that are expanding through their globalization activities. This puts increased pressure on the established organizations in these marketplaces, and in many cases, changes the entire nature of the competition.

In order to seize the opportunity of globalization, and withstand the impact of worldwide competition, corporations are seeking ever increasing levels of productivity improvement. These productivity requirements take many forms. They may involve better coordinated manufacturing and purchasing so as to make maximum effectiveness in economies of scale and production, local efficiencies of labor force, and efficient purchasing and warehousing of components. Furthermore, by increasing the responsiveness to market trends and the requirements of customers, it is intended to gain sales volume and minimize wasted energy.

Attaining such improvements in productivity would be a significant challenge under normal circumstances, but the process is made significantly more complex by the volatile environment that is becoming more evident and is expected to continue to increase into the 1990's. This volatile environment emerges not only from the business forces described above, but also through various governmental, sociological, and legal changes. As one example, the various steps being taken to form a boundary-free European community, targeted for 1992, represent a major change to the business map of the world, with implications both for business in Europe and businesses dealing with Europe, as well as rippling effects throughout the world.

As a result of various governmental and legal actions encouraging deregulation and privatization, as well as desires for corporations to rapidly accelerate their globalizations efforts, massive mergers, acquisitions, and divestitures are taking place. Thus, at the same time that one is attempting to coordinate for maximum efficiency the various plans and resources of a corporation, it is not clear which of those resources will be there tomorrow. Thus, if a division is sold off or divested, it is necessary to adapt both marketing and manufacturing activities to replace any critical capabilities that were being provided by those facilities. On the other hand, if other organizations or divisions are acquired, it is desired to assimilate these facilities as rapidly as possible into the total environment, eliminate unnecessary redundancies and duplication, and produce as efficient a marketing and production capability as possible. These two forces clearly represent a major challenge for the manager of the 1990's. In a highly stable environment, traditional approaches to optimization and productivity improvements can be used. But, in this highly volatile and unpredictable environment, novel approaches to attaining high levels of productivity must be pursued.

In almost every major industry, these business forces can be identified, and in many cases, are fairly advanced even at this time. Independent from these forces, there are various developments occurring in information technology that have a strong bearing upon the situation.

1990's Information Technology Opportunities

In reviewing the trends in information technology with key executives of the major vendors, as well as research colleagues, various trends can clearly be identified that are expected to continue and even increase into the 1990's. Continued advancement of both cost reductions and performance improvements in information technology is almost taken for granted based upon four decades of this process. As one colleague noted, although this progress seems continuous, in fact whenever a factor-of-ten improvement in cost and/or performance occurs, major discontinuities usually follow. The rapid infusion of personal computer technology throughout society in recent years is clearly one example, as is the current acceleration of local area networks (LAN). There are many simple explanations for this phenomena, but in general they revolve around the fact that a ten percent improvement merely allows existing applications and uses to be done somewhat more effectively or economically. A factor-of-ten change usually enables whole categories of applications to emerge that did not previously exist.

Although these advances in the individual components of information technology are quite important, an even more significant impact will occur as a result of new information technology architectures (i.e., new ways to organize and interconnect these components). Three particular trends in this direction can be identified: (1) Extensive communications networks, (2) Accessible distributive data bases, (3) Enhanced human interface work stations. These trends in architecture will be elaborated upon in later sections. In brief, the development of high-performance, high-reliability, comprehensive communication networks, both intraorganizationally and interorganizationally, are occurring at a rapid pace. At the same time, both hardware and software technology are evolving in ways that make it possible to maintain extensive amounts of information on line, and to be able to access this information in conjunction with the communication networks from almost any location. Furthermore, the increased capability of advanced personal computers, often referred to as "workstations", are providing many improvements in ease of use enabling people to work with these systems with much less formal training, yet be able to accomplish much more complex tasks. This is similar to the way that the automatic transmission was able to hide the detailed workings of the automobile's transmission system, involving multiple gears, with the actual selection of gear setting for optimal performance accomplished transparently to the user. In fact, these enhancements in the human interface have been one of the many important aspects, besides the reduction in costs, which have enabled personal computers to be absorbed into our society at such a rapid rate.

Increased Connectivity

At the intersection of the 1990's business forces and the 1990's information technology opportunities lies the need and ability to provide increased and more flexible connectivity. Here we mean connectivity in a very broad sense which can be manifest in many forms. Three of the most critical are identified below.

There have been dramatic increases in efforts to establish much more efficient and tightly coupled interorganizational business relationships. These take many forms. A major thrust is to increase productivity, reduce cost, and improve service by providing highly automated end-to-end electronic connectivity. In the most idealized case, all of the processes from the entry of an order at the customer site, through its processing at the manufacturer's site, and on to the request for replenishment of necessary supplies, are handled through direct electronic connections among customers and suppliers, with minimal manual intervention or paperwork.

In addition to these "electronic" arms-length transactions, there is an increasing emergence of "virtual corporations". This often occurs as the by-product of the agreement among two or more corporations to pool their resources for the purpose of pursuing a particular business opportunity. An example of this type of activity occurs in large-scale military projects, such as the development of the B-1 bomber. In that case, divisions or groups from upwards of two thousand separate corporations are essentially working together as departments within a "virtual corporation" to accomplish the design and manufacture of this product. This type of activity is being seen at an increasing rate in the civilian manufacturing sector as corporations strive to take advantage of manufacturing economies and expertise available in other organizations, or particular cost advantages of labor and raw materials. In these situations it is desirable to set up an efficient management structure and information flow comparable to that which would be expected in a large mature organization, but it is necessary to establish this infrastructure on a dramatically accelerated time-scale.

Looking within a single corporation, we find various similar forces at work, driving toward attempts to provide intraorganizational coordination for increased efficiency and effectiveness. Due to corporate culture or the past limitations of information technology, many corporations have run as a group loosely coupled mini-corporations, each running fairly autonomously from its siblings. Although this often simplifies the management structure, it has been found to incur various inefficiencies. For example, the designers are often removed from the concerns and needs of the manufacturing and purchasing groups. This has led some organizations to strive to develop task forces that consist of members from design, manufacturing and purchasing working together as a team on new products. Furthermore, in looking at the traditional regional separation or product separation within corporations, opportunities can exist to take advantage of the total corporation's resources, such as by pooling its purchasing activities to gain greater economies of scale, or pooling its warehousing and distribution capabilities so as to accomplish higher utilizations and reduced costs.

As a by-product of these inter- and intra-organizational needs, we must also add forces identified as part of the volatile environment of the 1990's. Significant adjustments to an organization may need to occur on rapidly accelerated timetables, leading to the need to develop approaches to support highly adaptable organizational structures. It is not uncommon to read about an organization that has switched from being structured based upon regional divisions to being organized along product lines. At the same time, another corporation, often in the same industry, is undergoing the exact reverse reorganization. Furthermore, in this environment of acquisitions and divestitures, it is not uncommon to find two large corporations merging to form a mega-corporation and spinning off various smaller divisions, either to operate as autonomous corporations, or themselves being absorbed into other corporations. There is a need to be able to rapidly restructure organizations in response to these changes, and to provide all of the necessary information technology infrastructure support. This is being made possible, to varying degrees, through the developments currently underway in information technology.

Cooperating Information Systems

A key generalization that arises out of our theme is that we can look for increased coordination and cooperation among information systems, whether they be humans or computers. As Figure 2 depicts, most situations involve an ensemble of humans and computers. Based upon current day examples, we can illustrate various forms of cooperation.

Electronic Mail

In an electronic mail application, information technology can provide a rapid and convenient means of communicating between two individuals. The connections might be fairly direct, as seen between individuals B and C, or may involve multiple electronic intermediaries, as seen between individuals A and C.

COOPERATING INFORMATION SYSTEMS

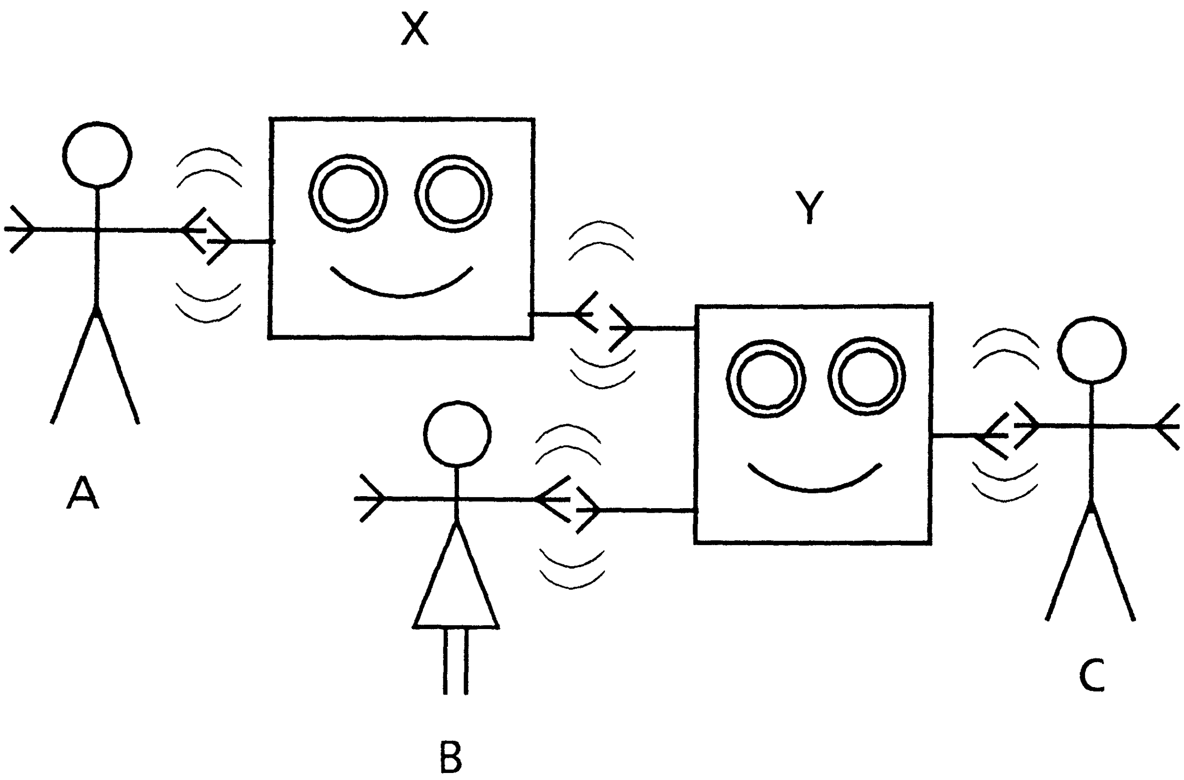

Figure 2

Electronic Data Interchange (EDI)

In electronic data interchange (EDI), information flows between computing systems may take place with minimal, if any, human intervention. For example, if computing system X detects that the inventory for certain items has dropped below its predefined automatic order point, it can initiate direct request for replenishment to its supplier's computer system, identified as computer system Y.

Electronic Conferencing

In looking at ways to improve the effectiveness of organizations, it is often necessary to combine and exploit the unique capabilities of both humans and computing systems. For example, in some experimental electronic conferencing or brainstorming systems, multiple humans interact. Rather than communicating in the form of simple electronic tranmissions, the computing systems help to organize, correlate, and structure the information flows so as to allow a larger number of individuals to participate in a more effective manner than would normally be possible in a noncomputer-enhanced environment.

As seen by these examples, we can improve the effectiveness of cooperating information systems consisting of a large array of humans and computers through more effective and increased capability for connectivity.

Information Technology Components

Before we attempt a detailed view of the trends in information technology capabilities that can be expected in the 1990's, it is necessary to identify the individual and critical components. Figure 3 depicts the four key components: (1) the workstation, (2) the shared distributive databases and knowledge bases, (3) the communications network, and (4) the specialized processors.

Workstations

For most users, the workstations are the most visible aspect of the information technology. This entry point may be a simple terminal, often referred to as a "dumb" terminal, or, as is becoming more common, a personal computer (PC). Although the definition of a workstation is subject to various interpretations, for our purposes it can be thought of as being a sophisticated personal computer. That covers significant territory since current day systems may range from a low-end personal computer costing under $1000 to an extremely sophisticated high-performance computer-aided design system costing in excess of $100,000. One key aspect that differentiates a workstation in our context from the common view of a personal computer is that we are focusing on workstations that can be interconnected both with other workstations and other information technology resources through communication networks as depicted in Figure 3.

Shared Access Distributed Databases and Knowledge Bases

The second key element is that of shared access distributive databases and knowledge bases. As the cost of on line storage technology, both in the form of magnetic disks and optical disks, continues to decline, and as the ease of use and sophistication of database management system software and knowledge-based systems software increase, the gathering and storage of vast amounts of information on line is increasing. Although Figure 3 depicts disk-type storage devices, these are actually controlled by computing systems that may range from traditional "mainframe" computers (running traditional database management systems software such as IBM's IMS), to newer local area network "file servers", or even the emerging "database machines" which are specialized computers optimized for maximum performance and minimum cost in supporting database and/or knowledge base processing.

Communication Networks

As noted earlier, a key element in our notion of a workstation is the fact that it can conveniently communicate with other work stations as well as other corporate resources, most notably the shared-access database and knowledge bases, through

COMPONENTS:

1. WORKSTATIONS
2. SHARED DISTRIBUTED DATA BASES & KNOWLEDGE BASES
3. COMMUNICATIONS NETWORK
4. SPECIALIZED PROCESSORS

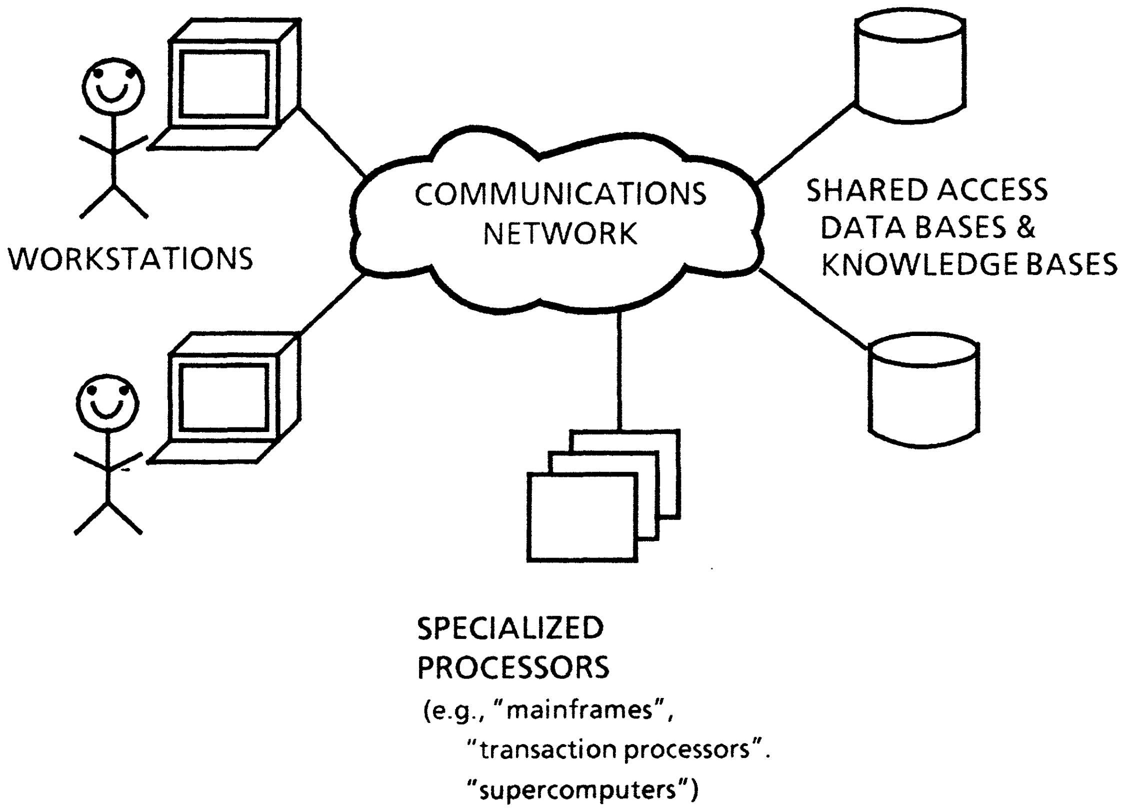

Figure 3

communications networks. It is important to note that these communications networks may take many forms. Three common forms are:

(1) Local area networks (LAN) which are communication networks internal to a corporation used to connect various systems within relatively close proximity, typically within a single building or plant location.

(2) Internal wide area networks, often called "corporate information networks", which are used to tie together the various geographically distributed parts of the corporation. These corporate networks often provide both voice (i.e., telephone) services as well as data transmission capabilities. These are referred to as internal networks because they are designed, managed, and, to a large extent, owned by that corporation to serve its needs. Most major corporations have either installed or have underway the development of such internal corporate communication networks. This has been driven by various forces, such as the ability to customize the network for maximum efficiency at minimum cost, by the ability to exploit specific technologies, and by the need and/or ability to circumvent various regulations that restrict or delay external public networks.

(3) External wide-area networks: We are all familiar with the traditional telephone network, which is an example of a public external network. That is, the equipment is owned by the telephone company, and is maintained and serviced by it, and services are supplied on a fee basis to customers, which might be individual residential users or large corporate users. In the same vein, there are various services provided both by traditional telephone companies, as well as specialized data communications corporations that provide similar data communications services.

Although these three types of networks have been listed separately, it is critical to be able to use them in combinations. For example, operating at a work station one may be able to access and communicate with other work stations in close proximity via a LAN. In order to access a remotely located corporate database it would be necessary to make use of a wide area network, either internal or external.

Furthermore, in order to communicate with work stations at a remote site, operating on a separate LAN, it is necessary to provide wide area communications capability to couple these individual LAN's together. In order to deal with members of other organizations that would not typically be connected to the internal communication network of your corporation, it would be necessary to operate through an external public network that connects individual corporate networks together, or, through prior arrangements, to have established appropriate agreements between the two corporations providing for direct means to link their internal corporate networks. Another example that has become increasingly popular is individuals who have work stations in their residence that communicate with their corporate counterparts and resources, accessed through external networks, typically using standard telephone services today.

Furthermore, there may be many variations of the three network types listed. As an example, "metropolitan area networks" (MAN) provide high performance capabilities for linking individual LAN's within a geographic area, such as a city. A MAN may be run as an internal service, an external service, or some intermediate form, such as linking all financial institutions in the city.

Thus, in many ways, Figure 3 greatly simplifies the issues that are involved in these communications networks. The depiction of a cloud is intended to indicate that the users of a network wish to have as little knowledge as possible about what actually goes on. In much the same way, a typical user of current day voice telephone service is unaware that a telephone call from Boston to Los Angeles may go through various technologies such copper wire, microwave, satellite, and fiber optics, and may undergo various transformations and switching at dozens of sites all along the way. Unfortunately, this transparency that we have come to accept on the voice telephone has been a major challenge to duplicate in complex communications networks. A major trend upon which we will elaborate later is the need and capability to provide increased transparency and functionality for these communications networks, so as to make the resources of the corporation and all of its partners accessible to the user's work station in a transparent easy-to-use manner.

Specialized Processors

The components described above represent a vision of the information technology capabilities and architectures for the 1990's. Examples reflecting this structure can already be found emerging in many organizations. One area that has attracted considerable controversy, and probably will represent differences in strategy, is represented in Figure 3 by the notation "specialized processors". There are several issues that are incorporated under this category.

One major issue involves the evolution of the traditional "mainframe". In many organizations today, the user is serviced through a "dumb" terminal, and all of the processing, both for servicing the end users interactions, the application, and the management and manipulation of the data, is accomplished on a single, shared, large-scale computing system. From a strictly technological point of view, there is often a comparable amount of computing power available in current day work stations (there are actually specialized work stations, such as for computer-aided design, which by certain measures provide more computational power to that single user than is available on most commercial large-scale mainframe computers). Furthermore, to increasing degrees, the amount of processing in support of the storage and management of large databases on specialized "file server" or "database server" systems, effectively eliminate much of the need for traditional mainframes. On the other hand, there are at least two reasons that the traditional mainframes are expected to continue to play a major role in many organizations.

First, there is a tremendous investment in software and procedures that has been based and centered on the mainframe concept. Thus, it is expected that there would be a slow evolutionary movement of certain capabilities, such as the front end user interface moving to the work station, and the data management moving to the shared access database systems. In many organizations, the key systems, such as accounting and inventory control, were first developed twenty years ago and have been updated through incremental developments. As we will discuss later, the ability to evolve such systems at a much more rapid rate will be necessary to deal with the business forces of the 1990's, as well as the ability to exploit the information technology architecture depicted in Figure 3.

A second factor involves various needs for centralized control. Although the combined computing power of the work stations and shared access database servers may be able to provide all the necessary processing ability, it is often desirable to have a central coordinating element. This is for processing that would not be particular to an individual work station, but in fact would coordinate the activities of many throughout the organization. For example, in determining the optimal manufacturing policies, or pooling and processing all of the purchasing needs, a system would be required that would scrutinize all of the appropriate data gathered on the shared access databases and make appropriate decisions for the organization. In addition, in certain areas such as financial transaction systems, there may be concern regarding both the integrity of the user and/or the software that is operating on the work station.

To meet these concerns, certain processors are designated as being transaction processors or coordinators. They may stand as intermediaries between the work station and the databases, that is, the transaction would be sent from the work station to the transaction processor that would validate the request and then perform the appropriate actions resulting in changes to critical database information. Alternatively, the work stations may directly interact with the shared database (which enforces certain restrictions and security restraints). The central coordinating processor would periodically examine the information that had been deposited in the shared data bases, take appropriate actions, and update the appropriate data in these databases.

In addition to the example of functionally-specialized processors to provide a coordination function as described in the previous paragraph, there is also the opportunity to make use of performance-specialized processors. Two particular examples would be ultra-high speed performance processors and ultra-high reliability processors. Although the processing power of work stations is continuously increasing, and is likely to match or exceed that of the "mainframe" computers of the 1980's, there will still be "top of the line" mainframes and "super computers" whose performance might be tens, hundreds, or even thousands of times faster than that of a single work station. Certain types of sophisticated mathematical calculation or the use of elaborate "expert system" software may require the use of such specialized processors. Through the use of emerging "network operating systems software", in conjunction with a transparent communications network, the user need not be aware of where processing is actually being performed, either on the local work station or by a remote specialized processor.

In this section, the key components that are expected in the 1990's information technology architecture have been identified. In the next section the key capabilities, especially those enablers that facilitate increased connectivity, will be elaborated upon, and areas of concern which might act as inhibitors to attaining these full benefits will be discussed.

Enablers and Inhibitors

Through discussions with the representatives of the sponsors of the "Management in the 1990's" program and the associated MIT researchers, as well as through a review of the related literature, a consensus has emerged as to key enablers and inhibitors that are expected to play a major role in the information technology environment of the 1990's. These key factors have been grouped into categories and summarized in Figure 4.

General Advances in Computer Hardware/Software

As noted earlier, and to be elaborated upon later, the continuation of advances in both the reduction in cost and improvements in speed and function of computer hardware and software are expected to continue through the 1990's. This force, on its own, will continue the trend of making information technology a ubiquitous force throughout all areas of business endeavor and social life. Certain specific developments are expected to play a pivotal role in increasing the impact and

OVERVIEW OF ENABLERS (+) AND INHIBITORS (-)

1. GENERAL
   + COMPUTER HARDWARE / SOFTWARE ADVANCES

2. COMMUNICATIONS NETWORKS
   + POWERFUL, TRANSPARENT, INTERNAL NETWORKS
   + EXTERNAL NETWORKS AND STANDARDS (EDI)
     - EFFECTIVE NETWORK MANAGEMENT

3. DISTRIBUTED DATABASE CAPABILITY
   + EASY ACCESS TO DISTRIBUTED DATABASES
     - EFFECTIVENESS OF SOFTWARE
     - DATA RESOURCE MANAGEMENT POLICY
     - INSTALLED BASE
     - CAPABILITY OF DICTIONARY / SEMANTIC RESOLVERS

4. WORKSTATIONS
   + UBIQUITOUS, NETWORKED, HIERARCHICAL WORKSTATIONS
   + AI / EXPERT SYSTEMS TO SIMPLIFY SYSTEM USAGE AND ACCESS TO NETWORK / DATABASE RESOURCES
     - INTERFACE TO DATA/KNOWLEDGE BASE RESOURCES
     - COGNITIVE SUPPORT / FACILITATIONS

5. IT ARCHITECTURE
   + INTEGRATE EXISTING OPERATIONAL SYSTEMS
     - APPLICATION ARCHITECTURE FOR ORG FLEXIBILITY

6. IT INFRASTRUCTURE / USAGE
   + IMPROVE GROUP PROCESSES & PERFORMANCE
   + IT LITERATE / IT CHAMPION MANAGEMENT
     - EFFECTIVE IT STANDARDS
     - GENERAL PRODUCTIVITY (INDIV, GROUP, FIRM)
     - SUPPORT OF MANAGEMENT PROCESSES
     - SYSTEMS INVESTMENT RATIONALE
     - USER ATTITUDE / READINESS
     - QUALITY OF IT WORKFORCE

Figure 4

importance of information technology, if certain associated inhibitors can be overcome.

Communication Networks

We are seeing an ever increasing ability to provide powerful transparent networks that bind together elements of the corporations (internal networks) and the world at large (external networks). Although these endeavors can be quite expensive and time-consuming, often taking up to a decade to be fully established, most major corporations have such efforts well under way and moving rapidly towards completion.

Another important development is the emergence of standards for both the development of new networks and the interfacing of existing network protocols. The Open Systems Interchange (OSI) effort, which has been progressing slowly, is expected to take on increasing importance in the 1990's.

In parallel with the development of the networks will be increasing applications making use of these networks to exchange information, both with internal divisions and with external customer and/or supplier relationships. Much of this activity is already evidenced through the increasing emergence of Electronic Data Interchange (EDI) activities. In some industries up to seventy per cent of all order-placing activities are being performed by direct EDI transactions between buyers and sellers. Due to the differences of structures among industries, the amount of EDI activity varies greatly. In all cases, it is expected that this activity will increase and become even more widely accepted as a basic way of performing business.

The technologies supporting the development of these communications networks, and the standards that provide for interconnection between these networks and related parties, such as buyers and sellers, can greatly increase the efficiency of many activities throughout industry. But there are also concerns that must be addressed which otherwise will inhibit the full potential of these communications networks to serve our needs. There is a major concern regarding our ability to develop effective network management. We are currently contemplating networks significantly larger, more comprehensive, and more complex than any that have ever been established before.

Even within a strictly internal network, there are technical problems with effectively monitoring and understanding the activity on this network and management problems regarding the operation and growth of such networks. When one looks further into the difficulties of building interorganizational networks, both the technical problems caused by the diversity of network technologies and strategies that may be deployed and the lack of a central network management organizational structure pose serious challenges. There are various ways that these problems may be mitigated, such as by the emergence of various network coordinating bodies (such as CCITT in regards to conventional public telephone networks) or de facto coordination by major network software and/or service providers (such as IBM with its Systems Network Architecture). Furthermore, advances in technology may help to reduce the apparent complexities and difficulties of managing such networks.

Distributed Database Capability.

A major driving force behind the emergence of such communications networks is the need to support and facilitate easy access to distributed databases. These databases may currently be distributed to various locations throughout the corporation, or may be centralized to a few locations, but with the need to be accessed from locations throughout the corporation. Furthermore, access and use of databases that traditionally operated in isolation will be particularly important in increasing our effective connectivity in support of better coordinated planning.

There have been significant advances in distributed database management systems software, including some currently announced products. Similar systems are expected from most of the major database management systems providers, including IBM, with widespread use accelerating from 1990 to 1995. With the proliferation of powerful communications networks supporting a distributed database capability, a user anywhere within the organization, or within cooperating organizations, would be able to access any authorized information transparent to its particular location.

Although, prototype examples of such systems have already been demonstrated, there are major concerns that may act as serious inhibitors to the effective deployment of this technology. The first of these concerns involves the rate at which this technology will evolve and the effectiveness of such distributed database systems. Most of the current implementations of such systems have many limitations, either in terms of their ability to provide complete transparency and/or their ability to provide adequate speed performance under all appropriate conditions. Increased experimentation and closer working relationships between the user community and the developers of such technology will be important in resolving these problems on a timely basis. If these function and speed problems are not overcome rapidly enough, the ability for effective distributed database management may be limited to that subset of circumstances that can be sufficiently addressed by the then available technology.

There are several other major concerns that, although identified as being inhibitors to a distributed database capability, are, to a large extent, problems inherent in attempting to run large complex organizations with increased coordination and cooperation. For example, a major concern involves the development of effective data resource management policies. In the past, each part of the organization or cooperating organization pursued both the functional definition and the information technology implementation of its databases fairly autonomously. To gain increased effectiveness in sharing data, it becomes important to develop common data resource management policies. These policies are needed to address a wide array of issues. One example might be the standardization of part numbers throughout the corporation or throughout the industry, in much the same way the Universal Product Code (UPC) represents a standardized part numbering scheme for products purchased through grocery stores. Although there has been considerable interest in the issue of data resource management, there are widely differing views as to how it might be best attained, considering both the technical limitations of systems, the organizational realities, and the complexity of the problems being addressed.

Another major concern involves the legacy of the installed base of data systems. These currently installed systems represent a major asset for the corporation, since both their development cost and the cost of acquiring and accumulating the data they now contain can be of enormous expense. Furthermore, despite possible deficiencies, the use of these systems represents knowledge that has been gained by many individuals in the organization, therefore, considerable disruption and retraining costs could be incurred if such systems are changed significantly. Thus, one is often faced with a tradeoff between a slow rate of change, which amortizes the value of the installed base over a longer period and minimizes the risks of technical or organizational disruptions, against the potential benefits of new systems that have been more effectively developed to support the need for increased connectivity and coordination. In almost every organization some form of evolutionary plan, possibly coupled with an overall data resource management policy, needs to be developed to deal with these opposing forces.

Our desire for an increasingly connected and coordinated organization, while having many potential benefits, does pose many challenges to our ability to manage such environments effectively. In particular, we are familiar with the old adage "divide and conquer". This adage in many ways represents the reality of how organizations and humans often function best. By limiting scope, the amount of information that one must process and comprehend can be kept within convenient limits. As we attempt to broaden this scope it becomes necessary to have capabilities to aid us in this process, in particular powerful global data dictionaries and semantic resolvers.

Global data dictionaries allow us to ascertain what information is in the system and how to request it. For example, we know that General Motors (GM) buys equipment from several of our divisions. If we want to find out the total volume of business with GM, how do we do it--which data systems, probably several, contain that information, what is it called, how is it accessed? The global data dictionary would assist us in this process.

Even more subtle issues involve the ability to perform semantic resolution. Within a specific part of an organization, or a specific culture, we make significant assumptions about our environment. For example, in the United States, if you were to ask a supplier the cost of a specific part and were told 16.45, it would be natural to assume that the price is expressed in U.S. dollars. When one starts to deal across parts of the organizations, or on a worldwide basis, traditional assumptions about the "semantics" or meaning of data needs to be made more explicit, and , often, transformed into compatible forms. For example, one supplier may express prices in terms of each individual part, whereas another supplier may express prices in terms of quantity of a dozen at a time. If one dealt with both suppliers on a continuing basis, the knowledge of these differences would be automatic assumptions, but as we expand our scope we will start to develop new relationships where these assumptions had not been explicit. The capability to have semantic resolver software represents a major challenge for information technology. Without such a capability, we would find ourselves faced either with making erroneous decisions due to an incorrect interpretation of data that we have retrieved or with the need to resort to much less efficient and extensive human intervention to validate the accuracy of the information and to perform appropriate transformations before that data can be used.

Workstations

An important enabling force toward making information technology usable throughout the organization is the emergence of the ubiquitous networked workstation. In much the same way that the telephone on our desk has become our access to the world of human communication through the telephone system, the workstation serves a similar role to accessing the new network of both humans and computer systems throughout these organizations.

An important enabling element is increased capabilities, especially through the deployment of Artificial Intelligence (AI) and Expert Systems (ES) technologies, that will allow these workstations to significantly simplify system usage and provide more transparent access to the network and database resources. This is an important area where the growing power of the workstation plays an important role. In a recent discussion with the developers of a sophisticated specialized information service, it was observed that the amount of software in the user's workstation to support an easy-to-use interface had increased from under 50,000 bytes in earlier models to now over 4,000,000 bytes. Thus, the advances in information technology are serving as the fuel making possible significantly more powerful and easier to use interfaces in much the same way that the addition of automatic transmission technology has made an automobile easier to operate.

There is concern that the goal for the workstations of the 1990's to provide a significantly easy-to-use capability that makes the complexities of the network and distributive data transparent will not materialize as fast as desired and thus will be an inhibitor to rapid connectivity.

There is also considerable effort to use these powerful workstations, in conjunction with AI and ES technology, to support the user's decision making and cognitive processing. This type of support could significantly enhance individual productivity, while increasing connectivity throughout the the organization will help to improve group productivity . Although these are reasonable goals to strive for, the capability of information technology to be able to fully support and facilitate such cognitive activities is an area of significant concern, and if not fully realized will be an inhibitor to their success.

Information Technology Architecture

The information technology architecture depicted in Figure 3 provides an evolutionary way to integrate existing operational systems to provide the basis for new informations systems that can help the disparate parts of the organization and multiple organizations to work cooperatively. This type of architecture is the focus of considerable attention by many leading edge organizations and needs to be pursued aggressively in order to provide truly connected and coordinated organizations.

As noted earlier, we must deal with a highly volatile environment. It is not only necessary to be able to integrate our operational systems together as a one-time activity, but we need to provide an architecture that provides for considerable organizational flexibility. Some suggestions for approaching this problem will be discussed later but if we are not able to attain this goal, the inflexibility and slow adaptability of the information systems structure itself may act as an inhibitor to the rate of change that must be accomplished to survive effectively in the highly volatile environment of the 1990's.

Information Technology Infrastructure and Usage

Although the above sections stressed the technology aspects, there are many organizational concerns that can either act as enablers or inhibitors to the effective use of information technology.

Two important enablers have been identified in our discussions. The first involves the ability to significantly improve group processes and performance through the various technologies supporting connectivity and coordination throughout the organization. Second, the increase in the number of managers that are both comfortable and literate in information technology will make it easier to identify and pursue these exciting opportunities, and the emergence of "information technology champions" with the vision and leadership to move the corporation aggressively in these directions is expected to increase as we move on into the 1990's.

These positive directions in the information technology infrastructure do not occur without significant obstacles that can significantly inhibit full attainment of our desired goals. Six specific concerns have been identified.

The first involves the development of effective information technology standards. The rapid development of new information technology capabilities, including hardware, software, and applications, while providing tremendous positive opportunities, also gives rise to wide diversities which can make connectivity and effective communication very difficult. The development of standards, or effective techniques for mapping between the alternatives, is a critical need.

The second factor involves our understanding of general productivity applied to both individuals, groups, and organizations. One of the many cliches regarding information technology is that it allows you to "perform the wrong actions hundreds of times faster". In order to gain the improvements in productivity that we seek, it is necessary to rethink the processes that are currently being used and, in many cases, transform the organization dramatically to gain the improvements in productivity that information technology makes possible. Although this point has been raised by many researchers in this field, the identification of a general and widely accepted approach to attaining these transformations still remains a major challenge.

Although the preceding paragraph primarily alluded to the products and services of the corporation, a very similar argument can be made regarding the decision making and general management processes of the organization. Due to their complexity, it has often been difficult to identify ways to apply information technology to improve these processes.

There is a need to reevaluate the systems investment rationale that has been traditionally applied to information technology initiatives. Often information technology investments are based upon direct cost savings accomplished through personnel reduction. In the environment of the 1990's, where we are exploring new ways to perform business and coordinate activities among organizations, a direct connection to personnel reduction may not always be possible nor fully reflect the benefits that are to be gained. Thus, we need new approaches to evaluate the information systems investment process in this environment.

There is also the duet of concerns regarding the quality of the information technology workforce and the user's attitude and readiness to accept and pursue the new options made available through advances in information technology. We will look at these two factors separately. Unless sufficient progress is made on either effective software that can hide the complexities of such systems or on the level of training and experience that will be generally available in the 1990's, there may be such a large mismatch so as to inhibit the widespread deployment of new technologies. Furthermore, even if there is sufficient transparency of the complexity and a sufficiently well-trained work force, there is concern as to how rapid a rate of change the users can endure, this includes both top management and middle management as well as the external partners to be interconnected and to work in a coordinated cooperative manner. Although the pressures to make such changes will be quite high, the organizational ramifications can be sufficiently onerous so as to act as major inhibitors to their adoption.

This section is intended to reflect the reality of the "good news/bad news" (or enablers and inhibitors) environment that will characterize the 1990's. To a large extent, the rate of progress and degree of success of the information technology advances anticipated will influence the degree to which the enablers will be significant and, to a lesser extent, the degree to which the inhibitors can be mitigated. But, many of the inhibitors represent management decisions and organizational changes which the information technology cannot completely minimize and will represent major challenges to the management of the 1990's.

Continuum of Proactive/Reactive Connectivity Environments

We have discussed the need to respond to the volatile environment of the 1990's by being able to support highly adaptable organizational structures. In looking at that volatility and the ways in which an organization can deal with it, we have identified two extremes which, although neither extreme is realistic, do represent pedagogically interesting perspectives as depicted in Figure 5.

The first category is referred to as proactive. As the extreme case under this environment, from the moment that the organization is formed all future requirements for connectivity are anticipated and planned for in advance. For example, if a corporation initially operates in Massachusetts in the United States, appropriate modes of operation including definitions of part numbering schemes, part number codes, financial reporting formats, will be determined. As this organization expands geographically and over time, these procedures are carried forward in unison throughout the entire organization. Under such a proactive environment, if it becomes desirable to coordinate the activities of a plant outside the USA with its counterpart in Massachusetts, this organizational consistency would make the exchange of data and processes quite easy, assuming that technical decisions such as choice of programming languages, operating systems, and

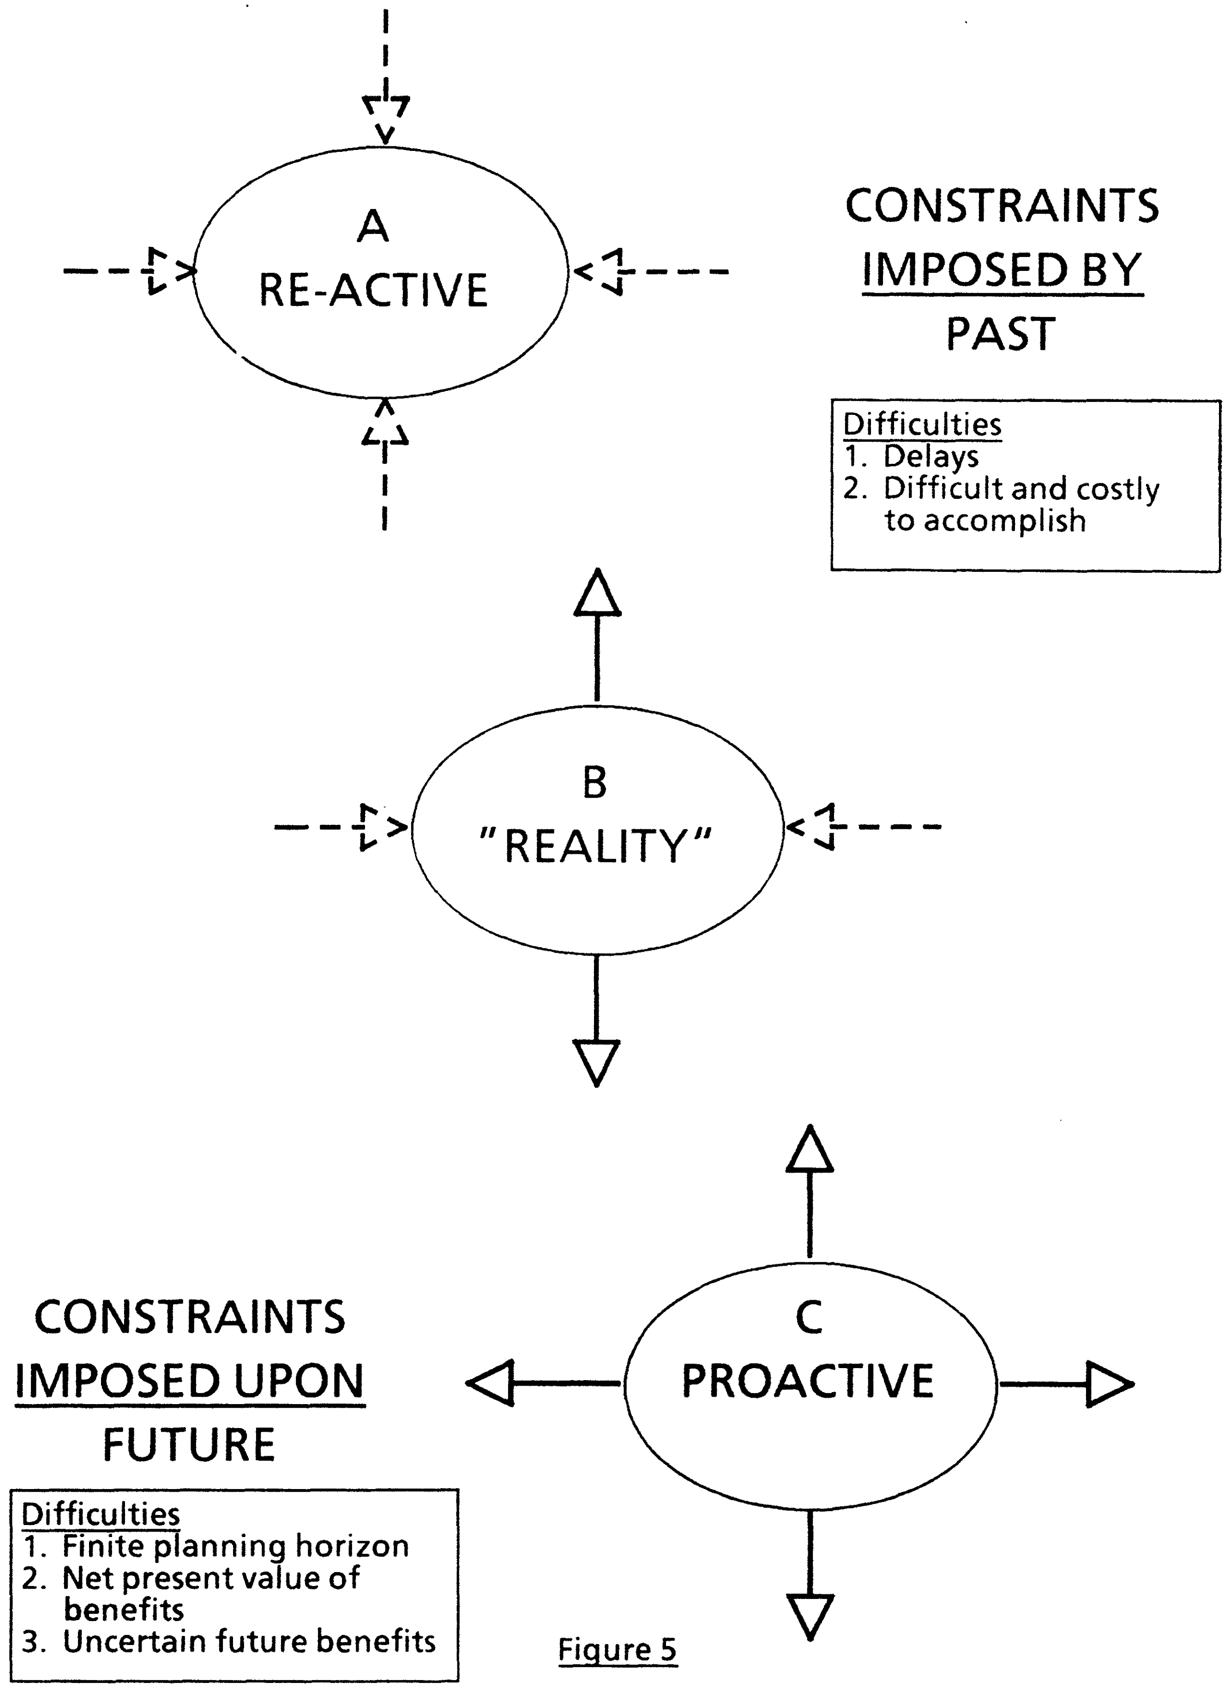
A
RE-ACTIVE
CONSTRAINTS
IMPOSED BY
PAST
Difficulties
1. Delays
2. Difficult and costly to accomplish
B
"REALITY"
CONSTRAINTS
IMPOSED UPON
FUTURE
Difficulties
1. Finite planning horizon
2. Net present value of benefits
3. Uncertain future benefits
C
PROACTIVE

Figure 5

communications networks have also been proactively planned amongst the divisions of the corporation.

The other extreme we can call the reactive environment. In its purest form, each separate part of the organization proceeds autonomously of the rest of the organization. Only when a specific need arises to connect and coordinate activity will such a concern be addressed. At that time either one part of the organization will alter its data and procedures to match the other, or appropriate transformation capabilities will be developed to allow them to interchange information effectively.

Obviously , both of these extremes are unrealistic. In the first case, to be completely proactive would require a level of effort and detail that could be an enormous cost to the corporation, although advances in information technology can help to mitigate this cost. Furthermore, it requires a planning horizon and stability that are contradictory to the volatile environment that has been assumed. For example, even if every division of the corporation has been coordinated so as to have exactly the same procedures, part numbers, customer codes, etc., this carefully planned structure would be thrown into disarray by acquiring a division from another corporation, which would presumably have its own different set of data and procedures, or by being acquired by a different corporation. Even if the intentions are to rapidly become restandardized, it is impossible for this task to be accomplished instantaneously, thus one must endure a period of reactive behavior, at least temporarily. Furthermore, especially when considering the globalization of corporations, one finds that the disparate customs, environment, and legal differences within different countries would impose forces that may be at odds with the goal of corporate standardization.

Likewise, an environment of total reactivity would be extremely costly and inefficient, essentially resembling a Tower of Babel. Thus, certain types of activities, such as financial reporting to headquarters to enable corporate-wide financial review and preparation of annual reports, would be a requirement of sufficient longevity and value that procedures would likely be instituted in each of the divisions to make such reporting convenient and minimize effort.

Although, from the descriptions that have been provided here, it would seem obvious that neither extreme of proactive or reactive environment would be feasible or desirable, many organizations behave as if they are at these extremes. In reality, of course, we exist in the continuum between these extremes. There are certain activities that are sufficiently critical and can be forecast with sufficient confidence, that a proactive plan may be feasible and appropriate (such as setting up procedures for electronic data interchange between all plants and the major suppliers and/or customers). At the same time, it must be realized that situations will arise that are beyond the forecasting window or that are unexpected events, such as a unique opportunity to acquire a complementary business. Under those circumstances, it is necessary to be reactive in as efficient and effective a manner as possible.

In such a heterogenous environment, one wishes to have information technology that can support both the proactive and the reactive strategies. In particular, as one colleague noted, in this uncertain world of the 1990's one wants to be able to "proactively plan to be reactive".

Information Technology for Both Proactive and Reactive Strategies

In studying the many information technology advances that have been identified as enablers in our research, it can be seen that they serve different roles and are of different levels of importance when seen in light of either proactive or reactive environments. We will use this differentiation to help organize the discussion in this section. Furthermore, there is another dimension that has been found helpful in organizing this material. That is the distinction between physical connectivity and logical connectivity. In physical connectivity we are focusing on the communication technologies and networks that handle the transmission of the data without regard to the contents or meaning of the data. This may be viewed as analogous to the postal system that, through a complex and elaborate process, delivers the mail from sender to receiver without the need to open the envelope or understand the contents. Logical connectivity refers to understanding the content of the letter. Thus, if an order is sent from a buyer in Germany to a manufacturer in the United States, the postal system will be able to deliver it if the addressing information on the envelope is adequate. But, when the plant in the U.S. receives the envelope and opens it, various difficulties may ensue. For example, if instructions need to be provided along with this order, either the sender in Germany will have to have translated the instructions into English, or alternatively the plant in the U.S. would have to have an interpreter who understood German. Furthermore, if the part numbering of the two organizations differs then either the sender will have to have translated its requirements into the part numbers of the supplier, or alternatively, the supplier will have to be knowledgeable about the part numbering system of its buyer and be able to perform the translation itself. Although we have used analogies of a postal system and of human buyers and sellers, we would like to have most of these activities performed through electronic technology, thus the problems both of physical connectivity and logical connectivity would need to be addressed.

Figure 6 provides a framework dividing the notion of connectivity into the categories of proactive versus reactive connectivity, and physical versus logical connectivity. Exploiting that framework, we can now organize the many enabling information technologies or inhibiting factors that have been identified as part of this project into one of the four parts of this matrix as illustrated in Figure 7. With few exceptions, this framework has been found very effective in organizing the over two dozen information technology enablers and inhibitors that have been identified. (Two issues, indicated by * and **, apply to entire rows of Figure 7 instead of individual quarters.) Furthermore, this framework also has the benefit of being prescriptive by highlighting which information technologies are most relevant to the particular form of connectivity being sought.

<u>Proactive Physical Connectivity</u>

We could start our discussion from any point in the framework of Figures 6 and 7. Arbitrarily we will start by looking at physical connectivity in the environment of proactive connectivity, the upper lefthand box.

Given the complexity of the task, a <u>homogenous network management</u> capability, whereby a consistent set of procedures and technologies for managing the network are employed throughout, is very important. Furthermore, by using <u>standard</u>

	PHYSICAL (COMMUNICATIONS)	LOGICAL (DATA)
PROACTIVE		
REACTIVE		

- **PHYSICAL VS. LOGICAL**
  - PHYSICAL = COMMUNICATIONS & NETWORKS
    (E.G., TRANSMISSION OF RAW DATA)
  - LOGICAL = SEMANTICS OF DATA
    (E.G., WHERE IS IT? WHAT DOES IT MEAN?)

- **PROACTIVE VS. REACTIVE**
  - PROACTIVE = SYSTEMS ARE DESIGNED SPECIFICALLY TO "FIT TOGETHER"
    (E.G., STANDARDS, HOMOGENEOUS NETWORKS AND DATABASES)
  - REACTIVE = REQUIREMENT AND PLANNING TO TIE SYSTEMS TOGETHER OCCURS AFTER SYSTEMS EXIST
    (E.G., HETEROGENEOUS NETWORKS AND DATABASES, SEMANTIC RESOLVERS, MULTI-VENDOR INTEGRATION)

Figure 6

	PHYSICAL (COMMUNICATIONS)	LOGICAL (DATA)
PROACTIVE [*]	• HOMOGENEOUS NETWORK MANAGMENT • STANDARD PROTOCOLS (ISDN, MAP, ETHERNET, ETC.) • STANDARDS FOR NETWORK INTERCONNECT • REGULATION OF WIDE AREA NETWORKS	• DISTRIBUTED HOMOGENEOUS DATABASES (TRANSPARENT) • EDI (APPLICATION SPECIFIC STANDARDS) • INTERNAL DATA STANDARDS • APPLICATIONS: GROUP DSS, E-MAIL, COMPUTER CONFERENCE • STANDARDS FOR FLEXIBLE APPLICATION ARCHITECTURE AND DATA ARCHITECTURE • PRE-PLANNED DATA RESOURCE MANAGEMENT
REACTIVE [**]	• MEDIA DIFFERENCES • PROTOCOL DIFFERENCES • GATEWAYS • HETEROGENEOUS NETWORK MANAGEMENT	• DISTRIBUTED HETEROGENEOUS DATABASES (TRANSPARENT) • MULTI-VENDOR SYSTEM INTEGRATION • DISTRIBUTED DATA DICTIONARY • DATA SEMANTICS • DATA RECONCILIATION • POST-PLANNED DATA RESOURCE MANAGEMENT • ORGANIZATIONAL FLEXIBILITY / INTEGRATION

[*] = MANAGEMENT OF STANDARDS PROCESS
[**] = MORE EFFECTIVE MANAGEMENT OF GEOGRAPHICALLY DISPERSED ORGS

Figure 7

protocols, the amount of diversity in the network can be minimized. There are many standardization efforts currently underway, such as the Integrated Services Digital Network (ISDN) which is primarily focused on the worldwide external network, the Manufacturing Automation Protocol (MAP) which is focused on local area networks for manufacturing factories, and Ethernet which is primarily focused on general purpose local area networks, such as in offices. Even though we may be able to minimize the number of different types of networks, there will be some variations due to the differences between the environment of a factory shop floor and the worldwide external network, thus interconnectibility standards between these networks is important. Each of these factors identified can be looked upon as enablers, if the prospects for their advancement proceed as intended, or may be inhibitors if we are unable to accomplish effective network management, develop standardized protocols, and resolve complexities and technical difficulties with efficient convenient interconnectibility amongst the networks.

Furthermore, most external networks, often referred to as wide-area networks (WAN), are regulated by government bodies. The regulatory factors in such WANs may provide enabling or inhibiting forces. If the individual government bodies of each country adopt differing and incompatible regulations, then producing worldwide public external networks can be quite difficult and, to a certain extent, these regulatory factors can impede worldwide corporate networks as well. On the other hand, if the various forces towards greater cooperation and standardization, both technical and regulatory, continues, then many of the current difficulties will be significantly reduced.

Reactive Physical Connectivity

Although diversity and disparity are not an explicitly sought after goal, the rapid changes in perceptions, breakthroughs in technology, and differences in philosophy make it possible for reasonable people to come up with different approaches to solving their communication network needs. These alternatives range from media differences, such as unshielded twisted pair versus shielded twisted pair versus optical fiber, etc., to differences in low level media access protocols, such as Ethernet versus Token Ring, to higher level protocols, such as IBM's Systems Network Architecture (SNA) versus Digital Equipment Corporation's DECNET.

Although these differences might be diminished through the forces of standardization described above, it is also likely that the striving to use new and better media technologies and protocols discoveries will lead to increased diversity. In the reactive environment of the "loosely-coupled" organization, we would expect the various parts of the organization to select the most appropriate network choices to meet their needs. Eventually certain of these groups will need to be connected. Thus, the ability to provide effective gateways that allow us to interconnect these networks will be quite important. Advances in information technology that make it possible to build general purpose gateways that can interconnect multiple networks, as well as development support technologies that allow rapid production of new gateway facilities, will mitigate this problem. Finally, heterogeneous network management capabilities, though much more difficult to accomplish than homogeneous network managment, will make it possible to monitor, control, and optimize these complex networks in much easier ways than are currently available today.

Proactive Logical Connectivity

In the past the demands for logical connectivity have been much less, both due to reality of the limitations of information technology and the lack of driving business needs. These situations are changing rapidly. Thus, organizations striving for increased logical connectivity are exploring new developments in information technology. One of the most important technologies is transparent distributed homogeneous database management systems. These systems allow one to treat data that is physically dispersed to multiple computers as if it were on a single system. The user interface and programming would appear similar to that of a traditional centralized site. By homogeneous we mean that the software is being supplied by a single vendor, or by software-compatible vendors, who designed the individual pieces to work cooperatively. There has been extensive research on this topic for several years which has led to the introduction of several commercial systems in recent times. It appears likely that most major vendors of database management systems software will be introducing offerings of this type by the early 1990's and the technology becoming well-established by the mid-1990's.

Merely being able to access distributive information may not be sufficient unless there is agreement on data standards. These data standards may come about from two different sources. First, there are industry-wide or application-specific efforts to establish electronic data interchange (EDI) data standards. Industry specific examples include extensive activities in the railroad industry and application-specific activities include EDI efforts on standardizing order entry applications data input across multiple industries. The advantages of these efforts are that they are focused on only that subset of data that needs to be interchanged between different organizations. There is still difficulty, of course, in the process of attaining agreement across a varied array of organizations and special interests.

The other activity involves internal data standards, whereby the organization attempts to coordinate all internal data, such as part number standardization, so that information may be exchanged even in cases that have not been explicitly anticipated in advance. The positive aspect of this effort is that since it is within a single organization, there may be opportunities to set up organizational structures to facilitate this process, although even within a single organization there may be competing goals and special interests. Furthermore, the magnitude of detail that one might be attempting to coordinate could be enormous. In one study for a major corporation, over 50,000 different tables were identified. If they averaged fifty columns to each table, then there would be approximately 2.5 million different data elements that might need to be coordinated (e.g., the pay grade codes used in the manufacturing plant in Hong Kong, and the pay grade codes used in the sales office in New York). A very promising approach, sometimes referred to as focused internal data standards, is where the critical needs of the corporation are identified and only those data elements found to be key to coordinating these activities are coordinated.

The value of having a environment supporting connectivity is most significant if there are ways to make use of this connectivity. Examples of connectivity applications include electronic mail encompassing the entire organization and its related parties, computer conferencing capabilities to allow collective agreement and activity to be reached, and group decision support systems (GDSS), by which multiple people, aided by computer-supplied information and modeling, arrive at key decisions. Research has been progressing on all of these areas, with some coming to be widely accepted and others still at the emerging stages.

In addition to exchanging information, it is also important to be able to exchange the processing (i.e. the programs) that are being used. This would make it possible to take a transaction validation capability that previously had been performed at the corporate headquarters computer in New York City and distribute it to the various branch operations to allow much faster and more effective feedback to the user. The development of standards for flexible application architecture represent an important area of attention in many organizations, especially those attempting to create a new proactive environment for the future.

Finally, we need to develop proactive data resource management which acts as a binding force for all of the other activities described above to assure that the appropriate technologies are identified and evaluated as to how they best serve important business needs and the organizational environment.

Reactive Logical Connectivity

In many ways, this is the most difficult environment to deal with, although it may also be the most important for many organizations. It is the area that has experienced rapid growth of importance and where the previous information technologies have most lacking. New developments in information technology show significant promise toward mitigating these problems.

Paralleling the discussion of proactive environment, an important development is that of transparent distributed heterogeneous database management systems.

These systems would allow one to access data that is geographically dispersed throughout the organization and that is running under the control of differing DBMS software from multiple vendors, as if it were a single centralized system. Although there have been prototypes demonstrated that have provided such capability in certain experiments, significant effort is needed to extend the functionality, speed of operation, and comprehensiveness of these systems. For example, although these prototypes are fairly effective at retrieving data, they are usually unable to perform updates to distributed data. Although these developments will lag behind those of the distributed homogeneous DBMS's, they are likely to become increasingly widely available by the mid-1990's.

Another area that has been gaining increasing attention is that of multi-vendor systems integration. It is becoming increasingly apparent that most organizations have, and will continue to have, equipment provided by multiple vendors. This provides both a need and a business opportunity for the systems integration industry, both from traditional information technology vendors that will support equipment from their competitors, as well as new business ventures that are strictly focused on the systems integration activity. The advances in this area will not only simplify much of the management of such a computer ensemble, but also will significantly simplify and improve the performance of the distributed heterogeneous database management systems.

As noted earlier, even though one may have the physical connection and database managements systems software that would allow access to any information, there is still a need to know what data exists and what it is called (even if the distributed heterogeneous DBMS software makes its geographic distribution transparent). Advances in distributed data dictionaries provide the opportunity to address this problem. Through these facilities one would be able to inquire of the "data about the data". One would be able to ask questions such as what are all of the attributes available regarding employees, or what are all the attributes available regarding warehouse capacity, etc. To a large extent, the advances in distributed data dictionaries represent a merging of the ongoing efforts at improving "traditional" data dictionary capabilities for centralized systems as well as the ability of the distributed database management systems software to make geographically distributed systems appear to be centralized. These capabilities will be particularly important when one needs to address connectivity requirements that have not previously been anticipated, as might be expected in a reactive connectivity environment.

Two additional important areas are data semantics and data reconciliation. Techniques in "data modeling" whereby the relationships between data elements are characterized and processed provide a starting point for data semantics efforts. More recent advances incorporating technologies of artificial intelligence and knowledge-based processing to further augment our ability to represent complex relationships amongst data. In the past, these efforts have primarily been focused on traditional centralized data systems. It is an even more complex and challenging problem to capture such meanings for distributed systems since that requires individuals throughout the organization to contribute knowledge about the meaning and interpretation of their information. Once one has captured the data semantics, the opportunity exists to be able to use this knowledge to resolve differences between data from multiple sources or restate information to meet the requirements of different individuals or parts of the organization.

Even though we have been assuming a reactive environment in this section, there is still need for data resource management. This need arises from two forces. First, we are never in a truly reactive environment, and thus we would wish to take the opportunity to be proactive and plan our data whenever possible. Second, in order to be as responsive as possible when new requirements occur, it is necessary to plan, develop, and maintain an appropriate infrastructure to support these activities, which might include procedures needed to continually update the distributed data dictionaries and the data semantics. This area has traditionally received much less attention than that of the proactive data resource management activities because the technologies have only recently emerged and the organizational problems that one would expect in an organization that is highly reactive can be very difficult.

Returning to the theme identified at the beginning of this section, we note that an important goal is that of organizational flexibilty and integration. This goal is not the exclusive property of information technology professionals but, in fact, represents the realities of the business environment of the 1990's. Thus, the interplay between the business forces and information technology must be paramount and the information technology must be as responsive as possible in meeting these evolving needs. The technologies that have been identified provide important advances that can make these goals more attainable.

Flexible Information Technology Architectures

Given that we wish to support both proactive and reactive Information Technology strategies, certain Information Technology architectures can provide considerable assistance. We want a flexible system architecture that can easily adopt to organizational changes, geographic shifts, and alternating forces of centralization and decentralization.

As part of our research, we have developed an example of such an architecture that consists of seven major functional components as depicted in Figure 8. These components are separated into five layers. This architecture attempts to mediate the conflicts between the goals of autonomy, integration, and evolution. All the layers of this architecture, except for message control and data control, lend themselves to as much, or as little, autonomy as desired -- and the balance can easily be changed.

Overview of Components

The components shown in Figure 8 are logical entities. They may each correspond to a separate computer or may all co-exist in a single computer. For purposes of an overview discussion, let us assume that each component is a physically separate computer. In this case, each group could acquire and manage its own resources including: (1) terminal/network gateways, (2) application processing computers and software, and (3) database computers and software. In the past these three decisions were often bundled together; with this architecture we can respond to the typical

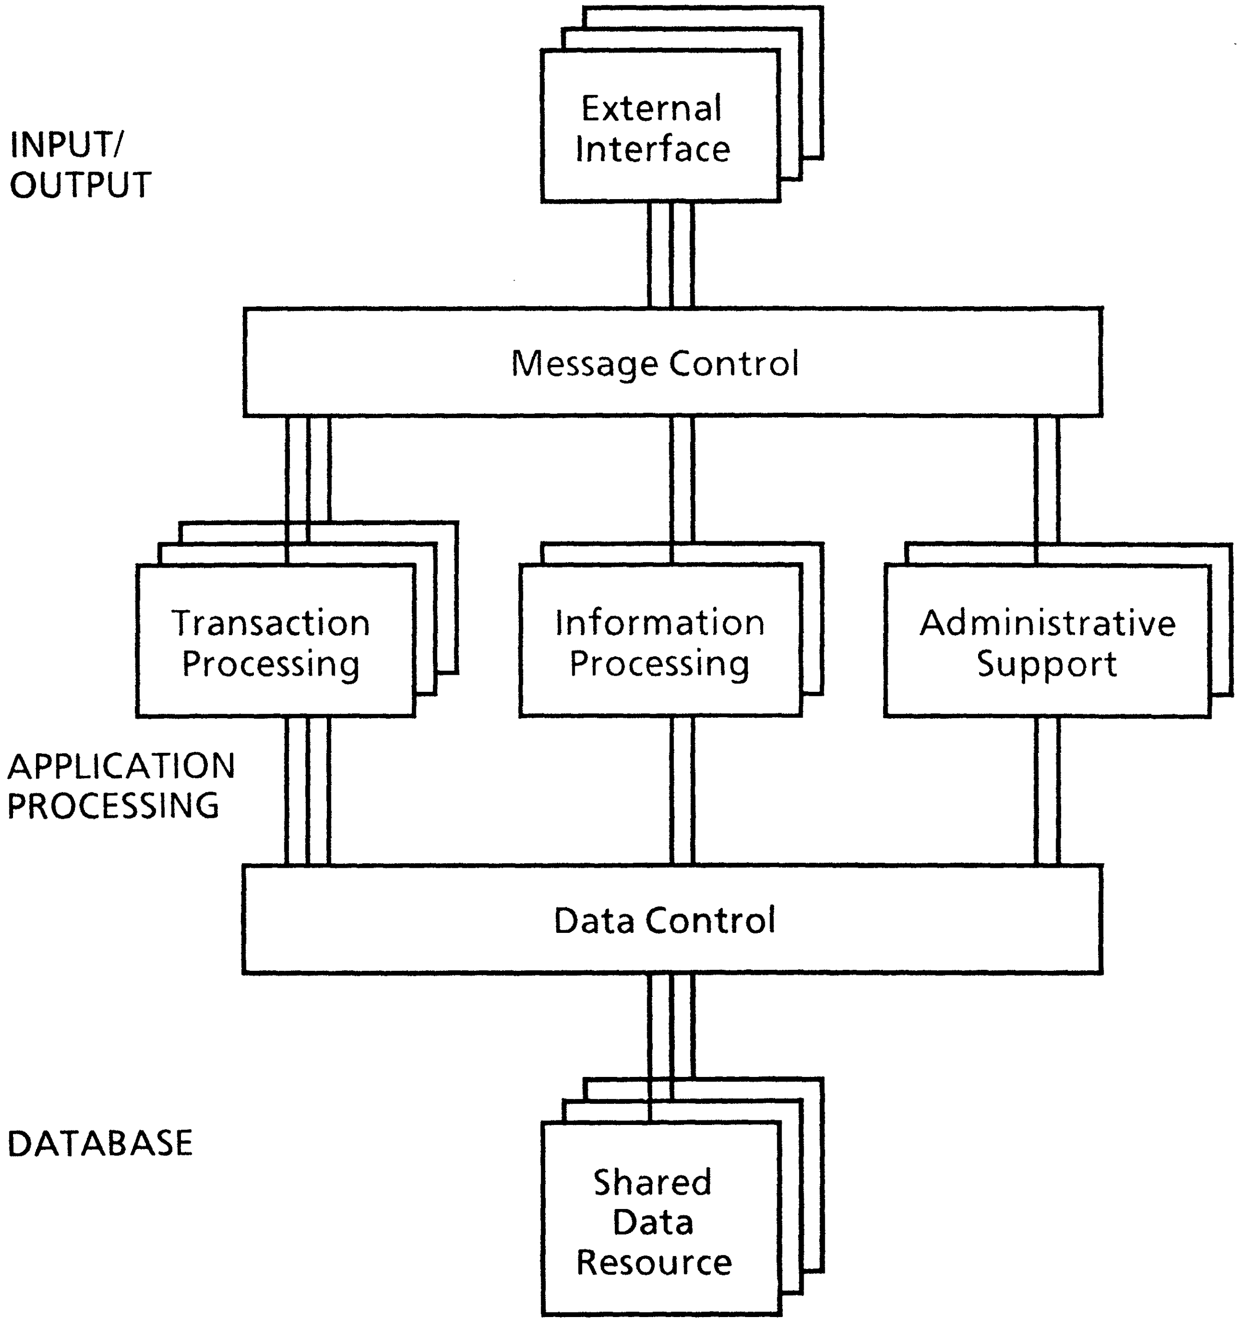
INPUT/
OUTPUT
External
Interface
Message Control
Transaction
Processing
Information
Processing
Administrative
Support
APPLICATION
PROCESSING
Data Control
DATABASE
Shared
Data
Resource

Figure 8

case where the highest degree of concern for autonomy involves the application processing, with a lesser concern for the database, and often a minimal concern for the diverse external interface. This is because it is usually in the application processing that the key functionality of a group's activities is manifest. It is important that the group be able to adapt to changes in environments, goals, or products rapidly.

Whereas the application processing may be tightly coupled to the activities of the group, the database requirements are more variable. In many cases there are separate databases for each application. But, it is expected that many of these will merge in response to the organization's need for integrated data access. The architecture of Figure 8 faciliates access to each of the individual databases as well as providing an evolutionary path for eventually integrating these databases, as the needs for integration intensify.

Integration and Evolution Aspects of the Architecture

The Message Control and Data Control components of the architecture address the issues of integration and evolution. They are the points at which all processing is coordinated. For example, in principle, any terminal can access any application (assuming appropriate permissions and authorizations). Similarly, any application subsystem can utilize the Shared Data Resource to manage and maintain data which is common to more than one component.

Message Control and Data Control are both conceptually *single* entities. The other components are *types* of processing functions. There may be many instances of each type (i.e., multiple external interfaces, multiple transaction processing systems, and multiple shared data resources).

Components of the Architecture

It is important to reemphasize that the components of Figure 8 are *logically* separate. They can be mapped to physical hardware in various ways, some examples of which are provided later.

External Interface: The External Interface provides the entry point to the system. In the case of banking systems, the external entities may fall into five categories: (1) payment networks, (2) communication networks, (3) customer terminals, (4) professional workstations, and (5) other intra- and/or inter-bank systems.

Message Control: Message Control coordinates the passage of messages between the processing components. This involves routing, translation, sequencing, and monitoring. *Routing* accepts a request for delivery of a message to a particular logical function and determines the currently appropriate physical address, thus routing can accommodate changes in the availability and location of functions. *Translation* maps a limited number of protocols from one standard to another. *Sequencing* determines the order in which messages are delivered to recipients on the basis of priorities. *Monitoring* determines the state of the messages within the system at any given time, this includes ensuring the integrity of the message from the time that it is presented by one component until it is accepted by another.

Data Control: Data Control coordinates access, format, and passage of data between application processing functions and the Shared Data Resources. It routes queries and updates to the appropriate component of the Shaared Data Resources, maintains concurring control over the shared data, and returns responses to the requesting applicaiton processing function.

Shared Data Resources: The Shared Data Resource is the component responsible for holding the information for one or more applications. Although this activity is logically centralized in the Shared Data Resource, it may conttain multiple separate components (e.g., each storing different segments of the shared data of different organizations of the shared data). The Shared Data Resource performs two functions: information management and storage management. *Information Management* determines what information must be accessed to satisfy the request, performs the data transformations necessary, and determines how the information is to be stored or retrieved. *Storage Management* determines the physical location of data on the storage devices and controls the actual data movement.

Application Processing: The integrating application-independent layers (external interface, message control, data control, and shared data resources) surround the application processing components. These application processing components can be divided into three classes: transaction processing, information processing, and administrative processing.

Transaction Processing refers to the applications that support specific operational activities, such as order-entry in a manufacturing environment or account balances in a financial environment. This includes many application-specific sub-functions such as validation which ensure that all information needed for processing is present and that the transaction does not violate limits, conditions, or policies that have been established and data update which records the impact of the transaction (e.g., update inventory balance) and initiates other transactions (e.g. billing) and/or manual instructions (e.g., a packing slip).

Information Processing refers to all application processing sub-systems that perform analysis, calculations, or restructuring of the data (e.g., produce a consolidated financial statement). Some example sub-functions include specialized user interaction (e.g. graphics), static reporting, ad hoc reporting, and coordination with external data resources.

Administrative Support provides facilities for the performance of office functions by administrative or managerial personnel. This activity is required to maintain organizational, procedural, or personal information. Example facilities include electronic mail, word processing, correspondence files, and scheduling.

Example Configuration

The hypothetical situation depicted in Figure 9 illustrates some of the various configurations possible using the above principles. The distributed system approach directly follows the "each component is a separate computer" route. These components could range from individual personal computers, to individual mini-computers, or even individual mainframe computers -- depending upon the performance requirements of each component. Appropriate communication facilities (depicted as vertical lines) interconnect the individual components, most likely a high-performance, possibly redundant, local area network.

The mainframe approach places all of the logical components on a single physical computer. There are explicit intra-computer communication procedures (depicted as vertical dotted lines) used to communicate among the components that parallel those used in the distributed system approach. Thus, in principle, identical, or almost identical, software modules can be used on both the distributed system approach and mainframe approach. Although Figure 9 may imply that this approach would be most appropriate for large sites, it may, in fact, be used in very small sites where a single "small mainframe" (e.g., a minicomputer or powerful personal computer system) hosts all of the logical components.

Two other variations, minimal installation and partial installation, are also depicted in Figure 9. In the minimal installation only certain layers are present (e.g., external interface and message control). Note that all of the message control and data control components are connected (depicted as horizontal dark lines). Thus, any transaction entered from Hong Kong can be forwarded by its message control component to the appropriate site for processing that transaction, which could be either New York or London in Figure 9.

In the partial installation all layers are present but not all components. For example, the administrative support components may be present but not the transaction processing components. Thus, certain activities may be completely processed locally, such as word processing, but other activities, such as order entry, will need to be forwarded to another site as in the minimal installation situation.

The minimal installation makes it possible to provide all of the company's information resources in a consistent compatible manner to all locations, no matter how small. The partial installation makes it possible to customize a configuration to best meet the high-volume high-performance requirements of a site yet still provide consistent compatible access to other resources that may be of lower volume and less critical. Furthermore, evolution from minimal to partial to full, or vice versa, can be

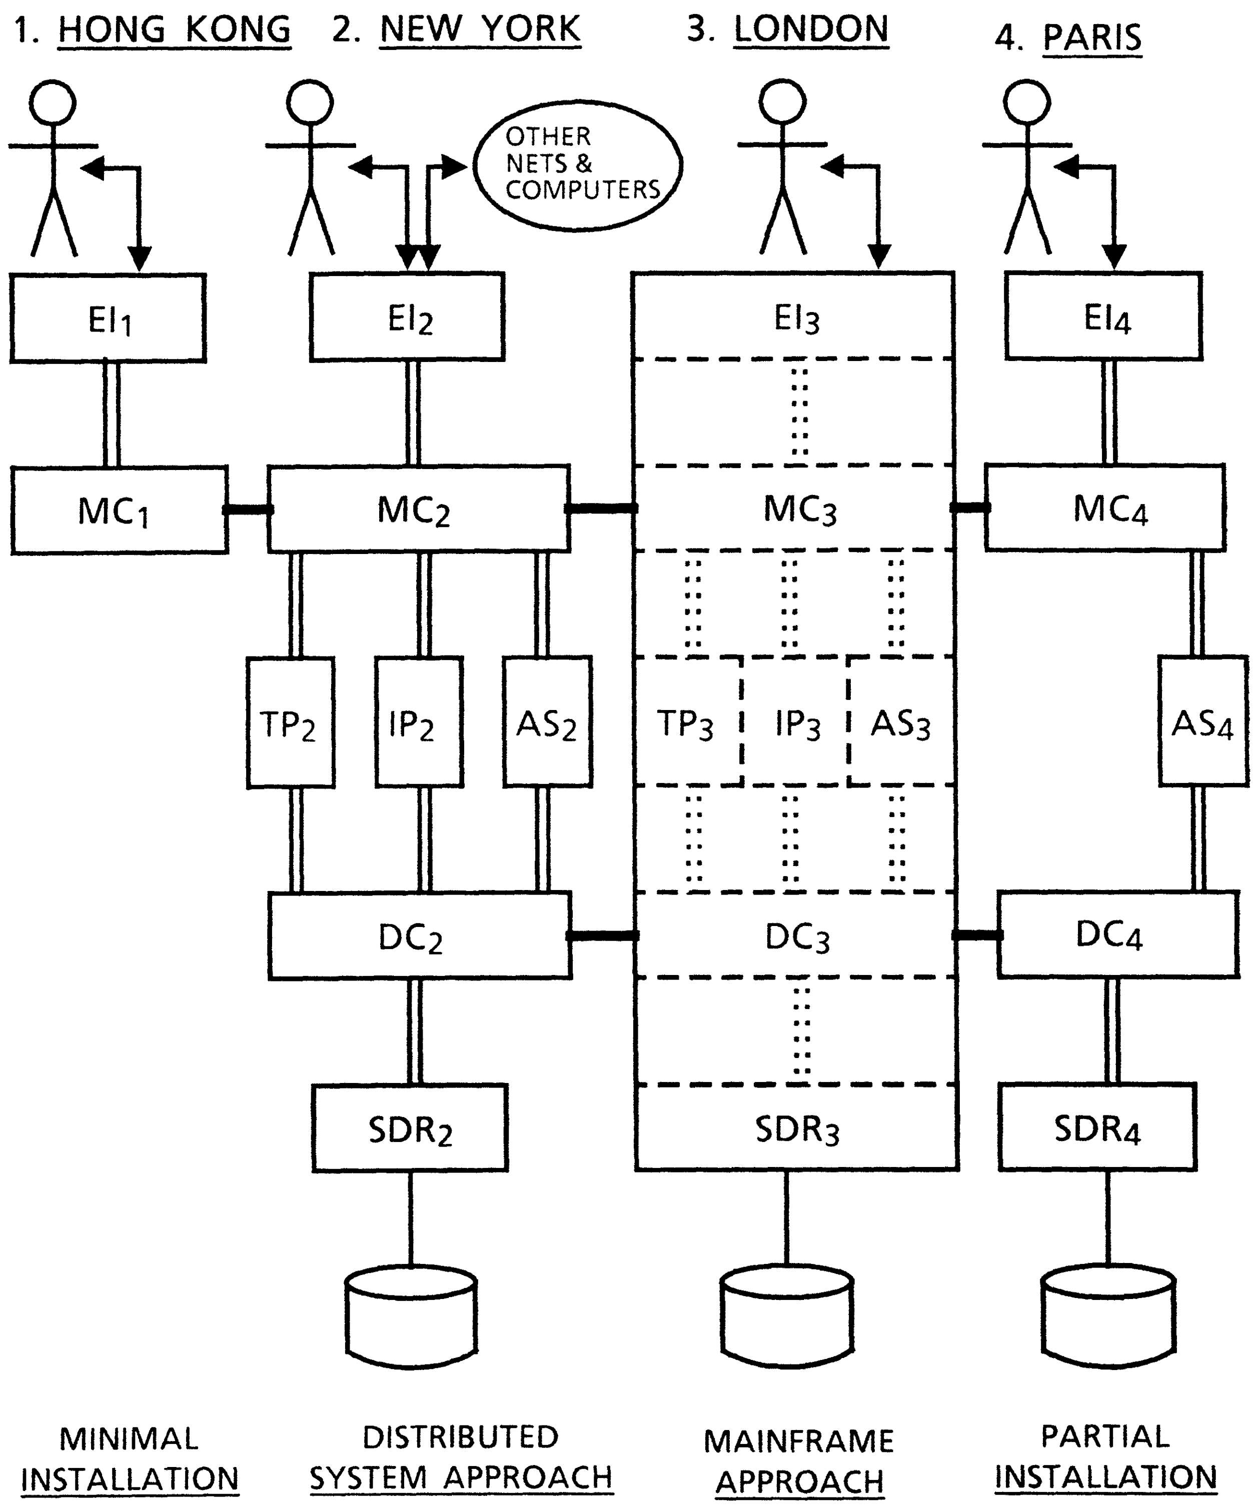
1. HONG KONG
2. NEW YORK
3. LONDON
4. PARIS
OTHER NETS & COMPUTERS
EI1
EI2
EI3
EI4
MC1
MC2
MC3
MC4
TP2
IP2
AS2
TP3
IP3
AS3
AS4
DC2
DC3
DC4
SDR2
SDR3
SDR4
MINIMAL INSTALLATION
DISTRIBUTED SYSTEM APPROACH
MAINFRAME APPROACH
PARTIAL INSTALLATION

Figure 9

accomplished in a smooth manner as the needs and structure of the organization change.

One more point is worth noting. Even the apparently full configurations, depicted as New York and London, may need to communicate with other sites for components that are not present locally. For example, if TP_2 and TP_3 support different types of transactions, New York will have to forward TP_3 type transactions to London.

To illustrate the cooperative natt.re of this achitecture, consider a transaction of type TP_3 that needs to operate upon data stored in shared data resource SDR_4 that is entered from a terminal connected to EI_1 in Hong Kong. The message control component MC_1 realizes that TP_3 is not local and forwards the request to its compatriot MC_3 in London for processing. TP_3 tells its data control component DC_3 that it needs certain information to process this transaction. DC_3 realizes that the data needed is stored on the shared data resource SDR_4 so it forwards the request to its colleague DC_4 who coordinates access to SDR_4. This path is then "rewound" as the data from SDR_4 is provided to TP_3 who processes it and in turn provides the appropriate response to the user in Hong Kong via EI_1. Although in reality we would expect most requests to be much simpler and only involve a single site, the situation could also be much more elaborate if the application required data from multiple shared data resources and/or it automatically generated one or more new transactions that circulated through the system to accomplish their functions.

The important thing to note about this architecture, and this simple example, is that it allows tremendous organizational flexibility and autonomy yet provides powerful capabilities for integration, cooperation, and coordination. Although this discussion has been somewhat abstract and idealized, given current-day technical shortcomings, there are many organizations moving in this general direction already, such as Citibank's Foundation Software Architecture and the Internal Revenue Service's Vision 2000 Technical Architecture.

Conclusions

The key message of this section is that there is significant good news to look forward to from information technology. The advances that are forecast provide for significant improvement through dramatic reductions in cost and improvement in speed performance and functionality of such a magnitude as to provide opportunities that cannot be fully identified at this time,. This is similar to the way that the advances in telephone technology made possible new ways to run organizations that were not foreseen at the time the telephone technology was first introduced.

Some of the most revolutionary information technology developments that we are expecting involve vastly improved and more convenient human interface to the systems, the ability to support the cognitive activities of the users, and the ability to access distributed information and make use of the data semantics to gain a more compatible and comprehensive view across the entire organization.

It is important to realize that many of these benefits are not self-applying. Attention and action on the part of management is critical. A parable I have found effective involves giving a power saw to a lumberjack. The lumberjack picks up this saw with a groan and attempts to move the saw back and forth in a cutting motion, commenting upon how much more difficult it is because it so much heavier and because the teeth on the chain keep moving as he moves the saw back and forth. The problem is that nobody explained that there is an ON switch that would enable the power saw to perform the cutting motion itself. In a similar vein, merely "dumping" new information technology in an organization does not necessarily mean that it will be used in a productive and effective manner. This repesents the major challenge for management of the 1990's.