DISTRIBUTED PROCESSING AND LIMITS TO ITS GROWTH

Gordon Bell
Vice President, Engineering, Digital Equipment Corp.
Professor, on leave, Carnegie-Mellon University

ABSTRACT

A fifth generation computer, can be fabricated on a very large scale integrated circuit (VLSI). Lower cost and increased use disperses computers in a manner analogous to the ubiquitous fractional horsepower motor. Distributed processing to interconnect dispersed computers is essential in order to avoid overloading people with information transmission and translation tasks.

The factors that affect and limit distributed processing are: physical technology and design complexity, ideas for new computer structures, basic tools to build applications, networking and other standards, useful applications, algorithms, and the human interface to the end user. A hierarchical, interconnecting model for distributing processing is based on established central and group level mini-computers, and evolving, personal computers.

DISTRIBUTED PROCESSING

Distributed processing matches computer systems to information processing needs (i.e. processing, memory, switching, transmission and transduction needs) on a geographical or organizational basis, and interconnects individual computers to form a single, integrated network so that related programs can share and transmit data among the computer nodes. The objectives are:

- to allow either local autonomy or central control of the various distributed parts;
- to provide an evolving open-ended system so that the development and installation of the parts can proceed in a quasi-independent fashion;
- to allow purchase and installation of hardware, taking advantage of timely, reduced hardware cost; and
- to build on and communicate with central systems, fully dispersed group-level mini-computer systems, and emerging personal computers.

Distributed processing is inherently hierarchical based on the principles that govern human organizational structures. In an organization, computers supplement their human, information processing counter-parts. As computers become better matched to people and organizations, and as people and organizations become more familiar with computers, an individual can interact directly with at least one computer and indirectly with group-level computers serving various functions of the organizational hierarchy. The opportunity of more egalitarian access to data provided by distributed processing may led to a change of the large organization from hierarchical to wider, functional matrix structures. Large organizations need to interconnect the hierarchy of computers for:

- communication among computer with dumb and intelligent terminals using large, central computers;
- organization of central, group and individual sites;
- a functional activity such as word processing or order processing; and
- a specialized computer-based function such as archiving, typesetting, message switching, and electronic mail.

FORCES CREATING DISTRIBUTED PROCESSING

Rapid evolution of semiconductor and magnetic recording technologies have forced computers improvements along paths of:

1. constant cost, with increased performance and productivity for evolutionary use;
2. reduced cost, with constant performance permitting new uses commensurate with the lower cost; and
3. higher cost and performance structures permitting radically new applications.

Costs for nearly all other forms of information processing are because they are labor intensive. Traditional storage, processing, and transmission in libraries and postal systems are increasingly soaring. Simple word processing computers that replace typewriters save the time-consuming process of correcting errors. When groups associated with information processing start using computers a positive feedback, learning curve effect begins further increasing computer markets and uses, and lowering costs.

The industry groups supplying these products and services include:

- computers - mainframe, minicomputers, personal computers and computer services;
- semiconductors - nearly all LSI components are either memory or a computer processor;
- communications - conventional voice and data, new packet networks and associated services;
- television and cable TV - stand-alone use with TV sets (eg. games, home computers) and as an alternative to conventional communication;
- office equipment - typewriters, copiers, and mechanical office equipment are increasingly electronic; and
- control - gears, cams and levers, and mechanisms for control will become electronic, limited only by transducers and sensors.

LIMITS AND PROBLEM AREAS OF DISTRIBUTED PROCESSING

Ultimately all information processing will be computer based. Presently the speed of the evolution is limited by two factors: technical solutions to distributed processing problems and user assimilation.

PHYSICAL TECHNOLOGY

Semiconductors and magnetic recording technology provide the basis for cost and performance improvements. Although, extrapolations too far into the future are generally dangerous, the following technological rates of change, based on the past ten years, will continue for at least five years:

TECHNOLOGY (PERFORMANCE)	YEARLY-RATE OF CHANGE FACTOR
semiconductor memory density	2.0
semiconductors, random logic	1.4-1.6
core memory density improvement	1.3
magnetic disk recording density	1.3-1.4
magnetic tape data-rate	1.25
magnetic tape density	1.2

TECHNOLOGY (COST)	YEARLY-RATE OF CHANGE FACTOR
memory price reduction	0.7
computer system cost reduction	0.8
crt terminal cost reduction	0.85
communication cost/bit transmitted reduction	0.9
packaging (cost/vol.) and power (cost/watt)	1.0
communication line cost <u>increase</u>	1.12
paper cost <u>increase</u>	1.12

Semiconductor technology, shared among several buyers groups, eg. consumer, communications, computers, has a faster rate of improvement than other technologies. Slower evolution has occurred in magnetic recording density because there is only one user, the computer industry. Widely used, well developed technologies, such as CRT's, previously improved for the mass television market are scarcely affected by their increasing use in computers. Costs of paper and communication lines increase with inflation.

Physical transducers that sense temperature, pressure and control power flow are slow to evolve, limiting computer use in automotive applications. Even the most widely used computer equipment, such as keyboards, printing devices and communications devices, evolve slowly by comparison with semiconductors.

COMPLEXITY OF SEMICONDUCTOR DESIGN

Gordon Moore of Intel, observed that the effort required to design semiconductors has doubled each 2-2/3 years since 1962, when a circuit only took 3 man months. 1979 circuits required 21 man years and 1982 circuits will take about 45 man years. While it is easy to conceive of organizing a team of 7 to complete a design in 3 years, the same time task by 15 people is difficult to imagine. Better management and design partitioning is required in order to avoid a drastic loss of productivity and quality that would increase the design effort even more. With one million circuits on a chip by 1982, new methodologies will be required to fully utilize VLSI's potential.

Because of the concern and numerous approaches being pursued, I am confident that it will only take another two semiconductor generations (six years) to solve the VLSI design complexity problem. Although we do not have a good measure of circuit complexity, a given circuit description is far less complex than the largest programs (e.g. a million bit, or 128 Kbyte program is not especially large).

IDEAS ABOUT WHAT TO BUILD

New directions in computer structures are difficult to predict by simply looking at conventional machines. Current limiting factors point to needed innovations. Applications involving two dimensional signal processing for pictures appear to require a different processor design, and speech signal analysis requires vector processing. A general purpose processor could emerge from these alternatives for one-and two-dimensional arrays:

- arrays of conventional microprocessors;
- application specific, functional processors;
- bit array processors to operate directly on the array data structures, including arrays, or associative processing;
- processing associated with memory; and
- data flow architectures.

We must get beyond the simple standards required for Packetnets and Internets to define protocols for passing high level messages, such as electronic mail, among computers. Office based applications, centered around text processing, electronic mail, user typesetting, office processing, and electronic filing, all require significant user level standards. Using only lower level communications protocol standards will cause a combinational explosion of high level protocol changing gateways. This leads to added overhead, extra development, delay, incompatibility, and often, misinterpretation of messages.

In the low priority area of intra-computer architecture, the U. S. Government has standardized on the existing defacto standard, the IBM Channel, as the means of interconnecting mass storage to computers. Unfortunately this act of standardization will limit change into newer systems architectures.

USEFUL APPLICATIONS AND DISTRIBUTING THEM

Decisions to use the major applications centered around office automation are very complex. Justifying an application generally requires an understanding of both computer systems (beyond that provided by manufacturers) and the organizational structure of individuals and group users. Although electronic mail seems right, measurements of increased productivity, decreased paper flow, better decision-making, efficiency of communication, and the creation of excess communication are hard to make. To my knowledge, they don't exist.

BASIC TOOLS TO BUILD APPLICATIONS

Coupling knowledgeable user needs to machine development produces more capable, yet harder to understand systems: a paradox in the attempt to build highly capable and easy to use systems. The popularity of the Bell Labs UNIX System is a testimony to a single, consistent, easy to use language, that is described in a small manual. The popularity of APL and BASIC systems can be similarly explained. Although one would expect that additional capabilities (memory) would make the user interface simpler, few good examples are known. The time to build a given application using the multitude of systems/databases/languages is highly variable, indicating a continued lack of understanding of the design process.

NETWORK AND OTHER STANDARDS

Because standards are evolving, the current situation of distributed processing among countries and vendor systems is a disaster. International protocol standards provided by manufacturers (Internets) and by various common carriers for Packetnets which are called by the same name, are fundamentally different and incompatible. Many standards mean no standards.

Given that few measures exist to rationalize, simple stand-alone applications, justifying a distributed network becomes a work of art. Tools have only recently become available for a system manager or developer to distribute the database, processing, and intercommunications over several systems. In the specific case of distributed processing for electronic mail, the results are encouraging but a general solution has not yet emerged.

An underlying difficulty of building applications beyond the generic office automation described above exists because problems are solved by patch-work. Usually programmers with computer science (computer engineering) training and a representative of a particular discipline (eg. accounting, mechanical engineering) put a solution together to get something started. This results in sub-optimal designs. In order to use the computer as a component of systems they design, rather than as a simple tool for problem solving, computer science must take on a pure role, like physics, and each of the disciplines take the responsibility for training people and engineering the systems within its own discipline.

ALGORITHMS

There are many cases of the adage: "It is better to work smarter rather than work harder". If always exponentially improving, technology will eventually permit solving a particular problem in a reasonable time, e.g. a 24 hour advanced weather forecast must be solved in less than 24 hours or an exponentially increasing machine population will be required. However, at a given time, algorithms limit when a problem can be solved and whether it is economically feasible.

HUMAN INTERFACE

The interface between the system and the final user is a barrier in the same way that a root system for building applications programs is a barrier to building applications. Adding more functions so that an application will perform better is generally accompanied by increased complexity requiring more documentation and training. The lack of standards at the user interface will limit getting the payoff inherent in a given system or set of systems, and may cause adverse user reaction. For example, word processing, electronic mail and user typesetting systems are all likely to have different syntax, semantics, manuals, training and procedures for dealing with the same text.

A DISTRIBUTED PROCESSING ENVIRONMENT

Proliferation of dispersed computing forces interconnection, hence distributed processing, so that human users don't have to become information carriers and translators between the different systems they use. Communication within and between organizations with common carrier networks is provided via an interconnected hierarchy.

INTERCONNECTING THE COMPONENTS

The three types of computers in a given organization will be connected via high bandwidth links in what may appear to be a hierarchical structure. In addition, clusters may be connected on a fixed basis. The alternative interconnect possibilities are:

- ethernets or rings to interconnect all terminals and computers with specialized terminal concentrators;
- evolution of phone circuit switches using digital techniques for both voice and data;
- packetnet switching; and
- direct interconnection among the computers with routing through each computer.

CENTRAL COMPUTERS

The top most computers of the hierarchy will evolve from the current, highly central computation facilities. These machines store most of the data and do most of the computing in today's organizations. Given the difficulty of migrating files and work from these machines, the emphasis within the centers will be interconnection among the machines within each center, creating in the short run, even larger data bases. The tight interconnection among the central computers will also permit trade-offs among cost, reliability, performance, and evolving performance, for a given application or set of applications. In order to get the economy of scale required to support the large human organizations that attend central computers, their functions will have to be specialized (e.g. front ends for handling many communications lines, and back ending for databases and archiving).

Central computing facilities will continue to be operated by large staffs whose emphasis is on knowledge of the operating systems and getting work done using highly specialized facilities such as CODASYL Databases. The casual user will be dependent on the central systems through the applications. Cost will be high for everything except the storage of very large files, where hardware provides an economy of scale. Programming costs at the center have to be the highest, because the facilities are general purpose and applications are most remote from the ultimate user. The role of central facility will be to provide:

- communications among all the other computers within the organization including gateways between various computer and telecommunications vendors;
- archival file storage;
- unique, sharable facilities such as very high speed computers and printing devices;
- computational functions for the entire organization e.g. electronic mail;
- operation of historical programs and data bases; and
- relatively high cost computing by having to provide generality and service for the worst case.

GROUP LEVEL COMPUTERS

Group level computers are based on the evolution of timeshared and real time minicomputers and cost roughly that of an additional person. Typically these machines support the single function of the group, (eg. order processing, engineering design and data base, laboratory data gathering and analysis, group word processing, single process control) running a single unattended program. Group level computers provide:

- relatively cost effective storage of the group data base;
- unique program(s) aligned with function of the group;
- relatively high performance processing; and
- cost-effective computing through sharing of a common function and specialization of work.

PERSONAL LEVEL COMPUTERS

Personal computers are emerging rapidly, and many believe that they will become the dominant form of computing. Since the only hardware technology for which economy of scale holds is mass storage, and given that all terminals already have embedded computers for control, it is easy to envision adding more primary memory and doing all the computation at the terminal instead of having computation done in any shared facility. A recent, Carnegie-Mellon University personal computer research proposal states:

> "The era of time-sharing is ending. Time-sharing evolved as a way to provide users with the power of a large interactive computer system at a time when such systems were too expensive to dedicate to a single individual...Recent advances in hardware open up new possibilities...high resolution color graphics, 1 mip, 16 Kword, 1 Mbyte primary memory, 100 Mbyte secondary memory, special transducers,...We would expect that by the mid-1980's such systems could be priced around $10,000."

Personal computers provide:

- personal data bases and security;
- more, average computing power, with better response time than shared systems;
- needed processing for the computationally intensive tasks like editing, and speech i/o;
- a program creation environment; and
- relatively higher costs than group level computing, unless the task is very specific and well-matched to the system.

Although both the novice and experienced user relish the independence that the personal computer provides, communications and support by the other levels is equally necessary. Given that we are substantially far from such distributed systems, there are surely additional problems, limits, and opportunities that are yet to be forecast.

DISTRIBUTED PROCESSING AND LIMITS TO ITS GROWTH

Gordon Bell
Vice President, Engineering, Digital Equipment Corp.
Professor, on leave, Carnegie-Mellon University

ABSTRACT

A fifth generation computer, can be fabricated on a very large scale integrated circuit (VLSI). Lower cost and increased use disperses computers in a manner analogous to the ubiquitous fractional horsepower motor. Distributed processing to interconnect dispersed computers is essential in order to avoid overloading people with information transmission and translation tasks.

The factors that affect and limit distributed processing are: physical technology and design complexity, ideas for new computer structures, basic tools to build applications, networking and other standards, useful applications, algorithms, and the human interface to the end user. A hierarchical, interconnecting model for distributing processing is based on established central and group level mini-computers, and evolving, personal computers.

DISTRIBUTED PROCESSING

Distributed processing matches computer systems to information processing needs (i.e. processing, memory, switching, transmission and transduction needs) on a geographical or organizational basis, and interconnects individual computers to form a single, integrated network so that related programs can share and transmit data among the computer nodes. The objectives are:

- to allow either local autonomy or central control of the various distributed parts;
- to provide an evolving open-ended system so that the development and installation of the parts can proceed in a quasi-independent fashion;
- to allow purchase and installation of hardware, taking advantage of timely, reduced hardware cost; and
- to build on and communicate with central systems, fully dispersed group-level mini-computer systems, and emerging personal computers.

Distributed processing is inherently hierarchical based on the principles that govern human organizational structures. In an organization, computers supplement their human, information processing counter-parts. As computers become better matched to people and organizations, and as people and organizations become more familiar with computers, an individual can interact directly with at least one computer and indirectly with group-level computers serving various functions of the organizational hierarchy. The opportunity of more egalitarian access to data provided by distributed processing may led to a change of the large organization from hierarchical to wider, functional matrix structures. Large organizations need to interconnect the hierarchy of computers for:

- communication among computer with dumb and intelligent terminals using large, central computers;
- organization of central, group and individual sites;
- a functional activity such as word processing or order processing; and
- a specialized computer-based function such as archiving, typesetting, message switching, and electronic mail.

FORCES CREATING DISTRIBUTED PROCESSING

Rapid evolution of semiconductor and magnetic recording technologies have forced computers improvements along paths of:

1. constant cost, with increased performance and productivity for evolutionary use;
2. reduced cost, with constant performance permitting new uses commensurate with the lower cost; and
3. higher cost and performance structures permitting radically new applications.

Costs for nearly all other forms of information processing are because they are labor intensive. Traditional storage, processing, and transmission in libraries and postal systems are increasingly soaring. Simple word processing computers that replace typewriters save the time-consuming process of correcting errors. When groups associated with information processing start using computers a positive feedback, learning curve effect begins further increasing computer markets and uses, and lowering costs.

The industry groups supplying these products and services include:

- computers - mainframe, minicomputers, personal computers and computer services;
- semiconductors - nearly all LSI components are either memory or a computer processor;
- communications - conventional voice and data, new packet networks and associated services;
- television and cable TV - stand-alone use with TV sets (eg. games, home computers) and as an alternative to conventional communication;
- office equipment - typewriters, copiers, and mechanical office equipment are increasingly electronic; and
- control - gears, cams and levers, and mechanisms for control will become electronic, limited only by transducers and sensors.

LIMITS AND PROBLEM AREAS OF DISTRIBUTED PROCESSING

Ultimately all information processing will be computer based. Presently the speed of the evolution is limited by two factors: technical solutions to distributed processing problems and user assimilation.

PHYSICAL TECHNOLOGY

Semiconductors and magnetic recording technology provide the basis for cost and performance improvements. Although, extrapolations too far into the future are generally dangerous, the following technological rates of change, based on the past ten years, will continue for at least five years:

TECHNOLOGY (PERFORMANCE)	YEARLY-RATE OF CHANGE FACTOR
semiconductor memory density	2.0
semiconductors, random logic	1.4-1.6
core memory density improvement	1.3
magnetic disk recording density	1.3-1.4
magnetic tape data-rate	1.25
magnetic tape density	1.2

TECHNOLOGY (COST)	YEARLY-RATE OF CHANGE FACTOR
memory price reduction	0.7
computer system cost reduction	0.8
crt terminal cost reduction	0.85
communication cost/bit transmitted reduction	0.9
packaging (cost/vol.) and power (cost/watt)	1.0
communication line cost increase	1.12
paper cost increase	1.12

Semiconductor technology, shared among several buyers groups, eg. consumer, communications, computers, has a faster rate of improvement than other technologies. Slower evolution has occurred in magnetic recording density because there is only one user, the computer industry. Widely used, well developed technologies, such as CRT's, previously improved for the mass television market are scarcely affected by their increasing use in computers. Costs of paper and communication lines increase with inflation.

Physical transducers that sense temperature, pressure and control power flow are slow to evolve, limiting computer use in automotive applications. Even the most widely used computer equipment, such as keyboards, printing devices and communications devices, evolve slowly by comparison with semiconductors.

COMPLEXITY OF SEMICONDUCTOR DESIGN

Gordon Moore of Intel, observed that the effort required to design semiconductors has doubled each 2-2/3 years since 1962, when a circuit only took 3 man months. 1979 circuits required 21 man years and 1982 circuits will take about 45 man years. While it is easy to conceive of organizing a team of 7 to complete a design in 3 years, the same time task by 15 people is difficult to imagine. Better management and design partitioning is required in order to avoid a drastic loss of productivity and quality that would increase the design effort even more. With one million circuits on a chip by 1982, new methodologies will be required to fully utilize VLSI's potential.

Because of the concern and numerous approaches being pursued, I am confident that it will only take another two semiconductor generations (six years) to solve the VLSI design complexity problem. Although we do not have a good measure of circuit complexity, a given circuit description is far less complex than the largest programs (e.g. a million bit, or 128 Kbyte program is not especially large).

IDEAS ABOUT WHAT TO BUILD

New directions in computer structures are difficult to predict by simply looking at conventional machines. Current limiting factors point to needed innovations. Applications involving two dimensional signal processing for pictures appear to require a different processor design, and speech signal analysis requires vector processing. A general purpose processor could emerge from these alternatives for one-and two-dimensional arrays:

- arrays of conventional microprocessors;
- application specific, functional processors;
- bit array processors to operate directly on the array data structures, including arrays, or associative processing;
- processing associated with memory; and
- data flow architectures.

We must get beyond the simple standards required for Packetnets and Internets to define protocols for passing high level messages, such as electronic mail, among computers. Office based applications, centered around text processing, electronic mail, user typesetting, office processing, and electronic filing, all require significant user level standards. Using only lower level communications protocol standards will cause a combinational explosion of high level protocol changing gateways. This leads to added overhead, extra development, delay, incompatibility, and often, misinterpretation of messages.

In the low priority area of intra-computer architecture, the U. S. Government has standardized on the existing defacto standard, the IBM Channel, as the means of interconnecting mass storage to computers. Unfortunately this act of standardization will limit change into newer systems architectures.

USEFUL APPLICATIONS AND DISTRIBUTING THEM

Decisions to use the major applications centered around office automation are very complex. Justifying an application generally requires an understanding of both computer systems (beyond that provided by manufacturers) and the organizational structure of individuals and group users. Although electronic mail seems right, measurements of increased productivity, decreased paper flow, better decision-making, efficiency of communication, and the creation of excess communication are hard to make. To my knowledge, they don't exist.

BASIC TOOLS TO BUILD APPLICATIONS

Coupling knowledgeable user needs to machine development produces more capable, yet harder to understand systems: a paradox in the attempt to build highly capable and easy to use systems. The popularity of the Bell Labs UNIX System is a testimony to a single, consistent, easy to use language, that is described in a small manual. The popularity of APL and BASIC systems can be similarly explained. Although one would expect that additional capabilities (memory) would make the user interface simpler, few good examples are known. The time to build a given application using the multitude of systems/databases/languages is highly variable, indicating a continued lack of understanding of the design process.

NETWORK AND OTHER STANDARDS

Because standards are evolving, the current situation of distributed processing among countries and vendor systems is a disaster. International protocol standards provided by manufacturers (Internets) and by various common carriers for Packetnets which are called by the same name, are fundamentally different and incompatible. Many standards mean no standards.

Given that few measures exist to rationalize, simple stand-alone applications, justifying a distributed network becomes a work of art. Tools have only recently become available for a system manager or developer to distribute the database, processing, and intercommunications over several systems. In the specific case of distributed processing for electronic mail, the results are encouraging but a general solution has not yet emerged.

An underlying difficulty of building applications beyond the generic office automation described above exists because problems are solved by patch-work. Usually programmers with computer science (computer engineering) training and a representative of a particular discipline (eg. accounting, mechanical engineering) put a solution together to get something started. This results in sub-optimal designs. In order to use the computer as a component of systems they design, rather than as a simple tool for problem solving, computer science must take on a pure role, like physics, and each of the disciplines take the responsibility for training people and engineering the systems within its own discipline.

ALGORITHMS

There are many cases of the adage: "It is better to work smarter rather than work harder". If always exponentially improving, technology will eventually permit solving a particular problem in a reasonable time, e.g. a 24 hour advanced weather forecast must be solved in less than 24 hours or an exponentially increasing machine population will be required. However, at a given time, algorithms limit when a problem can be solved and whether it is economically feasible.

HUMAN INTERFACE

The interface between the system and the final user is a barrier in the same way that a root system for building applications programs is a barrier to building applications. Adding more functions so that an application will perform better is generally accompanied by increased complexity requiring more documentation and training. The lack of standards at the user interface will limit getting the payoff inherent in a given system or set of systems, and may cause adverse user reaction. For example, word processing, electronic mail and user typesetting systems are all likely to have different syntax, semantics, manuals, training and procedures for dealing with the same text.

A DISTRIBUTED PROCESSING ENVIRONMENT

Proliferation of dispersed computing forces interconnection, hence distributed processing, so that human users don't have to become information carriers and translators between the different systems they use. Communication within and between organizations with common carrier networks is provided via an interconnected hierarchy.

INTERCONNECTING THE COMPONENTS

The three types of computers in a given organization will be connected via high bandwidth links in what may appear to be a hierarchical structure. In addition, clusters may be connected on a fixed basis. The alternative interconnect possibilities are:

- ethernets or rings to interconnect all terminals and computers with specialized terminal concentrators;
- evolution of phone circuit switches using digital techniques for both voice and data;
- packetnet switching; and
- direct interconnection among the computers with routing through each computer.

CENTRAL COMPUTERS

The top most computers of the hierarchy will evolve from the current, highly central computation facilities. These machines store most of the data and do most of the computing in today's organizations. Given the difficulty of migrating files and work from these machines, the emphasis within the centers will be interconnection among the machines within each center, creating in the short run, even larger data bases. The tight interconnection among the central computers will also permit trade-offs among cost, reliability, performance, and evolving performance, for a given application or set of applications. In order to get the economy of scale required to support the large human organizations that attend central computers, their functions will have to be specialized (e.g. front ends for handling many communications lines, and back ending for databases and archiving).

Central computing facilities will continue to be operated by large staffs whose emphasis is on knowledge of the operating systems and getting work done using highly specialized facilities such as CODASYL Databases. The casual user will be dependent on the central systems through the applications. Cost will be high for everything except the storage of very large files, where hardware provides an economy of scale. Programming costs at the center have to be the highest, because the facilities are general purpose and applications are most remote from the ultimate user. The role of central facility will be to provide:

- communications among all the other computers within the organization including gateways between various computer and telecommunications vendors;
- archival file storage;
- unique, sharable facilities such as very high speed computers and printing devices;
- computational functions for the entire organization e.g. electronic mail;
- operation of historical programs and data bases; and
- relatively high cost computing by having to provide generality and service for the worst case.

GROUP LEVEL COMPUTERS

Group level computers are based on the evolution of timeshared and real time minicomputers and cost roughly that of an additional person. Typically these machines support the single function of the group, (eg. order processing, engineering design and data base, laboratory data gathering and analysis, group word processing, single process control) running a single unattended program.
Group level computers provide:

- relatively cost effective storage of the group data base;
- unique program(s) aligned with function of the group;
- relatively high performance processing; and
- cost-effective computing through sharing of a common function and specialization of work.

PERSONAL LEVEL COMPUTERS

Personal computers are emerging rapidly, and many believe that they will become the dominant form of computing. Since the only hardware technology for which economy of scale holds is mass storage, and given that all terminals already have embedded computers for control, it is easy to envision adding more primary memory and doing all the computation at the terminal instead of having computation done in any shared facility. A recent, Carnegie-Mellon University personal computer research proposal states:

> "The era of time-sharing is ending. Time-sharing evolved as a way to provide users with the power of a large interactive computer system at a time when such systems were too expensive to dedicate to a single individual...Recent advances in hardware open up new possibilities...high resolution color graphics, 1 mip, 16 Kword, 1 Mbyte primary memory, 100 Mbyte secondary memory, special transducers,...We would expect that by the mid-1980's such systems could be priced around $10,000."

Personal computers provide:

- personal data bases and security;
- more, average computing power, with better response time than shared systems;
- needed processing for the computationally intensive tasks like editing, and speech i/o;
- a program creation environment; and
- relatively higher costs than group level computing, unless the task is very specific and well-matched to the system.

Although both the novice and experienced user relish the independence that the personal computer provides, communications and support by the other levels is equally necessary. Given that we are substantially far from such distributed systems, there are surely additional problems, limits, and opportunities that are yet to be forecast.