

15 empty rectangular boxes stacked vertically, likely for a mailing label.

**Large-Scale
Systems
Directions:
Large IBM and
Software-
Compatible
Mainframes**



INPUT[®]

43 Landings Drive, Mountain View, CA 94043. (415) 960-3990

U-IR3
1986
C.1



Digitized by the Internet Archive
in 2015

<https://archive.org/details/largescalesystem03unse>

LARGE-SCALE SYSTEMS DIRECTIONS:
LARGE IBM AND
SOFTWARE-COMPATIBLE MAINFRAMES

U-IR3
1986
C.1

AUTHOR

LARGE-SCALE SYSTEMS DIRECTIONS

TITLE

LARGE IBM & SOFTWARE COMPATIBLE MAINFRAMES

DATE
LOANED

BORROWER'S NAME

DECEMBER 1986



Published by
INPUT
1943 Landings Drive
Mountain View, CA 94043
U.S.A.

Information Systems Program (ISP)

Large-Scale Systems Directions: Large IBM and Software-Compatible Mainframes

Copyright ©1986 by INPUT. All rights reserved.
Printed in the United States of America.
No part of this publication may be reproduced or
distributed in any form or by any means, or stored
in a data base or retrieval system, without the prior
written permission of the publisher.

**LARGE-SCALE SYSTEMS DIRECTIONS:
LARGE IBM AND SOFTWARE-COMPATIBLE MAINFRAMES**

ABSTRACT

This report addresses the broad issues that need to be understood when an institutional entity attempts to quantify or qualify the value of its data/information/knowledge systems. The way in which the data/information/knowledge are distributed and how they are used are central to the realization of that value. Before one attempts to implement new hardware/software technology, one must have a thorough understanding of how that technology will be applied to benefit the productivity of the entire enterprise.

Residual values for selected large-scale IBM and software-compatible mainframes are updated based on recent announcements and other factors affecting that value.

This report contains 55 pages, including 10 exhibits.

**LARGE-SCALE SYSTEMS DIRECTIONS:
LARGE IBM AND SOFTWARE-COMPATIBLE MAINFRAMES**

CONTENTS

	<u>Page</u>
I INTRODUCTION.....	1
II SYSTEMS DEVELOPMENT FOR THE NEW "PERIPHERALS"	5
A. Distributed Systems Development	5
1. Development Center	5
2. Information Center	8
3. Departmental Systems	9
4. Personal Computers	10
5. Computer/Communications Networks	13
B. Disturbing Implications	18
1. Systems versus Peripherals	18
2. Top-Down versus Bottom-Up	23
3. Data Bases versus Information Flow	24
4. Quality versus Chaos	30
C. RISCs and Residual Values	33
1. Dinosaurs, Dragons, and the Ghost of von Newmann	33
2. Ephemeral "Solutions" and the Perennial Backlog	34
3. Hardware Chickens and Software Eggs	35
D. The Residual Cost of the Present	36
III RESIDUAL VALUE FORECASTS	39
A. Announcements	39
B. Review of Used Market Activity	42
C. Assumptions	43
D. Projected Used Market Prices and Residual Values	46

**LARGE-SCALE SYSTEMS DIRECTIONS:
LARGE IBM AND SOFTWARE-COMPATIBLE MAINFRAMES**

EXHIBITS

		<u>Page</u>	
II	-1	The Connectivity Problem	15
	-2	IBM's Network Strategy	20
	-3	Entropy Problems	28
III	-1	Projected Residual Value As a Percent of Vendor List Price As of January 1, 1987	48
	-2	Projected Used Market Retail As of January 1, 1987	49
	-3	IBM 308X Series Projected Residual Value As a Percent of Vendor List Price	50
	-4	IBM 3090 Series Projected Residual Value As a Percent of Vendor List Price	52
	-5	IBM 43XX Series Projected Residual Value As a Percent of Vendor List Price	53
	-6	Amdahl 58XX Series Projected Residual Value As a Percent of Vendor List Price	54
	-7	NAS AS Series Projected Residual Value As a Percent of Vendor List Price	55

I INTRODUCTION

- While this report series has been titled Large-Scale Systems Directions, INPUT's emphasis has always been on computer/communications networks--large-scale systems in the broadest sense of the words. There seems to be a propensity within the IS community to focus attention first on one level of processing and then on another: before 1981 everyone concentrated on mainframes, in 1983 PCs started to receive all of the attention, and now "departmental systems" are all the rage. The result has been a hodgepodge of computer hardware which is now supposed to be integrated into a "large-scale system" through the magic of connectivity.
- For those who have suffered through the evolution of computer/communications networks over the last 15 years, this report is dedicated to the bit of wisdom expressed by Ken Olsen (CEO of DEC) at the time VAXmate was announced: "The problem is that everyone's been going about this backwards--buying lots of computers and then trying to connect them together. We have to start thinking about the computers as peripherals. You start with the network, then you hang the computers on later."
- This point of view expresses very clearly a fundamental fact about computer/communications networks--the flow of information is substantially more important than the data locked up in host computers (or scattered on a bunch of PCs). Chapter II of this report will analyze some of the potential problems which result from going about network development "backwards" and the residual costs which may be inherent in starting to view large-scale

mainframes as "peripherals" if this is, indeed, the proper approach to networking.

- Chapter III of this report provides an update of IBM and software-compatible mainframe residual values. Chapter III also contains a brief summary of the IBM 3790 announcement, which though not a large-scale system is of significant importance to the overall computer/communications network hierarchy and, in the long term, may impact the residual values of mainframes and their associated peripherals.
- In fact, INPUT believes IBM is nearly ready to support distributed processing on something more than a superficial basis. For that reason, next year INPUT will present a new report series--Distributed and Office Systems Directions. This series will address Levels II and III of the processing hierarchy (mini-computers and intelligent workstations) while Large-Scale Systems Directions will continue to concentrate on Level I mainframes, major peripheral systems, and networking. The two report series will be companions in most senses of the word, but there will be a significant difference:
 - While the body of Distributed and Office Systems Directions will concentrate on IBM's strategy in departmental processors and intelligent workstations, residual value forecasts will not be a regular part of that report series. There will be specific analysis on the impact of office systems on productivity with special emphasis on measurement and performance analysis.
 - Next year, the three Distributed and Office Systems Directions reports will contain sections which will analyse performance at the following levels:
 - At the Human-Machine juncture an analysis will be made of the impact of intelligent workstations (personal computers) on the productivity of the various categories of office employees.

- . At the Work Unit level the impact of departmental processors (and LANs) on the productivity of identified work units will be analyzed.
 - . At the Institutional level the impact of information systems technology (computer/communications networks at all levels) on institutional (corporate) performance will be presented with special emphasis on cost justification.
- Therefore, the emphasis of Distributed and Office Systems Directions will be on the inherent value of existing and planned information systems. This is an area which has all too frequently been neglected.

II SYSTEMS DEVELOPMENT FOR THE NEW "PERIPHERALS"

A. DISTRIBUTED SYSTEMS DEVELOPMENT

I. DEVELOPMENT CENTER

- The development center is the bastion of the computer "professionals." Today these "professionals" fall into many different occupational categories and span a wide range of skills. Without attempting to establish any semblance of a professional pecking order (although every development center has its own caste system), the following list is representative of the types of personnel one is likely to find:
 - There are systems programmers who establish the operating systems environment in which applications are developed. They are the "bit fiddlers" and normally prefer to deal with binary, octal, and hexadecimal arithmetic rather than decimal.
 - The key to the operating environment is that systems are being developed in an environment of virtual machines, virtual memory, virtual terminals, virtual connections, virtual storage, and virtual access methods.
 - There is a price associated with this virtual environment when it comes time to buy the real machines, memory, terminals, etc., and the systems programmers do not know what the price is.

- There are systems designers and analysts in the development center who are concerned with the overall design of major applications systems to operate within the virtual environment. They generally do not understand the difference between virtual and real in terms of the hardware and operating systems with which they are dealing. Indeed, they have been told not to "concern" themselves with the physical limitations of the central host hardware/software or the network itself.
- In addition, it is frequently the case that the designers do not concern themselves with the physical (and mental) limitations of the programming staff which must implement the system or the quality of the data necessary to drive the system. Living in a virtual systems design world can have unfortunate consequences:
 - There can be emphasis on the design procedure itself which results in undue focus on design methodology and documentation rather than the practical reality of implementation.
 - Systems can be designed and specified which cannot be implemented (programmed) within any reasonable time schedule (or before the specifications change).
 - Systems can be developed which have unsatisfactory performance in the real hardware/software environment or which cannot be supported due to the quality of the data.
- There are programmers who are responsible for programming the major applications systems and who are several levels removed from both the hardware/software real environment and the requirements of the users of the system which is under development. Most of the productivity tools (such as 4GLs) to improve productivity are eschewed by "professional" programmers who prefer to use second or third generation

languages. Among programmers themselves there is a pecking order which is roughly as follows: assembly language programmers are the elite, PL/I and Fortran programmers look down on COBOL programmers, who in turn look down on those using DBMSs, and they all agree Basic is not a language for professional programmers.

- There are also a growing number of specialists who concern themselves with data. Depending on the particular organization, they may vary from clerical data base administrators to those involved with data modeling and entity relationships and the finer distinctions of relational normal forms. At the outer fringes of the data specialists are those concerned with the problems of distributing data bases over computer/communications networks. There are some problems associated with distributed data bases (such as security) which are beyond the current state of computer and mathematical science. These are not esoteric considerations; they are fundamental (and should be preliminary) to "building networks."
- While the organization and staffing of the development center may vary tremendously across institutions and there may be considerable internal conflict within the organization, the development center presents a united front against the outside world. The general attitude is one of defending a valuable asset (data base and information systems) against an "enemy." Unfortunately, the "enemy" is frequently the legitimate users of this asset, including executive management of the institution.
- In General Systems Theory, progressive centralization is described as the natural tendency for all systems to develop "leading parts" which control other levels of the system. The development center is the focus of a highly centralized approach to the development of information systems, and the large host mainframe is not considered "peripheral" to anything--it is the center of the IS universe.

2. INFORMATION CENTER

- Information centers came into being primarily because the development center could not (or would not) provide rapid turnaround for ad hoc reporting and "simple" applications. While this lack of responsiveness on the part of the central IS function had been at issue for a number of years, it was not until users started to install personal computers to satisfy their own reporting requirements that the concept of information centers came into being. Information centers were promoted by IBM as a means for the IS function to obtain some measure of control over the "foreign" hardware/software products which were beginning to proliferate in many companies.
- Information centers are considered "peripheral" to the mainstream information systems activities; and, in some ways, have been viewed by the development center as a means of buffering itself from "unreasonable" user requests. While the organization and functions of the information center vary greatly, the primary offerings include:
 - Advice and training on the acquisition of personal computer hardware/software.
 - Quick turnaround on ad hoc reporting and requests for "simple" applications.
 - Extracts from corporate data bases and advice and training in the use of these data bases.
 - Prototyping of new applications with end users.
- The tools of the information center--4GLs, relational DBMSs, and the whole assortment of personal computer software--are generally viewed with suspicion by the central IS function. There is fear that the "quick and dirty" systems will be inherited for maintenance and the contamination of corporate data bases will be the end result of end-user computing.

- The evolution of the information center represents another identified concept of GST--all systems demonstrate progressive differentiation in that some of the parts become more specialized. Thus, the information center is differentiated from the development center based on the types of systems (usually smaller) which are developed and the data (usually planning as opposed to operating) which are employed. Conceptually, the information center is a satellite of the central processing facility--separate, but very much under control of the "leading part."

3. DEPARTMENTAL SYSTEMS

- What used to be known as distributed processing has gone through several iterations of terminology, the latest of which is departmental systems. Originally distributed processing started with networked minicomputers. And, as INPUT has repeated so often over the last 10 years, IBM's primary networking strategy has been directed toward squeezing minicomputers out of the processing hierarchy. Earlier this year, Ken Olsen was quoted as stating: "If customers aren't buying our networking products, they don't understand what we are saying." What DEC is trying to say is that minicomputers have an essential role between mainframes and personal computers in computer/communications networks. Many IS departments have resisted either understanding or accepting minicomputers in the processing hierarchy.
- More recently, microprocessor-based office automation products have demanded "connectivity," which has, in turn, forced recognition that some intermediate level between desktop and mainframe may be desirable and perhaps even necessary. This leads to the term "departmental processor" in lieu of "minicomputer" which had become a bad (and misunderstood) term through nearly 15 years of SNA. It should be recognized that "connectivity" is a meaningless term for what is really happening. Once again GST states that all systems demonstrate progressive integration in which the parts become more dependent on the whole. Such integration is inevitable and, therefore, predictable.

- The early distributed processing installation of minicomputers was practically exclusively in scientific and engineering departments where the "open shop" programmers became reasonably proficient with both the hardware and software of the system. In fact, it is probable that the average minicomputer programmer understands computer hardware and software much better than does the average COBOL programmer. On the other hand, the departmental processors which are evolving from office systems and PCs will have casual users with less than rudimentary knowledge of operating systems, programming languages or even the simplest applications development tools. The software and support requirements for these two user sets are substantially different, and so is the way they view the central data processing facility.

- Whether departmental systems are driven by the distribution of processing and data from host mainframes or grow from within the departments, they have one disturbing attribute--they tend to restrict data sharing. Recognizing that data and information represent power, little data fiefdoms develop throughout the organization and data is used for personal or organizational enhancement rather than for the benefit of the institution.

- Departmental systems tend to be dependent on central data sources but relatively independent in their use of these data. The personnel and tools used for development of such systems are usually strong in the production of information (reports, screens, documents, etc.) and relatively weak in terms of data and information quality control. In other words, the fears of the professionals in the development center may be reasonably well founded concerning systems developed at the departmental level.

4. PERSONAL COMPUTERS

- It is a fact that personal computers are beginning to change the performance of office workers in all occupational categories. Unfortunately, there is little substantive evidence which indicates whether these changes are having a

positive or negative effect on office productivity in terms of cost benefits or the quality of information being produced. While it is beyond the scope of this study to analyze the true impact of personal computers on office productivity, it is necessary to correct what seem to be two popular misconceptions which are associated with the "PC revolution."

- Comparing performance of PCs and mainframes based on MIPS (with such statements as: "There are more MIPS installed on desktops than on mainframes") can be extremely misleading. This is true for several reasons:
 - Mainframe MIPS ratings do not consider the processing power embedded in the channels, controllers, etc.; therefore, the ability to move data between and among the various levels of the storage hierarchy is not considered when drawing such comparisons--throughputs of mainframes and PCs are not proportional to MIPS ratings.
 - The internal architectures of mainframes and PCs are so different (registers, interrupt handling, etc.) that performance on even computer-intensive applications cannot be equated based on MIPS comparisons.
 - Well over 99% of all potential PC MIPS are lost waiting for a human operator to hit a key, and this will always limit the effective use of all that power as long as the personal computer remains "personal" (interactive and user friendly) and involves a human component.

- While MIPS is clearly a meaningless measure of hardware performance across computer architectures, the current penchant for referring to PC software tools (such as spreadsheets and DBMSs) as "applications solutions" is incorrect. Buying the best spreadsheet or DBMS does not

solve any specific business problem or make any direct contribution to the organization's performance. Very specifically, even the availability of data from official data bases does not assure the quality of information generated using these so called "applications solutions." There is no such thing as a free ride in the development of applications systems.

- Once it is understood that a handful of PCs are unlikely to replace a 3090 mainframe and that you can pull together all of the PC software "applications" into one big integrated package and not have a "solution" to any of your problems, it is possible to isolate some of the very real potential advantages of microprocessor technology.
 - PCs get end users involved in the applications development process.
 - They can provide more responsive and cost-effective functions at appropriate levels in the applications system.
 - They can provide "user friendly" interfaces to otherwise complex systems through customization.
 - Users can develop their own personal systems when appropriate.
 - Fundamentally, PCs facilitate the fourth concept of GST which is the progressive mechanization of specific parts of the system. Mechanization is described as limiting the parts to a single function, in this case to a personal, specialized tool for an individual.
- Unfortunately, the history of personal computers has been one of users seizing the initiative from the central IS function, and these potential advantages have seldom been achieved. The cries for "connectivity" are coming from users who now want access to central data bases. There is little indication (or hope) that mass distribution of data bases to the desk top will:

- Result in properly designed distributed systems; users are not inclined to get involved with the central IS function because they see it as a threat to their newfound freedom.
 - Result in responsive, cost-effective systems; technical, professional, and managerial personnel will become inundated with data and applications programs which must be managed and maintained.
 - Provide user-friendly interfaces to those corporate data bases; PCs have established a new standard for "ease of use" which is not currently possible when interfacing with mainframe operating systems and data base management systems, and even 4GLs are not going to seem like productivity improvement tools to the great mass of end users.
 - Encourage the development of appropriate personal applications; end users are going to find that even the simplest of personal applications will require substantial maintenance and data base management, and they will rapidly find they are spending more time on systems activities than on their primary functions.
- Indiscriminate connectivity, data base access, and data base distribution would seem to be the antithesis of good information systems design, but that seems to be precisely the type of egalitarian network which has been prompted by the PC revolution.

5. COMPUTER/COMMUNICATIONS NETWORKS

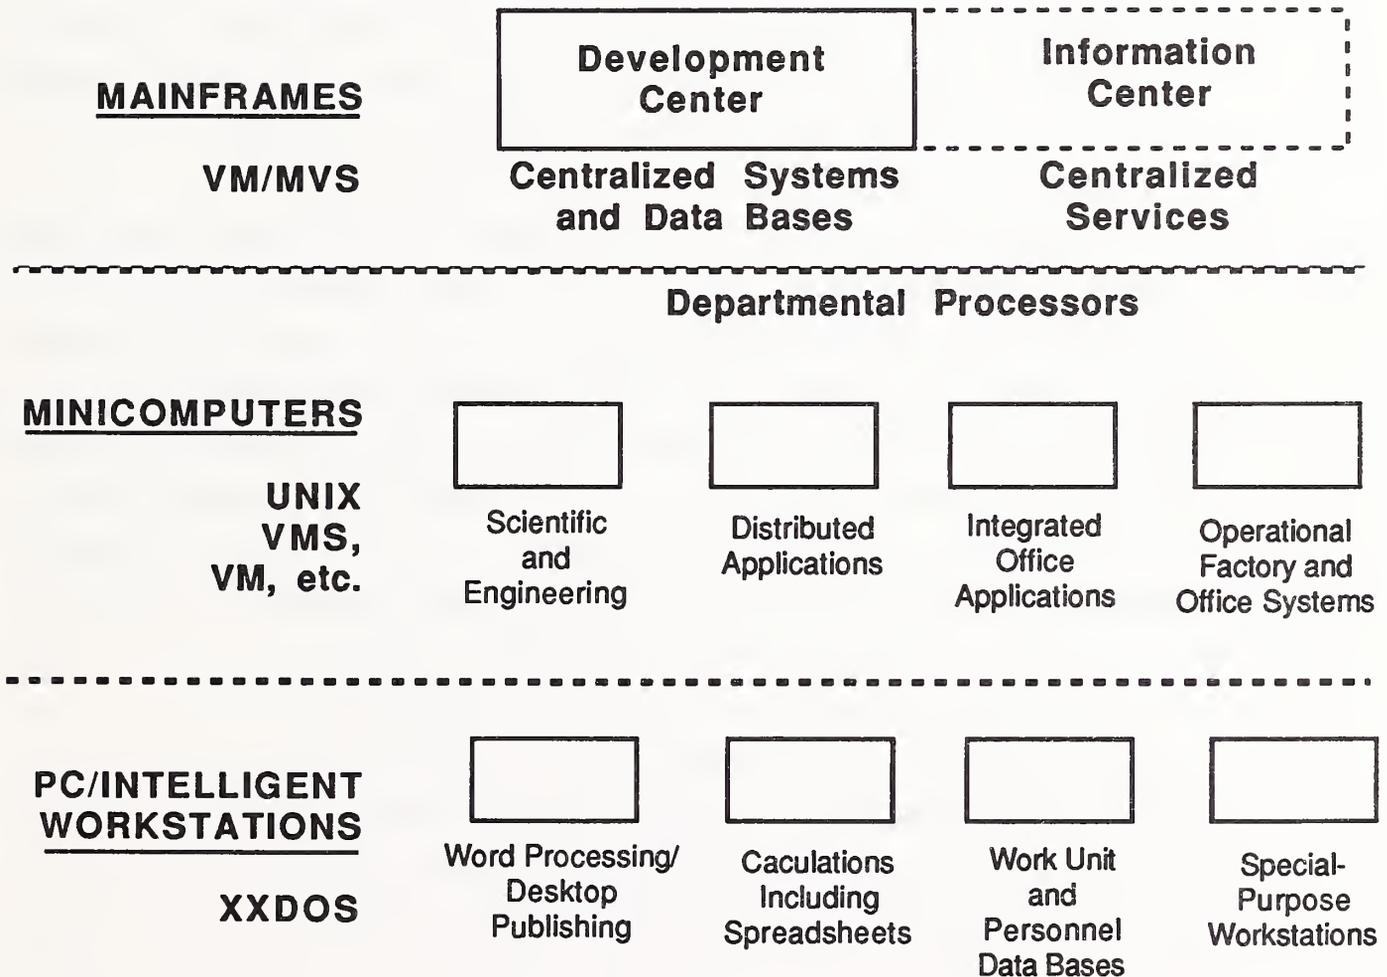
- Since we have been "going about this backwards" by installing a lot of computers at all levels of the processing hierarchy, we are now confronted with trying to connect them all together. Network design and implementation is normally left to communications specialists (if you are lucky enough to be able to find any). At the present time, anyone putting in a network is confronted with the distributed systems development activities which were

outlined above (see Exhibit II-1). There are conflicting objectives in this environment:

- The central development function wants to maintain tight control over major systems, central data bases, and what the end users are doing. Mainframe hardware/software has been designed for general purpose processing with little provision for differentiation and mechanization at lower levels in the processing hierarchy.
- The advocates of minicomputers want to differentiate by offloading data and applications from host mainframes to local processors and by integrating the personal computers and data bases within the work unit. In addition, minicomputers should be able to connect to each other on a peer-to-peer basis, and mainframes are considered to be just another node on the network.
- The users of PCs want to be able to communicate with all levels of the processing hierarchy and with each other. The PC is communications-oriented and so are its users. There is a very natural tendency to reach out for data since the other alternative is to key it in.
- There is a propensity on the part of vendors and users at the three levels to regard their particular systems as central and all others as peripheral.
 - . The central focus of the "host" mainframe is well established and the terminology itself certainly leaves little to the imagination.
 - . The departmental processors seek to occupy a central position in the network hierarchy (between mainframes and intelligent workstations) and, by serving as the switching systems, to control the network structure.

EXHIBIT II-1

THE CONNECTIVITY PROBLEM



- The intelligent workstations, as they are installed in higher organizational levels within the organization, view the network and its information resources as a service comparable to the central filing facility (albeit more difficult to use). And, who will argue with this when executives begin to use workstations?
- The conflicting views concerning the design of computer/communications networks represent a power struggle and always have. Information provides a competitive edge whether it is inside information about merger and acquisition activity or the control of the accounting information which reaches the company president.
- Most communications engineers are ill-equipped to deal with the complexity of the conflicting views of network design which are being presented by computer vendors. From its beginning, information theory has concerned itself with the transmission of bits and not with the content of the message. The fact that digital computers were one outgrowth of the theory has not changed that perspective very much. The world's largest and most reliable computers, called switches, are embedded in communications networks. (Unfortunately, these switches handle analog signals and ISDN will be a long time coming.)
- The hodgepodge of computer equipment, protocols, and network architectures which have been installed and proposed has been a real challenge to communications engineers, and they have been long suffering about the whole thing. They do not understand either the political (and technical) struggle among vendors or the applications systems requirements of the end users, and the biggest communications problem today is between computer systems people and those responsible for network development.
- When computer vendors talk about networks, they are talking about information flow which will position their particular products in the

most favorable position. SNA has been designed to squeeze mini-computers out of the processing hierarchy, and DECnet is designed to treat mainframes like peripherals.

- LANs have been used to confuse users, and a representative of IBM once publicly stated that the whole concept was a "tribute to advertising." This was followed with the statement that LANs solved the cabling problem but not the "connection" problem. So now IBM has given us "connectivity" which may solve the connection problem but does not address the distribution of data over the network hierarchy.
- So while it is fine to say that installing a bunch of computers is going about things backwards, it is not at all clear that starting with a network and "hanging computers on later" is a better solution. In fact, given the choice of a customer installing SNA now and hanging minicomputers on later or installing minicomputers now and connecting them to SNA later, it is not difficult to guess which DEC would recommend. The problem with network development is the same as that which has been experienced with data base systems.
 - If you build the necessary systems and applications software to support a data base system and do not give attention to the development of the data base you are not going to have a viable system.
 - If you concentrate on building a comprehensive data base to satisfy all the corporate requirements and then develop the ultimate system, the requirements will change before you ever get the system developed.
 - Most major corporations have experienced failure from both of these extreme approaches to the development of data base systems, and there is no reason to believe that the extremes of network development are any different. It is not a question of where you start (or even whether it is backwards or forwards), it is a question of knowing what you are doing.

- The exasperation of competent communications engineers with computer systems people is comparable to that of the central IS function with end users. It is expressed very simply as: "They don't know what they want, and when they think they do, they are sure to change their minds."
- An inimical attitude has developed over the years in the systems development process, and distributed systems development is the latest, and most disturbing, manifestation of that attitude. It is important to understand the implications of having systems developed by various groups with conflicting perspectives on how their systems interface and interact with other hardware/software systems.

B. DISTURBING IMPLICATIONS

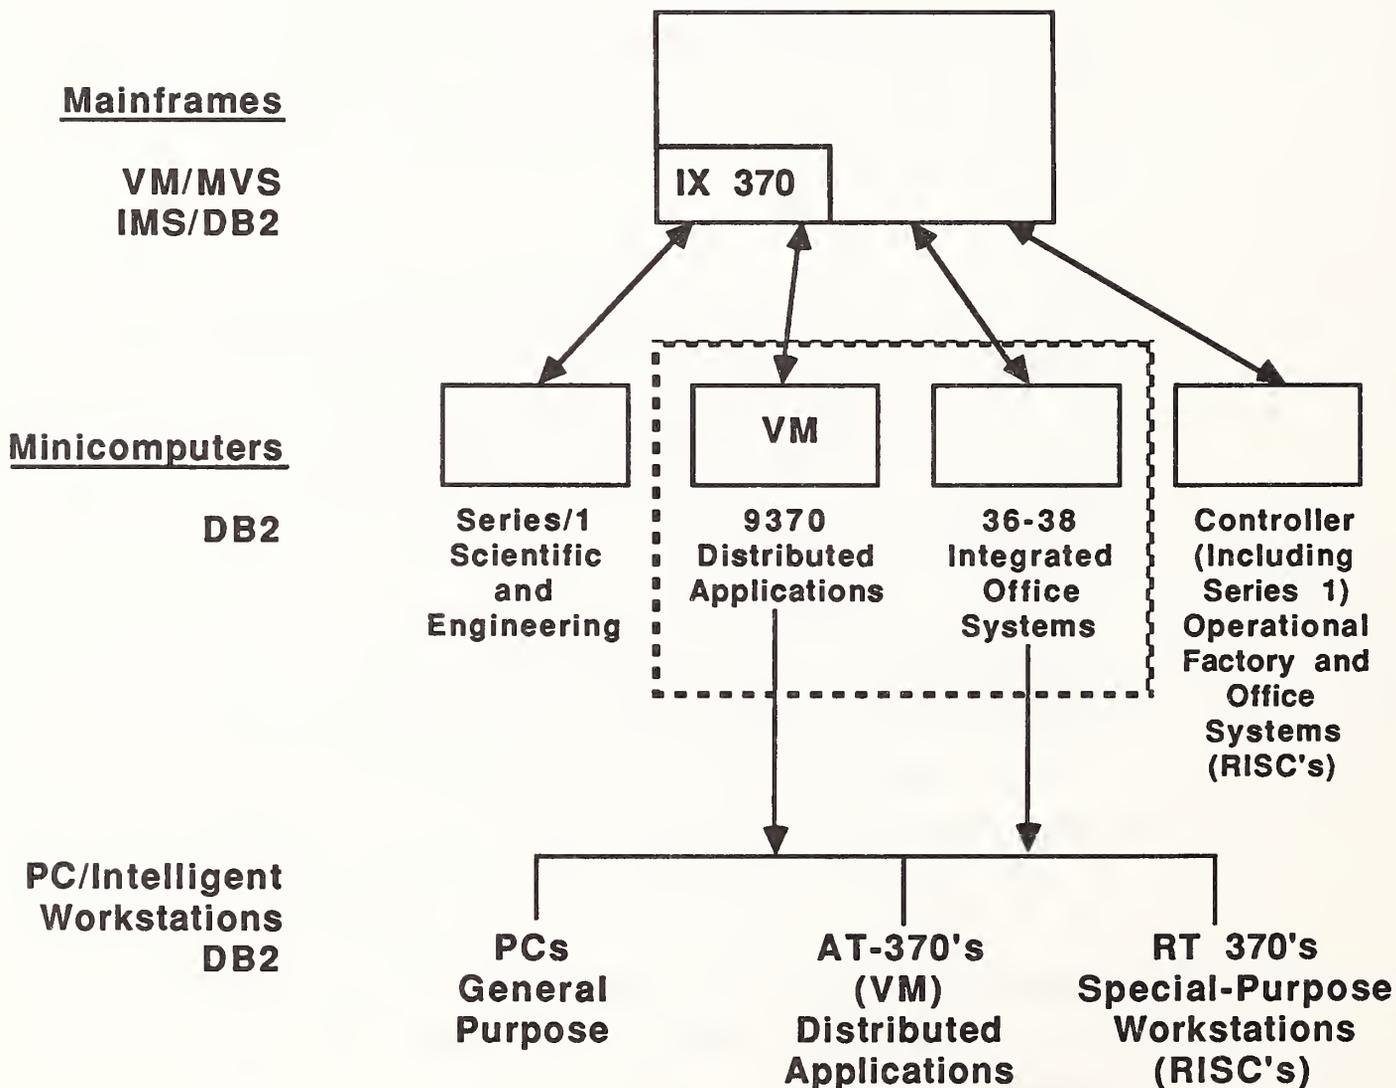
I. SYSTEMS VERSUS PERIPHERALS

- While INPUT has not endorsed IBM's networking strategy (as represented by SNA), the fact remains that IBM did define an evolutionary network architecture and then added computers later. The Large-Scale Systems Directions report series has described and analyzed that strategy in some detail. Two things are certain in that strategy--mainframe systems are not considered peripherals, and IBM and DEC do not agree on the architecture of computer/communications networks.
- While it is not our intention to review the checkered history of SNA, the recent announcement of the 9370 is significant enough to warrant at least an update of IBM's long and continuing battle against minicomputers. While the 3790 (is it chance that the same digits are used for the 9370, or does someone in IBM have a sense of history and humor?) and the 8100 were designed to limit the amount of processing which could be distributed from mainframes,

the 9370 indicates IBM may finally be getting serious about offloading mainframes. INPUT has long felt that this would be IBM's ultimate strategy, and indeed all of the pieces finally seem to have fallen into place (see Exhibit II-2).

- Despite the trade press's "discovery" of minicomputers in the form of departmental processors, IBM continues to emphasize that with rapid advances in microprocessor technology there will be no need for anything between the large mainframe and the desktop (see the last Large-Scale Systems Directions report). IS management in most large companies have adopted this point of view and will definitely wait to see what IBM will be doing with the 80386 before rushing out to install departmental processors.
- IBM's direct assault on the minicomputers being used for interactive scientific and engineering work (under UNIX) was described in the mid-year issue of Large-Scale Systems Directions, 1985. The use of the Series I as terminal controller (peripheral) and the absorption of the UNIX workload on a mainframe is clearly indicative of the highly centralized focus of SNA.
- When IBM thinks of distributed processing, it thinks of distributed data bases, and it has always been obvious to INPUT that the System 36 did not fit in the processing hierarchy because it did not have a data base system. It was equally obvious that the System 38 was not in the mainstream of IBM's networking plans because it was not compatible with the 370 architecture (and its supporting systems software). They were both standalone business systems which were pressed into service when integration of diverse office products and systems became necessary.
- The 9370 finally has the potential to address distributed data bases using IBM's preferred strategy which includes VM and DB2. The Large-

IBM'S NETWORK STRATEGY



Scale Systems Directions report series has been dedicated to the analysis of this evolving strategy. IBM's choice for a departmental processor is the 9370.

- Over most of the history of SNA, IBM has emphasized the use of terminal controllers for "distributed processing." Both the 3790 and 8100 fall into this category as do the more specialized boxes provided for point-of-sale and financial systems. It should be pointed out that controllers are reduced instruction set computers (RISCs) and ideally suited for systems differentiation and mechanization (in the GST sense of the terms). Therefore, IBM is now positioned to combat mini-computers in all of its major applications areas.
- What is important to IBM is a consistent and compatible hardware/software system from top-to-bottom in the network hierarchy, and that has theoretically been achieved with the 30XX, 9370, and AT 370 operating under VM and with DB2. And, whatever IBM does with the 80386, you can be sure it will enhance this distributed, data base-oriented structure. The ultimate objective will be to permit the user to have a single interface to a "system" with multiple levels of both processing power and data and with transparency as to where processing will take place.
- This has been the IBM technological dream for some time, and who can say when (if ever) it will become reality. At the rate SNA has progressed to date, it could take longer than most would be inclined to wait. However, there are three good reasons for IBM to pursue this course with all due diligence:
 - Some customers are beginning to understand what DEC has been talking about for all of these years and are beginning to demand something more than IBM has been willing to offer at Level II of the processing hierarchy (minicomputers). And, even though DEC's success in departmental systems has been vastly exaggerated by industry analysts (IBM profit margins remain higher than DEC's on a revenue

base which is approximately an order of magnitude larger), IBM management competes against its own past performance and remains sensitive to competition.

- IBM has a need to establish a complex hardware/software single-system view which can be used to exercise control over both its customer base and competition. While an open interface and connectivity may be the watch words of IBM's announced strategy, you can be sure that the mainstream IBM hardware/software systems are going to connect more promptly and more effectively than will competitors'. SNA (and systems software in general) is, first and foremost, a competitive weapon, and the fact that IBM has been relatively slow in deploying it (and seemingly reluctant to use it) is primarily a result of IBM not being threatened. The best thing that can happen for IBM competitors is for IBM to meet its business objectives--that way they will receive more benign attention. IBM's recent financial results have been bad news for everyone.

- IBM is not getting what it considers to be its appropriate share of the Level III (intelligent workstation) hardware/software market. The term "cooperative processing" is beginning to appear, and when real applications are truly split across the processing hierarchy, a consistent pattern of operating systems, DBMSs, and software development tools is required in order to produce applications which not only cooperate but are integrated into a coherent system. IBM must be the one to lead the way in the development of such integrated systems.

- SNA started with a mainframe-oriented strategy and has slowly evolved toward distributed processing. The architecture remains large system-oriented, and IBM would have difficulty changing this even if it wanted to--and it does not. The customer view is to be of one enormous, monolithic system. However, as INPUT has pointed out in this series of reports, when the functions which properly belong on a large mainframe are analyzed from the

perspective of the end user (or an objective observer), the host computer can be viewed as an enormous data base machine--in other words, a peripheral.

2. TOP-DOWN VERSUS BOTTOM-UP

- Distributed systems development, as defined earlier, leads naturally to a less structured development environment. Indeed, regardless of the justification for establishing (or permitting) such an environment, it can be fairly stated that the prevailing design philosophy is "bottom-up" rather than "top-down." One does not have to be a structured methodology purist in order to find this somewhat frightening.
- The integration of systems (or applications) developed in the information center or by end users on personal computers will frequently require more work than doing the entire application over again. Integration is the software equivalent of connectivity for hardware except it is substantially more difficult. Hardware boxes at least have the redeeming quality of being somewhat stable, whereas systems and applications developed in the DSD environment have the inherent attribute of being in a constant state of flux.
- The problems of top-down versus bottom-up systems design were discussed at some length in the last Large-Scale Systems Directions report. That computer/communications systems should be designed top-down is an intuitive decision reached by practically anyone who has ever designed a major system. When structured methodologies first appeared, the reaction of many old-timers to top-down design was ". . .is there any other way?". Building the ultimate corporate data base system just has not worked well for very many companies, and there is little indication that building the ultimate network will work any better. The problem seems to be that most "systems" design concentrates on computer hardware and software and not on the product of computer/communications networks--information.

3. DATA BASES VERSUS INFORMATION FLOW

- In INPUT's recent study on user applications of CD ROM, it was found necessary to present some practical definitions of data, information, and knowledge. Those definitions are as follows:
 - Data entry, data storage, data retrieval, data processing, data services, and all the rest refer simply to alphanumeric character strings. These strings are considered data from the point of view of the programmers, operators, and users of the computer.
 - Information and knowledge have a firm link, and the best way to define them is by distinguishing between them. The commonly accepted distinctions are as follows: (1) information is piecemeal, fragmented, and particular whereas knowledge is structured, coherent, and often universal; (2) information is timely, transitory, perhaps even ephemeral whereas knowledge is of enduring significance; (3) information is a flow of messages whereas knowledge is a stock, largely resulting from the flow in the sense that the "input" of information may affect the stock of knowledge by adding to it, restructuring it, or changing it in any way.
 - An additional fundamental distinction is that information is acquired by being told whereas knowledge can be acquired by thinking without new information being received.
- After 30 years of computer systems development, it is possible to draw some general conclusions about data, information, and knowledge:
 - Data, by definition, are stored in computer systems. However, it is possible to be more specific than that--data of institutional significance remain on host mainframes (for any but the smallest organizations), departmental processors are used primarily to concentrate data

for specific work units, and personal computers are used to generate paper documents (correspondence, reports, etc.).

- Information is transferred by voice (being told) or by paper documents. While the transfer of information by voice represents a substantially higher percentage of total office costs, the official communication of information remains on paper. If information of significance is generated in meetings or telephone conversations, it must normally be documented for purposes of validation, distribution, and storage. Paper remains the primary information media of organized human activity (business).
- Human minds remain the primary processors of information and the brain remains the primary storage device of knowledge (a very small percentage of individual human knowledge is ever documented). The best research efforts in artificial intelligence have resulted in precious little knowledge as to how either the mind or the brain works.
- Despite the change in terminology from DP to IS, it is a fair statement to say that the central development function is concerned primarily, if not exclusively, with data. Very few computer systems personnel have any interest in the vast flow of paper information within their organizations. If it can not be brought on-line it is ignored. Therefore, top-down design extends only to the printer. All of the problems of paper procedures--movement, filing, archiving, and disposal--are left to the "users," who as human beings are also the storage media for knowledge.
- The self-imposed isolation of most DP/IS organizations from the main information flow of the organization and from its knowledge bases (end users) has resulted in many systems which do not meet the needs of the organization. Paper-based systems and the essential organizational knowledge bases have not even been considered as "peripherals" to the computer/communications systems which have been developed. The objective of most "develop-

ment centers" has been to ignore information flow and knowledge. When one looks at what a so-called "knowledge engineer" is doing in building an expert system, it inspires one to ask: "Isn't that what systems analysts are supposed to be doing?"

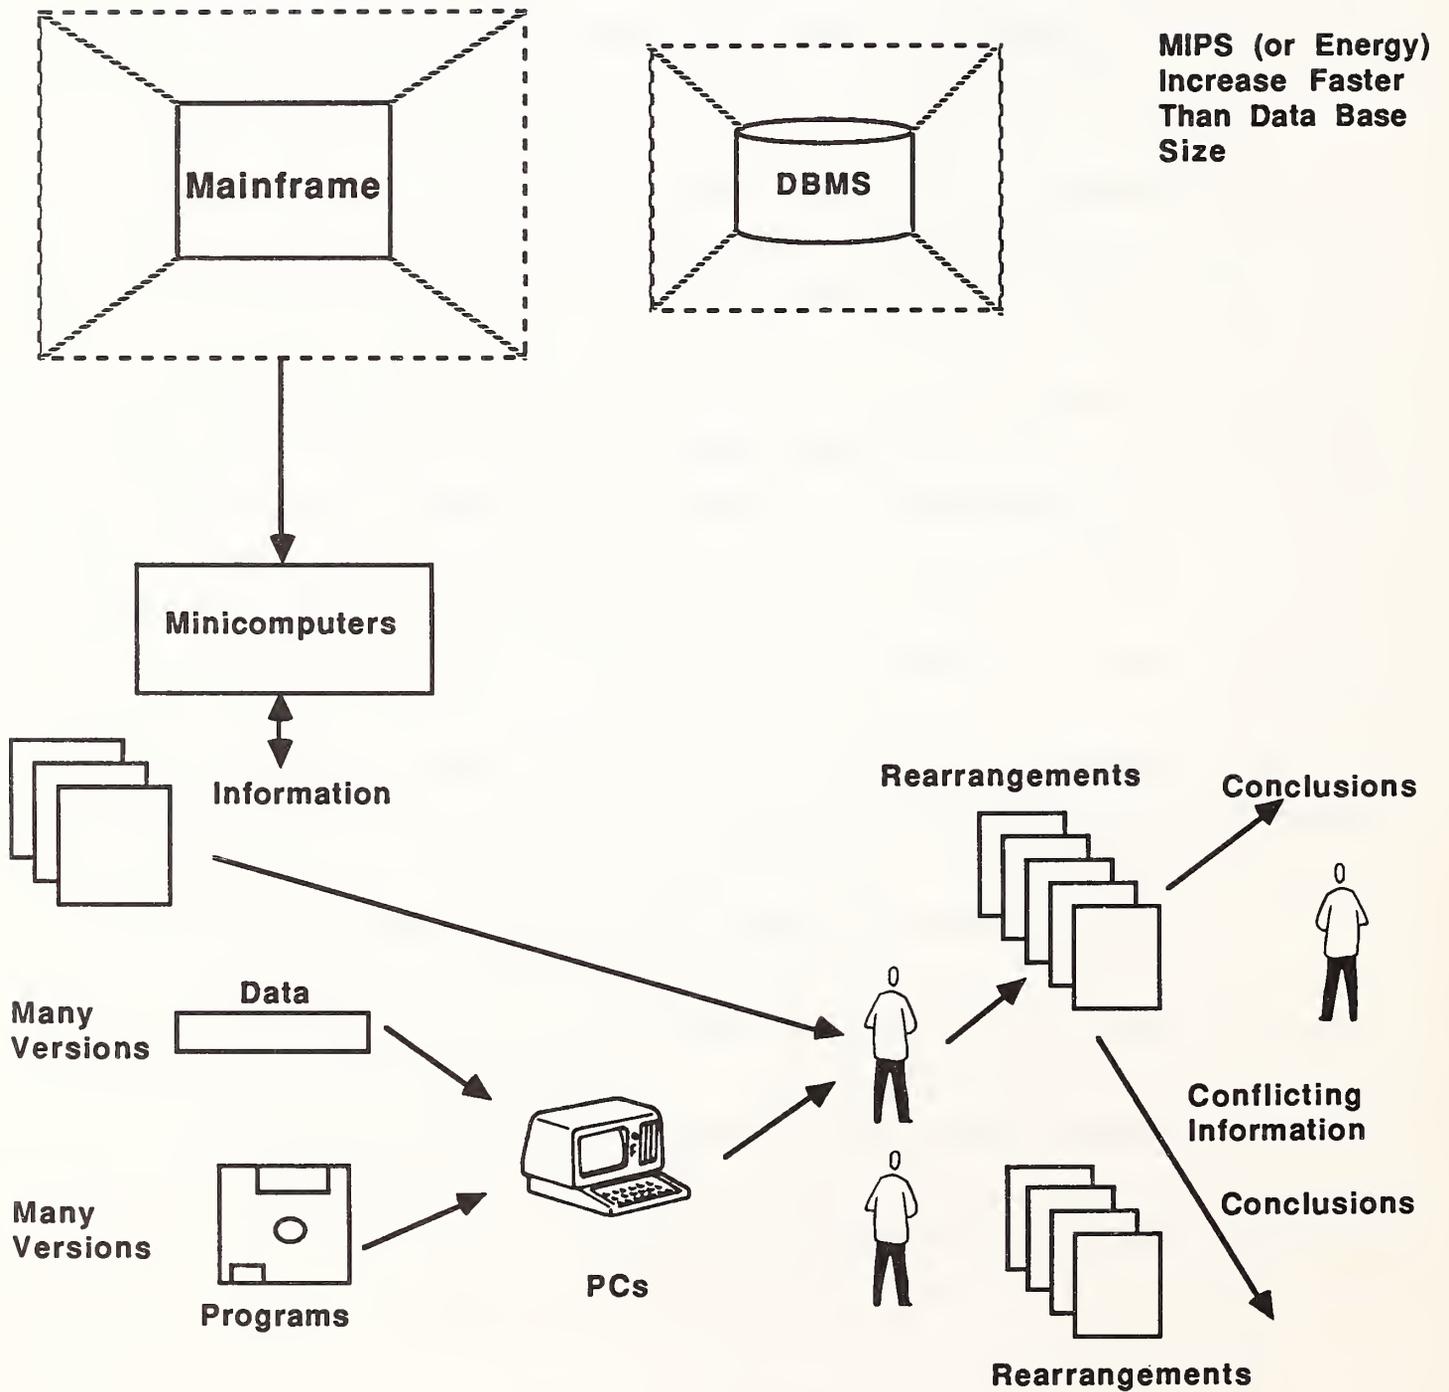
- On the other hand, the automation of offices has meant improving the production of paper documents with little regard for the quality of the contents. There is so much information flowing through offices that few office workers have time to devote to analysis; they are too busy generating pretty reports from their personal computers.
- It is important to recognize that knowledge can be generated without the receipt of additional information. In fact, it is probable that most significant knowledge comes only after turning off the flow of information; it is called thinking. But who has time to think on the information assembly line? Automated spreadsheets might permit the rapid calculation of numbers to support "what if" questions, but is there any indication that very many people are asking the right questions (much less getting the right answers from their spreadsheet package)?
- The central IS function has learned many valuable lessons about data processing and the importance of data quality, and end users have the necessary knowledge to improve information flow. However, neither top-down or bottom-up systems developed seem to hold very much promise of developing the systems which are needed to bridge the processing and data base hierarchy from corporate mainframes to the intelligent workstation.
- It is possible that the experienced minicomputer vendors such as DEC can provide the catalyst needed to bridge the gap between corporate data bases and the paper information flow in the office. Perhaps that is what is meant by building the network first and hanging the computers on later. Certainly, minicomputers have been highly successful on the shop floor in manufacturing environments. But, there is the nagging suspicion that most minicomputer

vendors are concerned only with the hardware network moving bits and bytes and have little appreciation of the complexities of managing large central or distributed data bases. In addition, they have generally given even less consideration to paper systems and procedures than have the mainframe vendors.

- It is becoming increasingly important to understand data, information, and knowledge and their contribution to productivity in the office environment. In the last Large-Scale Systems Directions report it was suggested that computer/communications network "software" be considered to include everything from SNA and systems software down through data/information/knowledge and the human beings connected to the network. In past reports it has also been recommended that white collar productivity be measured at three performance levels (see Exhibit II-3).
 - Performance Level I (large mainframe) is the conventional hardware/software network where software is viewed in the sense of computer programs. Its performance is measured by the relative cost of MIPS, cost per transaction, response time at the workstation, etc. While far from a science, Performance Level I at least has some metrics we can all argue about.
 - Performance Level II (human/machine dyad) is certainly on the periphery of the network, but it is dangerous to state that the human/machine dyad is of minor importance. Humans remain the primary "information processors" in terms of taking action whether it is merely pushing a button, making the decision to take over a multi-billion dollar company, or deciding to throw away a piece of "junk mail" in an electronic mail box. The measurement of performance at the human machine juncture (except for key strokes and routine clerical tasks) is extremely difficult to measure. Quality becomes more important than quantity at this level.

EXHIBIT II-3

ENTROPY PROBLEMS



- Performance Level III (work units) is the scene of a lot of current activity, but work units (departments, offices, etc.) vary so much that no generalized measures of performance seem to be appropriate. However, few quantitative measures exist other than meeting project and budgetary targets (both of which may have been inappropriate and even counterproductive in the first place). For example, a programming project team is a work unit, and we all know the problems with measuring productivity in the systems development process. Then, if a project is completed on time and within budget, it may be so costly at Performance Level I that it has a negative impact on performance at the institutional level.

- Performance Level IV (institution/enterprise) is theoretically where the benefits offset the cost of the computer/communications network. This has been so since the first computers were installed, and quite early in the evolution of computer systems the emphasis shifted from "saving money" to "making money." The exact contribution of computerization is extremely difficult to quantify in most cases. However, the mere availability of computers has increased reporting requirements so much that any net savings which might have appeared in the early days are extremely difficult to identify now. In fact, there seems to be a high correlation between the number of computers installed and an increasing number of higher-priced white collar workers.

- It is extremely important to identify more clearly what constitutes productivity in the office environment, and the contribution that computer/communications technology can make in these performance levels. (In fact, next year INPUT will begin to publish a companion report series--Distributed and Office Systems Directions--which will address these issues in considerable detail.) However, even now it is possible to state, almost categorically, that no computer/communications network can contribute to improved institutional performance unless the quality of data/information/knowledge on that network is maintained and enhanced.

4. QUALITY VERSUS CHAOS

- All INPUT research on the subject has indicated that the DSD environment has high potential for negative impact on the quality of data and information, and ever increasing cost for maintaining even the same level of quality. This problem manifests itself in IS management concerns about the distribution of processing power and data to minicomputers and intelligent workstations. These concerns have been repeated often--data base integrity and synchronization, privacy and security, mainframe performance, and conflicting reports to management. These problems seem to be recognized and accepted, and one is left with the feeling that while solutions are not readily available, at least the awareness will forestall any unpleasant surprises.

- However, INPUT believes that the real problem is data and information entropy, and no one seems to be very much concerned about that. Getting away from a technical discussion of information theory, information entropy can best be understood by stating that the more ways data or information can be arranged the higher entropy becomes. Entropy is the natural tendency toward energy loss (decreased information) which is inherent in all physical systems. As entropy increases more energy (processing power and human effort) is required to maintain order and preserve the information content. This is especially important in today's computer/communications networks for the following reasons:
 - As data bases increase in size and complexity, information entropy increases more rapidly; this in turn requires substantially increased energy (computer and human processing power) to maintain the same information content (quality). This is true whether we are talking about a 3090 running IMS and DB2 with ever increasing data base sizes or a PC running larger and larger spreadsheets. MIPS go up faster than bytes of on-line storage. (Earlier reports in this series explained this phenomenon.)

- Intelligent workstations with the power of yesterday's mainframes are data processors (or rearrangers). And the fact they can process and rearrange data in different ways will mean they will be used to do just that. The more ways the same data are rearranged, the higher the entropy of the system and the lower the overall quality of the information flow.
- To compound the problem, data are turned into information by software programs which are themselves information of varying quality (depending on both the original programmers of the tools as well as the users of the tools themselves). Then, of course, we have the connectivity problem itself. To permit various hardware and software systems to "talk" with each other, many rearrangements of information between the processors and systems software must take place, and there are additional problems associated the virtual environment which has been created.
- Exhibit II-3 presents a simplified version of the entropy problem. Mainframes may not be able to keep up with the increased demands which are being made on them by the increased size of on-line data bases, and rearrangements of essentially the same information (whether the conflicting conclusions are reached or not) inevitably lowers the quality of information content.
- While it is possible that the insertion of minicomputers between mainframes and the workstations may relieve the mainframe burden to the point where it becomes manageable, it is doubtful that there will be any decrease in information entropy as long as the quantity of paper documents continues to increase.
- Fortunately, the technology for controlling the proliferation of paper documents is becoming available in the form of optical memories. (See INPUT's report series on CD ROM.) Unfortunately, the quantity of

data which will then become available on the desktop will then threaten any control which is now being exercised from mainframes.

- An alternative, of course, is to reduce the number of human beings between data and decision makers. While current office automation products (primarily word processing) have created more paper documents to handle, the shift to optical media will encourage the reduction of paper output. The availability of virtually all information on-line will cause significant organizational changes which will have the following impacts:
 - Span of control will increase.
 - The vertical organizational hierarchy will be truncated. (Some levels of middle management will be eliminated.)
 - Organizational structures will become more flexible and even ephemeral. (Project teams and committees will be the rule and they will cut across traditional organizational structures.)
- This means that the design of a network to accommodate today's information flow will not be satisfactory tomorrow. Specifically, the work unit networks (LANs) and computers (departmental processors) have a high probability of being inappropriate for tomorrow's information flow.
- In the brave new world of electronic media it will also become apparent that information will have to be screened and filtered to extract what is new and timely from what is old and tired and to identify the information which is of lasting value (knowledge) from the "instant analysis" which is being prompted by the highly reactive environment which is being created. Call them expert systems or merely say it is good systems design, but we all are going to need some means of separating the "junk mail" from the information flow. And,

the more we can refine such systems to analyze both the content and the source of the message, the more effective the network will be.

- It is not at all clear that today's general purpose computers (whether mainframes or minicomputers) are appropriate for managing data and information flow or for serving as highly responsive "store and forward reservoirs" on the network. It is probable that differentiation and mechanization will progress very rapidly with the exponential increase in on-line information which will occur when optical memories begin to replace paper. This indicates that there are RISCs in your future whether you like it or not.
- In other words, it is all well and good to talk about building a network and hanging the computers on later, but it is highly probable that the requirements will change before either the computer/communications network is ever installed or the applications software is in place. Long-range planning is absolutely essential if order is to be maintained in today's high-entropy environment.

C. RISCs AND RESIDUAL VALUES

I. DINOSAURS, DRAGONS, AND THE GHOST OF VON NEWMANN

- We have often quoted Admiral Grace Hopper's statement of over 10 years ago that large-scale mainframes and their associated operating systems were becoming dinosaurs. However, it is difficult to recognize a dinosaur unless you have something to compare it with, and that is where personal computers have served a very useful purpose. Management is beginning to question whether it is necessary to pay for the care and feeding of a dinosaur when a mouse can do many things better (and a mouse is certainly more user friendly).

- Early in this report series the "von Neumann bottleneck" was described in some detail. The architecture of general purpose computer systems was determined; it is now possible to differentiate and mechanize many major computer functions and to explore new architectures which are better suited for processing arrays of data, handling relational tables, extracting information from paper documents, and a whole host of other functions. Large general purpose mainframes are becoming a bottleneck in maintaining order in today's dynamic environment.
- The "dragon" is the mainframe systems software which breathes fire at any one approaching the castle where the dinosaur is hidden. However, it is absolutely essential because it is responsible for scheduling the use of the dinosaur and telling it what to do. Whereas the dinosaur is very real, the dragon is imaginary; it is the virtual environment which has grown so forbidding that it is mistaken for reality. The last INPUT report described in detail how the dragon will disappear once the operating systems functions are properly distributed over the network.

2. EPHEMERAL "SOLUTIONS" AND THE PERENNIAL BACKLOG

- It has been virtually impossible to train (program) the dragon and dinosaur to do all the work we would like to have them do. Machine language gave way to assemblers, which gave way to FORTRAN and COBOL, which are in the process of giving way to 4GLs, which are currently giving way to applications generators, etc. At each step along the way the problem was supposedly solved; executives were supposed to be able to read COBOL listings and find out what the dinosaur was doing so they could instruct the beast's keepers in what they really wanted it to do.
- However, the dinosaur seems to have a mind of its own, and it keeps plodding along doing what it always has and growing bigger all the time. As it grows bigger more and more is expected of it, but it is very slow in responding to directions even when it is given instructions. Hidden deep in its subconscious

are patterns of past instructions dealing with 80 column cards and sequential batch processing. It is difficult to be responsive to the new instructions when the past patterns are so ingrained.

- The dragon's original role was clearly defined as being to make the dinosaur "easier to use." (This being the primary design point of OS/360.) What has developed is an additional barrier to getting the dinosaur (either to get useful work done or to kill it), and the dragon itself requires a special set of trainers known as systems programmers.
- Therefore, none of the solutions for getting the dinosaur to be more responsive have worked very well, and the want list continues to grow. Applications promised years ago remain in various states of implementation and promised results have only infrequently been delivered. However, it has only been with great reluctance that the keepers of the dinosaur have finally said: "Let them use mice."

3. HARDWARE CHICKENS AND SOFTWARE EGGS

- The productivity of the dinosaurs has been so low that they are occasionally replaced with an entirely new model (generation). When this occurs, there is a considerable period of adjustment before a new dragon can be installed to use all of the new capabilities of the new dinosaur. Once the new systems software is in place, the applications software must be written. There is an extremely long gestation period before the dinosaur produces an egg which gives promise of something really new for its keepers.
- Unfortunately, many of the eggs are not fertile and do not produce the long-awaited results. It is then decided that a new dinosaur is the answer, a new one is acquired, and the cycle starts all over again. The thing the dinosaur does best are those old patterns from the past--it really has that batch processing down pat. Fortunately, batch processing remains essential for the maintenance of those enormous data bases and distribution of extracts of

those data bases to those outside the castle walls. The dinosaur is assured of a home in the foreseeable future unless it dies from natural causes from overexertion before specialized data base machines can be hatched out of one of those eggs.

D. THE RESIDUAL COST OF THE PRESENT

- It is obvious that we are, at least temporarily, stuck with the care and feeding of the dinosaur regardless of what is happening outside the castle walls. However, the dragon is another story. It just is not capable of dealing with a bunch of mice making individual requests for data. (A thousand requests for a 100,000 bytes of data are substantially more work to handle than 10 requests for 10 megabytes.) Some of its functions can be distributed to other levels in the computer/communications network (remember the dragon is unreal so it can be recreated in different forms anywhere you like). The cost can be minimized if the operating systems can be simplified, and much of this cost savings will come from having fewer systems programmers.
- There is the residual cost of maintaining the dinosaur as long as it serves some useful purpose, and that useful purpose may go on for some time. The cost of maintaining the dinosaur will be in COBOL programmers to maintain existing applications systems and data base personnel to insure data quality of the central data bases. However, the development center should not train the dinosaur for new tasks. The development center should concern itself with relieving the poor monster of as much work as possible. The planned distribution of applications systems to minicomputers and/or microprocessors should be high on the list of priorities in any development effort. Otherwise, those big new applications systems for the dinosaur may become the residual costs of the future before they ever become operational.

- The true residual costs of the present are in the networks which have been installed without sufficient planning and in all those departmental processors which were installed (or are being installed) without regard for their proper applications role or how they will "hang on the network." Some top-down direction is needed from the development center before processing and data bases are distributed.
- The first generation of personal computers may have to be replaced, but the residual costs will not be too great since they can be recycled for home use. As far as the applications which have been developed for PCs are concerned, most of them are not worth saving or maintaining and they are the responsibility of the individuals who used them initially. The primary residual expense will be the decisionmaking process required to throw away many of the abortive attempts to develop applications at this level.
- However, unless the development center and the central IS function provide leadership and quality assurance, the DSD environment is going to create dinosaurs all over the place. So we agree that networks should be carefully planned first and the computers can be hung on later, but unless careful attention is given to data and information flow over the network (how applications and data are to be effectively distributed), the computer/communications networks which are being installed could have adverse impact on all four performance levels and the residual costs could be staggering.

III RESIDUAL VALUE FORECASTS

A. ANNOUNCEMENTS

- The most significant announcement since the last Large-Scale Systems Directions report was the IBM 9370. It is beyond the scope of this series of reports to analyse the announcement in detail, but INPUT believes it is significant to large-scale systems. This is true because we continue to believe that IBM's main thrust in distributed processing will be with 370-compatible systems; and, if IBM is getting serious about distributed processing, and we believe it is, it is time to start thinking about the distribution of applications from mainframes (offloading). This has obvious implications for long-range, large-scale systems planning and residual values.
- Only the briefest summary of the IBM 9370's characteristic will be presented in this report. However, as explained in the introduction, INPUT will start a new report series in 1987 (Distributed and Office Systems Directions) which will address both minicomputers and intelligent workstations.
- The essentials of the IBM 9370 announcement are as follows:
 - The 370 instruction set is extended down to the DEC VAX and even Micro VAX processor sizes, and the 9370 does have a bus architecture.

- There are four models (20, 40, 60, and 90) ranging in price from \$31,000 to \$190,000, which places them precisely in the minicomputer category as defined by INPUT over 10 years ago.
- Memory of 4, 8, and 16 Mb is available on the Model 20, and 8 and 16 Mb on the other models. IBM 1 Mb chips and a new air-cooled TCM are used in the packaging.
- DASD ranges from 368 Mb at the low end up to 39.6 Gb at the top of the line. A streaming tape drive is available for backup.
- The systems were designed for the office and do not require air conditioning. The Model 20 will operate on 110 v power, and the other models will run on "normal" 220 v. The systems are designed to conserve floor space with the largest system (Model 90) being about the size of two file cabinets.
- The systems are obviously aimed at the minicomputer market in general and DEC in particular. And, regardless of what IBM says, the 9370 will be the main competitor against the Systems 36 and 38. The 43XX line with the exception of the 4381 has obviously been replaced. While IBM does not like to use MIPS (a policy with which INPUT agrees quite heartily), some idea of relative performance can be obtained from the following statements which were made:
 - . There is a 100 times performance range of 370-compatible processing power between the 9370 Model 20 and the 3090 Model 400.
 - . There is a 5 times performance range within the 9370 line itself with the Model 20 being rated at 5 transactions per second and the Model 90 at a 25 transaction processing rate (TPR). (By contrast, the 4381 has a top performance of 80 TPR.)

- On-line terminals supported range from 64 on the Model 20 to 380 on the Model 90 (but it is doubtful that these are really practical at the present time).
 - There are four operating systems for the 9370 (when IBM goes for a multi-operating systems environment it goes all out). These are as follows:
 - VM/SP.
 - VIE/SP.
 - MVS/SP (the 9370 Model 90 and 4381 are both bridge systems to the 30XX environment).
 - IX/370 which is IBM's version of UNIX (IBM finally got around to running it on a minicomputer rather than on the mainframe).
 - IBM has not forgotten connectivity either; the communications options include having the 9370 serve as an:
 - ASCII Controller.
 - Token Ring Subsystem Controller.
 - Ethernet LAN Subsystem Controller.
 - Telecommunications Subsystem Controller (SDLC, BiSync, or X.25).
- For the first time since IBM announced SNA, it appears it is serious about distributed processing. The architecture of computer applications systems is going to change.

B. REVIEW OF USED MARKET ACTIVITY

- The 9370 obviously has impact on the 4300 series used market. At the present time, supplies of most 4300 series processors are plentiful and demand is moderate. However, IBM's announcement of a trade-in program (i.e., a 20% discount off the list price of a 4381 with a trade-in of a 4341 or 4361) essentially produces a list price which is equal to the used market price. This may have the effect of increasing used market demand for purposes of trade in.
- The 3090 has had its inevitable impact on the 308X market. At the present time, the supply of 3083s for sale is low, as is the demand. The 3083E used market price is essentially that of the memory and channels. Larger 3083s are holding their value reasonably well because at present levels they are price/performance-competitive with the 3090. Lead time for purchase can be from 60 to 90 days.
- The 3081 models are in moderate supply which is reasonably well balanced with demand. There is a sufficient trailing edge customer base to keep prices from dropping sharply at the present time. However, with general delivery of the 3090-400s, anticipated MVS-XA, and increased channel speed announcements, significant price reductions for 3081s are anticipated during 1987. At the present time the supply of 3084s is quite low and there is still moderate demand for the systems. However, they will suffer the same price erosion as the 3081s during 1987.
- There has been continuing price erosion in the used 3380 disk market, and lots of "bargains" are available for those who are not suffering severe space constraints. Remaining in the "trough of technology" has always been an attractive alternative for large DASD users. The theory goes that if you can hold off on purchasing the double density drives, IBM may announce new technology in 1987 (maybe a quadruple density drive), then the price of the 3380 E Model drives will begin to drop in the used market, and the cycle goes on.

C. ASSUMPTIONS

- INPUT's residual value forecasting methodology is proprietary and has been continually refined over the years. The assumptions underlying our forecasts fall into three categories: general, specific, and proprietary.
- The general assumptions underlying INPUT's residual value forecasts are as follows:
 - IBM is always operating against a plan which will maintain its traditional growth in revenues and its traditional profit margins.
 - IBM is essentially large-scale systems-oriented and will resist significant offloading of mainframes to minicomputers and/or microprocessors.
 - The means of control of the distribution of processing and data is through systems software (SNA, operating systems, and DBMSs), and software development will always lag hardware capability.
 - IBM will continue to be successful in controlling the distribution of processing because there will be no serious breakthroughs in competitive systems software development which IBM cannot effectively counter.
 - Large-scale hardware/software will continue to evolve pretty much on IBM's schedule, and there will not be any drastic changes in product cycles (a majority of customers are not going to decide to skip a generation).
 - Mainframes and associated peripherals will remain at the heart of IBM's strategy through the 1990s, and there will be a continuing used market for such equipment during that period.

- IBM has the administrative systems in place to facilitate product announcements and pricing changes which virtually give it control of residual values. (IBM's increased flexibility in both product introduction and pricing have become apparent over the years, and the importance of these improved internal systems should not be underestimated.)
- There are certain specific assumptions which are directly related to current residual value forecasting. These assumptions are as follows:
 - IBM will not deviate radically from historic patterns of price/performance improvement for large-scale processors and magnetic disk storage systems. (INPUT identified these patterns over 10 years ago, and they have proven to be remarkably accurate.)
 - Therefore, it is assumed that price/performance will improve at a rate of between 10% and 16% per year (depending on the particular product), and these rates are used to compensate for list price reductions over the product life cycle. (The specific methodology employed is proprietary.)
 - IBM will be able to delay the impact of optical memories on large-scale magnetic disks beyond the range of this year's forecasts (1991).
 - Modest impact of optical memories on large-scale tape systems will begin to be felt during this period, and this impact is built into the forecasts.
 - IBM is assumed to have been reasonably forthright in its large-scale systems presentation, as presented in our last report, and there do not appear to be any competitive technological developments which will force premature (from IBM's point of view) deviation from the highly

centralized architectural focus which has been described in this series of reports.

- As predicted earlier, IBM will start to release new versions of MVS/XA which will support the 3090 architecture (and any enhancements such as new channels). It is assumed that this announcement will be a clear signal that MVS/XA for the 308X has reached the end of the line.
- Alternate operating systems, such as UNIX and Amdahl's Aspen, will neither have serious impact on the 3090 sales nor serve to extend the life of 308X processors.
- IBM's dual data base strategy will prove successful--DB2 will become highly popular (and a de facto standard) and IMS will live on well past the forecast period. As data bases are "distributed," demand for archival storage on mainframes will continue to represent a strong growth area for magnetic disks. (In other words, data bases from minicomputers and intelligent workstations will be archived on the host mainframes which will in turn be archived on other storage media such as tape.)
- Privacy and security is going to become an increasingly important subject during the next five years, and the IBM solution is going to obsolete a lot of current hardware and software at all levels in the processing hierarchy. Secure, certified data bases are the key to large-scale systems growth in the 1990s.
- Large-scale printer technology is virtually frozen, and new technological developments will be concentrated on distributed printer systems. However, centralized printing facilities will remain viable and necessary during the forecast period.

- The announcement of the IBM 9370 may cause modification of one of INPUT's general assumptions--specifically, that IBM will resist the offloading of mainframes. While this report clearly states that IBM now has its preferred minicomputer offering in place, it is not felt that this will cause any fundamental change in strategy--IBM remains large-systems-oriented. However, a complete analysis of the potential impact of the 9370 on mainframes will be presented in the first Distributed and Office Systems Directions report which will be issued early in the second quarter of 1987. The research of that report series will determine the degree of impact which may be felt on large-scale systems, and assumptions will be adjusted accordingly.

D. PROJECTED USED MARKET PRICES AND RESIDUAL VALUES

- It is important to understand that at any given time, three price levels exist in the used market.
 - Retail price is the amount that an end user would pay for the equipment.
 - Dealer price is the amount that a dealer would pay another dealer to acquire equipment to complete a contracted sales obligation.
 - Wholesale price is the amount a dealer would pay to acquire equipment for resale.
- The dollar spread between levels is a function of the total value of the transaction. For large processors the wholesale price will typically be 80% to 95%, and for peripheral equipment, 70% to 90% of the retail price.

- Exhibit III-1 presents the projected residual value as a percent of vendor list price as of January 1 for selected IBM and software-compatible mainframes.
- Exhibit III-2 presents the projected used market retail prices based on those projected residual values.
- Exhibits III-3 through III-7 present the ranges of projected residual values for those mainframes.

EXHIBIT III-1

PROJECTED RESIDUAL VALUE
AS A PERCENT OF VENDOR LIST PRICE
AS OF JANUARY 1, 1987

VENDOR	PROCESSOR MODEL	CURRENT LIST PRICE	PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE AS OF JANUARY 1, 1987						
			1987	1988	1989	1990	1991	1992	
IBM	4331-J2	\$79,500	5%	3%	2%	1%	0%	0%	
	4331-L2	88,500	15%	8%	4%	2%	0%	0%	
	4341-L1	202,900	2%	1%	0%	0%	0%	0%	
	4341-M2	315,400	4%	2%	1%	1%	0%	0%	
	4361-K5	126,900	41%	25%	14%	7%	4%	2%	
	4361-M5	163,101	46%	31%	13%	8%	5%	3%	
	4381-M1	373,131	60%	47%	20%	11%	7%	4%	
	4381-P13	538,311	84%	67%	46%	33%	23%	15%	
	4381-M12	332,731	84%	66%	45%	31%	20%	13%	
	4381-P14	740,462	88%	70%	51%	39%	27%	18%	
	3083-CX	802,731	11%	6%	2%	1%	0%	0%	
	3083-E	1,237,731	7%	3%	1%	0%	0%	0%	
	3083-EX	932,731	11%	6%	2%	1%	1%	0%	
	3083-B	2,032,731	13%	5%	2%	1%	0%	0%	
	3083-BX	1,257,731	23%	9%	4%	2%	1%	0%	
	3083-J	2,607,731	16%	7%	3%	1%	0%	0%	
	3083-JX	1,667,731	29%	13%	6%	2%	1%	0%	
	3081-G	3,162,731	15%	7%	3%	1%	0%	0%	
	3081-GX	2,132,731	26%	11%	5%	2%	1%	0%	
	3081-K	3,902,731	17%	9%	4%	2%	0%	0%	
	3081-KX	2,642,731	32%	14%	7%	4%	1%	0%	
	3084-Q	6,750,462	27%	13%	7%	4%	1%	0%	
	3084-QX	4,800,462	42%	21%	11%	6%	2%	1%	
	3090-150	1,708,900	92%	74%	51%	20%	5%	3%	
	3090-180	2,608,900	93%	76%	53%	22%	6%	4%	
	3090-200	4,508,900	90%	80%	57%	27%	8%	5%	
	3090-400	8,515,785	92%	85%	65%	31%	12%	8%	
	AMDAHL	5840	1,270,000	41%	17%	6%	2%	1%	0%
		5850	1,534,000	39%	16%	6%	3%	1%	0%
		5860	1,928,000	39%	17%	6%	4%	2%	0%
		5868	2,850,000	40%	23%	9%	4%	3%	0%
		5870	3,236,000	39%	22%	9%	4%	3%	0%
		5880	3,516,000	40%	23%	9%	4%	3%	0%
5890-200		3,825,000	92%	82%	55%	18%	5%	2%	
5890-300		4,500,000	91%	78%	52%	22%	7%	3%	
5890-600		8,500,000	93%	85%	64%	29%	9%	4%	
NAS		AS/6630	341,500	18%	8%	5%	2%	1%	0%
	AS/6660	475,000	21%	12%	7%	4%	2%	0%	
	AS/8023	475,000	34%	16%	8%	3%	1%	0%	
	AS/8083	2,271,900	42%	23%	14%	6%	3%	0%	
	AS/9050	2,202,000	10%	6%	3%	2%	1%	0%	
	AS/9070	3,477,000	13%	7%	3%	2%	1%	0%	
	AS/XL50	3,050,000	90%	79%	54%	16%	4%	2%	
	AS/XL60	4,228,000	90%	80%	56%	18%	5%	3%	
	AS/XL90	9,789,000	0%	0%	60%	27%	8%	5%	

EXHIBIT III-2

PROJECTED USED MARKET RETAIL
AS OF JANUARY 1, 1987

VENDOR	PROCESSOR MODEL	CURRENT LIST PRICE	PROJECTED USED MARKET RETAIL as of January 1						
			1987	1988	1989	1990	1991	1992	
IBM	4331-J2	79,500	\$4,000	\$2,500	\$1,500	\$1,000	\$0	\$0	
	4331-L2	88,500	13,500	7,000	3,500	1,500	0	0	
	4341-L1	202,900	4,000	2,000	1,000	500	0	0	
	4341-M2	344,731	15,000	8,000	5,000	2,000	1,000	0	
	4361-K5	121,131	50,000	30,000	17,000	9,000	5,000	3,000	
	4361-M5	163,101	75,000	50,000	22,000	13,000	8,000	5,000	
	4381-M1	373,131	225,000	175,000	75,000	40,000	25,000	15,000	
	4381-P13	538,311	450,000	360,000	250,000	180,000	125,000	80,000	
	4381-M12	332,731	280,000	220,000	150,000	103,000	67,000	43,000	
	4381-P14	740,462	650,000	515,000	375,000	290,000	200,000	130,000	
	3083-CX	802,731	90,000	52,000	20,000	8,000	2,000	0	
	3083-E	1,237,731	88,000	41,000	15,000	6,000	1,000	0	
	3083-EX	932,731	102,000	60,000	18,000	10,000	5,000	0	
	3083-B	2,032,731	265,000	100,000	45,000	15,000	7,000	0	
	3083-BX	1,257,731	285,000	115,000	52,000	20,000	9,000	0	
	3083-J	2,607,731	425,000	190,000	85,000	30,000	10,000	0	
	3083-JX	1,667,731	485,000	215,000	97,000	35,000	12,000	0	
	3081-G	3,162,731	475,000	220,000	95,000	30,000	12,000	5,000	
	3081-GX	2,132,731	555,000	235,000	105,000	35,000	20,000	7,000	
	3081-K	3,902,731	665,000	350,000	150,000	80,000	16,000	8,000	
	3081-KX	2,642,731	850,000	370,000	175,000	100,000	28,000	10,000	
	3084-Q	6,750,462	1,850,000	900,000	500,000	250,000	100,000	25,000	
	3084-QX	4,800,462	2,000,000	1,000,000	550,000	275,000	115,000	30,000	
	3090-150	1,708,900	1,575,000	1,265,000	870,000	340,000	90,000	50,000	
	3090-180	2,608,900	2,425,000	1,980,000	1,385,000	575,000	160,000	98,000	
	3090-200	4,508,900	4,050,000	3,610,000	2,575,000	1,225,000	375,000	220,000	
	3090-400	8,515,785	7,850,000	7,250,000	5,525,000	2,650,000	1,025,000	675,000	
	AMDAHL	5840	1,270,000	520,000	210,000	75,000	25,000	10,000	3,000
		5850	1,534,000	600,000	250,000	90,000	50,000	20,000	5,000
		5860	1,928,000	750,000	325,000	125,000	75,000	30,000	6,000
5868		2,850,000	1,150,000	650,000	250,000	110,000	75,000	8,000	
5870		3,236,000	1,275,000	725,000	290,000	125,000	90,000	13,000	
5880		3,516,000	1,400,000	800,000	325,000	150,000	100,000	17,000	
5890-200		3,825,000	3,500,000	3,150,000	2,100,000	700,000	175,000	60,000	
5890-300		4,500,000	4,100,000	3,500,000	2,350,000	1,000,000	300,000	125,000	
5890-600		8,500,000	7,900,000	7,200,000	5,450,000	2,500,000	750,000	300,000	
NAS		AS/6630	341,500	61,000	27,000	17,000	7,000	3,000	0
	AS/6660	475,000	100,000	55,000	33,000	19,000	9,000	1,000	
	AS/8023	475,000	162,000	78,000	39,000	13,000	6,000	0	
	AS/8083	2,271,900	950,000	520,000	315,000	145,000	60,000	7,000	
	AS/9050	2,202,000	215,000	125,000	72,000	35,000	18,000	2,000	
	AS/9070	3,477,000	450,000	245,000	120,000	62,000	30,000	3,000	
	AS/XL50	3,050,000	2,750,000	2,400,000	1,650,000	475,000	125,000	55,000	
	AS/XL60	4,228,000	3,800,000	3,375,000	2,350,000	750,000	195,000	110,000	
	AS/XL90	9,789,000	--	--	5,850,000	2,650,000	785,000	490,000	

EXHIBIT III-3

IBM 308X SERIES
PROJECTED RESIDUAL VALUE AS A
PERCENT OF VENDOR LIST PRICE

PROCESSOR MODEL		PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE as of January 1					
		1987	1988	1989	1990	1991	1992
3083-CX	High	18%	12%	8%	5%	3%	2%
	Expected	11%	6%	2%	1%	0%	0%
	Low	7%	3%	1%	0%	0%	0%
3083-E	High	11%	7%	3%	2%	1%	1%
	Expected	7%	3%	1%	0%	0%	0%
	Low	4%	1%	0%	0%	0%	0%
3083-EX	High	15%	10%	5%	3%	2%	2%
	Expected	11%	6%	2%	1%	1%	0%
	Low	6%	3%	1%	0%	0%	0%
3083-B	High	18%	11%	7%	4%	2%	1%
	Expected	13%	5%	2%	1%	0%	0%
	Low	9%	3%	1%	0%	0%	0%
3083-BX	High	27%	17%	11%	6%	3%	2%
	Expected	23%	9%	4%	2%	1%	0%
	Low	12%	5%	2%	0%	0%	0%
3083-J	High	23%	15%	10%	4%	2%	1%
	Expected	16%	7%	3%	1%	0%	0%
	Low	9%	4%	1%	0%	0%	0%
3083-JX	High	32%	20%	13%	7%	5%	2%
	Expected	29%	13%	6%	2%	1%	0%
	Low	14%	5%	3%	1%	0%	0%
3081-6	High	21%	13%	6%	4%	2%	1%
	Expected	15%	7%	3%	1%	0%	0%
	Low	11%	4%	1%	0%	0%	0%
3081-6X	High	33%	20%	10%	5%	3%	2%
	Expected	26%	11%	5%	2%	1%	0%
	Low	18%	6%	2%	1%	0%	0%
3081-K	High	25%	15%	6%	5%	3%	1%
	Expected	17%	9%	4%	2%	0%	0%
	Low	12%	5%	1%	0%	0%	0%
3081-KX	High	40%	21%	14%	9%	5%	2%
	Expected	32%	14%	7%	4%	1%	0%
	Low	23%	9%	4%	2%	0%	0%

EXHIBIT III-3 (Cont.)

**IBM 308X SERIES
PROJECTED RESIDUAL VALUE AS A
PERCENT OF VENDOR LIST PRICE**

PROCESSOR MODEL		PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE as of January 1					
		1987	1988	1989	1990	1991	1992
3084-Q	High	32%	19%	12%	9%	5%	3%
	Expected	27%	13%	7%	4%	1%	0%
	Low	23%	7%	3%	1%	0%	0%
3084-QX	High	45%	25%	17%	11%	7%	4%
	Expected	42%	21%	11%	6%	2%	1%
	Low	38%	12%	5%	2%	0%	0%

EXHIBIT III-4

IBM 3090 SERIES
PROJECTED RESIDUAL VALUE AS A
PERCENT OF VENDOR LIST PRICE

PROCESSOR MODEL		PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE as of January 1					
		1987	1988	1989	1990	1991	1992
3090-150	High	95%	80%	60%	35%	10%	6%
	Expected	92%	74%	51%	26%	5%	3%
	Low	87%	65%	40%	15%	2%	0%
3090-180	High	96%	82%	63%	38%	12%	9%
	Expected	93%	76%	53%	22%	6%	4%
	Low	89%	66%	44%	16%	3%	0%
3090-200	High	95%	85%	65%	40%	15%	12%
	Expected	90%	80%	57%	27%	8%	5%
	Low	86%	70%	45%	18%	4%	1%
3090-400	High	96%	90%	70%	50%	22%	15%
	Expected	92%	85%	65%	31%	12%	8%
	Low	88%	74%	52%	21%	6%	2%

EXHIBIT III-5

IBM 43XX SERIES
PROJECTED RESIDUAL VALUE AS A
PERCENT OF VENDOR LIST PRICE

PROCESSOR MODEL		PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE as of January 1					
		1987	1988	1989	1990	1991	1992
4331-J2	High	8%	5%	3%	2%	1%	1%
	Expected	5%	3%	2%	1%	0%	0%
	Low	2%	1%	0%	0%	0%	0%
4331-L2	High	20%	13%	8%	5%	2%	1%
	Expected	15%	8%	4%	2%	0%	0%
	Low	10%	4%	1%	0%	0%	0%
4341-L1	High	5%	3%	2%	1%	1%	1%
	Expected	2%	1%	0%	0%	0%	0%
	Low	0%	0%	0%	0%	0%	0%
4341-M2	High	7%	4%	3%	2%	1%	1%
	Expected	4%	2%	1%	0%	0%	0%
	Low	2%	0%	0%	0%	0%	0%
4361-K5	High	46%	32%	20%	12%	8%	5%
	Expected	41%	25%	14%	7%	4%	2%
	Low	35%	17%	5%	3%	1%	0%
4361-M5	High	50%	35%	20%	10%	7%	5%
	Expected	46%	31%	13%	8%	5%	3%
	Low	40%	25%	10%	5%	1%	0%
4381-M1	High	68%	58%	30%	20%	12%	8%
	Expected	60%	47%	20%	11%	7%	4%
	Low	57%	40%	12%	5%	2%	0%
4381-P13	High	86%	75%	55%	40%	26%	21%
	Expected	84%	67%	46%	33%	23%	15%
	Low	80%	60%	38%	25%	15%	5%
4381-M12	High	87%	73%	51%	37%	27%	20%
	Expected	84%	66%	45%	31%	20%	16%
	Low	78%	57%	36%	23%	12%	8%
4381-P14	High	90%	79%	57%	45%	31%	25%
	Expected	88%	70%	51%	39%	27%	18%
	Low	83%	65%	42%	28%	17%	7%

EXHIBIT III-6

AMDAHL 58XX SERIES
PROJECTED RESIDUAL VALUE AS A
PERCENT OF VENDOR LIST PRICE

PROCESSOR MODEL		PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE as of January 1					
		1987	1988	1989	1990	1991	1992
5840	High	47%	23%	12%	5%	3%	1%
	Expected	41%	17%	8%	2%	1%	0%
	Low	35%	10%	4%	0%	0%	0%
5850	High	44%	22%	11%	6%	3%	1%
	Expected	39%	16%	6%	3%	1%	0%
	Low	33%	9%	3%	1%	0%	0%
5860	High	45%	22%	13%	8%	6%	1%
	Expected	39%	17%	8%	4%	2%	0%
	Low	35%	11%	3%	1%	0%	0%
5866	High	47%	29%	15%	7%	6%	2%
	Expected	40%	23%	9%	4%	3%	0%
	Low	34%	14%	5%	2%	1%	0%
5870	High	45%	30%	17%	11%	7%	2%
	Expected	39%	22%	9%	4%	3%	0%
	Low	33%	16%	6%	2%	1%	0%
5880	High	44%	27%	15%	10%	6%	3%
	Expected	40%	23%	9%	4%	3%	0%
	Low	35%	15%	4%	1%	0%	0%
5890-200	High	96%	88%	63%	37%	12%	5%
	Expected	92%	82%	53%	18%	5%	2%
	Low	87%	69%	41%	14%	1%	0%
5890-300	High	96%	85%	61%	33%	14%	6%
	Expected	91%	78%	52%	22%	7%	3%
	Low	86%	63%	40%	15%	2%	1%
5890-600	High	97%	89%	70%	45%	18%	7%
	Expected	93%	85%	64%	29%	7%	4%
	Low	86%	70%	47%	19%	4%	1%

EXHIBIT III-7

**NAS AS SERIES
PROJECTED RESIDUAL VALUE AS A
PERCENT OF VENDOR LIST PRICE**

PROCESSOR MODEL		PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE as of January 1					
		1987	1988	1989	1990	1991	1992
AS/6630	High	25%	13%	10%	6%	5%	2%
	Expected	18%	8%	5%	2%	1%	0%
	Low	14%	4%	2%	0%	0%	0%
AS/6660	High	27%	17%	12%	8%	5%	2%
	Expected	21%	12%	7%	4%	2%	0%
	Low	15%	8%	3%	1%	0%	0%
AS/8023	High	42%	21%	13%	9%	6%	3%
	Expected	34%	16%	8%	3%	1%	0%
	Low	29%	9%	3%	0%	0%	0%
AS/8083	High	48%	30%	20%	11%	6%	3%
	Expected	42%	23%	14%	6%	3%	0%
	Low	35%	17%	6%	2%	0%	0%
AS/9050	High	18%	11%	7%	5%	3%	2%
	Expected	10%	6%	3%	2%	1%	0%
	Low	7%	2%	0%	0%	0%	0%
AS/9070	High	20%	13%	8%	5%	3%	2%
	Expected	13%	7%	3%	2%	1%	0%
	Low	9%	2%	0%	0%	0%	0%
AS/XL50	High	94%	84%	62%	30%	18%	11%
	Expected	90%	79%	54%	16%	4%	2%
	Low	85%	71%	42%	10%	1%	0%
AS/XL60	High	95%	87%	65%	34%	17%	13%
	Expected	90%	80%	56%	18%	5%	3%
	Low	86%	73%	42%	12%	2%	0%
AS/XL90	High			70%	41%	22%	16%
	Expected	*	*	60%	27%	8%	5%
	Low			45%	15%	4%	1%

*Delivery Scheduled Early 1988.

About INPUT

INPUT provides planning information, analysis, and recommendations to managers and executives in the information processing industries. Through market research, technology forecasting, and competitive analysis, INPUT supports client management in making informed decisions. Continuing services are provided to users and vendors of computers, communications, and office products and services.

The company carries out continuous and in-depth research. Working closely with clients on important issues, INPUT's staff members analyze and interpret the research data, then develop recommendations and innovative ideas to meet clients' needs.

Clients receive reports, presentations, access to data on which analyses are based, and continuous consulting.

Many of INPUT's professional staff members have nearly 20 years' experience in their areas of specialization. Most have held senior management positions in operations, marketing, or planning. This expertise enables INPUT to supply practical solutions to complex business problems.

Formed in 1974, INPUT has become a leading international planning services firm. Clients include over 100 of the world's largest and most technically advanced companies.

Offices

NORTH AMERICA

Headquarters

1943 Landings Drive
Mountain View, CA 94043
(415) 960-3990
Telex 171407

New York

Parsippany Place Corp. Center
Suite 201
959 Route 46 East
Parsippany, NJ 07054
(201) 299-6999
Telex 134630

Washington, D.C.

8298 C, Old Courthouse Rd.
Vienna, VA 22180
(703) 847-6870

EUROPE

United Kingdom

INPUT
41 Dover Street
London W1X 3RB
England
01-493-9335
Telex 27113

Italy

Nomos Sistema SRL
20124 Milano
Viale Vittorio Veneto 6
Italy
228140 and 225151
Telex 321137

Sweden

Athena Konsult AB
Box 22232
S-104 22 Stockholm
Sweden
08-542025
Telex 17041

ASIA

Japan

ODS Corporation
Dai-ni Kuyo Bldg.
5-10-2, Minami-Aoyama
Minato-ku,
Tokyo 107
Japan
(03) 400-7090
Telex 26487

INPUT[®]
Planning Services For Management

