Market Analysis: Data Base Management Systems

INPUT®

MARKET ANALYSIS:
DATA BASE MANAGEMENT SYSTEMS

MARKET ANALYSIS:
DATA BASE MANAGEMENT SYSTEMS

CONTENTS

MARKET ANALYSIS:
DATA BASE MANAGEMENT SYSTEMS

EXHIBITS

I INTRODUCTION

A. REASONS FOR PREPARING REPORT

- Last year INPUT prepared a major study--Market Impacts of IBM Software Strategies. That report provided a framework in which to view general software trends and market opportunities, and it provided, for the first time, a realtively clear picture of the complex interrelationships between various levels in the software hierarchy. This report will analyze the market for data base management systems (DBMS) within the framework of IBM's strategy and projected technological trends.

- There are two major trends which promise to have substantial impact on the market for DBMS:

  - Advanced office automation systems which integrate voice, data, and images (electronic filing).

  - The extension of decision support systems toward expert- or knowledge-based systems.

  - These trends present both opportunities and challenges for vendors currently competing in the DBMS market.

- IBM's DBMS strategy and industry trends toward information centers, prototyping, and micro-mainframe links portend a complex set of problems for DBMS users. These problems were defined in Market Impact of New Software Productivity Techniques, INPUT, 1984. Essentially, the problems relate to:

  - Performance of generalized DBMS in highly centralized (very large) data base environments.

  - Data base integrity and synchronization in a distributed data base environment.

  - Security, protection, and privacy considerations in both environments.

  - The potential deterioration of data and information quality as a result of these problems poses a substantial threat to the general purpose DBMS market unless they are resolved, but, conversely, represents opportunities for those who can find solutions.

B. RESEARCH BASE

- In addition to the two reports mentioned above, several other INPUT reports (and their supporting research bases) were primary sources for this report. Prominent among them were:

  - Impact of Upcoming Optical Memory Systems, 1983.

  - Methods of Cost/Benefit Analysis for Office Systems, 1983.

  - Relational Data Base Development, 1983.

  - New Opportunities for Software Productivity Improvement, 1984.

- Integrated DBMS-Application Software, 1984.

- Artificial Intelligence and Export Systems, 1985.

- In addition to focusing extensive past INPUT research on the market requirements for DBMS, it was also necessary to broaden the theoretical perspective of the DBMS market to include information and knowledge.

C. SCOPE AND METHODOLOGY

- This study will focus on the market requirements for tools, aids, and facilities to manage data, information, and knowledge quality in an environment which currently tends to emphasize ease of access (and use) and to assume quality. In other words, it will analyze the market from the perspective of the data base administrator, the information systems professional, and the corporate information officer, not the end user.

- Today's rapidly changing technological environment (hardware, software, and communications networks) is constantly probing technological and theoretical boundaries and exceeding practical developmental limits. There is an expanding divergence in the DBMS market between wants, needs, expectations, and current and/or potential products. This report will attempt to:

  - Distinguish between wants, needs, and expectations.

  - Analyze current products in light of these "requirements."

  - Identify rough technological and theoretical boundaries (problems) which currently exist in satisfying these "requirements."

- Isolate true market requirements which can be satisfied without technological or theoretical breakthroughs.

- The methodology employed in preparing this report was as follows:

  - A critical analysis of past INPUT and external research projects was made, with emphasis upon the reports and publications listed in the section which described the research base.

  - As new conclusions were reached (or old conclusions were reaffirmed), selected experts and authorities (both user and vendor) were interviewed by telephone to discuss and refine INPUT's findings.

- INPUT considers the management and control of data, information, and knowledge to be fundamental to improved productivity in the systems development process. To the degree that this report redefines the structure of the market for DBMS, it will have significant market impacts upon language developments and applications development tools. These impacts are analyzed in two companion MAPS reports:

  - Market Analysis: Fourth Generation Languages, 1985.

  - Market Analysis: Applications Development Tools, 1985.

- INPUT believes these reports will provide a framework for software product and software market planning into the 1990s.

II EXECUTIVE SUMMARY

- This executive summary is designed in presentation format in order to:
  - Help the reader review the key research findings.
  - Provide an executive presentation to facilitate group communications.
- The key points of the entire report are summarized in Exhibits II-1 through II-7.
- It is recommended that the entire report be read in order to make effective use of the summary presentation, and supporting INPUT reports should be reviewed to support any product planning activities.

A. DBMS: KEY TO IBM CONTROL AND GROWTH? (SNA/DDP PERIOD)

- INPUT projects that IBM's primary emphasis for the remainder of the 1980s will be on the following:
  - Continued emphasis upon highly centralized host control through SNA, operating systems, and DBMS.
  - Control of distributed processing and data bases.
  - This is the IBM "SNA/DDP strategic period."
- During the SNA/DDP strategic period, IBM will:
  - More than double its total revenue.
  - More than triple its software revenue.
  - Increase its DBMS revenue over five times.
- While DBMS revenue will still represent only 3% of IBM's total revenue, INPUT believes DBMS will be the key to IBM's account control and to its all-important mainframe and magnetic storage sales.
- It is anticipated that the effective DBMS market--that market left after IBM takes its share--will approach $3 billion and be highly competitive. Exhibit II-1 summarizes INPUT's market projections for the IBM software/DBMS markets through 1990.

DBMS: KEY TO IBM CONTROL & GROWTH?

MARKET	$ Billions	
	1984	1990
• Total IBM Revenue	$46.0	$100.0
- IBM Software Revenue	2.4	10.3
• Total DBMS Market	1.0	6.4
- IBM DBMS Share	.6	3.6
- Non-IBM DBMS Share	.4	2.8

B. AN IBM CONTROLLED ENVIRONMENT

- IBM has awesome tools and sound technical reasons for centralizing control of distributed data base synchronization, assuring central data base integrity, and centralizing data protection and security.

- Interfacing with IBM hardware/software systems is becoming increasingly complex and expensive.

  - Not only will IBM continue to shift functions between hardware, firmware, and software, but the same shifting of functions will take place between SNA, operating systems, and DBMS. INPUT projects IBM will tend to move functions toward their DBMS.

  - The withholding of source code will obviously make the task of interfacing more difficult for other software product vendors.

- IBM's schedule for micro-mainframe linkages and LANs leaves many open questions for both users and vendors. This keeps the market in suspense, delaying the development of available markets for users and other vendors and requiring them to synchronize their plans with IBM's schedule of announcements.

- IBM will exercise control over distributed data base development by a confusing array of workstations and by emphasizing the very real, unresolved problems associated with distributed data bases.

- Exhibit II-2 graphically portrays this user/vendor dilemma.

AN IBM CONTROLLED ENVIRONMENT

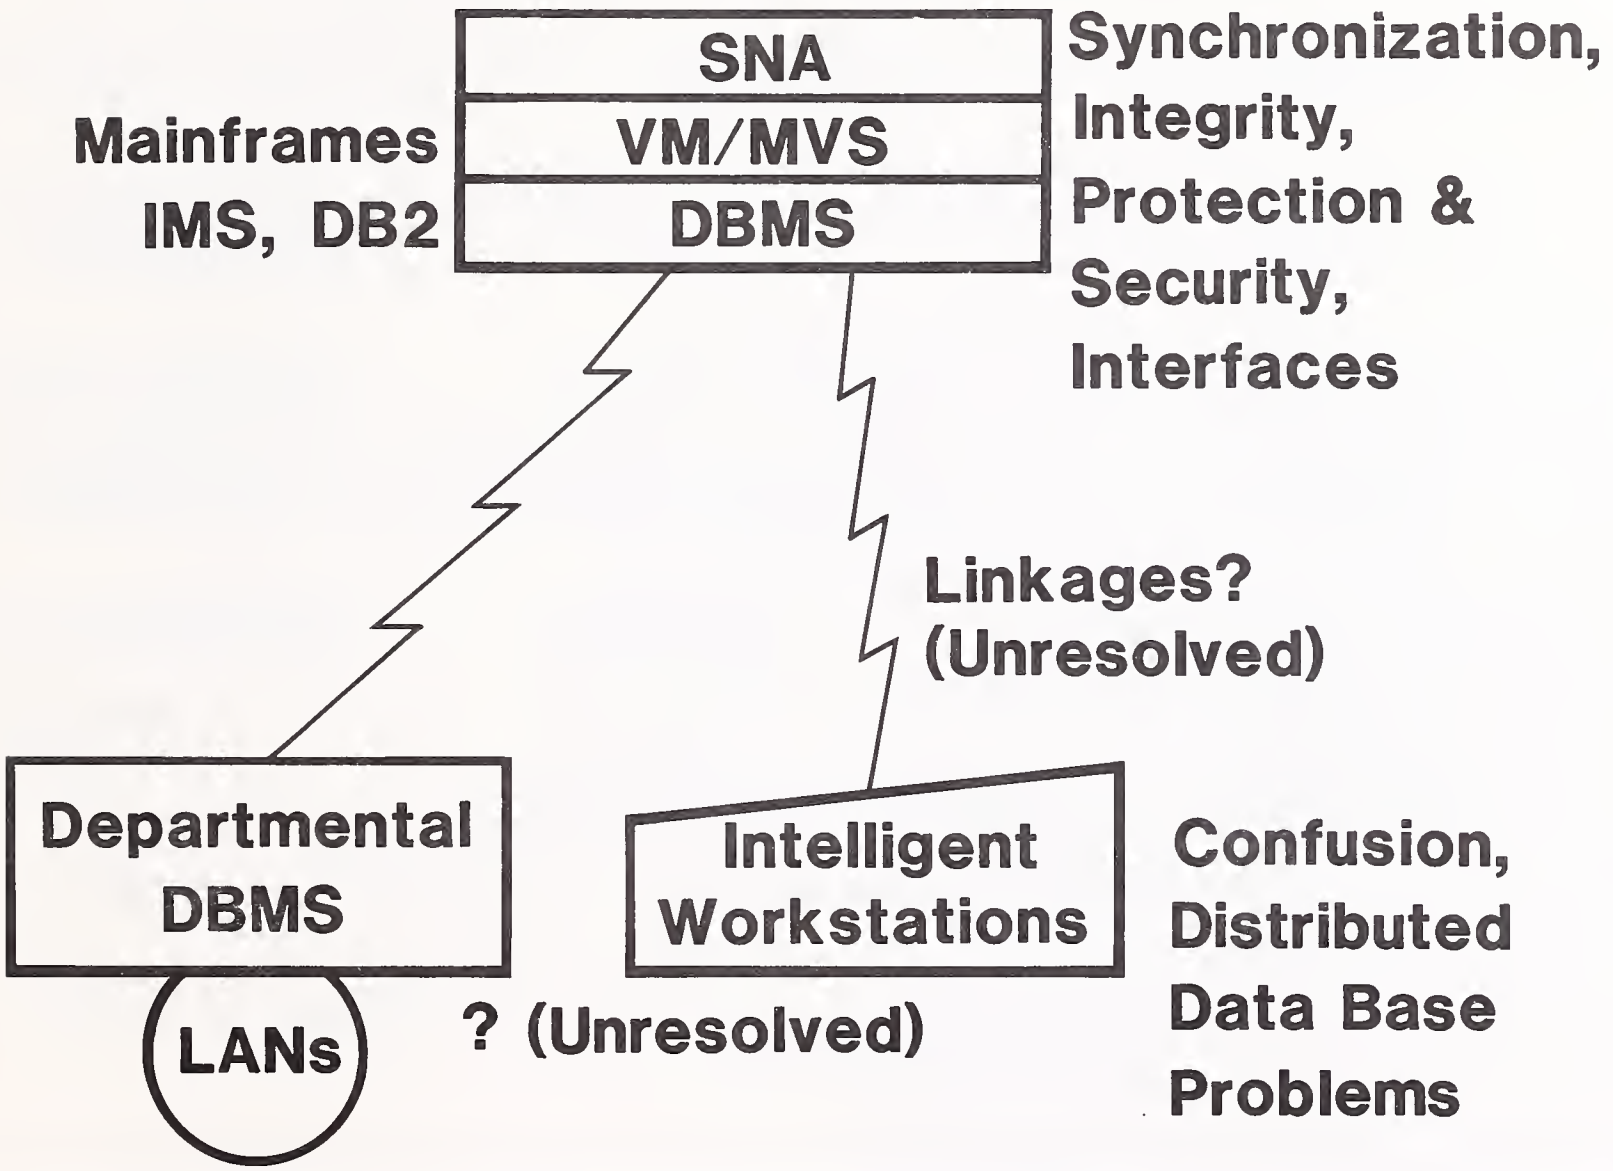

C. MEDIA INTEGRATION

- There are fundamental changes taking place in the storage of data, information, and knowledge.

  - Currently, data are generally stored in processable form on magnetic media, information is stored on paper or micrographic media, and knowledge is stored in human brains or on paper media.

  - The trend is toward integration of data, information, and knowledge bases, but magnetic media remains too costly to store significant amounts of information and knowledge.

- Optical memories offer the promise of dramatic advances in storage capacity and cost. As a result, optical memories will facilitate the integration of current data, information, and knowledge bases.

- It is probable that the following will result:

  - Optical memory will replace magnetic media for storage of archival and raw data, although magnetic media may remain during processing.

  - Significant amounts of paper information will be replaced by substituting optical disks for paper files.

  - While humans will remain the primary repository for knowledge, optical memories will permit storing (and ready maintenance and access) of knowledge bases currently contained in text books and encyclopedias.

- Exhibit II-3 synthesizes these trends.

INPUT®

MEDIA INTEGRATION

Base Type	Present		Future
Data Bases	Magnetic	→	Optical
Information Bases	Paper (Files) Micrographic	→	Optical Paper
Knowledge Bases	Humans Paper (Books & Files)	→	Humans Paper Optical

D. PROCESSING INTEGRATION

- The processing of data, information, and knowledge bases is also undergoing significant changes; at present:

  - Data bases are being processed primarily on general purpose mainframe computers.

  - Information processing is accomplished through human interaction and document preparation.

  - Knowledge processing is accomplished in the individual human brain or through human interaction.

- Developments in computer technology and networking will permit the integration of distributed data bases with voice, image, and text information bases, and the integration of both of these with human beings who will continue to be the primary processors of data and information into knowledge.

- The technical developments supporting this integration will result from an increasing trend toward differentiation and mechanization of processing functions into data base machines, distributed minicomputers and microprocessors, and specialized processors for the implementation of artificial intelligence-based systems (see Exhibit II-4).

INPUT®

PROCESSING INTEGRATION

Processing Type	Present	Future
Data Base Processing	Mainframes	DBM, Mini/Micro
Information Base Processing	Human Interaction PC/WP, Graph.	Human Interaction Integrated Networks
Knowledge Base Processing	Human Interaction	Human Interaction Human-Computer Interaction

E. OPPORTUNITIES FOR DATA, INFORMATION, AND KNOWLEDGE MANAGEMENT

- Current general purpose mainframes can be supplemented with optical memories (e.g., Reference Technology Inc.) and back-end data base machines (e.g., Britton-Lee). This will provide needed performance improvement and storage capacity for the emerging distributed data base environment.

- Technology is currently available to facilitate the development of integrated office systems which will:

  - Effectively integrate data, image, and voice in a cost effective processing and storage hierarchy.

  - Reduce paper handling by providing electronic filing and library access.

- The experience gained through the development of DBMS must be applied to office automation and expert systems development where integration of data, information, and knowledge is necessary, but poorly understood.

- The windows of opportunity for the application of advanced processor and storage technology, and for the integration of data (D), information (I), and knowledge (K) will continue while IBM concentrates on centralized control in the SNA/DDP period (see Exhibit II-5).

- INPUT estimates this strategic period will extend through the remainder of the 1980s.

OPPORTUNITIES FOR D/I/K MANAGEMENT

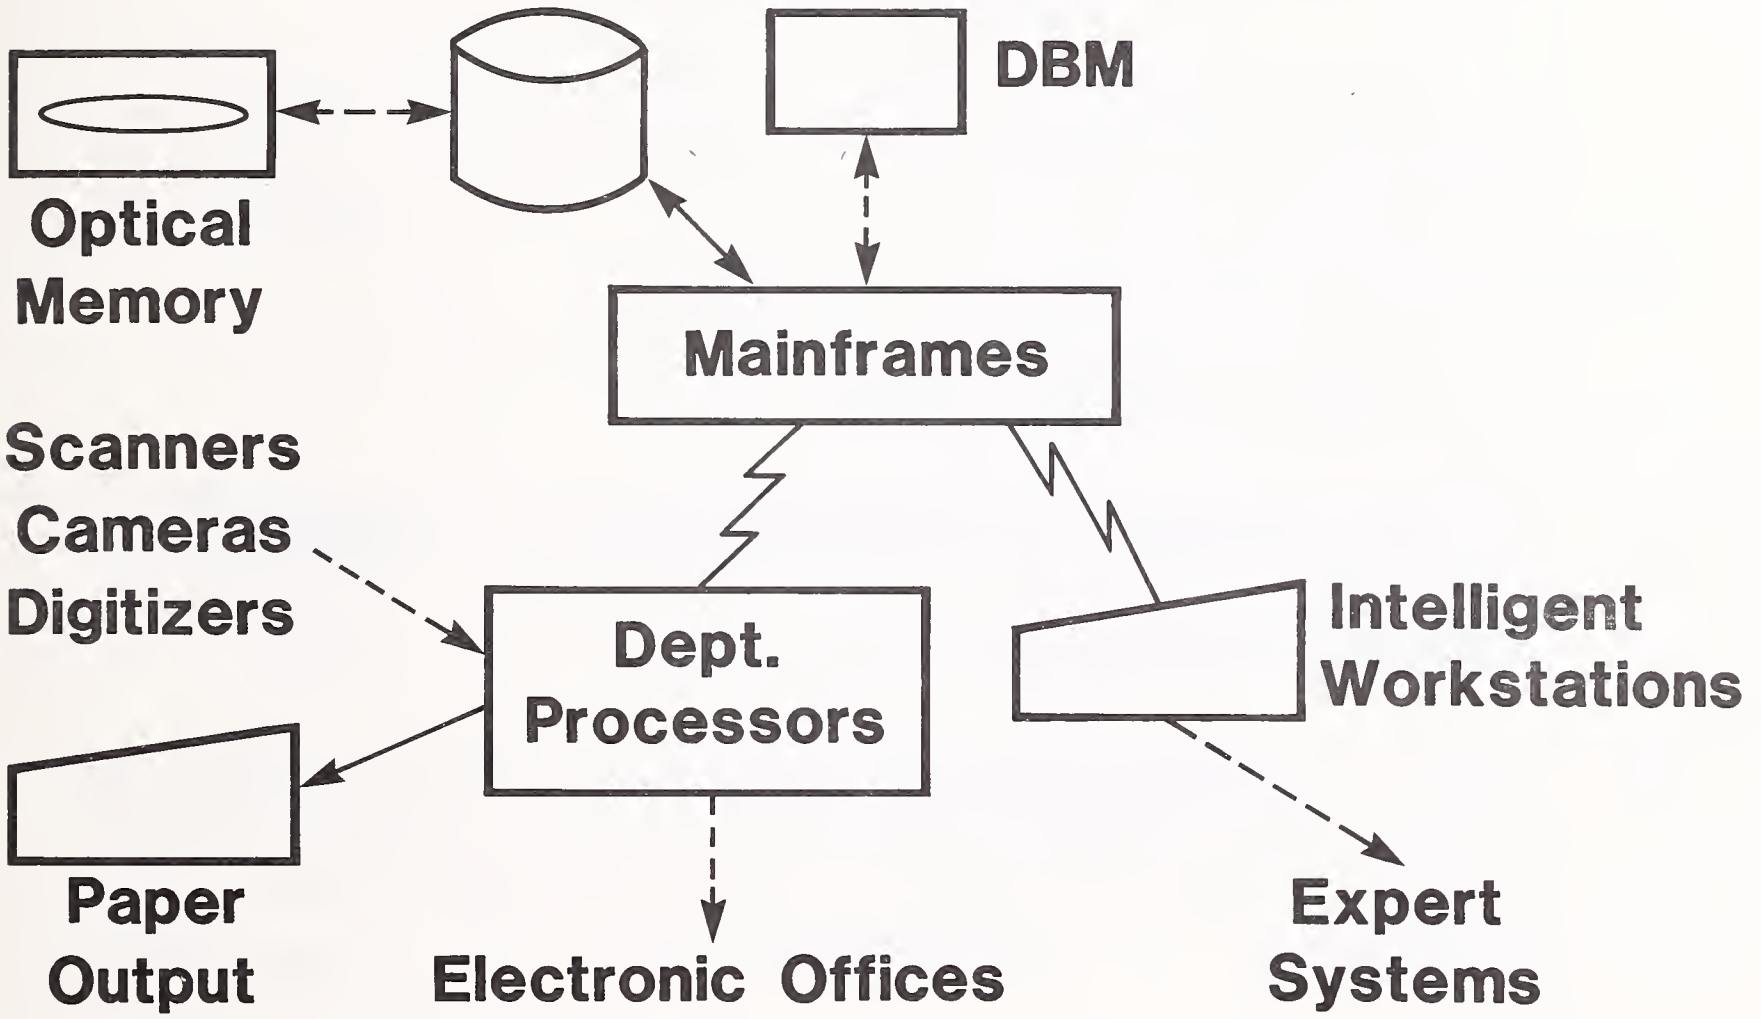

F. TECHNICAL CONSIDERATIONS (CENTRAL HARDWARE)

- All of the opportunities for the extension of DBMS into new areas assume that the outstanding problems of distributed data bases will be addressed. However, when selecting opportunities of promise, it is important to understand specific technical considerations:

  - Optical memories will require new data (information) models to accommodate storage and retrieval of potentially large, variable length records.

  - Effecting implementation may require storage management facilities with larger page sizes.

- Data base machines will require knowledge of advanced computer and memory architecture, plus the ability to assure a clean mainframe interface with IBM hardware and DBMS.

- Support of scanners, cameras, and digitizers will require knowledge of data/information compression and pattern recognition techniques. Both artificial intelligence and information theory will be important in addressing these problems.

- In each area a successful vendor must not only understand and control the techniques necessary for addressing the market, but have an understanding of how users' needs may be practically served by the new capabilities. This is a far from easy task (see Exhibit II-6).

INPUT®

TECHNICAL CONSIDERATIONS
(Central Hardware)

Optical Memories	**- Data & Information Structure**
	- Variable Length Records (Pages)
	- OS Paging Bottlenecks
Data Base Machines	**- Alternative Architectures**
	- Data & Information Models
Scanners, Cameras, Digitizers	**- Pattern Recognition**
	- Compression

G. TECHNICAL CONSIDERATIONS (SYSTEMS)

- At the systems level, these problems are compounded.

- Electronic offices will require a highly advanced communications and systems integration capability to service new storage devices, processor complexes, I/O devices, and workstations. Local area networks for integrated data, information, and knowledge bases will bring to the surface performance problems which will probably require general knowledge of operations research techniques (such as queuing networks) for analysis and operation (see Exhibit II-7).

- Expert systems are in their infancy, and it is important to consider the requirements for both development and maintenance of knowledge-based systems as well as the hardware/software tools (such as LISP development engines) which will be necessary for their successful implementation and integration.

- Experts on expert systems are also in short supply (in 1984 it was estimated that the total U.S. population of expert systems engineers was only 3,000), and in-house capabilities may have to be developed within an organization contemplating entering this market.

- In summary, the next five years promise to be ones of extensive upheaval and rapid technical development affecting the very core of the U.S. information systems installed base.

INPUT®

TECHNICAL CONSIDERATIONS (Systems)

Electronic Offices	**- Integration of Data, Information & Voice**
	- LAN Structure
	- Queuing (Performance) Problems
Expert Systems	**- Knowledge-Base Structure**
	- Development Engines (Workstations)

III CURRENT MARKET ANALYSIS

A. INDUSTRY TRENDS

1. RELATIONAL EMPHASIS AND DATA MODELS

- In 1983 INPUT published a report, Relational Data Base Developments, which summarized the advantages and disadvantages of relational data base systems (RDBS) as follows:

  - The primary advantages normally associated with RDBS are flexibility and ease of use. These advantages apply to both professional systems developers and end users; they are considered of special importance because of current trends toward prototyping, information centers, and decision support systems. (Today, we might add micro-mainframe links.)

  - An advantage which is not frequently mentioned, but which INPUT considers to be especially important, is the communicability of the simple relational structures (tables) among users, programmers, and data base administrators, especially in the current environment where end-user development is encouraged. In fact, bridging the communications gap between the corporate data base designers and end users could be the most important contribution of the RDBS.

- The primary disadvantages of relational systems has been and continues to be performance. There are costs associated with flexibility and simplicity even when an RDBS is used properly. When misused, an RDBS can result in a prohibitively expensive system.

- While there is a pronounced and very real trend toward relational systems, even experienced systems personnel (much less end users) do not understand the relative technical advantages and disadvantages of various data models. And, as any marketing manager is already aware, you do not sell DBMS talking about data structure types, relational algebra, and normal forms. In fact, the general attitude in the marketplace is best summarized by an IS manager who responded to an INPUT questionnaire by saying, "I don't even like to talk about data models--the whole subject bores me."

- Therefore, this report will not delve into the technical aspects of various DBMS except to state the following:

  - While the world may appear to be going relational, there remain various user sets who "think about" data in different ways, and whether it is a simple sequential file, tables in third normal form, a hierarchy, a network, or a seven dimensional array, the mere fact that data are viewed in different ways is sufficient reason to believe various data models are needed.

  - In addition, information and knowledge bases (to be defined later) will present a different set of problems from those encountered with the management of encoded data elements. Specifically, information elements are of varying length and format, and knowledge bases require complex structures and navigation schemes.

  - Therefore, it can safely be assumed that not only will current hierarchical, network, and relational models remain, but new models will be required in support of new applications.

2. DISTRIBUTED DATA BASES

- There have been substantial theoretical discussions of distributed data bases and their potential--unresolved problems for years. However, there is little question that the predominant trend in the DBMS market is in that direction.

- It is important to recognize that there are many ways to distribute data bases, and that there is no one right way. In fact, INPUT has long seen fit to distinguish between "geographic" and "architectural" distribution of processing and this implies the distribution of data. Exhibit III-1 presents a diagram of some of the ways data may be distributed.

  - The current emphasis upon micro-mainframe links is essentially a distributed data base problem--not a communications problem. Various solutions are possible, with a simple and popular one being to distribute data in report format. (A natural solution since standalone personal computer users frequently had been keying in their data from computer reports.)

  - Then, of course, there is hierarchical distribution by tying micros to minis (or small mainframes) to large-host mainframes.

  - When processing power or data base size become problems on the central host, users frequently find it necessary to distribute data bases across mainframes. In fact, it is sometimes desirable to duplicate data bases across mainframes (as in the case of information and/or development centers), and this can be perceived as just another form of distributed data bases.

  - Of course, the use of multiple data models and/or DBMS represents another form of distribution which is being selected by (or forced upon) users because of data base size, transaction rates, or performance.

EXHIBIT III-1

DISTRIBUTED DATA BASES

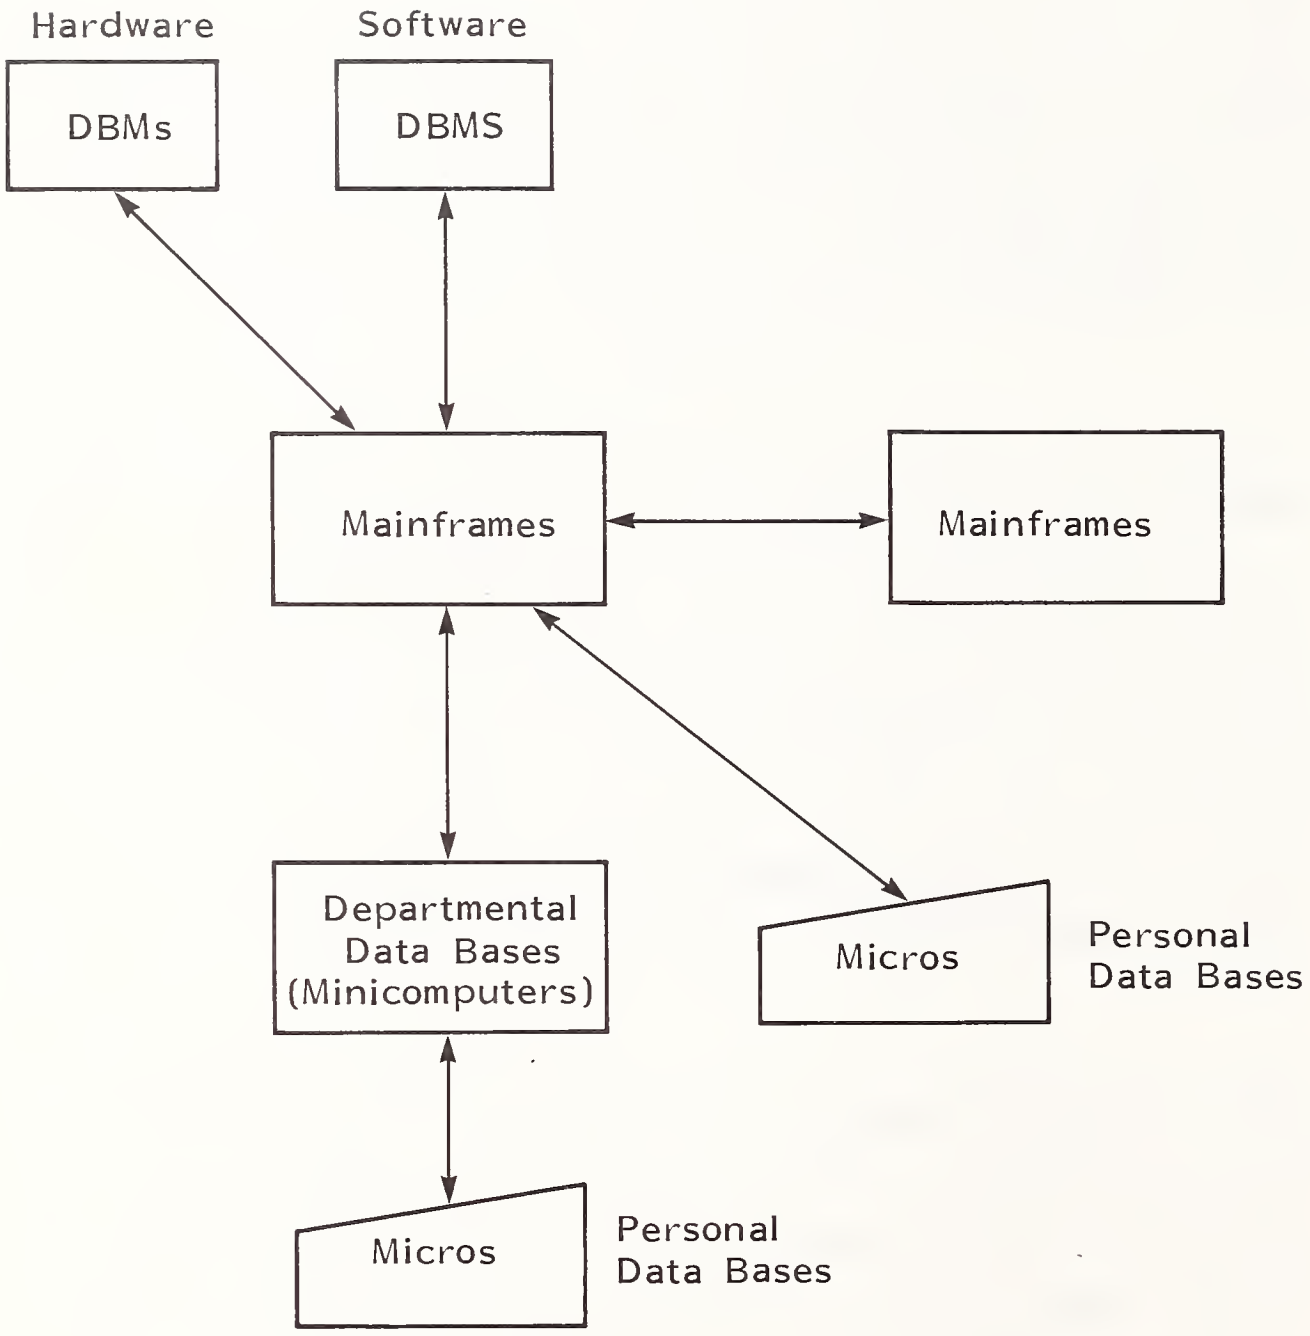

(Another view of distribution, not shown on the diagram, is based on activity and can be classified as high response, batch, and archival data bases.) In addition, back-end data base machines (DBM) are being selected to solve certain performance problems of software DBMS.

- While the form of distributed data bases varies considerably, they create and share a common set of problems (which will be discussed later in this report). It is INPUT's belief that these problems will become more complicated for the following reasons:

  - Networks are becoming more complex and data is flowing more frequently from data base to data base and location to location whether on international or local area networks (LANs).

  - Data, information, and knowledge bases are interdependent, and representatives and media vary substantially as the boundaries are crossed. Computer scientists and information systems specialists, who have been struggling with the problems of data base management, are being confronted with a new set of problems which will be several orders of magnitude more difficult to solve. In fact, most of the new problems are not generally recognized, much less understood.

3. DATA/INFORMATION/KNOWLEDGE DEFINED

- While there are fine nuances of what data, information, and knowledge really are (and even experts cannot agree at the points of interface), INPUT finds the following definitions convenient:

  - Data are facts (admittedly of varying quality) and for our purpose it will be assumed they are encoded for computer processing.

  - Information is the "action of informing or the fact of being told something." In other words, it is communication and requires a sender and

receiver. It may be transmitted directly between humans or be broadcast to many humans, but it doesn't exist until the connection is made. (If I write a book or generate a report from a data base, it isn't information until somebody read it--an important point.)

- Knowledge and information are closely linked, with information generally being referred to as "the transfer of knowledge." However, several ways of distinguishing knowledge from information are necessary:

  . Information is piecemeal, fragmented, and particular, whereas knowledge is structured, coherent, and often universal.

  . Information is timely, transitory, and perhaps ephemeral, whereas knowledge is of enduring significance.

  . Information is a flow of messages, whereas knowledge is a stock (pool, base, etc.) which results from the (information) flow.

  . In addition, there is an additional important distinction. Information is the result of being told, but knowledge can be created by thinking (in other words, without information being received).

  - It is also important to recognize that these definitions do not relate to practical usefulness--it is possible to have useless data, information, and knowledge. In fact, quantity versus quality of data, information, and knowledge is the most important challenge to the information systems industry.

- For all practical purposes, it is possible to say that data are stored encoded and processable on magnetic media, information (including that generated from computer systems and telephone conversations) is stored on paper in file cabinets, and knowledge resides in libraries and human brains. At the present time, the following trends can be discerned:

- Data are being widely distributed for processing (information generation) by lower level computer systems and human beings in hierarchical networks of varying complexity (as previously described). Responsibility for central data bases and their distribution resides in the information systems function.

- Information at the office level is increasingly being communicated over local area networks (LANs), but the quantity of paper documents continues to increase more rapidly than it can be absorbed (or effectively used). End users and a relatively inexperienced cadre of "office automation experts" are responsible for the efforts to apply advanced computer and communications technology in the office environment.

- Knowledge required for operating, planning, and decision making remains highly specialized in individual and small work units. Efforts to encode knowledge in computer-processable form (expert systems) is in its infancy and depends upon a select group of "knowledge engineers" spawned out of academic research in artificial intelligence.

- It is obvious that data, information, and knowledge are interdependent in terms of quality--good information depends upon good data, knowledge required for decision making depends upon a continuous flow of quality information, and human knowledge gained must be communicated back into data bases if computer systems are to be improved. Unfortunately, the difficult lessons learned over years of development of data bases and DBMS seem to have been lost in the current systems development environment which appears to emphasize quantity (responsiveness) over quality. In addition, there is a general spirit of competition among those developing information systems, office systems, and expert systems.

4. DEVELOPMENT TRENDS

- Current development efforts which focus on information centers, prototyping, micro-mainframe links, and decision support systems (evolving into expert systems) are all dependent upon a source of high quality data. Sometimes the developers of information- and knowledge-based systems are unaware of their dependence, much less of the complexities and problems associated with maintaining data quality (especially with distributed data bases, which is what information and knowledge bases require). However, the dependence is inevitable.

- Specifically, the necessity for integrating voice, data, and images in office automation systems has been apparent for some time. However, only now is the technology becoming available to drive such integration. Electronic filing is already being referred to as "image base management," and it is probably only a question of time before "audio-base management" systems emerge. Such distinctions will only serve to slow integration.

- Current knowledge base development tools do not interface with current DBMS, and the developers (rather than being embarrassed) cannot decide whether to build an interface or build their own DBMS.

- The market for DBMS is now going to be confused with the entry of "information base management systems," "image base management systems," and "knowledge base management systems," all of which are going to have substantial areas of overlap.

5. ADDITIONAL LEVELS OF INTEGRATION

- Data, information, and knowledge base (DIK) management systems (in the strict sense that they will be addressed in this report) do not include FGLs and/or other program or systems development tools. However, FGLs, applications generators, systems generators, and visual programming aids are all dependent upon some type of DBMS or DIK management system.

- Whether the market is considered to be the DBMS-FGL market, FGL-DBMS market, the DSS market, the knowledge-based systems market, or the expert systems market is somewhat immaterial--there is going to be a substantial amount of overlap and confusion. INPUT believes DBMS are the foundation and common thread for all of these various integration efforts and that is why specific tools for controlling the development and maintenance of high quality data bases will be emphasized in this report.

- Language development and markets will be analyzed in Market Analysis: Fourth Generation Languages, and the markets for other development tools will be analyzed in Market Analysis: Applications Development Tools.

B. IBM'S STRATEGY (DBMS)

1. GST AND STRATEGIC PERIODS

- In Market Impact of IBM Software Strategies, INPUT, 1984, IBM's software strategy was broken down into four strategic periods based on the emphasis placed upon the general systems theory concepts of centralization, integration, differentiation, and mechanization. (Please refer to that report for detailed definitions and descriptions.) These strategic periods were defined as follows:

  - The SNA/DDP (System Network Architecture/Distributed Data Processing) Period would extend through the 1980s, and IBM would emphasize centralization and control through concentration on SNA, operating systems, and DBMS.

  - The Electronic Office Period would extend from 1990-1995, and IBM would emphasize integration of languages/decision support systems,

industry turnkey systems, and applications packages under the centralized control established in the SNA/DDP period--in other words, under IBM's SNA, operating systems, and DBMS.

- The Expert Systems Period would extend from 1995-2000, and IBM would emphasize differentiation of the integrated systems established during the electronic office period. This would be accomplished by providing data/information/knowledge-based systems (expert systems) in specific domains--general DSS systems would be differentiated into expert systems.

- The Custom Products Period would extend beyond 2000, and IBM would emphasize mechanization through the narrowing of hardware/software/DIK product and service offerings down to the individual.

- IBM's software focus, GST emphasis, and strategic periods are summarized in Exhibit III-2. IBM will naturally give attention to all of the software areas during all four strategic periods, and GST trends always progress in parallel as complex systems develop--IBM's strategy represents emphasis. Essentially, this chart predicts:

  - It will be business as usual in the 1980s with emphasis upon conventional data processing applications and data bases.

  - The early 1990s will see emphasis upon information bases which will substitute electronic for paper media.

  - In the late 1990s, IBM's emphasis will be on DIK (knowledge-based) products and services within expert domains.

  - After the year 2000, custom-tailored products and services will be provided to the individual.

EXHIBIT III-2

SUMMARY OF IBM'S STRATEGY

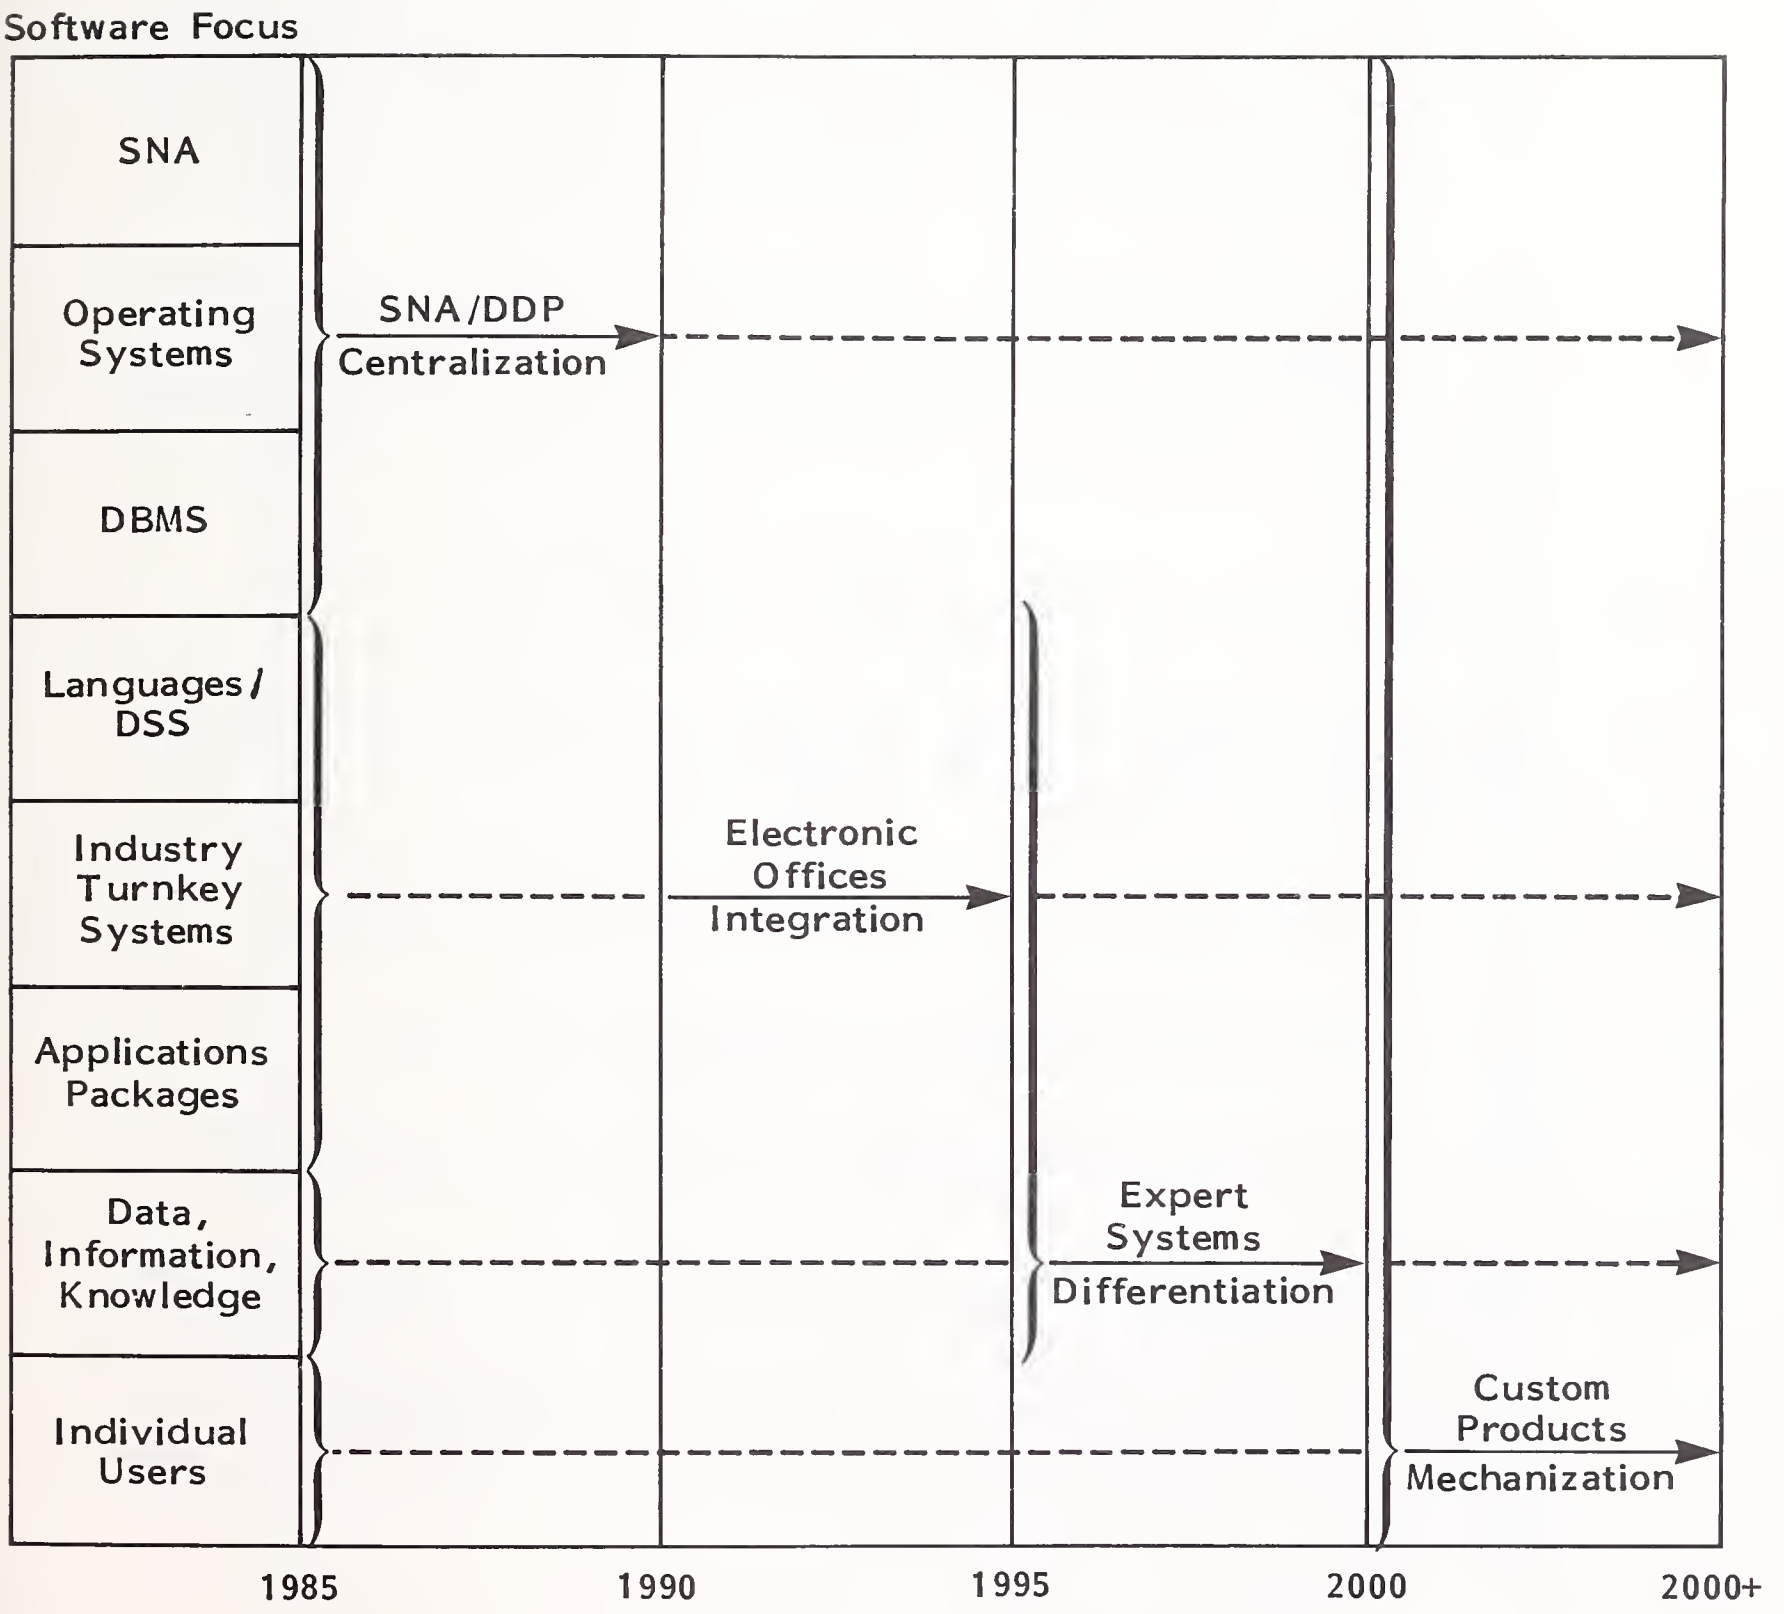

- While the primary emphasis of the SNA/DDP is upon centralization of control on large host mainframes, this applies primarily to operating systems functions such as storage management, protection and security, resource allocation, and system structure. It was specifically pointed out in Market Impacts of IBM Software Strategies that IBM would have to emphasize integration of DBMS during the SNA/DDP period because DDP means distributed data bases (and a variety of DBMS and data models) and these would require integration.

- It should also be remembered that IBM's primary strategy is to sell hardware, and when IBM thinks of hardware it thinks big. IBM's data base strategy has been extremely effective in selling hardware in the past, and it appears it will remain that way in the future.

2. LARGE MAINFRAMES AND MAGNETIC STORAGE

- Early in 1985, IBM made a public statement about the anticipated growth for its "average" large customer. The targets are for:

  - Processing power to increase from 100 MIPS to 1,000 MIPS.

  - System memory to increase from 200 MB to 10,000 MB.

  - Direct access storage (DASD) to increase from 300 GB to 6,000 GB.

- These growth rates are projected in Exhibit III-3, and it is apparent DBMS vendors are going to have a lot of DASD to work with and a lot of MIPS to drive their systems. In fact, in its Large-Scale Systems Directions series of reports, INPUT has referred to IBM mainframes as "large host data base machines."

EXHIBIT III-3

LARGE-SCALE SYSTEMS GROWTH TO 1990

(IBM Projections)

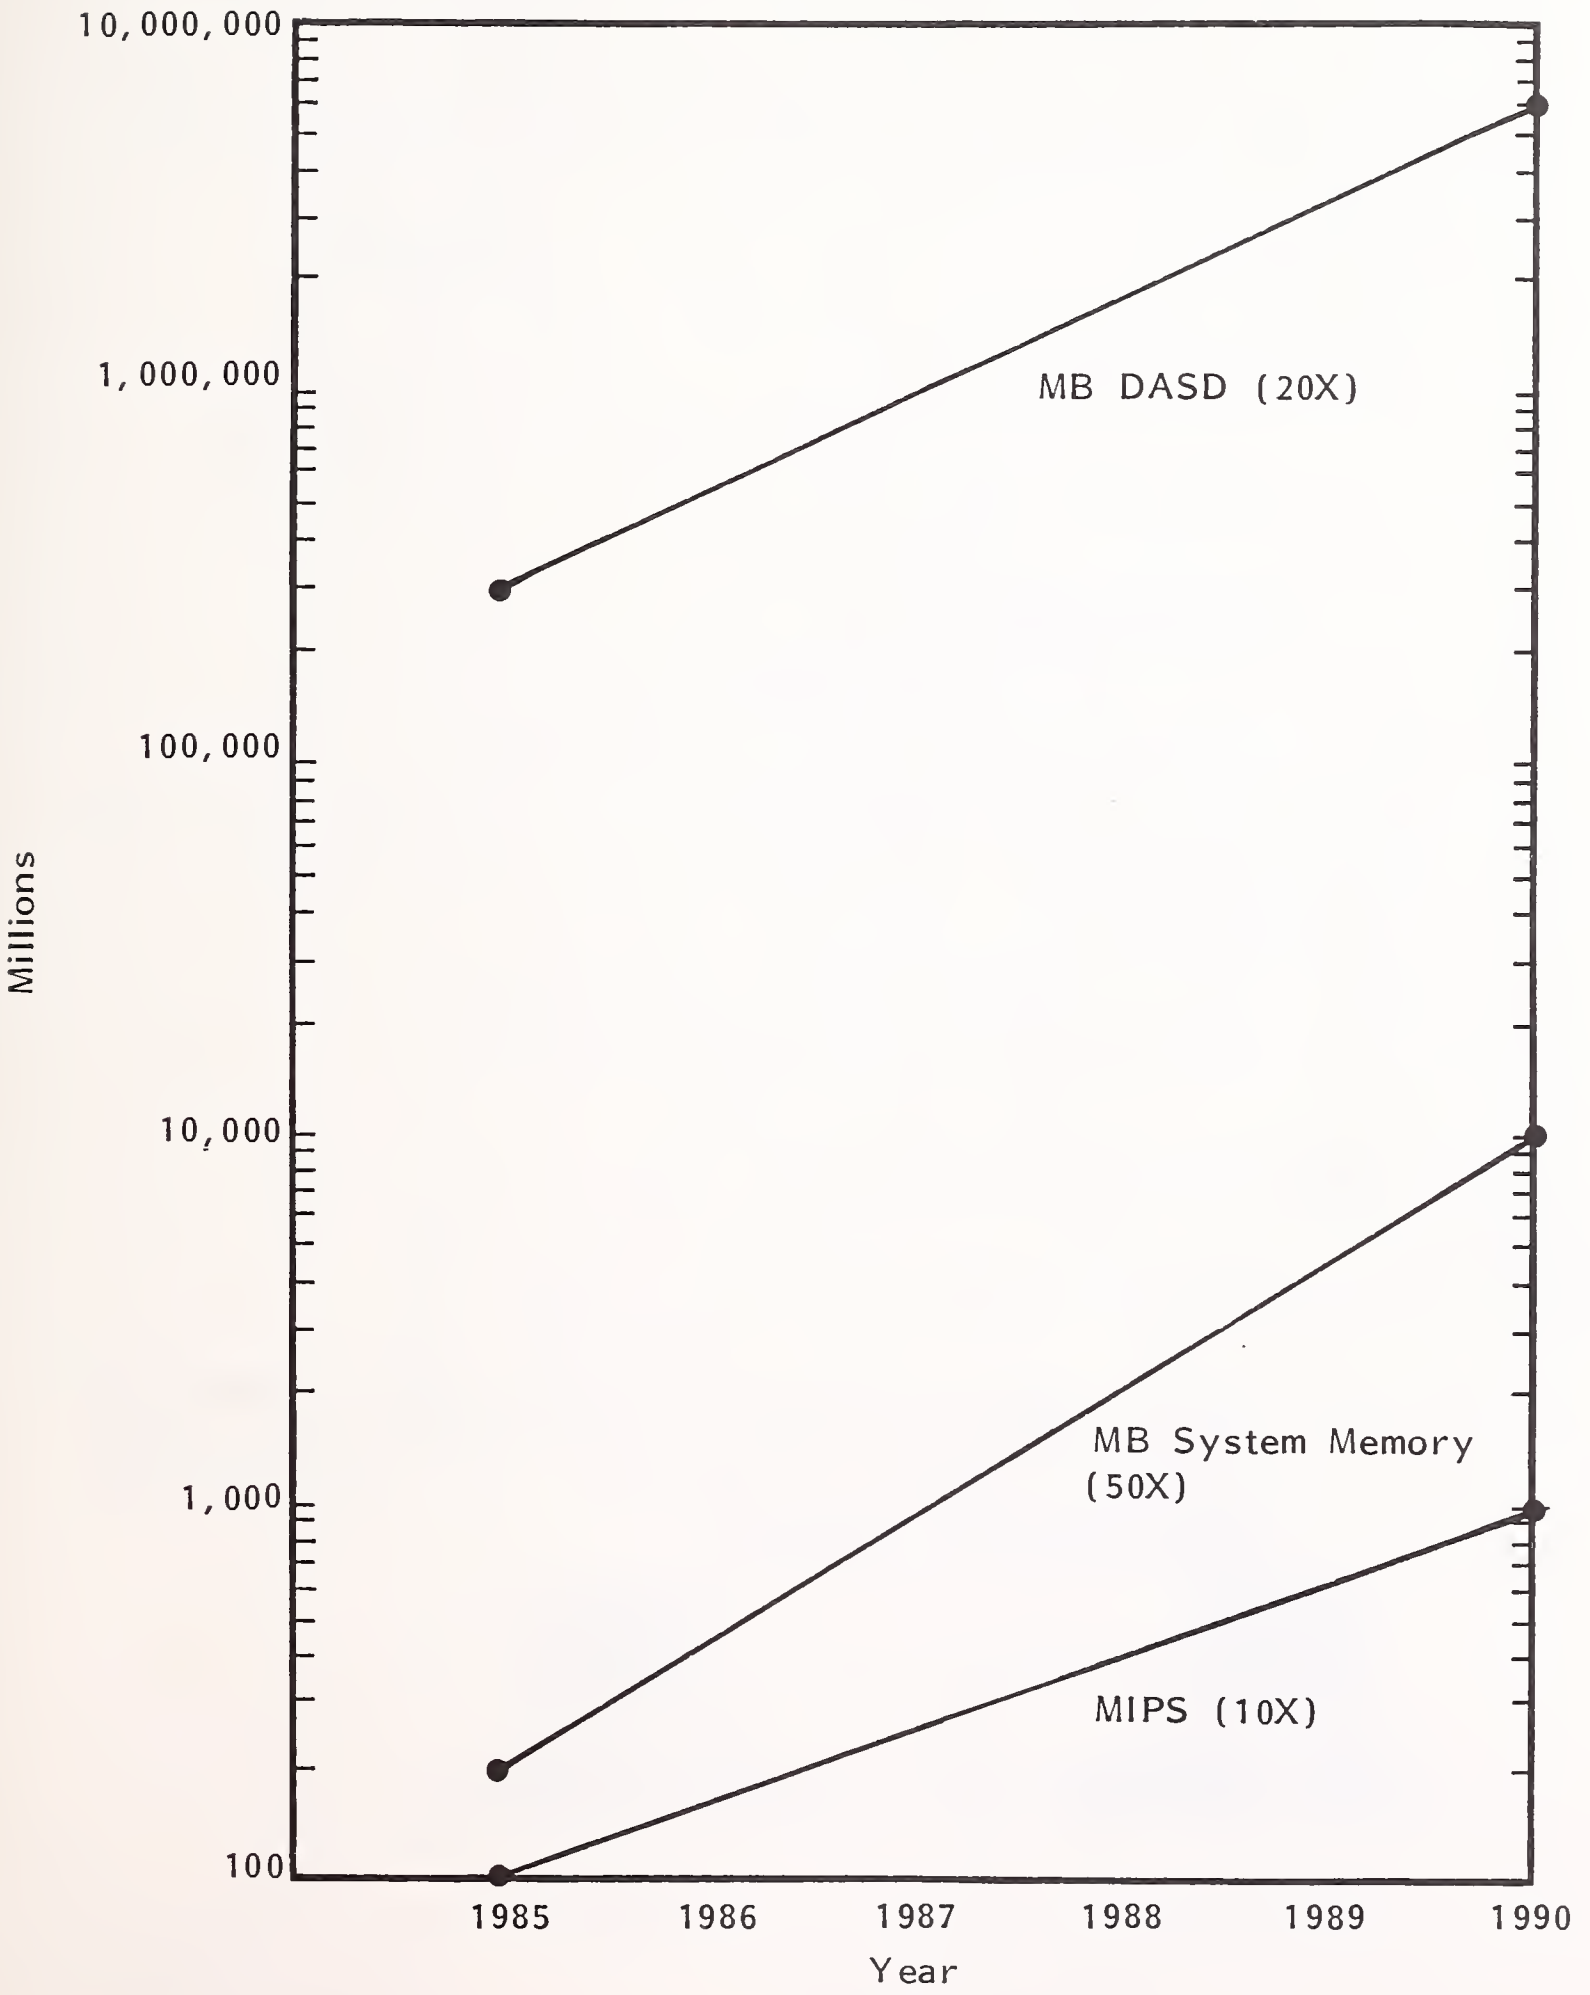

3. ANNOUNCED PRODUCTS AND IMPLICATIONS

- IBM strategy has been relatively clear since it announced DB2 for mainframes in 1983. It is the implementation of the highly centralized strategy which will be emphasized during the SNA/DDP period. From IBM's point of view, it is an ideal strategy--it even takes advantage of the batch orientation of IBM's operating systems (see Exhibit III-4). In this environment, the large host data base machine would:

  - Schedule batch requests for:

    . Extracts from files and other data bases to build DB2 tables.

    . Updates of data bases.

    . Backup of host and distributed data bases.

    . Extraction and transmission of data bases (primarily DB2 tables).

    . Heavy computations (they will be required).

  - Service interactive processing of:

    . Transactions against IMS data.

    . Queries against DB2 tables using the Query Management Facility (now shown, to be discussed under Market Analysis: Fourth Generation Languages).

  - Centralized protection and security functions such as:

    . Controlling file and data base access.

EXHIBIT III-4

LARGE HOST DATA BASE MACHINES

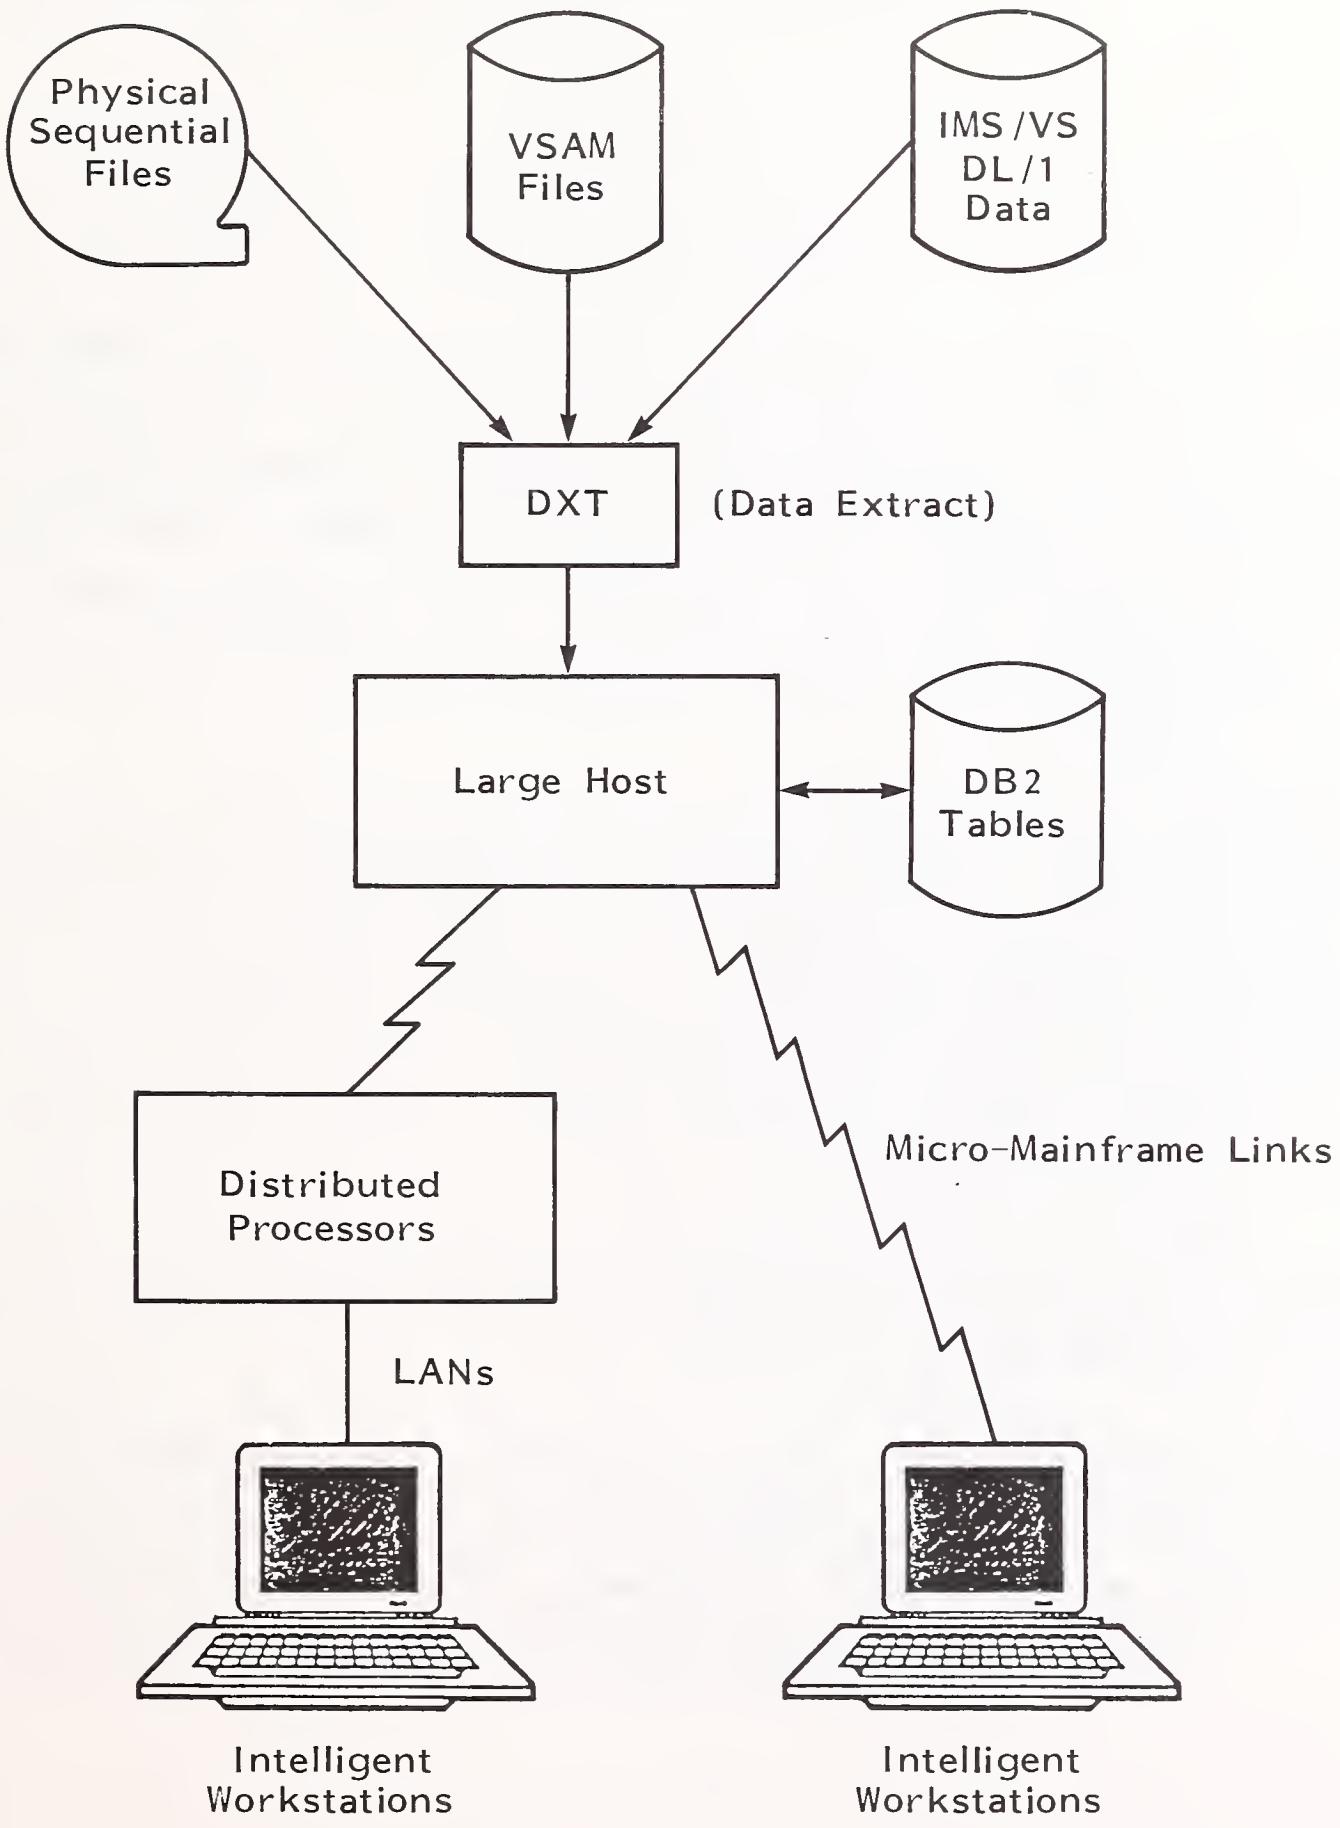

- Control of data base updates.
- Assure file and data base synchronization and integrity.
- Provide "certification" of files and data bases.
- Monitor information flow.
- Control encryption.

- The emphasis upon centralization of control and integration of existing user data bases is obvious in IBM's strategy, and the replication of data bases implicit in this strategy will drive IBM's projected demands for increased storage and MIPS. It clearly indicates the importance of IBM's DBMS strategy to its business strategy. The market is not the current estimated $600 million from DBMS software sales, but the tens of billions in increased hardware sales.

- During the SNA/DDP strategic period, IBM will not place much emphasis upon differentiation and mechanization of DBMS functions unless it becomes necessary to support the growth of large host data base machines. At present, there is general differentiation between large transaction-oriented data bases (IMS) and planning data bases (DB2), but the general approach remains toward bringing everything under large host control using slowly-evolving systems software. It is unlikely that information and knowledge bases will be differentiated from data bases during the SNA/DDP period, and any mechanization of DBMS functions will be limited.

- IBM's approach will be to emphasize some very real problems associated with rapid advances in DBMS technology (both hardware and software) much as they have with LANs. In other words, IBM will attempt to control the acceptance of new technology as they have in the past. There will be windows of

opportunity for IBM competitors who can successfully integrate cost-effective new hardware ahead of IBM's schedule and for those with imaginative solutions to unresolved data base management problems.

C. DEVELOPMENTS OF SIGNIFICANCE

- Optical memories have the cost, capacity, and performance characteristics which are necessary to effect substantial changes in both data and information base management. It is not necessary to wait for erasability to make imaginative use of this technology. As it now exists, optical disks can be extremely valuable for archival storage (magnetic tape and micrographics replacement), electronic filing (image processing), and for supplementing and complementing magnetic disk storage in a variety of ways. (The significance of optical memories has been discussed in many INPUT reports, the most comprehensive being Impact of Upcoming Optical Memory Systems, 1983.)

- Data base machines (DBM) in a variety of capacities and architectures are beginning to appear as alternatives to general purpose processors with conventional DBMS software. While impact on the DBMS market has not been great up to this point, the recognition of the central host as being fundamentally a data base machine will prompt direct comparison against DBMS. (There are already reports of DBMs handling an IBM 3084 workload at one-tenth the cost.)

- Rapid advances in microprocessor technology are facilitating improved document scanners and voice digitizers, but regardless of the efficiency of compression algorithms, image and voice information bases are not handled very effectively by current DBMS which are designed to handle encoded data. As a matter of fact, page sizes in virtual storage operating systems are not handy for image or voice storage either. At the very least, DMBS must interface relatively smoothly with these coming information bases.

- Microprocessor developments are also permitting rapid advancement in cost-effective pattern recognition of symbols/text, voice, and graphics. This will facilitate the updating of encoded data bases from sources other than conventional keyboard devices, but also represents another level away from the disciplined data entry environment assumed by most DBMS. Indeed, pattern recognition is a key element in tying together data, information, and knowledge (consider voice and handwritten annotation of conventional computer reports as part of the DBMS feedback loop), but the inherent quality problems will inevitably result in "fuzzy sets" which may have to be analyzed before entry into conventional data bases.

- All of the above developments are of importance in the DBMS market, and they represent not only opportunities but technical challenges because they tend to exacerbate some unresolved problems associated with the current data base environment.

D. UNRESOLVED PROBLEMS AND UNPLEASANT SURPRISES

- The problems IS management perceives in the current distributed systems development (DSD) environment were analyzed in detail in New Opportunities for Software Productivity Improvement, INPUT, 1984. These problems are summarized in Exhibit III-5.

  - Problems related to data base integrity, data base synchronization, and data security and protection were rated "very serious" by nearly half of the respondents.

  - Approximately 80% of the respondents felt there would be misunderstandings of data from central data bases and conflicting management reports would be generated.

EXHIBIT III-5

HOW USERS RATE DSD PROBLEMS

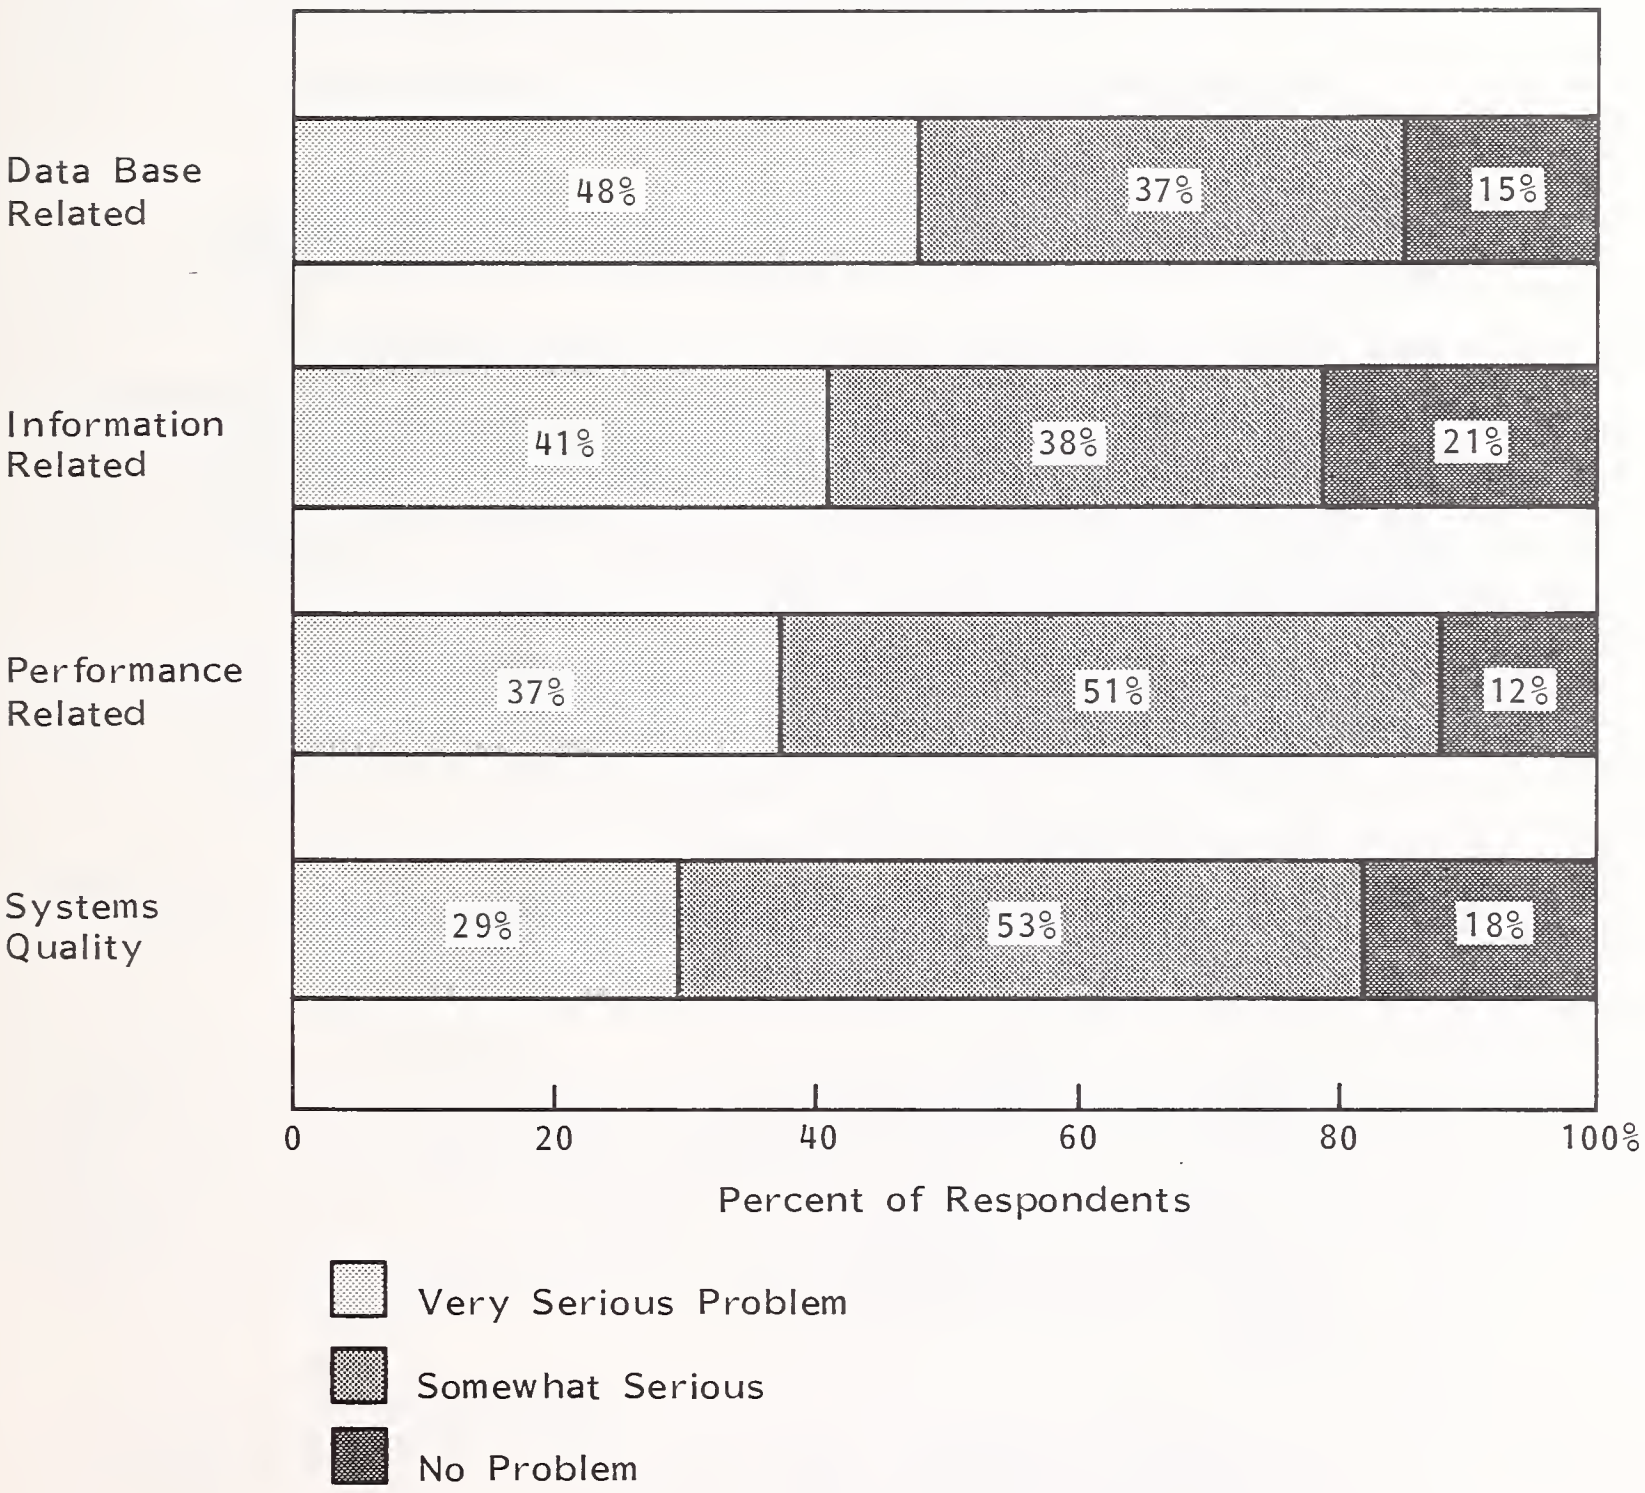

- Nearly 90% felt that performance of large mainframes would be impacted by demands for distributed data systems and expert systems.

- While over 80% felt that overall systems quality would suffer, it was rated as a very serious problem by less than 30% of respondents. INPUT's conclusion was that the impact on overall systems quality was underrated and that it would, in fact, be the most serious of all of the DSD problems.

- The basis for this conclusion was that there is a natural tendency toward chaos (entropy) associated with both data and information and the effort to maintain order (integrity) increases more rapidly than does data base size. Data and information entropy are not generally understand and can result in unanticipated performance and quality problems. Preliminary investigation by INPUT reveals:

  - That MIPS tend to increase more rapidly than DASD in very large installations.

  - Human effort associated with data base administration increases more rapidly than data base size in these installations.

  - If these preliminary conclusions are correct, and we believe they are, there will be major impacts on IBM's highly centralized DBMS strategy. For example, IBM is projecting installed DASD to increase 20 times by 1990 and installed MIPS to increase only 10 times (see Exhibit III-3).

  - Anticipating unpleasant surprises associated with IBM's strategy points to opportunities in the DBMS marketplace.

IV MARKET OPPORTUNITIES AND CHALLENGES

A. PROJECTED MARKET SIZE

- In Market Impact of New Software Productivity Techniques, INPUT forecasts for the systems software markets in the U.S. were presented from 1984 through 1989. INPUT breaks systems software down into three major areas: applications development, systems control, and data center management. The applications development area includes two major categories of products: program development and production tools, and data base management systems. The major subcategories included are:

  - Applications generators.
  - Assemblers.
  - Automatic documentation.
  - Compilers.
  - Debugging aids.
  - Languages (all generations).
  - Systems development control.

- Retrieval systems.
- Translators.
- DBMS.
- Data dictionary.
- Others.

- It was forecast that the total market for the applications development area would grow from $1.8 billion in 1983 to $10.3 billion in 1989, and the market for products associated with new productivity techniques (FGLs plus the quality control tools outlined in Market Impact of New Software Productivity Techniques) would grow from $0.75 billion in 1984 to $5.2 billion in 1989. However, INPUT has not made market forecasts of DBMS in the past and that, obviously, ignores a major segment of the applications development area.

- This report will forecast the DBMS market from 1984 through 1990, and two comparison reports (Market Analysis: Fourth Generation Languages and Market Analysis: Applications Development Tools) will provide forecasts and complete coverage of the applications development area for the first time.

- Exhibit IV-1 shows that the overall market for DBMS will grow from approximately $1 billion in 1984 to $6 billion in 1990, and IBM's share of that market will remain relatively constant at 60%. Therefore, the effective market for other competitors will grow from approximately $0.4 billion in 1984 to $2.4 billion in 1990.

- This DBMS forecast differs from several which have been reported recently in two important regards: 1) the market is growing more rapidly than most analysts have been projecting, and 2) IBM is retaining its market share (most

EXHIBIT IV-1

MARKET FORECAST DBMS
(USA 1984 - 1990)

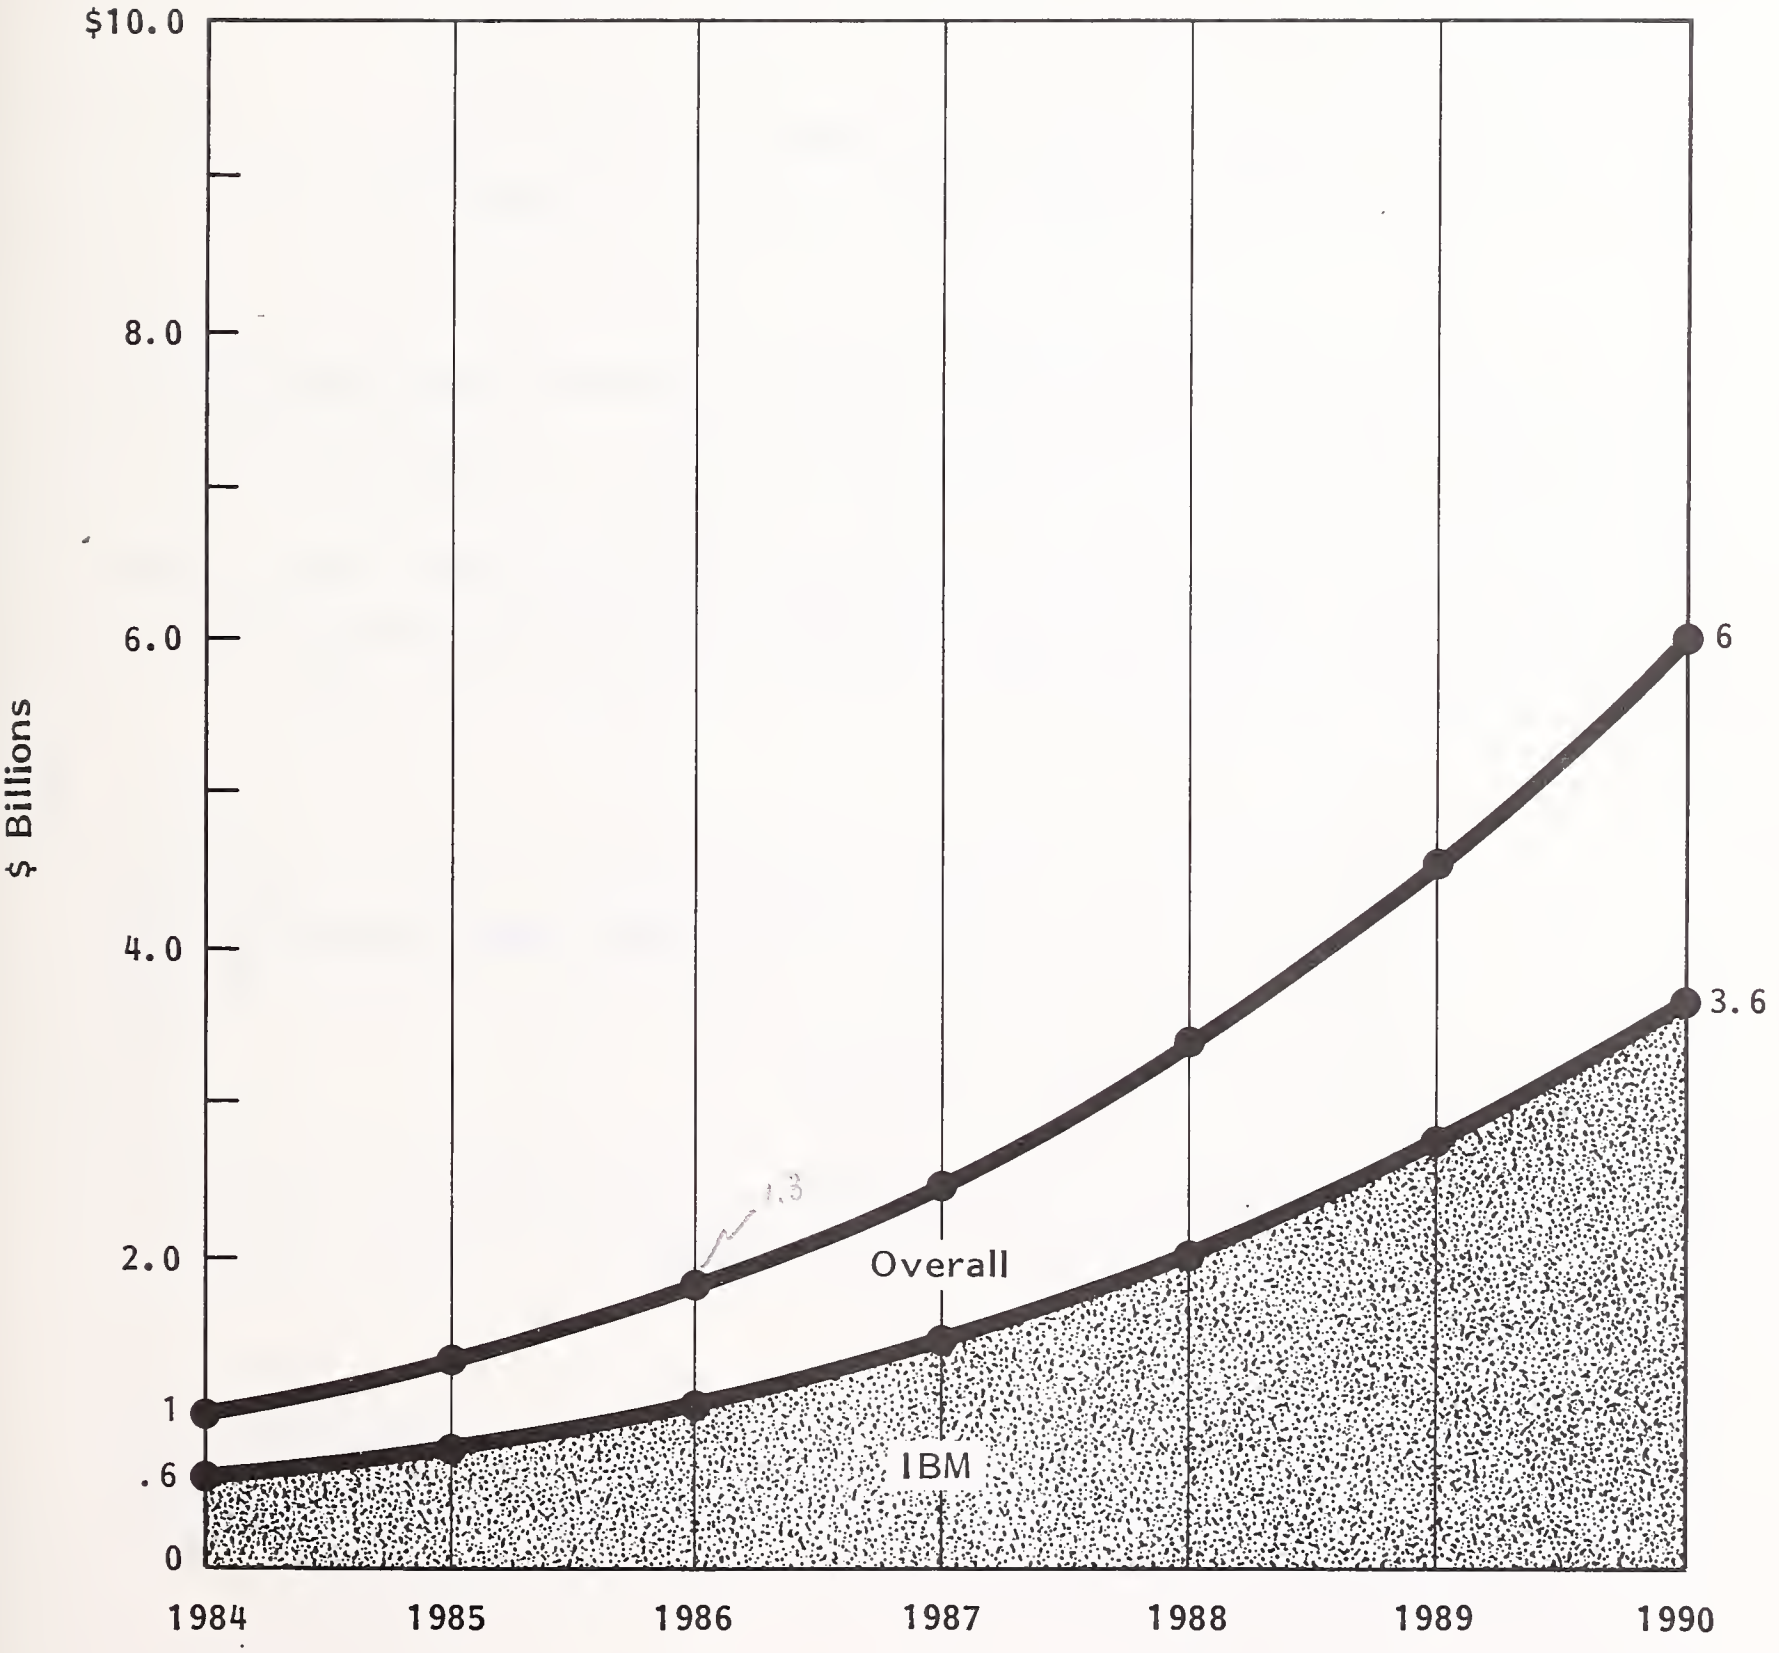

market research studies indicate IBM will lose its DBMS market share). INPUT's forecast is based on the following:

- The critical importance of DBMS in IBM's business plan for both hardware and software during the SNA/DDP strategic period (see Exhibit III-2).

- The power of IBM's highly centralized DBMS strategy to force both user and vendor conformance in many critical areas. For example:

  . Data base synchronization and integrity requires central control, and IBM can make a compelling argument that all competitive data base systems must be subordinate in terms of data flow.

  . Protection and security are severe problems, and whatever IBM solution turns out to be, both users and vendors will be hard put to ignore it.

  - IBM's continued integration of SNA, operating systems, and DBMS is assuring not only increased interdependence but also the transfer of functions among these systems. The dynamics of these changes will impact competitors comparable to IBM's transfer of functions among hardware, firmware, and software.

- An example of structural change is SNA's Logical Unit Type 6.2 (LU 6.2) or Advanced Program-to-Program Communication which is currently receiving a great deal of attention.

- LU 6.2 is the recognition that distributed processing is really a data base rather than a conventional communications problem, and the linkage is being transferred from SNA (communications) to applications (DBMS).

- It is nice to think that "peer-to-peer" communications may permit a micro to address a mainframe as a peripheral and extract data from an IMS data base, but you can be sure than IBM's implementation will assure that some peers are more equal than others.

  - Then, the question of an open systems architecture and the availability of source code is always in the background. It will become increasingly easy to justify expanding the point of interface beyond IBM's DBMS as protection and security of data weaves a complex and changing thread through SNA, operating systems, and DBMS.

- Simply put, IBM's entire software strategy during the 1980s is aimed at not only maintaining control of central data bases but extending that control. Considering the technological trends which are apparent and the tools IBM has at its disposal to establish interfaces and standards, it would be naive to believe that IBM will lose DBMS market share. In fact, unless competitors select targets of opportunity carefully, IBM could dominate the DBMS market by 1990.

B. DIRECT COMPETITION

- IBM's strategy during the SNA/DDP period emphasizes centralization on large host mainframes, and it can be anticipated that direct competition will occur for any mainframe-oriented DBMS.

- IBM has not solved the technical problems (including performance) of distributed data bases, and DBMS of higher quality have been available for some time. For the established leaders in DBMS technology, there remain opportunities to compete effectively by staying ahead of IBM. However, the cost of competing against IBM on the large host data base machines is going to increase as the software environment and interfaces become increasingly complex.

- For any new entrants in the central host DBMS market, the costs may prove to be prohibitive. It is not only conceivable but probable that establishing and maintaining interfaces to IBM systems software will exceed the development cost of the DBMS itself. In fact, even those providing lower level DBMS for minicomputers and microprocessors may find the cost of interfacing with IBM's central DBMS to be prohibitive.

- However, IBM's highly centralized approach to DBMS and its dependence upon hardware sales to achieve its growth objectives do leave significant windows of opportunity for competitive vendors who can anticipate technological trends which run ahead of IBM's SNA/DDP strategy. (Please see Market Impacts of IBM Software Strategies, INPUT, 1984, for more comprehensive analysis of both IBM's strategy and potential opportunities.)

C. INTEGRATION OPPORTUNITIES

- IBM has traditionally been slow to support new hardware advances, and this will be especially true for optical memories, which pose a significant threat to magnetic storage products upon which IBM remains dependent for revenue growth. The effective integration of optical memories into the storage hierarchy is especially attractive because it is a necessary prerequisite for the electronic office strategic period which follows (see Exhibit III-2).

- Exhibit IV-2 presents one of seven potential systems incorporating optical memories which were defined in Impact of Upcoming Optical Memory Systems, INPUT, 1983. The ramifications for data base management and the coming electronic office strategic period are obvious, and the hardware technology is currently available.

EXHIBIT IV-2

INTEGRATED IMAGE PROCESSING SYSTEM

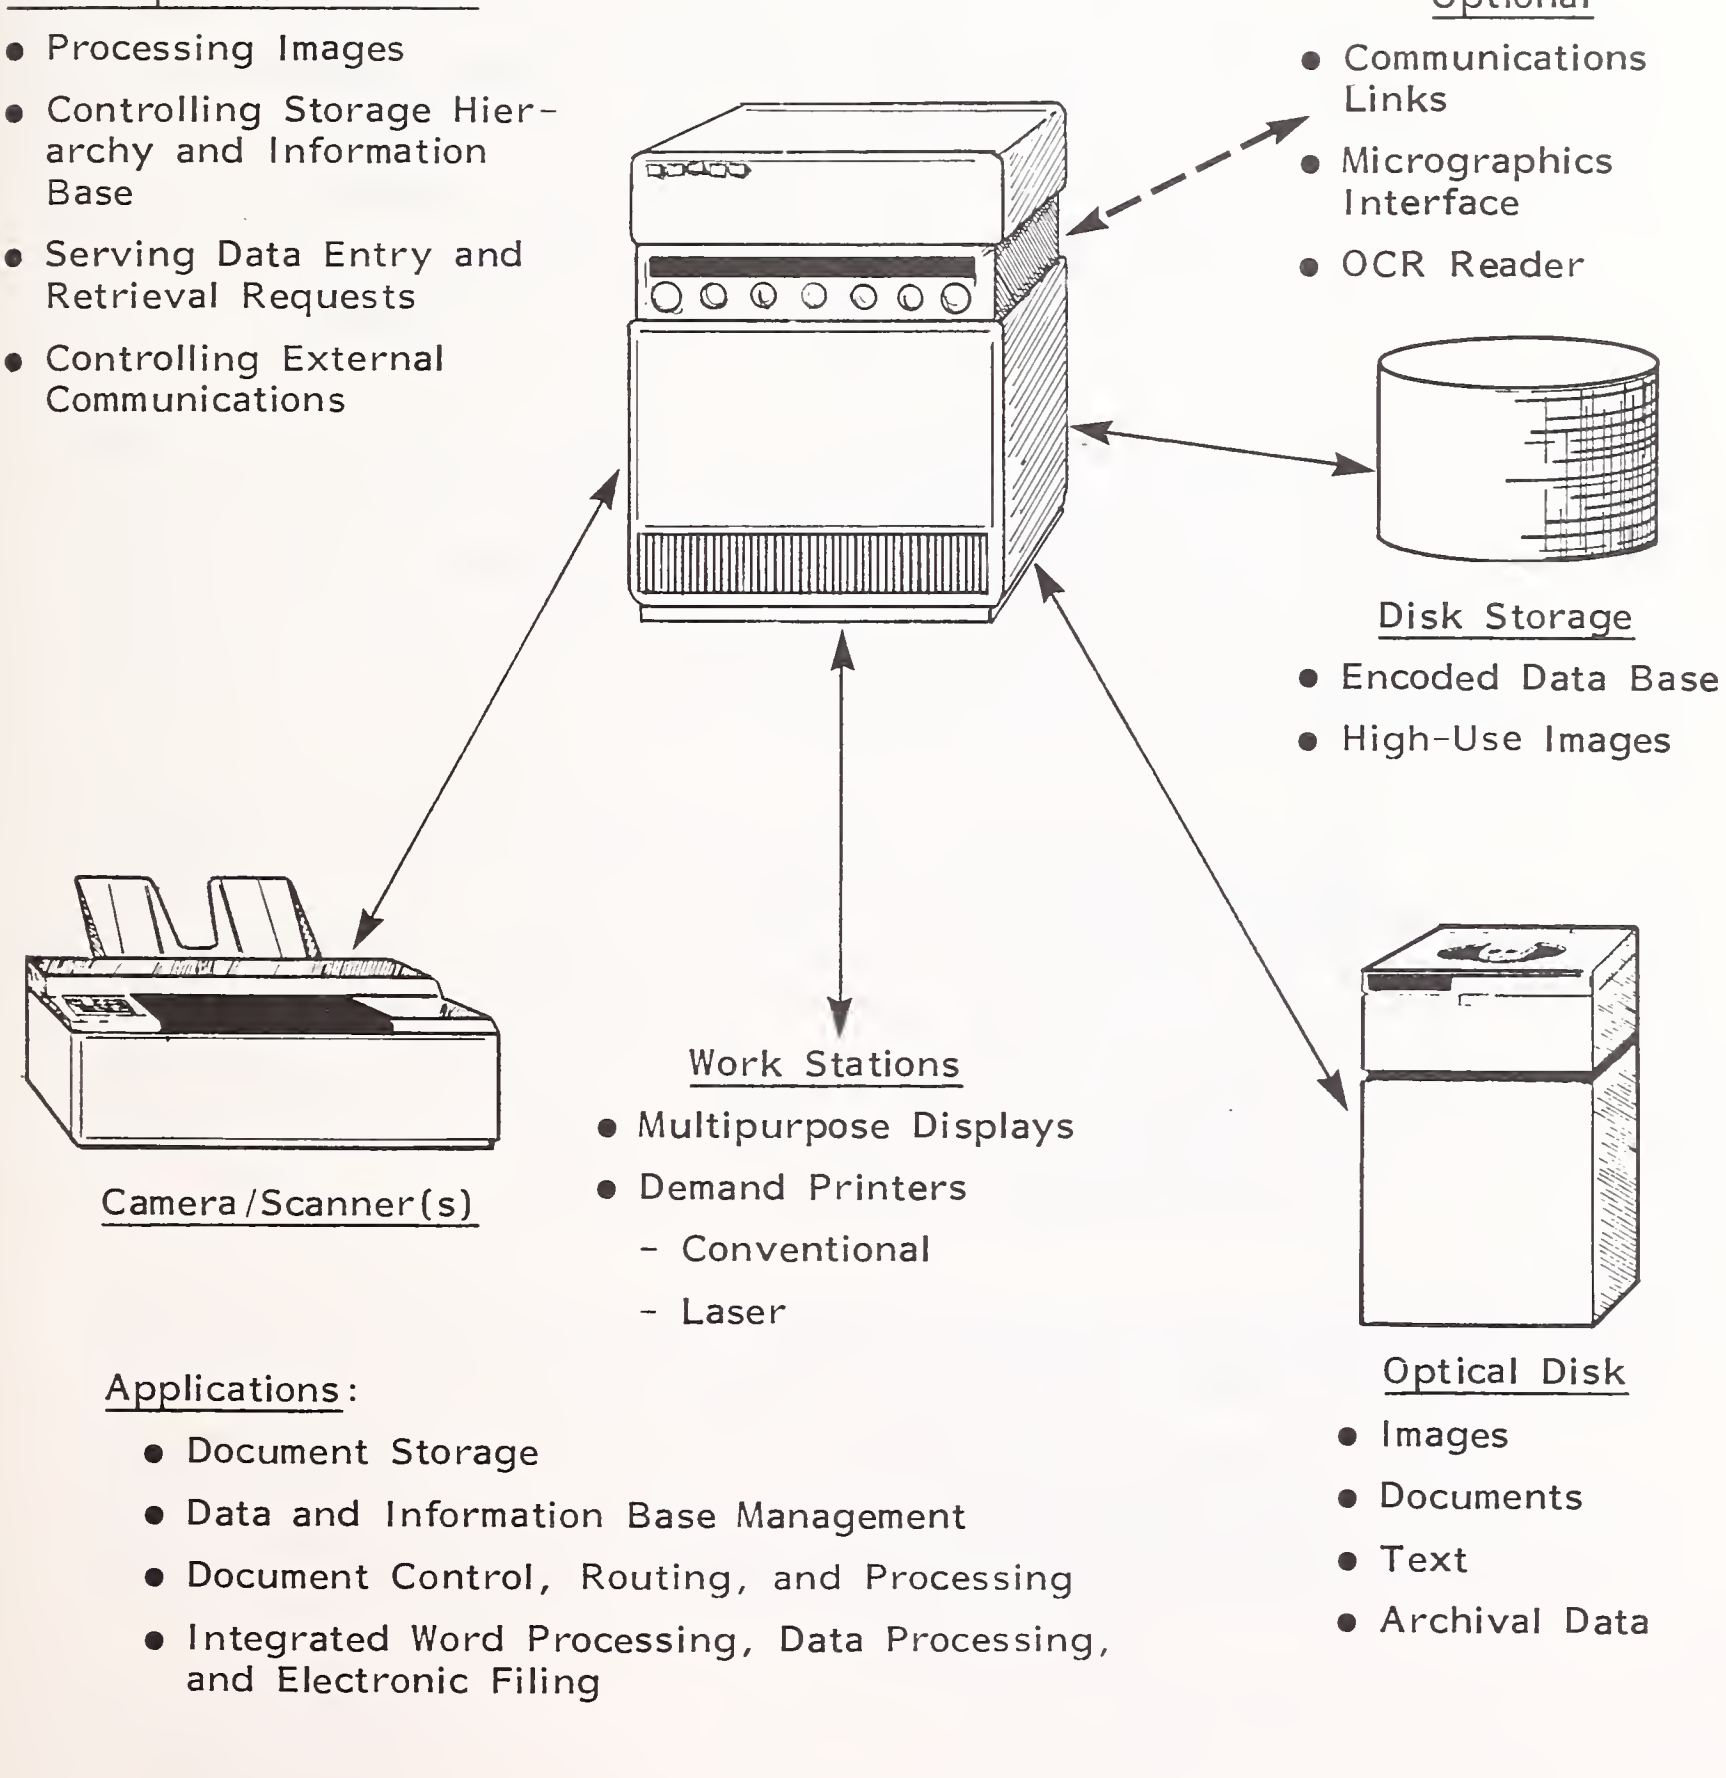

Applications:

- Document Storage
- Data and Information Base Management
- Document Control, Routing, and Processing
- Integrated Word Processing, Data Processing, and Electronic Filing

- The ability to have backup of data bases (and files) and even archival data on-line (or at least on more easily handled media) has obvious advantages in an environment where replication of data bases at various levels in the processing hierarchy seems inevitable.

- The need for integration of encoded data and information in various formats (images and text) is obvious, and it should also be apparent that current DBMS will have to be extended in terms of data/information models, structure, and access methods in order to be effective.

- Such a fundamental change in media (from paper to optical or magnetic storage) represents a substantial technical challenge, but the opportunities are more than directly proportional to the challenge.

- It should also be apparent that decision support systems, as they evolve toward expert systems, will require access to not only encoded data but text, images, and documents. Thus, the integration of data and information are a prerequisite to any significant penetration of knowledge-based systems in the commercial market.

- One of the other optical memory-based systems proposed in Impact of Upcoming Optical Memory Systems was as a "network store-and-forward reservoir" which would serve to facilitate data collection, information store and forward, and network archiving. Considering that optical storage would be used for encoded data, images, digitized voice messages, and video "messages," the impact on data, information, and knowledge-based management is apparent.

- The requirements for advanced languages to facilitate access to current and integrated data, information, and knowledge bases represent an extremely attractive opportunity, especially since IBM is not projected to place major emphasis upon such development until the electronic office strategic period. These opportunities will be analyzed in Market Analysis: Fourth Generation

Languages, which will be published shortly after this report. However, backward integration from languages into specialized data base systems represents a significant opportunity.

D. DIFFERENTIATION OPPORTUNITIES

- Even assuming the necessity to interface with central IBM DBMS, there will be abundant opportunities for product differentiation as data bases are distributed. In fact, IBM's current support of large host data bases (see Exhibit III-4) would seem to actually encourage (or at least facilitate) the development of alternative DBMS solutions.

  - By providing the data extract facility to build relational tables from IMS and DL/I data bases, VSAM, and sequential files, IBM has not only provided its "blessing" of the relational model, which is especially appropriate for lower levels in the processing hierarchy, but has provided the data format for transfer to other DBMS. Even the fact that there is currently no provision for update from DB2 to IBM's other DBMS should be viewed as an opportunity.

  - The reasons behind IBM's opening of the market for distributed DBMS are probably a combination of the following:

    . Simple customer demand for both access to central data bases and for "relational-like" compatibilities.

    . The recognition that the unresolved problems of distributed data bases should be left open for customer and competitive solution.

    . The recognition that distributed data bases are going to sell a lot of IBM hardware at various levels in the processing hierarchy regardless of whose DBMS is used.

. Reasonable confidence that there are a variety of tactics available within its general hardware/firmware/software strategy to establish control of the DBMS market as the SNA/DDP strategic period progresses.

- Regardless of IBM's motivation and/or strategy, there are ample opportunities for alternative DBMS for distribtuion of data bases on both an architectural and geographic basis (see Exhibit III-1). Essentially, the alternative implementations should provide solutions to the existing problems which were discussed earlier. There is the need for DBMS products which are differentiated by providing:

  - High performance alternatives to IBM's large host data base systems--both IMS and DB2.

  - Data and information quality assurance by solving some of the problems of data base integrity and synchronization as well as conflicting reports to management (information flow).

  - High levels of security for data bases and communications at all levels in the processing hierarchy.

- In addition, there are requirements for industry-specific data and information management systems in support of the electronic office. The elimination (or truncation) of paper document flow implicit in the integrated image processing system depicted in Exhibit IV-2 does not readily lend itself to general purpose DBMS solutions across industries. Just as the airline industry required a specialized operating system, it is a safe assumption that major industries will provide opportunities for DBMS product differentiation as office automation progresses toward the electronic office.

E. MECHANIZATION OPPORTUNITIES

- The performance problems associated with the relational model (and multiple data models) offer numerous opportunities for mechanization of data base functions. It behooves vendors of conventional DBMS to consider hardware implementations. A variety of data base machines are beginning to appear, and advances in LSI technology will make associative memories increasingly attractive. Even a simple sort/merge box (as advocated by INPUT several years ago and implemented as part of the Japanese fifth generation effort) can provide significant DBMS performance enhancement.

- As data and information management systems become integrated, there is a requirement for specialized input/output facilities to handle images and digitized audio information. Mechanizations of compression and expansion in I/O devices (scanners, digitizers, and displays) and controllers will become highly desirable, and the transfer of the resulting information to and from the information base requires attention. Specifically, variable length records are back with us and they do not necessarily lend themselves to the page sizes in IBM operating systems.

- With the availability of optical memories, automatic archiving and backup of data bases is becoming possible and economically justified at various levels in the processing hierarchy. Indeed, it may be that this is the simplest and most practical DBMS application of the new technology.

- In addition, expert systems domains will require highly specialized data/information/knowledge bases along with new hardware (LISP or other development engines), and knowledge bases will be the most sensitive information in most enterprises. The mechanization of the knowledge-based systems for purposes of development, maintenance, and security should be independent of the data and information bases upon which the expert systems may depend, and probably from other expert systems domains. In other words, there is a

need for highly specialized knowledge-based management systems to support knowledge engineers in the development of expert systems.

F. DATA/INFORMATION/KNOWLEDGE SERVICES

- It is INPUT's belief that the opportunities for data, information, and knowledge-based integration, differentiation, and mechanization described above create an opportunity for providing centralized data, information, and knowledge services. In fact, it is our opinion that many knowledge-based systems will not be economically justified without the availability of centrally maintained data and information services.

- Therefore, it is anticipated that the owners of proprietary data and information bases will increasingly provide hardware/software delivery systems to support their data and information management services. And, the developers of the hardware/software data/information/knowledge-based management systems described above will see fit to provide services. In addition to data and information itself, such services could include:

  - Backup of major data, information, and knowledge bases for purposes of both archival storage and disasters.

  - The provision of certified, secured storage and delivery networks for highly sensitive information.

- The provision of centralized data and information services will become closely integrated with knowledge bases as expert systems become an extension of decision support systems. Early indications are that expert systems in even relatively narrow domains require extensive maintenance by scarce (and ill-defined) "knowledge engineers." It is probable that many expert sytems will be impractical to maintain except on an interorganizational basis. In other words, there will be opportunities for major service offerings in carefully selected knowledge-based systems.

V CONCLUSIONS AND RECOMMENDATIONS

A. CONCLUSIONS

- IBM's strategy during the SNA/DDP period (1985-1990) will continue to emphasize highly centralized control of data bases and information flow. Considering the unsolved technical problems of distributed data bases and information flow, this "go slow" strategy can now be justified on both a technical and business basis by IBM.

- It is INPUT's opinion that IBM intends to integrate SNA, operating systems, and DBMS so closely during this period that the point of interface for most other vendors will become IBM's DBMS. Considerations of data base integrity, synchronization, and security will force many customers to wait for IBM's solution. The current flurry of speculation concerning LU 6.2 and peer-to-peer communications sometimes ignores the fact that the problems of distributed data bases remain (and are even compounded). It is inevitable that IBM's implementation of LU 6.2 will emphasize centralization of data base control (the definition of centralization in general systems theory is the development of a "leading part"), and it is probably that the control will be exercised by IBM's DBMS.

- While software remains the key to IBM's strategy of account and technological control, IBM's business plan is dependent upon hardware sales. During the SNA/DDP period, the key element in IBM hardware sales remains large-scale

mainframes and magnetic disk storage, and IBM's DBMS strategy is designed to increase such hardware sales.

- To the degree that hardware/software technological developments seem contrary to IBM's clearly indicated strategy of using large-scale mainframes as data base machines, there will be opportunities for competitive vendors. Specifically, INPUT believes the following technological devleopments represent threats to IBM's SNA/DDP strategy:

  - Optical memories offer new solutions to many data and information handling and storage problems. They will be the key to office automation and the electronic office. IBM's revenue growth is heavily dependent upon the sale of magnetic storage devices, and IBM will be slow in taking advantage of the technology.

  - Back-end data base machines which offload (or replace) mainframes are both desirable and necessary to improve performance of DBMS, particularly those employing the relational model. The architecture of IBM mainframes is not conducive to the handling of the emerging DBMS environment. In fact, INPUT believes that unanticipated performance problems will arise in IBM's highly centralized DBMS environment and force the distribution of processing (both architectural and geographic).

  - In fact, a good case can be made that IBM mainframe hardware/software architecture is not especially suitable for the emerging information- and knowledge-based environments where processing of images and variable length messages becomes combined with complex symbolic (or list) processing.

  - In addition, many of the problems associated with distributed data bases and information flow will prove to be especially difficult to solve in the complex IBM hardware/software environment which has evolved

over the last 20 years. This means that initial solutions to these problems will probably be easier for those who have more tactical flexibility than exists within the IBM hardware/software strategy.

- It is INPUT's opinion that numerous opportunities exist for competitors which are supplementary and complementary with IBM's DBMS strategy and will not necessarily be in total conflict with IBM's ultimate objectives. It is also our conclusion that major conflict with IBM in its DBMS strategy is especially ill-advised.

B. RECOMMENDATIONS

- Vendors must anticipate and understand IBM's strategy rather than react to specific tactices (IBM hardware/software announcements or lack thereof). Market Impacts of IBM Software Strategies provides a general framework for such anticipation and understanding, and we encourage you to use it in placing the DBMS market in proper perspective.

- Unless vendors are already established in the mainframe DBMS market, INPUT does not encourage entry (and direct confrontation with IBM). The cost of interfacing with IBM hardware and software systems is already high and is going to increase because of IBM's DBMS emphasis during the SNA/DDP period. If you are already established, be prepared to compete against IBM by:

  - Emphasizing improved performance and offloading of mainframe processing burden. This implies either the support back-end data base machines or the distribution of processing burden (and data bases) to minicomputers (departmental processors) and intelligent workstations.

- Understanding and solving the problems associated with large central data bases, distributed data bases, and information flow as defined in this report. INPUT research has revealed that users feel vendors do not understand these problems and are not being responsive to their requirements. There is a clear need to emphasize quality and control of data and information in the distributed environment.

- Security, protection, and privacy of data and information at all levels in the processing hierarchy are of increasing importance. They must be addressed soon or there is the potential for severe impact on the acceptance of existing products. All DBMS vendors must be responsive to these problems, and there will be a substantial market for those who can provide leadership in solving them.

- Understand the difference between data, information, and knowledge and be prepared for their integration. A major window of opportunity exists for such integration prior to IBM's electronic office period (1990-1995), and there are specific things which can be done.

  - Take advantage of optical memories and extend DBMS to include image- and voice-based management.

  - Recognize that structures of data, information, and knowledge are not compatible, and develop different models for storage and handling--including retrieval.

  - Concentrate on both facilitating and controlling the flow among data, information, and knowledge. For example:

    . Enhanced pattern or voice recognition can be used to faciitate the updating of encoded data bases from information.

  - Flexible archiving facilities would permit either documents (images) or supporting data and programs to be stored and retrieved.

  - As data are accumulated, they can eventually change knowledge bases in unanticipated ways. Information concerning these changes must be generated in order to evaluate the impact on existing knowledge-based systems. (In other words, tools must be provided to automatically report significant changes in DBMS content, quality, etc.)

- In addition, various industries have particular paper-based systems and procedures which will require specialized data, information, and knowledge-based systems. There will be increased opportunities for the development of such systems for vendors who elect to invest the necessary resources.

- Then, of course, there are substantial opportunities for the development of imaginative tools and aids to facilitate the use of data, information, and knowledge bases. The integration of such tools and aids will be the key factor in the marketability of advanced data, information, and knowledge-based systems. The markets for languages and applications development tools will be analyzed in two follow-on reports--Market Analysis: Fourth Generation Languages and Market Analysis Applications Development Tools.

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

Detroit
220 East Huron
Suite 209
Ann Arbor, MI 48104
(313) 971-0667

New York
Park 80 Plaza West-1
Saddle Brook, NJ 07662
(201) 368-9471
Telex 134630

Washington, D.C.
11820 Parklawn Drive
Suite 201
Rockville, MD 20852
(301) 231-7350

EUROPE

United Kingdom
INPUT, Ltd.
41 Dover Street
London W1X 3RB
England
01-493-9335
Telex 27113

France
La Nacelle
Procedure d'abonnement 1-74
2, rue Campagne Premiere
75014 Paris
France
322.56.46
Telex 220064 X5533

Italy
Nomos Sistema SRL
20127 Milano
Via Soperga 36
Italy
Milan 284-2850
Telex 310352

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

Singapore
Cyberware Consultants (PTE) Ltd.
2902 Pangkor
Ardmore Park
Singapore 1025
734-8142

INPUT®
Planning Services For Management

INPUT®
1943 Landings Drive, Mountain View, CA 94043, (415) 960-3990