Proceedings

DISTRIBUTED COMPUTING

COMPCON '80-Fall
21st IEEE Computer Society International Conference



SEPTEMBER 23-25, 1980 TWENTY-FIRST IEEE COMPUTER SOCIETY INTERNATIONAL CONFERENCE CAPITAL HILTON HOTEL, WASHINGTON, D.C.

IEEE Catalog No. 80CH1598-2C Library of Congress No. 80-83217



THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Additional copies available from:

IEEE Computer Society 10662 Los Vaqueros Circle Los Alamitos, CA 90720

IEEE Service Center 445 Hoes Lane Piscataway, NJ 08854

Copyright © 1980 by The Institute of Electrical and Electronics Engineers, Inc., New York, N.Y., Printed in USA

IEEE Catalog No. 80CH1598-2 Library of Congress No. 80-83217

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy articles in this volume beyond the limits of U.S. copyright law without fee if for private, noncommerical use of patrons. Instructors are permitted to photocopy isolated articles for noncommerical classroom use without fee. For other copying, reprint or republication permission, write to Director, Publishing Services, IEEE, 345 E. 47 St., New York, NY 10017. All rights reserved. Copyright © 1980 by the Institute of Electrical and Electronics Engineers, Inc.

GENERAL CHAIRPERSON'S MESSAGE



It has been a pleasure to head the team which has brought you Compcon 80 Fall. I think this has turned out to be an outstanding conference, with thirty nine technical sessions plus the plenary session and evening panel debate. I would expect these proceedings to be well thumbed in the years to come.

I must tell you that everything went smoothly, even through changes of employment and conference committee membership. Everyone carried his/her share of the load, but no one was overtaxed. Based on previous years' experience, we established a realistic schedule and followed it. Problems arose and were dealt with. My point is to suggest that if you've never participated in Computer Society committee work, you might consider working on a conference.

I'm especially pleased with the program balance. There is good treatment of each of several levels of involvement with distributed processing: casual interest, beginning involvement, and highly experienced (two years, perhaps?). The pre-conference tutorials continue this multi-level presentation, including entry level as well as specialized presentations. My only regret is that there are too many simultaneous sessions I want to attend—and I haven't felt that way about a conference in years.

Of course this was a team effort. I salute the conference and program committees. Special tribute is due David Nelson, program chair, who organized an outstanding group of people, who in turn produced this outstanding program.

The subject of this year's conference, Distributed Computing, has been selected because of its emerging importance. It is sufficiently broad to encompass most of computer-land's specialities, while at the same time being sufficiently focused to promise specific benefit to attendees.

I trust that interactions at the conference will be as much a benefit as these proceedings. Your suggestions for, and participation in, future conferences is invited.

Marshall D. Abrams General Chair COMPCON FALL '80

PREFACE



Man's desire to interconnect—railroads, highways, telephones, and now computer systems—is resulting in significant advances in distributed computer systems. These advances are as significant for the 1980s as timesharing was to the 1970s and batch processing was to the 1960s. Just as batch processing yielded to timesharing, so will the latter yield, in many respects, to distributed designs, as their advantages become understood and verified through experience.

We are entering a new phase in distributed processing. We finally have actual experience with several systems, and one gets a feeling similar to when a computer program becomes too large and complex—a desire for simpler solutions, a desire to discard unworthy ideas, and a desire to reimplement the good ones. Through all of this, we see the emergence of ideas which are more elegant, integrative, and in better compliance with both application and technological requirements.

This year's conference reaches a very broad set of topics, all of which directly bear on the design of distributed systems. The Program Committee is equally broad with respect to member's interests, affiliations, and backgrounds. We hope you find the program valuable and interesting.

David L. Nelson Program Chair COMPCON FALL '80

TABLE OF CONTENTS

SESSION 1: MANAGEMENT ISSUES Key Decisions in Distributed Processing	3	Human Cognitive Factors in Front-End Interaction A.N. Badre	119
H. Lorin		Human Factors Experiments for Refining Interactive	
The Economics of Distributed Systems D.J. Sykes	8	System Designs	123
Distributed Systems Management—A Key to Success L.W. Smith	16	SESSION 7: SHORT NOTES	
SESSION 2: USER INTERFACE TO DISTRIBUTED PROCESSING SYSTEMS		SESSION 8: USER-ORIENTED DESIGN APPROACHES TO HUMAN COMMUNICATION SYSTEMS	
Command Interaction in Distributed Systems K.A. Lantz	25	Towards a Methodology for User-Driven Design of Electronic Office Systems	
Considerations for a Logical Virtual Terminal		D. Tapscott, M. Greenberg, D. Henderson and	danie i
L. Greninger and R. Roberts	33	M. Collins An Experiment in the Design of a Public Message	
SESSION 3: NETWORK ARCHITECTURES		System	137
A Distributed Virtual Machine on a Tightly Coupled		N. Naffah	Hosel
Multiprocessor	43	Domesticating the Computer Terminal	
A Distributed File System in EPOS	50	SESSION 9: FUTURE DIRECTIONS OF UNIX— A PANEL	
Parallel Processing Simulator for Network Systems		SESSION 10: COMPUTER NETWORKS	
Using Multi-Microcomputer	55	A Computer Aided Logic Design System for Custom	
H. Takenouchi, M. Hatada, K. Hiyama, H. Inamori and A. Toda		LSI's S. Nakamura, M. Terai, T. Ogihara, C. Tanaka and	151
SESSION 4: COMPUTER-BASED MESSAGE	Distrib	S. Murai	E off
SYSTEMS: IMPLEMENTATIONS		Computer Network Coupled by 100 MBPS Optical Fiber	
Solving Communication Problems—The Case for Electronic Mail	188,18	Ring Bus: System Planning and Ring Subsystem Description	150
J.B. Holden	65	K. Ikeda, T. Fujima, Y. Ebihara, T. Nakamura,	159
Office Automation at TI	69	M. Ishizaka and K. Nakayama	
Future Message System Design: Lessons from the		SESSION 11: TEXT PROCESSING LANGUAGE STANDARDS—A PANEL	
T.H. Myer	76	Computer Languages for the Processing of Text: A	ieas
SESSION 5: NETWORK ARCHITECTURE FOR DISTRIBUTED COMPUTERS	mmoJ mmoJ *	Status Report by a Panel of Experts from the Standards Development Project	169
An Assessment of the Status of Network			
Architectures	87 "	SESSION 12: LOCAL COMPUTER NETWORKS I A Fail-Soft Distributed Processing System	179
Distributed Machines	95	D.J. Paulish A Local Network for Distributed Laboratory	
Backup and Recovery for Distributed Interactive		Microcomputers	185
Computer Systems	101	J.I. Powell, R. Fico, W.H. Jennings, E.R. O'Bryan and A.R. Schultz, Jr.	
Network Management in DECNET	108	A Local Distributed Microprocessor Net with Decentralized Access to an Optical Bus	191
SESSION 6: HUMAN FACTORS IN DISTRIBUTED PROCESSING SYSTEMS—A PANEL	le autor	A.M. Sauer and H.G. Schwartzel	
The Application of Human Factors to Distributed		SESSION 13: DISTRIBUTED DATA DESIGN METHODOLOGY	
Systems	117	A Design Methodology for Distributed Data Bases	100
M.L. Schneider		P.S. Fisher, P. Hollist and J. Slonim	199

System Design Selection for Distributed Data		Estimation of Intrinsic Parallelism in High Level	
Systems	203	Programs Using Data Flow Execution	313
A User Demand Model for Distributed Database		Abstract Dataflow Protocol for Communication in	
Design J.L. Dawson	211	V.P. Srini and B.D. Shriver	321
SESSION 14: VLSI CHIPS		SESSION 19: RESEARCH DIRECTIONS	
The TMS 9918 Video Display Processor: A Brief		An Overview of Research Directions in Distributed	
Overview K. Guttag	219	Processing	333
The 16-Bit Family Roadmap Shows the Way J.F. Stockton	224	Centralized Operation and Control of a Remote Computer Using a Microprocessor	341
Logic-Enhanced Memories: An Extensible VLSI		A.A. Desrochers, M. Heydari and J.B. Selep	
Architecture (Increased Performance Without Increased Complexity)	228	SESSION 20: SWITCHING TECHNIQUES On the Building of a Packet Switched Network of	
SESSION 15: WIRELESS NETWORKS DEVELOPMENTS		Microcomputers	351
Distributed Routing and Relay Management in Mobile		Packet Switching Using Concurrent Pascal in a Network	
Packet Radio Networks	235	A.M. van Tilborg and L.D. Wittie	
Wireless Terminal Communications Using Spread-		Auditing and Message Reconciliation in a Datagram	
Spectrum Radio	244	Store-and-Forward Packet Switching Network K. Talati and C.G. Gray	366
Performance of Some Satellite Communication		- SESSION 21: DISTRIBUTED DATA BASE	
Protocols	249	SOFTWEAR ARCHITECTURES Alternative Architectures for Distributed Data Sharing:	
Minimization of Response Time by Optimum Assignment of Packet Repetition Rates in a Broadcast		Functional Issues	
Environment		Data Architectures for Distributed Data Bases J.A. Larson and T.B. Wilson	378
SESSION 16: LOCAL COMPUTER NETWORKS II		The Multiple-Schema Architecture of DSEED: A	
The Extended Ethernet: EE-Net	261	Distributed CODASYL Prototype System F. Germano, Jr.	. 383
Functional and Logical Description of a New Fiber-		SESSION 22: PROCESSOR ARCHITECTURES	
Optic Contention Bus Network. S.A. Kahn, R.L. Stewart, S.G. Tolchin and S.J. Healy	268	Single Chip Pascal Processor: Its Architecture and Performance Evaluation	395
A Survey of Local Network Hardware	273	K. Tanabe and M. Yamamoto	400
K.J. Thurber SESSION 17: SCHEMA-MERGING RECOM-		A 3 Microwatt C-MOS 4 Bit Microcomputer	400
POSITION PROCEDURES		Communication Control System for Distributed Data	
Automatically Merging Databases	279	Processing Services	406
Design of Independent Database Schemes for Distributed		Integrated Disk Cache System with File Adaptive	
Processing Using Multivalued Dependencies	287	T. Tokunaga, Y. Hirai, and S. Yamamoto	412
A Normal Form Based on Theta-Join and Projection for Relational Data Bases	295	SESSION 23: DTC/DCE INTERFACES—A PANEL	
Towell, R. Leas, W. M. Johnson, E. R. O. Bryan.		SESSION 24: COMMUNICATIONS SYSTEMS	
SESSION 18: DATAFLOW AND CENTRAL CONCEPTS FOR DISTRIBUTED PROCESSING		Channel Extension at Very Long Distances D.J. Schibonski and J.E. Thornton	421
	301	Resource Management Aspects of Ringnet: A De-	
FP and Its Use as a Command Language L.A. Stabile	301	centralized Local Network	424
A Model for Decentralized Control in a Fully Distributed Processing System	307	ENCOMPASS: Tandem's High Performance Relational Data Base Management System	428

9	SESSION 25: SYNCHRONIZATION, RECOVERY,		0	SESSION 31: ISO/ANSI REFERENCE MODEL	
	AND RELIABILITY ISSUES IN THE DISTRIBUTED DATA BASE ENVIRONMENT			Overview and Status of the ISO/ANSI Reference Model of Open Systems Interconnection	. 55
	Overview of Synchronization and Recovery Problems in			R. desJardins	
	Distributed Databases	433		Open Systems Interconnection Upper Layers Activity L.L. Hollis	
	Reliability Issues for Completely Replicated Distributed		1	Middle Layers of Open Systems Interconnection: Session	1
	Databases	442		J.P. McGovern and D. Basu	. 56
	Analysis of Update Synchronization Schemes in Distributed Databases	450	t	Lower Layers of Open Systems Interconnection Archi- tecture: Network, Data Link, and Physical Layers G.W. White	. 57.
	SESSION 26: PERFORMANCE EVALUATION		, 5	SESSION 32: DISTRIBUTED OPERATING	
	An Integrated Design System for Distributed Database			SYSTEMS	
	A.R. Hevner and G.M. Schneider	459	I	The Series/1 Distributed Operating System: Description and Comments	
	On Fairness in Packet-Switching Networks			W.D. Sincoskie and D.J. Farber	579
	Modeling and Analysis of Single and Multiple-Bus Multi-Microcomputer Systems	471		Experiences With the Series/1 Distributed System	585
	D.A. Protopapas and E.J. Smith	4/1	A	an Interprocess Communication Facility for Distributed	
	SESSION 27: A CRITICAL ASSESSMENT OF SOFTWARE COST TRENDS IN		A	D.E. Britton and M.E. Stickel	590
	TELECOMMUNICATIONS—A PANEL		S	ESSION 33: MULTICOMPUTER	
	SESSION 28: NETWORK OPERATING SYSTEMS			ARCHITECTURE II	
	Implementation of a Dial-Up Network of UNIX TM		0	n a Distributed-Processor Communication	
	D.A. Nowitz and M.E. Lesk	483	A	rchitecture	599
-	SDS/NET—An Interactive Distributed Operating		G	eneralized Multistage Interconnection Network	
	C.J. Antonelli, L.S. Hamilton, P.M. Lu, J. J. Wallace, and K. Yueh	487		GEMINE)	606
	GCOS 8: A Distributed Processing Model	40.4	A	Novel Distributed Processing Scheme for Computer	
	G. Colliat	494	A	nalysis of Motion D.P. Agrawal and R. Jain	614
	SESSION 29: MULTICOMPUTER ARCHITECTURE I		SI	ESSION 34: APPLICATIONS AND	
I	Distributed Control of Ring Networks Using A "Play			IMPLEMENTATION	
7	Through" Protocol	507	PI No	an-Net: A Secure, Fault-Tolerant Packet Switching	
F	Functional Distribution Via Control Sharing	516		C.G. Gray and K. Talati	625
A	A Functionally Distributed Architecture for		A	Large Distributed Processing System for Real-Time viation Applications	600
(Communications Processors	522		P.D. Bergstrom and T.B. Fowler	629
	H. Oku, T. Matsuyama, F. Shimomura, and M. Yamashita		A	Case Study of a Distributed System Design L.E. Deimel, Jr., R.J. Fornaro, D.F. McAllister and	636
S	ESSION 30: SPECIFYING SYSTEM COMPOSITION		Di	T. Sanford stributed Computing in a University Environment	740
0	On the Decomposition of Distributed Systems into			R.T. Thomas	643
N	A. Silberschatz	531	SE	SSION 35: ANALYSIS OF DISTRIBUTED	
A	Protocol for Distributed Execution and Consistent			SYSTEMS	
	esource Allocation	535	AS	Study of Load Levelling	647
A	T. Christopher, O. El-Dessouki, M. Evens,	543	Ass	signment of Tasks and Resources for Distributed occasing	655
	W. Kabat, and S. Wagle			S.B. Wu and M.T. Liu	000

SESSION 36: DISTRIBUTED OPERATING SYSTEM KERNELS		Integration Testing of Distributed Software T.S. Chow	706
An Operating System Kernel for a Hierarchical Multicomputer	665	A Table-Driven Data Validator	712
D. Bhatt, P. Sadayappan, R.B. Kieburtz and D.R. Smith		SESSION 39: PROGRAMMING LANGUAGES FOR DISTRIBUTED PROCESSING	
A Kernel for Supporting a Distributed File System J. Mukerji and R.B. Kieburtz	673	A Survey Note on Programming Languages for Distributed Computing	719
A Secure Archival Storage System	679	A. Silberschatz A Language Model for Fully Distributed Systems	723
SESSION 37: USE OF MUMPS IN DISTRIBUTED		A.B. Maccabe and R.J. LeBlanc	729
SYSTEMS The Logic of DDP and Its Impact on Programming	:35.	The Starmod Distributed Programming System	127
N.F. Hirst	685	CSP/80: A Language for Communicating Sequential Processes	736
Distributed Database With ISM-11 P.T. Ragon	691	M. Jazayeri, C. Ghezzi, D. Hoffman, D. Middleton and M. Smotherman	
The Functionally Distributed Database: ANS MUMPS As An Implementation Tool		LP. Smyd, J.W. Wong, and J.A. Held elling anti Analysis of Single and Multiple-Eus L-Microccaputar Systems 471	
SESSION 38: NETWORK TOOLS		LATE PAPER diamed L.S bear escapotors, A.G.	
Debugging Commands for a Distributed Processing System J.A. Stankovic	701	Digital TV Communication Network as a Usage of Microcomputers in Japan	741

Treough Protocol.
T.C. Wilson and Q.S. Silio, Jr.

Sessions	PAGE:
dan and the same of the same o	1
Management issues	
Jser Interface to Distributed Processing Systems	41
Network Architectures	
Computer-Based Message Systems. Implementations	85
Network Architecture for Distributed Computers	
Human Factors in Distributed Processing Systems	
Short Notes	
Future Directions of UNIX	147
Computer Networks	149
Text Processing Language Standards	167
Local Computer Networks I	
Distributed Data Design Methodology	
VLSI Chips Wireless Networks Developments	
Wireless Networks Developments	259
Local Computer Networks II	
Schema-Merging Recomposition Procedures	299
Dataflow and Central Concepts for Distributed Processing	
Research Directions	
Switching Techniques	
Distributed Data Base Software Architectures	. 393
Processor Architectures	. 417
DTC/DCE Interfaces	. 419
Communications Systems	. 413
Synchronization, Recovery, and Reliability Issues in the Distributed Data Base	. 431
Environment	
Performance Evaluation	
A Critical Assessment of Software Cost Trends in Telecommunications	
Network Operating Systems	. 481
Multicomputer Architecture I	. 505
Specifying System Composition	
ISO/ANSI Reference Model	. 551
Distributed Operating Systems	. 577
Multicomputer Architecture II	
Applications and Implementation	
Analysis of Distributed Systems	
Distributed Operating System Kernels	. 663
Use of MUMPS in Distributed Systems	. 683
Network Tools	. 699
Programming Languages for Distributed Processing	. 717

SESSION 1: MANAGEMENT ISSUES

W. Kerns U.S. Air Force

SESSION I: MANAGEMENT ISSUES

W. Kerns U.S. Air Force

KEY DECISIONS IN DISTRIBUTED PROCESSING-

Harold Lorin

IBM Systems Research Institute, New York, New York

The establishment of "distributed processing" involves a number of key decisions. These decisions address the dispersion of responsibility, hardware, and software across various operational units. This paper comments on the relationship between organizational decisions and systems structures.

PLACEMENT OF APPLICATION DEVELOPMENT SKILLS

The development of computer applications requires the skills to recognize an application opportunity, to conceptualize a solution consistent with the characteristics of information processing devices, to develop a computer oriented design, to create programs, to debug and test programs, to maintain and evolve programs in view of evolving business needs.

The life-cycle of application development contains some reasonably well defined points of transition from conception to production. In the most rigorously centralized organizations the perception of an application possibility may lie with the data processing department which then imposes applications upon business units or staffs. However, it is at least equally used for a business unit or staffs unit to recognize the potential for the "computerization" of some procedure and approach the data processing department with a proposal. This proposal is then subjected to a benefits/cost analysis which may result in its being undertaken as a project. The role of the business unit at this point becomes the role of a purchaser of professional services from the data processing department.

The vesting of design and implementation skills in a data processing department is justified for a number of reasons. Important among them are:

1. The professional management expertise of the data processing department that gives it the methodologies to exert controls on the development process.

2. The economies of scale that may accrue from an ability to manage populations of expensive professional people more flexibly and to share certain general systems skills which may be needed to develop an application of any complexity.

3. The ability to provide career paths, co-ordinate education, and maintain technological competence.

4. The ability to coordinate applications proposals in order to reduce redundancy or insure proper communication between applications which may be related.

There are alternate views of the degree to which these benefits may be obtained in various situations in various business environments.

The movement of skills from the data processing department to business units may or may not constitute an extreme form of decentralization and it may or may not be complete. It is possible, for example, to put systems justification and design skills into a business unit while maintaining the approval rights and the implementation skills in the data processing department. There are various levels of assistance or enforcement rights which may remain residual within a data processing department even when the skills of a business unit to recognize applications is enhanced.

The actual "exportation" of professional computer talent in the form of applications programmers is an important step that will have major impact on the data-processing styles of an enterprise. It provides a business unit with an independent resource which it can use to generate needs for computational power. It is a key decision in data processing. The decision may be made independently, as we will see, from decisions about the physical placement of computing devices or decisions about the extent to which identified sets of hardware units are dedicated to a particular business unit's applications.

The application development talent which we address here is beyond the level of personal computing or the

use of canned application packages. The image is one of professional programmers who are employed by a business unit and who are primarily concerned with the development and maintenance of computer applications software. A closely related skill, naturally, is the evaluation and selection of available applications packages for use by the business unit.

This export is an important decision for at least three reasons. It implies new management and organizational responsibilities for the business unit, it implies possible new structures within which business units can coordinate and cooperate with each other, and it lays the basis for some hard questions about the relationship between applications development and systems planning. These questions are often difficult to answer because there is commonly a fuzzy intersection between systems development and applications

Before determining to vest professional programming talent in a business unit the following questions must be considered:

1. Is the set of applications associated with business units clearly partitionable so that analysts and designers associated with an applications area can be assigned in such a way that their current skills are directly applicable?

2. Is there a corresponding specialization within the programmers and, if not, will some benefit in programmer productivity or program quality be achieved by dedicating subsets of programmers to the applications of a particular business unit?

3. How much additional structure must be built within the business unit to provide the management skills associated with successful programming projects? How much incremental expenditure may be necessary to replicate necessary unit?

4. Is there some specialized systems environment associated with the applications of the business unit? Will the business unit then find it necessary to take over some systems software responsbilities as well? Will the applications developers coming on board have these skills?

5. Should the business unit undertake to act as a kind of mini-data processing department within its own structure, providing a relationship between "end users" nuch like the relationship between

the business unit and the data processing department, or should the progammers be integrated with end users on a project basis? To what extent will the relationship between end users and programming staffs truly be made easier?

6. Can the business unit offer career paths to its computer people? Would there be a clear line of advancement within the department, or a clear line of advancement out of the department into systems management areas which may be residual in the data processing department?

7. Is the decision to export programmers really being made on its own merits or is the decision entangled in ways that may not be necessary with issues of where machines are placed, how systems are chosen, etc.?

CONTROL OF DATA

Data may be looked upon as a corporate resource and resource and responsibility for data management planning, database organization, data structure, data security, integrity, accessability, vested in a centralized data management organization. At the other extreme data may be considered an application or a program resource and the definition of new data files, structures, etc., left to application development projects.

In planning for distributed processing a policy about data control must be formulated. Between a pure central control approach and the approach that considers data a convenient adjunct to processing, there are a number of operative points of view.

One approach is to vest control of business unit data with business unit data data with business unit data data denoted the concentration of the organization has responsibility and authority to decide what data is to be kept, how data is to be structured, how it is to be protected, etc. Some external standards for integrity, security, recoverability may be imposed by a central data processing standards function, and some requirements for external reporting may be imposed, but the decision about what data to keep and how to keep it is left essentially to the business unit. The concept of corporate wide data management and control is not formed in such a context. The positions of Database Administrator, Database planner, Database designer do not exist at a corporate function level. Attempts to form highly integrated data structures and minimize replication at an enterprise level

are not undertaken. This decision has potentially very wide impact both on application development style and on the relationship between operational units and a data processing staff.

The reasons commonly given to define a corporate wide data management function are:

1. Data is a valuable corporate resource whose definition and integrity, timeliness, quality, accessibility, lies at the center of both tactical and strategic management. It is a unique and valuable resource and must be managed at a level which has an enterprise wide point of view.

2. Software packages which provide appropriate functions of data description, integrity, recoverability are commonly rather complex to use and install. As a consequence advanced systems skills are required to achieve the benefits of such software and these skills are appropriately placed in a professional staff.

3. The use of such software packages is desirable not only from a data control point of view but because of their potential impact on program preparation productivity. Advancing technology in program preparation methodologies holds out the promise that program preparation effort can be reduced by the use of data management facilities associated with highly structured program development language interfaces.

The decision to vest data control with business units need not be total. As with application development skills, partitionings of responsibility and authority may be undertaken. Classes of data may be defined which may or may not meed enterprise wide control, and levels of authority over data may be defined.

In determining what policy should be formed for data control the following factors must be considered:

1. Does the distribution of authority over data suggest a change in the software systems which are used to provide data residence and access?

2. Does the distribution of authority suggest a fundamental change in the hardware structure of systems which will be used to house the data? A physical model of centralized control data with some amount of distribution of processing involves dependent applications processors communicating with a data utility. A dispersed data control physical model involves cooperating

peer processors with local data and mechanisms for data shipment from one to another.

3. Are there facilities for enhanced program development in the context of systems whose use would be implied by the decision to disperse data control?

4. Is there an implication that the physical placement of hardware systems will change as a result of the decision? Does control over data imply where the data is physically placed, or whether specific machine systems will be dedicated to the business unit data?

5. What is the nature of additional data management skill which will be necessary to achieve standards of integrity, etc., by employees of the business unit?

6. To what extent may patterns of business unit cross-referencing emerge so that there is important data referencing across functional business units? If this develops who is responsible for developing such a capability, who is responsible for controlling it? Does the data-processing department undertake to implement and support cross-referencing capability or can each functional unit develop systems which can refer to the data of other units?

The above list of factors suggests that the data control decision is closely related to systems structure design, application autonomy, and physical location of equipment. This is an area where special effort is necessary to discover the implications of various approaches to data control.

PHYSICAL PLACEMENT OF EQUIPMENT

Evolution in data-processing may involve the interconnection of computer systems which are already in place. In such situations a determination about where to physically locate equipment is not necessary. It is possible, however, that because of changes in organization, business style, and technology, a decision about where to locate equipment must be made. It is important to realize that this decision is a separate and distinct decision. It must not be confused with the decisions regarding the degrees of dedication of specific systems with specific departments or applications.

Prequently the idea of distribution is closely associated with the presence of computer equipment physically within a business unit. It is not necessarily true, however, that this dispersion of equipment is

necessary in order to achieve the benefits that application or workload partitioning will bring. An interesting exercise, well worth undertaking, is for business unit management to gather in a meeting with data-processing management in order to discuss what benefits might accrue with various forms of distribution of function and equipment.

A list of what is expected from dedicating equipment to business units may be formed. This list will contain such items as increased availability, increased stability of performance, increased flexibility faster response time, etc. The next step is to determine exactly what kind of difference in benefits will occur if the computer equipment is not located with the functional unit but is retained in a data center. Often it will become obvious that the equipment should be placed in the business unit. Sometimes it will be recognized that the benefits which are sought may be achieved by partitioning and dedication without a physical dispersion of equipment.

A next phase is for data-processing personnel to explain in some detail what is involved in the "ownership" of computing equipment of various kinds. The need for professional staffs with various skills, etc., and the nature of systems management may discourage some whose desire for equipment is only a surrogate for a desire for more responsive service from a data center.

The factors which should be considered in determining whether and to what degree computing equipment should be physically dispersed are:

1. Which of the benefits of dedication of equipment require that the equipment be physcially placed in the business unit? Stability of performance, local programmability, no systems sharing may be accomplished without physical dispersal of the system. Non-dependence on communications equipment and absolute physical security require local siting.

2. How much equipment should be placed in the local facility? Is enough equipment to support some failsoft status sufficient or need the entire applications area operate in the local facility?

3. Are there operational costs which will be different in important ways depending on the physical location of the equipment? Will systems programming or operational costs go up or down? Will local professional staff be needed or will certain aspects of systems control be retained in the data center even if

the equipment is dispersed?

4. Are communications costs going to be a serious factor indicating where equipment should be actually physically placed?

5. Will back-up systems cost and design effort be greater or less if the equipment is dispersed?

ORGANIZATIONAL LEVEL OF DISPERSION

An important decision is to what depth in the organization computing equipment and associated function may be dispersed. Within a functional area there may be some number of organizational levels and each may have a reasonable separable set of applications. At what level in the organizational chart should equipment be dispersed, what will be the equipment/application ratios, at how many discrete sites should equipment be placed?

In considering what levels of functional management constitute an appropriate level of dispersion, the following things should be considered:

1. What is the level at which the organization is prepared to take responsibility for enforcing standards in data control, programming project control, operational control?

2. To what extent is there intensive and continuing interaction between organizational structures within the same functional grouping.

3. To what extent does systems design, communications design, applications design become more complex as a function of more granular partitioning?

4. To what extent has management at lower levels expressed an interest in having dedicated and/or locally sited equipment?

5. At what level of structure is the organization stable so that dispersion of hardware will not reduce the flexibility of subsequent re-organization.

DETERMINATION OF SYSTEMS EDGES

This is an issue that must be carefully worked out between data-processing professionals and functional units who are acquiring a voice in the direction of applications development, systems development and communications design.

The decision to disperse applications programming, for

example, must necessarily involve some questions about what the intersection is between applications development and systems development. Does the capability of developing an application involve the capability to choose what operating system or data manager the application will use? Does local control of data formatting and design imply the right to choose what data manager system will be used to support the data?

In different systems software environments the functional edges between applications programs and systems programs differs widely. Some systems provide functions that must be programmed into the application when other systems are used. A clear understanding of who is responsible for determining systems environments must be achieved before distribution decisions can be put into place.

Similarly, the edges of the network must be defined and discovered. It is good practice for a professional staff to maintain control of a "backbone" network. This backbone network is designed to allow the computer systems of various functional units to interact with each other or with higher level management systems. The design of logical and physcial networks to do this, the establishment of standards and protocols is often considered an important professional activity even when computer power is dispersed. Avoiding the proliferation of applications specific networks, achieving efficient network usage, the flexibility of cross-application access from sets of terminals is seen as a major advantage.

The definition, however, of just what is the "backbone network" is something to which explicit attention must be paid. One approach is to have a centralized staff define the business unit computer to business unit computer network, but leave the definition of subnetworks of terminals or satellite processors which talk to business unit systems to the operational departments. Thus the department A system talks to the department B system across the enterprise network, but terminals talking to A use an isolated network which may be defined by department A.

In considering how much global responsibility the backbone network staff should have, the following factors are important:

 How much cross system access is expected from any particular terminal? Now? Within Five Years?

2. How much terminal specialization is involved in supporting particular applications? Will the nature of the terminals suggest interest in communicating only with one system or will general purpose terminals which may grow in functional use be more common?

3. To what extent will back-up and failsoft situations involve the use of the backbone network?

4. How stable is the functional organization which is used as the basis for distribution? Will it tend to change so that access to systems not currently needed become a requirement?

5. Is there a real possibility that more and more local function will be placed in local units as the intelligence of local units increase? Will this tend to change the systems access patterns of local units?

6. To what extent does technology suggest that distinctions between large capacity, long haul, or low capacity short haul (or various combinations of speed, capacity and distance) will fade in the future and a single technology will be chosen for all unit to unit interaction?

CONCLUSION

We have proposed that the decisions to disperse programming power, data control and physical equipment are key issues in distributed processing. Some organizational decisions have little impact on systems structure. Some, like the decision about data control, seem to imply a great deal about how the computing systems of an enterprise will be shaped.

References

1. Lorin, H., Aspects of Distributed Computing Systems, John Wiley & Sons, New York 1980

2. Lorin, H. Distributed Processing: An Assessment, IBM Systems Journal, Vol 18, No 4, 1979

Processing, IBM Systems Journal, Vol 17, No 4, 1978

seed lewens of a grant to a get a seed lewens of a grant to a get a seed lewens of a grant to a get a seed lewens of a grant to a get a seed lewens of a grant to a get a seed lewens of a get a