Report 77-27
Stanford -- KSL

Scientific DataLink

Overview and Bibliography of Distributed Databases.
Hector Garcia-Molina,
1977

card 1 of 1

HPP-77-27

OVERVIEW AND BIBLIOGRAPHY OF DISTRIBUTED DATA BASES

Hector Garcia-Molina[1]

Department of Computer Science
Stanford University
Stanford, California, 94305

1 INTRODUCTION

The goal of this paper is to provide an overview and a critical bibliography for the area of distributed data bases. Because of the recent technological advances in computer networks and communications, and because of the cost reduction of computer hardware, there has been a great interest in distributed data bases including some attempts at actual implementations. In this paper, we will first define what we mean by a distributed data base. Then we will give some of the reasons why people are so interested in this new field. After classifying the different types of distributed data bases, we will describe the current areas of research. Finally, we will give an annotated bibliography that lists the most important papers in this area.

2 DEFINITIONS

One of the most serious problems in the fast growing area of distributed data bases is that there is no well defined vocabulary: the same words are given different meanings and diverse names are used for the same thing. This is especially true for the term "distributed data base" itself, so in this section we will try to define it.

The first step is to define what we mean by a standard (non-distributed) data base: A data base is a collection of related data that is accessible by a computer. Usually the data is shared by several users with diverse objectives. A data base must also have a set of procedures for handling the data. The operations on the data include storing, updating, searching and retrieval of the data items. The system that handles the data base is called the data base management system.

For our definition of distributed data bases, we will try to give

---

[1] This work was partially supported by a fellowship from Loving Grace Cybernetics. Computer facilities were provided by the SUMEX-AIM facility at Stanford University under National Institutes of Health grant RR-00785. A number of useful suggestions were provided by Kjell G. Knutsen, Reid G. Smith, and Gio Wiederhold.

the most general one possible. This way we will be able to encompass all of the types of distributed data bases. Our definition is the following:

> A distributed data base is a system that allows integrated access to a collection of logically independent data bases.

Notice that we did not mention the term "network of computers" in our definition. There will be a computing facility, called a node, associated with each data base in the distributed data base and there will be communication mechanisms between the nodes. However, a distributed data base does not necessarily have to be spatially distributed, nor is it necessary to have a different computer for each data base. Two or more of the data bases may be physically located in a single computer. In this case, the communication mechanisms are straightforward (e.g. through shared memory). To differentiate a single non-distributed data base in a single computer from a set of data bases on a single machine, it is important that the set of data bases be logically independent. That is, it is necessary that from a logical point of view they could as well be located on separate machines. (In the remainder of this paper, we will still use terms like "network" or "remote site" to simplify explanations.)

By integrated access we mean that a query entered at a node can access data in any one of the data bases. This is a minimum requirement for integrated access; in particular note that being able to update or add data at a remote node is not a requirement.

A typical example of a distributed data base would be a system used by a large manufacturing company. The company has several sites and at each site there is a computer. All the computers are interconnected through a network. The data bases at each of the sites might contain data on the local raw materials and finished products inventory, the planned production at the site, as well as some data on the employees that work there (e.g. name, address, shift, extra hours, etc.). The data base at the company's headquarters might have data on all of the employees (e.g. name, salary, site where employed, etc.) and data on the purchase and sale orders. Typical operations with the data base could be to find out the address of a given worker, to update the inventory, to shift production from one site to another, or to give employees of a certain classification (at any site) a raise.

3 ADVANTAGES OF DISTRIBUTED DATA BASES

Distributed data base systems are by no means the final solution to all data management problems; they are only an alternative to the more common centralized data base systems. Not all data bases should be designed as distributed systems. Only by understanding the particular objectives of a given system and by knowing the advantages offered by a distributed system, will it be possible to decide if distribution pays off. It is also important to understand the advantages and disadvantages

of a distributed solution because emphasizing certain advantages will result in widely varying systems.

The potential advantages of distributed data base systems are the following. We could also call this list the reasons for choosing a distributed system over a centralized one.

1) PERFORMANCE. By taking advantage of the available parallelism and of the increased compute power, we can speed up operations in the data base system. Our gains can be of three types:

a) Response times for searches can be decreased.

b) If the data of immediate importance to the user is kept locally, then this data can be kept more up to date since local updates can be done faster.

c) Larger data bases can be handled without degrading performance.

2) RELIABILITY. By having duplicate data at different nodes, the system will be more reliable. If one node goes down we can still access data from another node. By not having centralized control, the system will be able to operate even if some nodes are down.

3) CONTROL and QUALITY of data. If the data is distributed among the users, they will have direct control of their own data while still being able to share it with other users. When a user is in charge of his own data, he will be responsible for it and will take better care of it. Therefore, the data in the system will be of higher quality.

4) SHARING of geographically distributed data. If the data base already exists and is geographically distributed, then a distributed data base system will interconnect the data bases and allow the sharing of the data.

5) ECONOMY. If the data base users are geographically distributed and if their interactions exhibit strong "locality", then it might be less expensive to do the processing locally. That is, the tele-communications costs can be higher than the tele-processing costs.

6) LOAD DISTRIBUTION. The distributed data base will allow us to move programs and/or data from overloaded nodes to nodes with available capacity.

7) MODULARITY. A distributed data base system can be modular and therefore easier to expand.

8) SECURITY. Distributed data base systems have a potential for greater security because the data bases can be kept in completely independent computers with access from other nodes in the network carefully controlled.

Now we will give a list of the potential disadvantages of data distribution.

1) COMPLEXITY. The main problem with distributed data base systems is that they are considerably more complex than centralized systems. In addition to some of the common issues related to standard data bases, there is an entire set of

questions that are related to the data distribution. (These problems will be treated in Section 5.) The higher complexity implies greater design costs and more sources for error.

2) HARDWARE COSTS. A good distributed data base system inherently has more hardware than a centralized system. Some of the sources of extra hardware are the communication mechanisms, the replicated processors and the extra storage needed for redundant data. Although hardware prices are changing rapidly, most distributed alternatives will be more expensive.

3) LACK OF EXPERIENCE. There are currently only a few experimental and limited distributed data base systems being designed. So there is none of the security implied by a tested and widely used idea.

4) LACK OF CENTRAL CONTROL. It is commonly stated that centralized control is an advantage of centralized data base systems. J. Fry and E. Sibley state that centralized control "is necessary for efficient data administration" [2]. However, the truth of this statement is debatable. If by "efficient" we refer to hardware efficiency (e.g. no wasted resources), then a centralized system would be advisable; but if we are talking about efficient service for the users, then a distributed system might be better. Paul G. Comba uses the following argument against centralized control [15]: "A large complex enterprise does not stand still long enough for the data base administrator and his staff to understand the information needs of every user and integrate them into a complete data base specification. ... The only sensible way to proceed is for the users to participate directly in the specification and development of those parts of the data base system that are intended to facilitate their work; and for the design/implementation process to proceed interactively."

4 CLASSIFICATION OF DISTRIBUTED DATA BASES

The term distributed data base encompasses many diverse types of systems. In order to narrow down our view to a particular system or design, it is necessary to describe some of the characteristics of such systems. Thus we propose the following criteria to classify a distributed data base system:

1) NUMBER AND TYPE OF COMPUTING FACILITIES. Distributed data bases can be categorized by the number and type of nodes. Nodes vary according to their computing power and their storage capacity. For example, some nodes can be back-end data machines [58] which are special purpose computers dedicated to data management. Other nodes can be multi-functional nodes which perform a variety of tasks.

2) LOGICAL NETWORK ORGANIZATION. In some systems the nodes are organized in a "horizontal" fashion, that is, all the nodes are functionally equivalent. On the other hand, some networks

can be "vertically" organized. In this case, some nodes are functionally subordinated to others (e.g. a hierarchy). (Distributed data bases could also be classified to the physical network used, but this classification is not as important. From the distributed data base point of view, it is almost irrelevant how the nodes are physically connected.)

3) SPATIAL DISTRIBUTION. Distributed data bases can also be grouped according to the degree of spatial distribution. Just like it is possible to have the whole system in a single room, it is also possible to have a geographically distributed system.

4) DATA MODEL. Distributed data base systems, together with centralized systems, can be classified by the data model used (e.g. relational, network, hierarchical, etc.). However, in distributed systems, it is possible to use more than one data model in a single system. (See next item below.) For information about the data models, see references [1-6].

5) HOMOGENEITY. Another way to categorize distributed data bases is by the uniformity of the data model and data languages used throughout the system. In a homogeneous system, all the nodes will use the same data model and data language. Therefore it will not be necessary to have either data or command translation for inter-node communications. In a heterogeneous system, translation between nodes will be necessary.

6) DEGREE OF CENTRALIZED CONTROL. The amount of centralized control allowed can drastically change the design of a system. Centralized control can include a system directory, a deadlock prevention mechanism, a synchronization center or a locking mechanism. In a centralized control system, one or more special nodes are in charge of the control functions, while in a distributed control system, the control of the system operation is shared by all the nodes. As a general rule, a distributed control system is more complex but more reliable.

7) AMOUNT AND TYPE OF DATA REDUNDANCY. Data can be duplicated for two reasons: performance and reliability. By having several copies of the data, access time to it can be reduced (performance) and if one of the copies is lost, the data is still accessible (reliability). In some systems, only certain parts of the data base will be stored redundantly at some nodes, while in other systems the complete data base may be replicated at every node.

8) COMPLEXITY OF THE INTERACTIONS. Since not every system allows the most general types of interactions, distributed data bases can be grouped by the complexity of the interactions permitted. Some common restrictions are to limit the number of nodes a query can span, to limit the number of intermediate steps for query processing, to permit only updates that have been generated at the same node (local updates), etc.

9) DEGREE OF SEMANTIC COUPLING. Distributed data bases can also be classified by the amount of semantic coupling between the data bases. In a closely coupled system, the data bases are

related through the actual meaning of the data. For example, a data base containing data on employees is more closely coupled to a payroll data base than to a data base with data on the weather. (The coupling is closely related to how the total data is distributed among the data bases, so that this coupling classification could also be called a classification by how the data is distributed in the system.) The type of data base coupling can affect the types of queries that are received as well as the algorithms used for processing the queries.

10) AMOUNT OF DYNAMIC RECONFIGURATION. Distributed data base systems can be categorized by the amount of dynamic reconfiguration that is allowed. In a static system, the location of the data and directories, the names and numbers of files (or relations) and the system structure are all fixed at creation time. In a dynamic system, data can migrate and be duplicated, files (or relations) can be created and destroyed, and nodes can be added to the system, all without interrupting the operation of the system. Clearly, dynamic systems are more complex but more flexible.

5 CURRENT AREAS OF RESEARCH

The field of distributed data bases is a complex one where there are still a lot of unresolved issues. It is a relatively young field where some research has been done but where much more is needed. We will now list and briefly describe the current areas of research and some of the main problems of distributed data bases. (A more detailed description of these areas can be found in some of the overview papers [1,14,16,17,18,19].)

1) PROGRAM AND DATA ALLOCATION. The problem here is to find the optimal location and the optimal number of copies of the program and data files in the distributed data base. What is to be minimized are the combined storage, communication and processing costs. By choosing different sets of assumptions and variables, several solutions of varying complexity have been obtained [20-26].

2) MAINTENANCE OF DUPLICATE COPIES. Since copies of the data may exist at different nodes, it is necessary to have algorithms that make sure that all copies are updated properly. Special mechanisms are needed to know where the duplicate copies, if any, exist. Several algorithms for different types of distributed data bases have been suggested and work is in progress for proving the algorithms correct [31,32,43].

3) CONCURRENCY CONTROL. In a distributed data base system, several users may be attempting to read and/or update a set of data. In order to always provide users with a consistent view of the data, it is necessary to have concurrency control. This control, which can either be centralized or distributed, should include synchronization and locking mechanisms. Some work has

been done defining the basic concepts [30] and analyzing the available options [34,37,39,40,42].

4) DEADLOCKS. Just like in any system where multiple users compete for access to a set of finite resources, in distributed data base systems there is a possibility for deadlocks. There are two ways to deal with deadlocks: deadlock prevention and deadlock detection. Both of these alternatives have been analyzed for general systems [67] and in particular for distributed data base systems [29,33,34,35,37,40].

5) QUERY PROCESSING. This area involves the design of algorithms that process queries into strings of data manipulation commands. These algorithms are more complex than the usual algorithms for centralized data bases. First of all, the local knowledge at the node where the query is processed might not be enough to understand the query, so that help from other nodes is needed. Then there is the problem of locating the relevant data. If the data is duplicated, we must choose the copy to use. Finally, one must decide how to actually manipulate the data. There are three options for this: transmit commands and transmit results back; move all of the necessary data to a node and work there; or a mixture of these two methods (i.e. filter data before moving). These problems are discussed in [45,56,59].

6) MANAGEMENT OF NAME SPACE. A directory contains a description and the location of files (or relations) in the system. Directories can be global or local, distributed or centralized and they can have one or many copies. The tradeoffs involved with the different options are being analyzed [38,51]. If the distributed data base is dynamically changing, it is necessary to have mechanisms to add or delete names to the directory.

7) DATA AND PROGRAM TRANSLATION. In a non-homogeneous distributed data base, it is mandatory to have translation mechanisms between the data bases. Since it would be very inefficient to design an interface for every possible pair of dissimilar data bases, it is necessary to design general procedures for translating data and programs. These procedures can include languages for describing the data and program formats plus definition of a common intermediate format [27,28].

8) PRIVACY. There has been very little work done in the area of data privacy. It is necessary to design good ways of identifying users, both local and remote. It would also be nice to be able to restrict access not only by who the user is but by what his application is. For example, a user might not be allowed to access a particular employee's salary, but he may be permitted to look at the average salary of a group of employees.

9) RECOVERY. If a distributed data base system is to be reliable, procedures for detecting errors and recovering from failures are required. It is important that when some data base fails, the rest of the data is left in a consistent form. When a node comes up after a failure, it is indispensable to get its

data base up to date. Recovery can become extremely hard if failures cause the network to partition (i.e. to split up into several isolated pieces). These problems are treated in [36,41].

6 BIBLIOGRAPHY

The following is a bibliography for distributed data bases. Additional references to papers in the field can be found in the survey papers [7,9,13,14,16,18,19]. The papers marked with a "+" form a basic reading set for the distributed data base newcomer. These papers can provide a more complete and thorough introduction to the field. The papers that provide important relevant results or ideas are marked with a "*".

Symbols used:

- + Recommended reading.
- * Important paper.
- @ Institution where author is located.
- / Source.
- = Short description or comment.

Categories:

- I. Non-distributed data bases.
- II. General.
- III. Research issues.
- IV. Systems being designed.
- V. Other related subjects.

------ I. Non-distributed data bases. ----------

1. +* C. J. Date
@ IBM (UK) Laboratories
"An Introduction to Database Systems"
/Addison-Wesley Publishing Co., January 1976
= This is a good introductory textbook for data bases. It defines the three basic data models currently used (relational, network, and hierarchical). This book also has some material on data base security and integrity. The book is biased towards the relational model.

2. J. P. Fry and E. H. Sibley
@ U. of Michigan and U. of Maryland
" Evolution of Data-Base Management Systems"
/ ACM Computing Surveys, Vol. 8, Num. 1, March 1976, pp. 7-42
= This is an introductory paper. It deals with the history and definitions common to data base technology. It describes the objectives of data base management systems. Other relevant concepts are also discussed.

3. D. D. Chamberlin
@ IBM Research Laboratory, San Jose, Ca. 95193
" Relational Data-Base Management Systems "
/ ACM Computing Surveys, Vol. 8, Num. 1, March 1976, pp.43-66
= Introductory paper for data bases based on the relational model. Discusses the essential concepts, relational languages and implementation issues.

4. R. W. Taylor and R. L. Frank
@ IBM Research Laboratory, San Jose, Ca. 95193
@ and Computer Science Dept., U. of Utah
" CODASYL Data-Base Management Systems "
/ ACM Computing Surveys, Vol. 8, Num. 1, March 1976, pp. 67-103
= Introductory paper for data bases based on network model. Presents in tutorial fashion the concepts, notation and data base languages that were defined by the CODASYL Data Description Language and Programing Language Committees.

5. A. S. Michaels, B. Mittman and C. R. Carlson
@ Bell Telephone Labs, Inc., Naperville, Illinois 60540
@ Vogelback Computing Center, Northwestern U., Evanston, Ill. 60201
@ and Dept. of Computer Science, Northwestern U., Evanston, Ill. 60201
" A Comparison of the Relational and CODASYL Approaches to Data-Base Management"
/ ACM Computing Surveys, Vol.8, Num. 1, March 1976, pp.125-151
= Introductory paper. Compares and contrasts the relational and network data models.

6. * G. Wiederhold
@ Computer Science Department, Stanford U., Stanford, Ca. 94305
" Database Design "
/ McGraw Hill Book Co., June 1977
= This textbook provides comprehensive coverage of most of the issues in data bases.

------ II. General. ----------

7. +* M. E. Deppe and J. P. Fry
@ Graduate School of Business Administration, U. of Michigan,
@ Ann Arbor, Mi 48109
" Distributed Data Bases-- A Summary of Research "
/ Computer Networks, Vol. 1, Num. 1, March 1976, pp. 130-138
= A very good introductory paper. Defines some of the basic concepts; describes the areas of current and future research. Has an extensive bibliography.

8. G. M. Booth
@ Honywell Information Systems, Phoenix, Arizona
" The Use of Distributed Data Bases in Information Networks "
/ Proceedings of the First International Conference on Computer Communications: Impacts and Implications, October 1972, pp. 371-376

= This paper does not say much; Booth's 1976 paper is better.

9. G. M. Booth
@ Honywell Information Systems, Phoenix, Arizona
" Distributed Information Systems "
/ National Computer Conference 1976, pp. 789-794
= Not very deep paper; provides some general definitions and some examples of distributed data bases.

10. H. Merten
@ Siemens AG, Munchen, F. R. of Germany
" Communication with Data Bases"
/ Proceedings of the 3rd International Conference on Computer Communications, 1974, pp.61-66
= A very limited paper. Analyzes requirements for data communication between computers.

11. J. C. Chupin
@ CII Scientific Center, I.M.A.G., Grenoble, France
" Control Concepts of a Logical Network Machine for Data Banks "
/ Proceedings of IFIP congress 1974, pp. 291-295
= Describes the control concepts of a Logical Network Machine to be used for a data base application using the Socrate system. Its hard to understand what the author is saying in the paper.

12. * D. N. Streeter
@ I. B. M. Yorktown Heights, N. Y.
" Centralization or Dispersion of Computing Facilities "
/ IBM Systems Journal, No. 3, 1973, pp. 283-301
= Describes advantages and disadvantages of distributed systems (including data bases). Introduces simple equations for analyzing cost and performance.

13. * R. G. Canning
@ Canning Publications Inc., Vista, Ca. 92083
" Distributed Data Systems "
/ EDP Analyzer, Vol. 14, Num. 6, June 1976
= Provides a good overview of distributed data bases from the "business" point of view. Describes a few systems that are being implemented in some banks and companies.

14. * F. Aschim
@ Central Institute for Industrial Research, Oslo, Norway
" Data Base Networks - An Overview "
/ Management Informatics, Vol. 3, Num. 1, 1974, pp. 13-28
(The magazine is now called Management Datamatics.)
= Provides a general overview of the area. Classifies some of the available alternatives. Has a good bibliography.

15. * P. G. Comba
@ IBM Corporation, Cambridge, Ma. 02139

" Needed: Distributed Control "
/ Proc. of the International Conference on Very Large Data Bases, Farmingham, Ma., September 1975, pp. 364-373
= States that for very large data bases, distribution is the "only" way to go. Lists the political, technical, organizational, etc. advantages to distributed data bases.

16. * E. Grapa
@ Dept. of Computer Science, U. of Illinois at Urbana-Champaign
" Characterization of a Distributed Data Base System "
/ U. of Illinois Thesis and Report UIUCDCS-R-76-831, October 1976
= Describes some of the current work done in the area. Describes and compares several methods for sharing and updating distributed data.

17. E. F. Severino
@ Auerbach Publishers Inc.
" Databases and Distributed Processing "
/ Computer Decisions, Vol. 9, Num. 3, March 1977, pp. 40-42
= A very brief and informal overview of the distributed data base area.

18. +* F. J. Maryanski
@ Computer Science Department, Kansas State U.,Manhattan, Kansas 66506
"A Survey of Developments in Distributed Data Base Management Systems"
/ Technical Report CS77-01, Kansas State U., January 1977
/ To appear in Computer
= A very complete survey of the field of distributed data bases. Lists the main areas of research and the important work done in each area. Has a very complete bibliography.

19. * P. S. Fisher and and F. J. Maryanski
@ Computer Science Dept., Kansas State U., Manhattan, Kansas 66506
" Design Considerations in Distributed Data Base Management Systems "
/ Technical Report CS 77-08, Kansas State U., April 1977
= Also provides a survey of the field of distributed data base systems.

------ III. Research issues. ----------

20. * S. Mahmoud and J. S. Riordon
@ Carleton University, Canada
"Optimal Allocation of Resources in Distributed Information Networks"
/ ACM Transactions on Database Systems, Vol. 1, Num. 1, March 1976, pp. 66-78
= The problems of file allocation and capacity assignment in a fixed topology distributed computer network are examined.

21. +* K. D. Levin and H. L. Morgan
@ The Warton School, Philadelphia, Pennsylvania
" Optimizing Distributed Data Bases-- A Framework for Research "

/ National Computer Conference 1975, pp. 473-478
= Discusses the file location problem at the program and data sharing level, i.e. dependencies between programs and files are taken into account.

22. * W. W. Chu
@ Computer Science Department, UCLA, Los Angeles, Ca.
" Optimal File Allocation in a Multiple Computer System "
/ IEEE Trans. on Computers Vol. 18, Num. 10, October 1969, pp.885-889
= A model is developed for allocating information files required in common by several computers. The model is formulated into a nonlinear integer zero-one programing problem.

23. * R. G. Casey
@ IBM Research Laboratories, San Jose, Ca.
" Allocation of Copies of a File in an Information Network "
/ Proc. of the Spring Joint Computer Conference, 1972, pp.617-625
= Demonstrates several properties of the optimal assignment of files to nodes in a network with multiply-located files. A test useful in determining the optimum configuration is derived.

24. R. G. Casey
@ IBM Research Laboratories, San Jose, Ca.
" Design of Tree Networks for Distributed Data "
/ Proceedings of National Computer Conference 1973, pp. 251-257
= Networks are restricted to tree structures. A procedure for finding the close to optimal network (file location, line capacities and network topology) is described.

25. * P. A. Alsberg
@ Center for Advanced Computation, U. of Illinois at Urbana-Champaign
" Data Distribution Strategies "
/ Proc. (First) Berkeley Workshop on Distributed Data Management and Computer Networks, May 25-26, 1976
= This short paper describes the ways that data can be distributed in a computer network.

26. G. G. Belford and J. D. Day
@ Center for Advanced Computation, U. of Illinois at Urbana-Champaign
" Cost Analysis of Data Distribution "
/ Data Base Engineering, Vol. 1, Num. 1, March 1977, pp.5-12
= This paper is a summary of a detailed report. The key question addressed is: When is it more economical to store data at a remote site and bring it over the network when needed?

27. +* A. G. Merten and J. P. Fry
@ Department of Industrial and Operations Engineering, U. of Michigan, Ann Arbor, Michigan, 48104
" A Data Description Language Approach to File Translation "
/ Proceedings ACM SIGMOD Workshop, May 1974, pp. 191-205
= Describes the design of a system that translates data created under

a particular system so that it can be read and processed by another system which may be on different hardware. The system uses a Stored Data Definition Language and a Translation Definition Language to specify the translation.

28. * K. Yamaguchi and G. Merten
@ Department of Industrial and Operations Engineering, U. of Michigan, Ann Arbor, Michigan, 48104
" Methodology for Transfering Programs and Data "
/ Proceedings ACM SIGMOD Workshop, May 1974, pp. 141-155
= Presents an overview of current and proposed solutions to the problems of incompatible programs and data between dissimilar computer systems.

29. * D. D. Chamberlin, R. F. Boyce and I. L. Traiger
@ IBM Research Laboratory, San Jose, Ca.
" A Deadlock-free Scheme for Resource Locking in a Data Base Environment "
/ Proceedings of IFIP Congress 1974, pp. 340-343
= This is an interesting paper. It discusses some of the problems in locking records in a data base. An algorithm is proposed whereby the system guarantees a consistent snapshot of the data base to each process which issues a lock request.

30. +* K. P. Eswaran, J. N. Gray, R. A. Lorie and I. L. Traiger
@ IBM Research Laboratory, San Jose, Ca 95193
" On the Notions of Consistency and Predicate Locks in a Data Base System "
/ IBM Research Report RJ 1487, December 1974
/ Also: CACM November 1976, pp. 624-633
= This paper discusses and defines the concepts of transaction, consistency, locking and scheduling. An implementation of predicate locks is suggested. This paper is "required reading" for a good understanding of these fundamental concepts.

31. * C. A. Ellis
@ Xerox Palo Alto Research Center
" A Robust Algorithm for Updating Duplicate Databases "
/ Xerox PARC Report; No date, received March 1977
/ Also: Proc. 2nd Berkeley Workshop Dist. Data Management, pp.146-158
= A decentralized algorithm used to maintain multiple copies of a data base at various nodes is discussed. The algorithm maintains consistency, avoids deadlocks, and is robust.

32. +* P. R. Johnson and R. H. Thomas
@ Bolt Baranek and Newman Inc., Cambridge, Ma. 02138
" The Maintenance of Duplicate Databases "
/ Network Working Group RFC 677, January 1975
= Duplicate data base problems are avoided by using timestamps for updating and creating data.

33. +* W. W. Chu and G. Ohlmacher
@ UCLA, Los Angeles, Ca.
" Avoiding Deadlock in Distributed Data Bases "
/ Proceedings of the ACM National Symposium, Vol. 1, November 1974 pp. 156-160
= Deadlock prevention mechanisms and a deadlock protection mechanism for distributed data base systems are presented.

34. * A. Shoshani and A. J. Bernstein
@ Dept. of Electrical Eng., Princeton U., Princeton, New Jersey and
@ Research and Development Center, General Electric Co.,
@ Schenectady, New York
" Synchronization in a Parallel-Accessed Data Base "
/ CACM, Vol. 12, Num. 11, November 1969, pp. 604-607
= Discusses the rules necessary for accessing a shared data base without producing deadlocks.

35. * P. F. King and A. J. Collmeyer
@ Xerox Corporation, El Segundo, Ca.
" Database Sharing - An Efficient Mechanism for Supporting Concurrent Processes "
/ Proceedings of National Computer Conference 1973, pp. 271-275
= A model based on graphs for a shared set of files is presented. A deadlock detection and recovery mechanism is introduced.

36. +* P. A. Alsberg and J. D. Day
@ Center for Advanced Computation, U. of Illinois at Urbana-Champaign
" A Principle for Resilient Sharing of Distributed Resources "
/ 2nd International Conference on Software Engineering, San Francisco 1976, pp. 562-570
= A technique is presented which permits distributed resources to be shared in a resilient maner. One of the hosts is declared a priori the primary and the others the backups.

37. L. C. Sekino
@ Bell Telephone Laboratories, Holmdel, N. J. 07728
" Multiple Concurrent Updates "
/ Proc. of the International Conference on Very Large Data Bases, Farmingham, Ma., September 1975, pp. 505-507 (Abstract only.)
= Contains little new material. A method for updating a shared data base is proposed. Deadlocks are prevented by ordering lock units.

38. * W. W. Chu and E. E. Nahouraii
@ Computer Science Department, UCLA, Los Angeles, Ca., 90024
@ (Nahouraii is currently at IBM System Development Division,
@ Palo Alto, California.)
" File Directory Design Considerations for Distributed Data Bases "
/ Proc. of the International Conference on Very Large Data Bases, Farmingham, Ma., September 1975, pp. 543-545 (Abstract only)
= The centralized, the localized and the distributed file directories are studied. Mathematical models are formulated to study the tradeoffs

in cost and response time.

39. * J. N. Gray, R. A. Lorie and G. R. Putzolu
@ IBM Research Laboratory, San Jose, Ca.
" Granularity of Locks in a Shared Data Base "
/ Proc. of the International Conference on Very Large Data Bases, Farmingham, Ma., September 1975, pp. 428-449
= Proposes a locking protocol which associates locks with sets of resources. This protocol allows simultaneous locking at various granularities by different transactions.

40. * J. E. Shemer and A. J. Collmeyer
@ Xerox Corporation, El Segundo, Ca.
" Database Sharing: A Study of Interference, Roadblock and Deadlock "
/ 1972 ACM SIGFIDET Workshop on Data Description Access and Control, Denver, Colorado, December 1972, pp. 147-163
= A shared data base locking mechanism is simulated. The results and a description of the simulation is presented.

41. * P. A. Alsberg
@ Center for Advanced Computation, U. of Illinois at Urbana-Champaign
" Synchronization and Resiliency in Network Data Access "
/ Proceedings (First) Berkeley Workshop on Distributed Data Management and Computer Networks, May 25-27, 1976
= Defines and discusses the concept of resiliency for a distributed data base system. Discourses on the inherent cost of synchronization.

42. * D. J. Rosenkrantz, R. E. Stearns and P. M. Lewis
@ General Electric Research and Development Center
"A System Level Concurrency Control for Distributed Database Systems"
/ Proceedings Second Berkeley Workshop on Distributed Data Management and Computer Networks, May 1977, pp. 132-145
= Proposes two algorithms for distributed concurrency control in a particular distributed data base model. Discusses some of the implications of maintaining global data base consistency.

43. * B. Thomas
@ Bolt Beranek and Newman Inc., Cambridge, Ma. 02138
" On the Maintenance of Duplicate Databases "
/ Working Paper, no date, received May 1977
= Presents an update algorithm where an update is accepted by having all nodes vote. The algorithm is deadlock free, robust, has distributed control and is supposed to be provable correct.

------ IV. Systems being designed. ----------

44. M. Stonebraker, E. Wong, P. Kreps and G. Held
@ UC Berkeley and Tandem Computers Inc. (Held)
" The Design and Implementation of INGRES "
/ ACM Trans. on Database Systems, Vol. 1, Num. 3, September 1976, pp. 189-222

= Describes implementation of INGRES, a non-distributed relational data base system. This paper is useful for understanding the distributed INGRES paper.

45. +* M. Stonebraker and E. Neuhold
@ Electronics Research Laboratory, UC Berkeley and
@ U. of Stuttgart
" A Distributed Data Base Version of INGRES "
/ Electronics Research Laboratory Report, UC Berkeley, September 1976
/ Also: Proc. 2nd Berkeley Workshop Dist. Data Management, pp. 19-36
= Discusses some of the design choices made for a system that will be implemented. The paper is not well organized, but has some good ideas.

46. E. A. Ozkarahan, S. A. Schuster and K. C. Smith
@ U. of Toronto, Toronto, Ontario, Canada
" RAP-- An Associative Processor for Data Base Management "
/ National Computer Conference 1975, pp. 379-387
= This paper describes a very different type of distributed data base. A small processor is assigned to each disk track. The processors work in parallel for solving a query.

47. J. D. Foster
@ Advanced Systems, Bank of America, San Francisco, Ca.
" The Development of a Concept for Distributed Processing "
/ Digest of Papers, COMPCON Spring 1976, pp. 28-30
= Describes the reasons why the Bank of America is going to a distributed data base. Also gives a brief description of what the system will look like.

48. R. Peebles and E. Manning
@ Dept. of Computer Science, U. of Waterloo
" A Computer Architecture for Large (Distributed) Data Bases "
/ ACM Symposium on Very Large Data Bases, 1976, pp. 405-427
= The operation (software) of a distributed transaction system is described.

49. E. Chang
@ Computer Science Dept., U. of Waterloo, Waterloo, Ontario N2L 3G1
" A Distributed Medical Data Base "
/ Computer Networks, Vol. 1, Num. 1, 1976 pp. 33-38
= Describes the software design aspects of a distributed medical data base to be implemented on a homogeneous minicomputer loop network.

50. B. Thomas
@ Bolt Beranek and Newman Inc., Cambridge, Ma. 02138
" A Resource Sharing Executive for the Arpanet "
/ National Computer Conference 1973, pp. 155-163
= Describes the RSEXEC system. RSEXEC is a distributed executive system which functions to integrate the operation of the Arpanet Tenex hosts. RSEXEC supports a distributed file system.

51. * W. R. Sutherland
@ Bolt Beranek and Newman Inc., Cambridge, Ma. 02138
" Distributed Computation Research at BBN "
/ BBN Report 2976, Volume III, December 1974
= Describes the RSEXEC system. Also has some comments on distributed data bases (pp.42-46).

52. * F. J. Maryanski
@ Dept. of Computer Science, Kansas State U., Manhattan, Kansas 66506
"Design Considerations for a Distributed Data Base Management System"
/ Technical Report CS-76-14, Kansas State U., September 14, 1976
= Design considerations for a system that will be implemented for the Army. System will be based on CODASYL specifications; will use a network of mini-computers. In particular, this paper discusses a memory management algorithm as well as a deadlock prevention algorithm for the system.

53. F. J. Maryanski
@ Dept. of Computer Science, Kansas State U., Manhattan, Kansas 66506
" Language Specifications for a Distributed Data Base Management System "
/ Technical Report CS-76-13, Kansas State U., May 1976
= The schema data definition, sub-schema data definition, data manipulation, and device media control languages for the system are described. This paper does not focus on the distribution problems.

54. V. E. Wallentine and F. J. Maryanski
@ Dept. of Computer Science, Kansas State U., Manhattan, Kansas 66506
" Implementation of a Distributed Data Base System "
/ Technical Report CS-76-03, Kansas State U., February 1976
= Provides an overview of the system being designed. The paper focuses on the inter-task (across processors) communication system.

55. +* J. B. Rothnie and N. Goodman
@ Computer Corporation of America, Cambridge, Massachusetts 02139
" An Overview of the Preliminary Design of SDD-1: A System for Distributed Databases "
/ Proceedings Second Berkeley Workshop on Distributed Data Management and Computer Networks, May 1977, pp. 39-57
= Presents an overview of a very interesting system that will be implemented. System uses relational model, allows duplicate data, allows distributed directories.

56. +* E. Wong
@ U. of California, Berkeley and
@ Computer Corporation of America, Cambridge, Massachusetts 02139
" Retrieving Dispersed Data from SDD-1: A System for Distributed Databases "
/ Proceedings Second Berkeley Workshop on Distributed Data Management and Computer Networks, May 1977, pp. 217-235
= Presents an algorithm for efficient retrieval from the SDD-1 system.

The queries are decomposed into data transfers and local processing at the nodes of the system.

57. B. D. D'Ambrosio
@ U. of California, Berkeley, California
" Design Studies for an On-line U. C. Library Catalog Network "
/ Proceedings Second Berkeley Workshop on Distributed Data Management and Computer Networks, May 1977, pp. 257-263
= Discusses design considerations for a planed state-wide computer network for the combined U. C. library catalog.

58. * R. H. Canaday, R. D. Harrison, E. L. Ivie, J. L. Ryder and L. A. Wehr
@ Bell Telephone Laboratories Inc., Piscataway, New Jersey
" A Back-end Computer for Data Base Management "
/ CACM, Vol. 17, Num. 10, October 1974, pp. 575-582
= Describes the implementation of a back-end data base management computer, the XDMS. Lists the advantages and disadvantages of the back-end concept.

59. * H. Garcia-Molina and G. Wiederhold
@ Department of Computer Science, Stanford U., Stanford, Ca. 94305
" Application of the Contract Net Protocol to Distributed Data Bases "
/ HPP Report 77-21, Stanford U., March 1977.
= The paper applies the contract net formalism to develop a model for interaction in a distributed data base.

------ V. Other related topics. ----------

60. * T. Marill and D. Stern
@ Computer Corporation of America, Cambridge, Massachusetts
"The Datacomputer-- A Network Data Utility "
/ National Computer Conference 1975, pp. 389-395
= The datacomputer is not a distributed data base system but rather an alternative to a distributed system.

61. * D. J. Farber and F. R. Heinrich
@ U. of California, Irvine, Ca.
" The Structure of a Distributed Computer System-- The Distributed File System "
/ Proceedings of the 1st International Conference on Computer Communications, October 1972, pp. 364-370
= Describes a reliable fail-soft distributed file system for the computer network that exists at UC Irvine. A distributed file system is not really a distributed data base.

62. J. J. Passafiume and S. Wecker
@ Digital Equipment Corporation, Maynard, Massachusetts
" Distributed File Access in DECNET "
/ Proceedings Second Berkeley Workshop on Distributed Data Management and Computer Networks, May 1977, pp. 114-129

= Describes a file transfer and access protocol for DECNET. This is not really a distributed data base system.

63. Y. K. Dalal
@ Digital Systems Laboratory, Stanford U., Stanford, Ca. 94305
" Distributed File Systems "
/ Proceedings (First) Berkeley Workshop on Distributed Data Management and Computer Networks, May 25-26, 1976, pp. 296-307
= Examines the logical structure of file systems and shows how they could be distributed. Adaptive distributed algorithms are proposed for achieving automatic file migration.

64. S. E. Madnick
@ MIT, Cambridge, Massachusetts
" INFOPLEX-- Hierarchical Decomposition of a Large Information Management System Using a Microprocessor Complex "
/ Proceedings NCC 1975, pp. 581-586
= Describes how micro-processors could be used to hierarchically decompose a large memory system. The decomposition can be functional (e.g., query language, n-ary relations, binary relations, virtual storage) or physical(e.g., cache, main, secondary, mass).

65. * A. J. Collmeyer
@ Xerox Data Systems, El Segundo, California
" Database Management in a Multi-access Environment "
/ Computer, November-December 1971, pp. 36-46
= A tutorial on the problems of shared data bases with concurrent access. Defines the parts of a data base management system; identifies the different types of file accesses and briefly discusses integrity and security.

66. S. I. Saffer, D. J. Mishelevich, S. J. Fox and V. B. Summerour
@ U. of Texas Health Science Center at Dallas
" NODAS-- The Network Oriented Data Aquisition System for the Medical Environment "
/ Proceedings NCC 1977, pp. 295-299
= Describes an elementary network of computers used in medical laboratories. The network is a simple star with a PDP-10 host and PDP-11 satellites.

67. * E. G. Coffman, M. J. Elphick and A. Shoshani
@ Pennsylvania State U., U. of Newcastle upon Tyne and
@ System Development Corporation, respectively
" System Deadlocks "
/ Computing Surveys, Vol. 3, Num. 2, June 1971
= A survey on system deadlocks from both the theoretical and practical points of view.

-- End of bibliography. --