

8-25-1995

# Decision Support for Distributed Database Fragmentation and Allocation Schema Design

Goul M.  
*Arizona State University*

Follow this and additional works at: <http://aisel.aisnet.org/amcis1995>

---

## Recommended Citation

M., Goul, "Decision Support for Distributed Database Fragmentation and Allocation Schema Design" (1995). *AMCIS 1995 Proceedings*. 65.  
<http://aisel.aisnet.org/amcis1995/65>

This material is brought to you by the Americas Conference on Information Systems (AMCIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in AMCIS 1995 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact [elibrary@aisnet.org](mailto:elibrary@aisnet.org).

# Decision Support for Distributed Database Fragmentation and Allocation Schema Design

M. Goul, R. D. St. Louis, and R. A. You  
Department of Decision and Information Systems  
College of Business  
Arizona State University  
Tempe, AZ 85287-4206  
Phone: (602) 965-6350

## ABSTRACT

*If peer to peer distributed database systems meet current expectations, they are likely to replace virtually all centralized database systems over the next decade. One impediment to the proliferation of peer to peer distributed database systems is the lack of proven and established normative methodologies for designing distributed database fragmentation and allocation schemas. The literature discussed here serves as the basis of research-in-progress for designing, implementing, and empirically evaluating a support system to aid distributed database fragmentation and allocation schema decision making. Future manuscripts are planned that describe the prototype decision support system and our empirically-based experiences.*

## Introduction

A distributed database schema differs from a centralized database schema in that relations can be horizontally and/or vertically divided into a disjoint 'fragmentation schema.' Relations and/or partial relations are then selected to be resident at specific nodes of a computer network, where each node is running the distributed database system kernel. Some relations and/or partial relations can be selected to reside at more than one node, and such relations and/or partial relations are called 'replications.' The manner by which replicated and non-replicated relations and/or partial relations are assigned to nodes is called an 'allocation schema.' Taken together, the formulation of a fragmentation schema and an allocation schema represent a set of database

administrator decisions that are not associated with normative centralized database schema design concepts and methodologies such as functional dependencies, normal forms, dependency preserving decompositions, and lossless join decompositions.

The focus of our research is on 'logical' fragmentation and allocation schema design as opposed to decision making with respect to the 'physical layer' (the bottom-most layer of the ANSI/SPARC proposed database architectural framework as described by Apers and Scheuermann, 1991). The most widely cited distributed database design methodology is

Ceri, et al.'s [1983] top-down approach, which includes the phases of requirements collection, requirements specification, view analysis and integration, distribution design, and physical database design for each database residing at a specific network node. Fragmentation and allocation schema decision making are the foci of the distribution design phase of Ceri, et al.'s approach. The two most important elements of the distribution design phase are the fragmentation and allocation decisions.

## **Fragmentation**

A fragmentation schema is a division of the global schema of relations into smaller relations. A 'horizontal' fragment is a partial relation containing all the attributes of the original relation, but with fewer tuples, all of which match on a specific attribute. A 'vertical' fragment is a partial relation containing a subset of attributes of a relation and a key to the original relation upon which vertical fragments can be re-joined to form the original relational. Please note that fragmented relations can themselves be further fragmented, but the fragmentation schema must be disjoint with respect to all relations and extensions of the global schema.

## **Allocation**

Once a fragmentation schema has been decided upon, each fragment must be assigned to reside at one or more nodes in the distributed database management system. This assignment of fragments to nodes is called an allocation schema. The problem is similar to that of file allocation [Dowdy and Foster, 1982], but has several significant differences [Apers, 1988]. Allocation basically involves finding the best possible distribution of fragments to sites, where minimal cost and maximum performance represent two of a host of potentially competing multiple objectives to be weighed.

## **Related Work**

Virtually all of the existing research on fragmentation and allocation schema decision making involves algorithm development that is focused on a particular area, i.e., fragmentation or allocation. In other words, there is very little research that considers the possibility of integrating fragmentation and allocation decision making under the auspices of one decision making context. Following is a discussion of recent literature in each area.

Navathe, et al. [1984] developed a two-phase approach for determining fragments given a global database schema. The gist of the approach is to rely on the apriori collection of cost factors that can then be used to generate fragments that can then be allocated so as to maximize query locality. Cornell and Yu's approach [1990] to fragmentation takes into consideration the access plan selections made by a distributed database system's query optimizer [1990]. Their approach incorporates a query analysis step that uses a binary partitioning algorithm based on an integer linear programming solver. Meghini and Thanos [1991] developed algorithms for determining fragments as well as an algorithm for checking the completeness of a fragmentation schema. While the above literature is

by no means exhaustive, it is representative of the types of approaches traditionally taken in addressing the fragmentation schema design problem.

Allocation schema decision making, when viewed in isolation of fragmentation schema decision making, basically involves finding the 'optimal' allocation of fragments to sites. The major differences in existing research approaches are related to definitions for the term 'optimality.' For example, Dowdy and Foster [1982] define optimality with respect to two measures: minimal cost and maximum performance. Ozsu and Valduriez [1991] view the allocation problem as a cost minimization problem with response time constraints. Bell and Grimson offer a more creative approach; they adopted a semantic model that doesn't even consider quantitative data [1993]. Each of these approaches are representative of other approaches included in the literature.

## Summary

Decision support for database administrators implementing distributed database systems needs to be improved in the areas of fragmentation and allocation schema design. Current approaches view fragmentation and allocation schema design and selection as loosely coupled decision processes. The heart of existing research approaches appears to be the objective of maximizing the locality of data in areas where that data is most frequently used. One problem with this approach is that as organizations evolve, so does the manner in which their data is utilized. For example, Hac's [1989] research focuses on algorithms for automatically shifting data and processes in distributed database systems to achieve the objective of balancing overall system load. Since load balancing helps to keep down overall distributed database system costs, this factor should be considered in fragmentation and allocation schema design. The decision support system we are developing takes other factors such as this into account.

## References

Apers, P.M.G., "Data Allocation in Distributed Database Systems." *ACM Transactions on Database Systems*. v.13, no.3, pp. 263-304, September, 1988.

Apers, P.M.G., "Schema Architectures and Their Relationship to Transaction Processing in Distributed Database Systems." *Information Sciences*. v.15, no.6, pp. 41-66, 1991.

Bell D. and J. Grimson. *Distributed Database Systems*. Addison-Wesley, 1992.

Ceri, S., S. Navathe, and G. Wiederhold. "Distribution Design of Logical Database Schemas." *IEEE Transactions on Software Engineering*. V. 9, no. 4, pp. 97-106, 1993.

Cornell, D.W. and P.S. Yu, "An Effective Approach to Vertical Partitioning for Physical Design of Relational Databases." *IEEE Transactions on Software Engineering*. v. 16, no.2, pp. 248-258, 1990.

Dowdy, L.W. and D.V. Foster. "Comparative Models of the File Assignment Problem." *ACM Computing Surveys*. v.14, no.2, pp.287-313, June 1982.

Hac, A. "A Distributed Algorithm for Performance Improvement Through File Migration." *IEEE Transactions on Software Engineering*. v. 15, no.11, pp. 1459-70, 1981; reprinted 1989.

Meghini, C, and C. Thanos. "The Complexities of Operations on a Fragmented Relation." *ACM Transactions on Database Systems*. v.16., no.1, pp. 56-87, March 1991.

Navathe, S., S. Ceri, G. Wiederhold, and J. Dou. "Vertical Partitioning Algorithms for Database Design." *ACM Transactions on Database Systems*. v. 9, no.4, pp. 680-710, December, 1984.

Ozsu, M. T., and P. Valduriez. *Principles of Distributed Database Systems*. Prentice Hall, 1991.