


COMPUTER CORPORATION OF AMERICA
CAMBRIDGE A DISTRIBUTED DATABASE
MANAGEMENT SYSTEM COMMAND AND
CONTROL APPLICATIONS.

BIBLIOTHEQUE DU CERIST



IST
680

A Distributed Database Management System for Command and Control Applications: Semi-Annual Technical Report 3

C 68

SCH. E. 076

BIBLIOTHEQUE DU CERIST

Technical Report
CCA-78-10
July 30, 1978




16 SEP. 1979

CENTRE DE RECHERCHES SCIENTIFIQUES
ET TECHNIQUES ET DE TRANSFERTS
TECHNOLOGIQUES
B.P. 315 ALGER - GARE

Computer Corporation of America,
575 Technology Square
Cambridge, Massachusetts 02139

A Distributed Database Management System
for
Command and Control Applications
SEMI-ANNUAL TECHNICAL REPORT III
January 1, 1978 to June 30, 1978



This research was supported by the Defense Advanced Research Project Agency of the Department of Defense and was monitored by the Naval Electronic System Command under Contract No. N00039-77-C-0074. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

BIBLIOTHEQUE DU CERIST

Table of Contents

1. Introduction	2
2. SDD-1 Design	6
2.1 Overview	6
2.1.1 Introduction	6
2.1.2 Data Modules	7
2.1.3 Transaction Modules	9
2.1.4 Distributed Data Organization	11
2.1.5 Directory Management	14
2.2 Analysis of Concurrency Control in SDD-1	15
2.2.1 Introduction	15
2.2.2 A Formal Model of SDD-1	17
2.2.2.1 Introduction	17
2.2.2.2 Database Designs and Execution Histories	18
2.2.2.3 Admissible Execution Histories	22
2.2.3 Serializable Histories	23
2.2.3.1 Consistency and Serializability	23
2.2.3.2 Equivalence	24
2.2.3.3 Serializability	27
2.2.4 Well-Behaved Execution Histories	31
2.2.4.1 Time-ordered Serialization Graphs	31
2.2.4.2 The SDD-1 Synchronization Mechanism	34
2.2.4.3 <u>Proof of Serializability</u>	39
2.2.5 Serializability of Logical Transactions	51
2.3 Initial SDD-1 Implementation	54
2.3.1 System Architecture	55
2.3.1.1 User Interface	61
2.3.1.2 Relational Query Generation	61
2.3.1.3 Fragment Transformation	62
2.3.1.4 Access Planner	64
2.3.1.5 Move Manager	65
2.3.1.6 Final Query Processing	67
2.3.2 Distributed Query Processing	67
2.3.3 The Demonstration Setup	71
2.4 Reliable Broadcast	87
3. Enhancement of Datamodule 1	113
3.1 Study Procedure	114
3.2 Initial Model Definition	115
3.3 Initial Measurements	119
3.4 Future Work	126
References	127



BIBLIOTHEQUE DU CERIST

Project Staff Members:

B. BERKOWITZ

P. BERNSTEIN

S. FOX

N. GOODMAN

M. HAMMER

T. LANDERS

C. REEVE

J. ROTHNIE

S. SARIN

D. SHIPMAN

1. Introduction

This report summarizes the third six month period of a project entitled, "A Distributed Database Management System for Command and Control Applications" which has been undertaken by CCA and sponsored by ARPA-IPTO. The primary focus of this effort is to design and implement a distributed database management system called SDD-1 (System for Distributed Databases). SDD-1 is specifically oriented toward command and control applications and will be installed in phases and tested in the Advanced Command and Control Architectural Testbed (ACCAT) at the Naval Ocean Systems Center (NOSC) in San Diego.

The motivation behind building a distributed database management system like SDD-1 is to take advantage of the decreasing cost of distributed processing environments and at the same time respond to the increasing data handling needs of geographically distributed organizations. SDD-1 permits data to be managed on a network of computers in an integrated environment that presents the user with the illusion that he is dealing with a centralized DBMS.

SDD-1 is designed to achieve other goals that are only possible in a distributed system. These goals are:

1. distributed access -- Data is accessible over the network from many sites.
2. fast response/low communication cost -- Through intelligent database design, data can be stored geographically near where it is most often used so that the majority of accesses are essentially local.
3. reliability/survivability -- Redundancy of processors, communications and data can be used to achieve very high levels of system reliability. The loss of one processor can be compensated for with others in the system and using the redundant data, the user can still run transactions that require data stored at the disabled site.
4. modular upward scaling -- Incremental addition of new sites to the system can be used to enhance system capacity without major reconfiguration of existing sites.

BIBLIOTHEQUE DU CERIST

In the process of designing and implementing SDD-1, three key technical problems have been identified. They are:

- synchronizing update transactions
- distributed query processing
- handling failures of processors and communications channels

Solutions to these problems were designed and reported during the first year of this project ([CCA a], [CCA b], [BERNSTEIN et al b], [ROTHNIE and GOODMAN] and [WONG]). These design results have been further refined and exploited with the following results:

1. A new and simpler proof of the correctness of the update synchronization algorithm was developed.
2. The distributed query processing algorithm has been implemented in an initial version of SDD-1. This version of the system has been demonstrated to ARPA with the aid of a graphic display of the algorithms behavior.
3. The set of reliability mechanisms have been further refined and an initial implementation has begun.

BIBLIOTHEQUE DU CERIST

Section 2 of this report presents an overview of the SDD-1 design and summaries of the above design results including a description of the initial working version of the system.

In addition to SDD-1 design and implementation, CCA began a study of the enhancement of datamodule 1. This project involves studying possible enhancements to the Datacomputer [MARILL and STERN] in order to make its performance more compatible with SDD-1 and other command and control applications. Section 3 summarizes the study techniques and some initial results.