

AD-A082 548

CALIFORNIA UNIV BERKELEY ELECTRONICS RESEARCH LAB  
INGRES - A RELATIONAL DATA BASE SYSTEM.(U)  
SEP 79 M STONEBRAKER, E WONG

F/6 5/2

DAA629-76-6-0245

UNCLASSIFIED

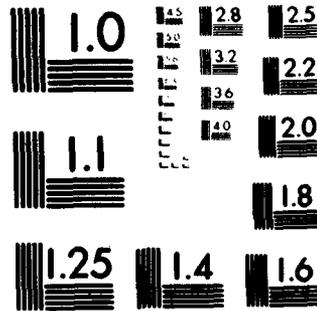
ARO-13585.14-EL

NL

1 of 1  
As  
Microfilm



END  
DATE  
FILMED  
5-80  
DTIC



MICROCOPY RESOLUTION TEST CHART  
NATIONAL BUREAU OF STANDARDS 1963-A

ADA 082548

ARO 13585.14-EL

18

19

12

6

INGRES - A RELATIONAL DATA BASE SYSTEM,

9

FINAL REPORT

1 Jun 76 - 10/31/79

34 Sep 79

10

M. Stonebraker

E. Wong

12 9

11

Sep [redacted] 79

DTIC ELECTE  
MAR 31 1980

U. S. ARMY RESEARCH OFFICE

ARO CONTRACT DAAG29-76-G-0245

15

A

Department of Electrical Engineering and Computer Sciences  
and the Electronics Research Laboratory  
University of California, Berkeley, California 94720

APPROVED FOR PUBLIC RELEASE;  
DISTRIBUTION UNLIMITED

DDC FILE COPY

80 - 3 - 26 - 070

127550



Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle)  INGRES - A relational data base system		5. TYPE OF REPORT & PERIOD COVERED final 6/1/76 - 9/30/79
7. AUTHOR(s)  M. Stonebraker		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS  Electronics Research Laboratory University of California Berkeley, CA 94720		8. CONTRACT OR GRANT NUMBER(s)  DAAG29-76-G-0245
11. CONTROLLING OFFICE NAME AND ADDRESS  U. S. Army Research Office Post Office Box 12211 Research Triangle Park, NC 27709		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE
		13. NUMBER OF PAGES
		15. SECURITY CLASS. (of this report)  Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report)  Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)  NA		
18. SUPPLEMENTARY NOTES  The findings in this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)  Data Base Management                      Data Base Design Data Base Machines                         Concurrency Control Distributed Data Bases		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)  The report summarizes the progress made during a three year period on research in data base management. The primary effort has been the design development of a major data base management system of the relational type, INGRES. In addition, a number of new directions, such as distributed data bases and data base machines were initiated.		

## DESCRIPTION OF RESEARCH FINDINGS

Under the support of the U.S. Army Research Office the INGRES project has made major advances on several different fronts during the period June 1, 1976 - September 30, 1979. We briefly indicate each in turn.

### I. Development of a Usable Relational Data Base System.

Considerable effort has gone into producing a relational data base system that can be used by others to test the viability of the relational approach. This has included making INGRES acceptably fast, reasonably free from bugs and containing needed services such as crash recovery, concurrency control and protection. It also included producing good user level documentation.

To date there are about 130 INGRES installations around the world including several military ones. These are installations at the National Security Agency, the Army Computer Systems Command, the Navy Ocean Systems Center, the Navy Personnel Data Center, the Department of Defense, the Air Force Data Services Center, the Navy Personnel Research and Development Center, the Naval Postgraduate School, the Army Research Office - AIRMICS, and the Defense Advanced Research Projects Agency

It is our conclusion that the development of INGRES as a one machine system has had considerable impact. However, there is little reason to pursue this development further; that should be left to commercial venture. A summary and critique of the design experience is contained in [STON 79].

### II. Distributed Data Bases

In late 1976 we embarked on extending INGRES to function on multiple computer systems loosely coupled together by a communications network. The problems to be overcome included:

- 1) storage and redundancy of system catalogs

- 2) query processing in a distributed environment
- 3) concurrency control
- 4) crash recovery
- 5) control of multiple copies of data

Initial thoughts in these areas were presented in [STON 77] and more recent results in [STON 79a, EPST 78]. We have embarked on building a prototype distributed data base system, and it is nearing completion at the current time. It should be mentioned that most of the difficulties we have encountered concern the communications system and the operating system software to support it; the data base code has been relatively straightforward. As such no performance data is currently available but we hope to have some soon.

### III. Data Base Machines.

In 1977 we began investigating the feasibility of a data base machine. Moreover, we examined all the existing proposals and discarded them as not particularly viable as a mechanism to increase the speed of our existing software. Hence we proposed a design in 1978 [STON 78] and a new design in 1979 [STON 79b]. The latter design is in the process of study refinement and initial coding is commencing.

The basic point we are exploiting is that data base machines are really nothing other than distributed data base systems from a software point of view. Rather increased speed is obtained by a fast network and specialization of function.

### IV. Data Base Design.

It has become evident that designing data bases (regardless of what data base system is available) is very hard. Moreover, doing designs in such a way that migration from CODASYL network system to relational systems or vice versa is harder yet. Consequently, we developed a representation

scheme and a set of rules for this scheme that preserve the possibility of mapping to either CODASYL or Relational systems. This set of rules also guarantees that many of the ambiguities that can happen with poor designs are avoided completely. This work is reported in [KATZ 79].

#### V. Concurrency Control

One of the key problems in a data base management system is to control multiple, concurrent, possibly conflicting updates. This problem is invariably solved by creating some lockable object in the data base and then establishing a lock management protocol. However, the problem remains: "How large should this lockable object be?" If it is chosen too small, there is excessive overhead getting and releasing locks. Alternately, if it is chosen too large, possible parallelism is not exploited.

In [RIES 77, RIES 79] we have completely answered this question.

#### VI. Programming Language Interface to a Data Base System

We have investigated how a data manipulation language should be embedded in a general purpose programming language [STON 77b] and some of the issues in designing a new programming language oriented toward data manipulation [PREN 78].

- STON 79  
Stonebraker, M. "Retrospection on a Data Base System" to appear  
in ACM Transactions on Data Base Systems.
- STON 77  
Stonebraker, M. and Newhold, E. "A Distributed Data Base Version of  
INGRES" Proc. 2nd Berkeley Workshop on Distributed Data Bases and  
Computer Networks, Berkeley, CA., May 1977.
- STON 79a  
Stonebraker, M. "Concurrency Control Crash Recovery and Consistency  
of Multiple Copies of Data in Distributed INGRES" IEEE Transactions  
on Software Engineering, May 1979.
- EPST 78  
Epstein, R., Stonebraker, M. and Wong, E. "Distributed Query Processing  
in a Relational Data Base System" Proc. 1978 ACM-SIGMOD Conference  
on Management of Data, Austin Tx., June 1978.
- STON 78  
Stonebraker, M. "A Distributed Data Base Machine" ERL Technical  
Report # M78/23, Berkeley, CA, May 1978.
- STON 79b  
Stonebarker, M. "MUFFIN: A Distributed Data Base Machine" Proc. 1st  
International Conference on Distributed Computing, Huntsville, Al.,  
October 1979.
- KATZ 79  
Katz, R. and Wong, E. "Logical Design and Schema Conversion for  
Relational and DBTG Data Bases" Proc. Internation Conference on the  
Entity Relationship Approach to Systems Analysis and Design, Los  
Angeles, CA., December 1979.
- STON 76  
Stonebraker, M. et. al. "The Design and Implementation of INGRES"  
ACM Transactions on Data Base Systems, September 1976.
- RIES 77  
Ries, D. and Stonebraker, M. "Effects of Locking Granularity in a  
Database Management System" ACM Transactions on Data Base Systems,  
September 1977.
- RIES 79  
Ries, P. and Stonebraker, M. "Lock Granularity Revisited" ACM  
Transactions on Data Base Systems, March 1979.
- STON 77b  
Stonebraker, M. and Rowe, L. "Observations on Data Manipulation  
Languages and Their Embedding in General Purpose Programming Languages"  
Proc. 1977 Very Large Data Base Conference, Tokyo, Japan, October 1977.

**PREN 78**

Prenner, C. and Rowe, L. "Programming Languages for Relational Data Base Systems" Proc. National Computer Conference, New York, New York, June 1978.

**VII. Publications.**

- a) "Observations on Data Manipulation Languages and Their Embedding in General Purpose Programming Languages" by Michael Stonebraker and Lawrence Rowe, Proc. 1977 Very Large Data Base Conference, Tokyo, Japan, October 1977.
- b) "Distributed Query Processing in a Relational Data Base System" by Robert Epstein, Michael Stonebraker and Eugene Wong, Proc. 1978 ACM-SIGMOD Conference on Management of Data, Austin, Texas, June 1978.
- c) "Query Processing in a Relational Data Base System" by Karel Youssefi and Eugene Wong, Proc. 1979 Very Large Data Base Conference, Rio de Janeiro, Brazil, October 1979.
- d) "Retrospection on a Data Base System" by Michael Stonebraker, to appear in ACM Transactions on Data Base Systems.
- e) "The Design and Implementation of INGRES" by Michael Stonebraker et. al. ACM Transactions on Data Base Systems, September 1976.
- f) "Effects of Locking Granularity in a Data Base Management-System" by Dan Ries and Michael Stonebraker, ACM Transactions on Data Base Systems, September 1977.
- g) "Lock Granularity Revisited" by Dan Ries and Michael Stonebraker, ACM Transactions on Data Base Systems, March 1979.
- h) "Programming Languages for Relational Data Base Systems" by Charles Prenner and Lawrence Rowe, Proc. National Computer Conference, New York, New York, June 1978.
- i) "Logical Design and Schema Conversion for Relational and DBTG Data Bases" by Eugene Wong and Robert Katz, Proc. International Conference on the Entity Relationship Approach to Systems Analysis and Design, Los Angeles, CA, December 1979.

**VIII. Degrees Awarded**

Robert Epstein	Ph.D. Candidate in Philosophy	(November 1979) (August 1978)
Randy Katz	Candidate in Philosophy Ph.D. expected	(October 1978) (June 1980)

Douglas Tsui	B.S.E.	(June 1977)
Christopher Zegelin	B.S.E.	(June 1978)
Paula Hawthorn	Ph.D. Candidate in Philosophy	(September 1979) (July 1978)

**LATE  
LME**