CONTRASTING OLCP WITH OLTP: WHAT IS IT?

DATABASE WORLD '90, SANTA CLARA, CA.

David McGoveran

July 19, 1990

ABSTRACT

Many companies need OLTP and database management systems that can support it. Nonetheless, OLTP applications are only a special, albeit historically difficult, category of applications to support with a relational database management system. As the business community moves to respond to a world market with needs that change in times as short as minutes, companies will not be able to compete without what I call OLCP: On-line Complex Processing. The database management systems which support OLCP applications will have to be based on the relational model or some extension of it. OLCP may well replace the role played at present time by OLTP.

This seminar defines OLCP, discusses the requirements for an OLCP database management system, and documents a set of rules that can be used to measure how well a DBMS meets the needs of OLCP environments. These topics are discussed in contrast to OLTP. This list of objectives is not meant to be exhaustive, nor are they meant to provide a "scoreboard" for DBMS evaluations. However, they can be used to guide the DBMS selection process and to provide a foundation on which to build further evaluation criteria.

"WHY INTRODUCE OLCP VS. OLTP?" I

- OLTP DBMS REQUIREMENTS MAY CONFLICT WITH OLCP REQUIREMENTS
  - Colin White, "OLTP: What is it?", InfoDB, Spring, 1989
  - Database Product Evaluation Report Series
    - LESSONS LEARNED
- WHY IS HIGH PERFORMANCE AN RDBMS REQUIREMENT?
  - 12 DB/CR TPS SUFFICIENT FOR 90 PERCENT OF ALL OLTP APPLICATIONS
    -- THE ABERDEEN GROUP, BOSTON.
  - WHO NEEDS PERFORMANCE AND FOR WHAT?
    - OLCP COSTS ARE NOT ADDRESSED!

"WHY INTRODUCE OLCP VS. OLTP?" II

- A NEED TO CHARACTERIZE OLCP
  - D. McGoveran, "OLCP: Beyond OLTP", InfoDB, Summer, 1990
  - SUITABLE DBMSs
  - DATABASE TRANSACTIONS
  - DIFFERENTIATE OLCP FROM OLTP

"WHY INTRODUCE OLCP VS. OLTP?" V

- END-USER TRAINING COSTS

  1. TIME
  2. MONEY

- SMARTER USER INTERFACES

- SMARTER APPLICATIONS

  1. APPLICATION DRIVES THE END-USER

     NOT

  2. THE END-USER DRIVES THE APPLICATION

     ■ MINIMIZE AND SPECIALIZE HUMAN INVOLVEMENT

"WHY INTRODUCE OLCP VS. OLTP?" VI

- CONCLUSION: THE FOLLOWING WILL CONTINUE TO GROW RAPIDLY

  ■ PERFORMANCE REQUIREMENTS

  ■ TRANSACTION VOLUMES

  ■ TRANSACTION COMPLEXITY

  ■ AVERAGE DATABASE SIZES

  ■ AVERAGE DATABASE SCHEMA COMPLEXITY

ON-LINE COMPLEX PROCESSING: AN EXAMPLE III

- PASSIVE VS. ACTIVE MAINTENACE

  1. NORMALIZATION
  2. INTEGRITY
  3. PHYSICAL OPTIMIZATION

     - THE MANUFACURING ANALOGY

- HIGHLY COMPETITIVE

- NEW PRODUCT DESIGN

- RAPID EVOLUTION

ON-LINE COMPLEX PROCESSING: AN EXAMPLE IV

- VOLUME OF TRADE DRIVE PROFITABILITY

- MASSIVE AMOUNTS OF DATA

  1. COLLECTED
  2. QUERIED
  3. ANALYZED
  4. UPDATED
  5. AVAILABLE
  6. RECOVERABLE

- WITHOUT SACRIFICING INTEGRITY

DEFINING ON-LINE TRANSACTION PROCESSING I

WHAT IS OLTP?

- PERFORMANCE RANGE OF 5-1000 TPS
- SIMPLE STATEMENTS:
  - FEW TABLES AFFECTED
  - FEW COLUMNS UPDATED
  - FEW COLUMNS/TABLE
  - NARROW COLUMNS
- RECORD-AT-A-TIME UPDATES/QUERIES
- "THIN" TRANSACTIONS:
  - FEW STATEMENTS PER TRANSACTION

DEFINING ON-LINE TRANSACTION PROCESSING II

- HIGH AVAILABILITY ESSENTIAL
- SOPHISTICATED RECOVERY AND TRACKING
- A FEW LARGE (I.E. DEEP) VOLATILE TABLES
- BATCH REPORTING
- OPTIONAL BATCH UPDATING
- RELATIVELY STRAIGHT-FORWARD INTEGRITY CONSTRAINTS
  - MASTER/DETAIL OR MASTER/DETAIL/DETAIL
  - FEW REFERENTIAL CYCLES
- RELATIVELY STABLE SCHEMA AND DATA MODEL

DEFINING ON-LINE COMPLEX PROCESSING III

- ON-LINE DECISION MANAGEMENT (OLDM) REQUIREMENT
- HIGH AVAILABILITY
- COMPLEX INTEGRITY CONSTRAINTS
- HYBRID ENVIRONMENT
  - DECISION SUPPORT
  - INTERATIVE AD-HOC QUERY
  - OLTP
  - BATCH OPERATIONAL PROCESSING
  - BATCH END-USER PROCESSING

D1-13

OLCP OBJECTIVES OVERVIEW I

<u>OLCP RDBMS FEATURE OBJECTIVES</u>

- COMPLEMENT EXISTING OBJECTIVES
- CODD'S RELATIONAL FIDELITY RULES
  - LESS TOLERANT OF DEVIATIONS FROM RELATIONAL MODEL
    - INTEGRITY
      1. ENTITY
      2. DOMAIN
      3. REFERENTIAL
      4. USER-DEFINED
    - LOGICAL INDEPENDENCE
      1. BEYOND UPDATABLE VIEWS

OLCP OBJECTIVES: PERFORMANCE I

PERFORMANCE

- RULE 1. EFFICIENT CONCURRENCY SCHEME

  - *1.1 CONCURRENCY SCHEME SUPPORTS USER-CONTROLLED LOCKING GRANULARITY WHEN READING DATA:

    1. SYSTEM WIDE
    2. BY APPLICATION
    3. BY TRANSACTION
    4. BY STATEMENT

D1-15

  - *1.1' LOCK GRANULARITY WHEN READING DATA CAN BE:

    1. RECORD
    2. PAGE
    3. TABLE
    4. PREDICATE
    5. INDEX
    6. INDEX PAGE

OLCP OBJECTIVES: PERFORMANCE II

- 1.2 UPDATE TRANSACTIONS DO NOT AFFECT THE PERFORMANCE OF READ- ONLY TRANSACTIONS, AND VICE-VERSA:

  1. MULTI-VERSION READ:
  2. "DIRTY READ":

- *1.3 CONCURRENCY SCHEME SUPPORTS USER-CONTROLLED LOCK GRANULARITY WHEN UPDATING DATA:

  1. SYSTEM WIDE
  2. BY APPLICATION
  3. BY TRANSACTION
  4. BY STATEMENT

OLCP OBJECTIVES: PERFORMANCE V

- 1.7 SYSTEM SUPPORTS MULTIPLE LEVELS OF TRANSACTION CONSISTENCY (AS DEFINED BY ANSI/ISO SQL2):

  1. LEVEL 0:
  2. LEVEL 2:
  3. LEVEL 3:
  4. LEVEL 4 (REPEATABLE READ):

D1-17

- *1.8 THE USER MAY INFLUENCE THE CONCURRENCY CONTROL MECHANISM:

  1. OPTIMISTIC:
  2. PESSIMISTIC:

OLCP OBJECTIVES: PERFORMANCE VI

- RULE 2. EFFICIENT COMMIT LOGIC
  - 2.1 NO SYNCHRONIZED DATABASE I/O DURING COMMIT:
  - 2.2 SUPPORTS GROUP COMMIT:
  - 2.3 SYSTEM COORDINATES THE COMITTING OF DATABASE CHANGES AND COMMUNICATION MONITOR MESSAGES:
- RULE 3. EFFICIENT DATABASE MANAGEMENT
  - 3.1 SUPPORTS LOOK-ASIDE BUFFERING:
  - 3.2 SUPPORTS SEQUENTIAL PREFETCH:
  - 3.3 SUPPORTS DEFERRED WRITE BUFFERING:
  - 3.4 SUPPORTS CHAINED WRITE BUFFERING:
  - 3.5 SUPPORTS PARALLEL I/O OPERATIONS:
    1. FOR WRITE PROCESSING:
    2. FOR READ PROCESSING:

OLCP OBJECTIVES: PERFORMANCE IX

- RULE 5. EFFICIENT OPTIMIZATION
  - 5.1 OPTIMIZER USES A COST MODEL AND STATISTICS:
    1. *USES COST FUNCTIONS (NOT COST INDEXES OR HEURISTICS):
    2. *IS NOT SENSITIVE TO STATEMENT SYNTAX:

D1-19

OLCP OBJECTIVES: PERFORMANCE X

- 5.2 OPTIMIZER OUTPUT IS CACHED IN MEMORY FOR RE-USE BY THE SAME OR DIFFERENT TRANSACTION:
  1. RE-USABLE BY SAME TRANSACTION:
  2. RE-USABLE BY DIFFERENT TRANSACTION:
  3. *RE-USABLE BY DIFFERENT PROCESS:

- 5.3 SUPPORTS DBMS STORED PROCEDURES:
  1. *ACCEPTS PARAMETERS:
  2. *SUPPORTS ERROR HANDLING:
  3. *MAY BE NESTED:
  4. *SUPPORTS PROCEDURAL LANGUAGE:

OLCP OBJECTIVES: AVAILABILITY I

CONTINUOUS OPERATION AND HIGH AVAILABILITY

- RULE 6. ON-LINE UTILITIES

  - 6.1 DATABASE UTILITIES CAN BE RUN WHILE OLTP SYSTEM IS ACTIVE:

    1. LOAD:
    2. BACKUP:
    3. RECOVERY:
    4. REORGANIZATION:

D1-21

OLCP OBJECTIVES: AVAILABILITY II

- 6.2 AVAILABILITY OF TABLE DATA IS NOT AFFECTED BY UTILITY OPERATIONS:

  1. OTHER TABLES REMAIN AVAILABLE DURING TABLE LOAD:
  2. OTHER TABLES REMAIN AVAILABLE DURING TABLE BACKUP:
  3. TABLE DATA CAN BE READ DURING BACKUP OPERATIONS:
  4. TABLE DATA CAN BE UPDATED DURING BACKUP OPERATIONS:
  5. OTHER TABLES REMAIN AVAILABLE DURING TABLE RECOVERY:
  6. OTHER TABLES REMAIN AVAILABLE DURING TABLE REORGANIZATION:
  7. TABLE DATA CAN BE READ DURING INDEX BACKUP:
  8. TABLE DATA CAN BE READ DURING INDEX RECOVERY:
  9. TABLE DATA CAN BE READ DURING INDEX REORGANIZATION:

OLCP OBJECTIVES: AVAILABILITY V

- 8.3 SUPPORTS RESTARTABLE UTILITIES:

  1. LOAD:
  2. BACKUP:
  3. RECOVER:
  4. REORGANIZATION:

- 8.4 SUPPORTS PARALLEL SEARCHING OF DISK VOLUMES DURING QUERY PROCESSING:

  1. SINGLE TABLE READ:
  2. MULTI-TABLE JOIN:
  3. INDEX SCAN:

OLCP OBJECTIVES: AVAILABILITY VI

- RULE 9. EFFICIENT SYSTEM RECOVERY

  - 9.1 SYSTEM HAS EFFICIENT DISK RECOVERY LOGGING:

    1. DISK LOGGING:
    2. RECORDS PHYSICAL DATA CHANGED ONLY:

  - 9.2 FULL DISK LOG AUTOMATICALLY ARCHIVED TO ARCHIVE VOLUME WITHOUT IMPACT TO SYSTEM OPERATION:

  - 9.3 NO MANUAL INTERVENTION BY THE OPERATOR IS REQUIRED TO RESTART SYSTEM FOLLOWING SYSTEM FAILURE:

OLCP OBJECTIVES: AVAILABILITY IX

- RULE 11. EFFICIENT APPLICATION RECOVERY
  - 11.1 BACKS OUT TO LAST COMMIT POINT AFTER AN APPLICATION FAILURE:
  - 11.2 SUPPORTS APPLICATION RESTART FROM LAST COMMIT POINT AFTER AN APPLICATION FAILURE:

OLCP OBJECTIVES: AVAILABILITY X

- RULE 12. FAULT-TOLERANCE
  - 12.1 SUPPORTS DUPLEXED LOGS:
  - 12.2 SUPPORTS DUPLEXED DATABASES/TABLES:
  - 12.3 SUPPORTS "HOT" STANDBY PROCESSOR:
  - 12.4 DBMS SUPPORTS SHARING OF A DATABASE BETWEEN LOOSELY COUPLED PROCESSORS FOR AVAILABILITY:
    1. DEC VAXCLUSTER:
    2. IBM:
    3. OTHER:
  - 12.5 SUPPORTS DISK MIRRORING WITH ON-LINE MIRROR RECOVERY:

D1-27

OLCP OBJECTIVES: ARCHITECTURE III

- 14.5 UNLIMITED (PRACTICAL) NUMBER OF INDEXES ON A TABLE:
- 14.6 UNLIMITED (PRACTICAL) NUMBER OF CONNECTED OR CONCURRENT USERS:
- 14.7 UNLIMITED (PRACTICAL) BUFFERPOOL SIZE:
- *14.8 UNLIMITED (PRACTICAL) NUMBER OF ALLOWED JOINS:

OLCP OBJECTIVES: ARCHITECTURE IV

- *14.9 UNLIMITED (PRACTICAL) NUMBER OF TABLES REFERENCED IN A STATEMENT OR TRANSACTION:
- *14.10 UNLIMITED (PRACTICAL) NUMBER OF EXPRESSIONS IN A STATEMENT:
- *14.11 UNLIMITED (PRACTICAL) NUMBER OF SUBQUERIES IN A STATEMENT:
- *14.12 UNLIMITED (PRACTICAL) SUBQUERY NESTING LEVELS IN A STATEMENT:
- *14.13 UNLIMITED (PRACTICAL) NUMBER OF CURSORS IN AN APPLICATION OR IN THE SYSTEM:
- *14.14 UNLIMITED (PRACTICAL) NUMBER OF CHARACTERS IN A STATEMENT:
- *14.15 UNLIMITED (PRACTICAL) NUMBER OF COLUMNS REFERENCED IN A STATEMENT:

OLCP OBJECTIVES: SYSTEM MANAGEMENT CONTROLS III

- RULE 17. PERFORMANCE MANAGEMENT TOOLS
  - 17.1 PROVIDES OFF-LINE PERFORMANCE REPORTING:
  - 17.2 PROVIDES ON-LINE PERFORMANCE REPORTING:
  - 17.3 PROVIDES RESOURCE GOVERNOR:
    1. CPU:
    2. I/O:
    3. DISK STORAGE:
    4. MEMORY:
    5. RECORDS PROCESSED:
  - 17.4 CAN CONTROL NUMBER OF LOGGED-ON AND CONCURRENT USERS:

OLCP OBJECTIVES: FLEXIBILITY I

PROCESSING FLEXIBILITY

- *RULE 18. DATA DEFINITION LANGUAGE
  - 18.1 THE USER MAY CREATE TEMPORARY TABLES WHICH ARE AUTOMATICALLY DROPPED AT THE END OF THE TRANSACTION OR SESSION:
  - 18.2 THE USER MAY CREATE DATABASE TRIGGERS:
  - 18.3 BLOBS, TEXT, AND IMAGE ARE SUPPORTED DATA TYPES:
  - 18.4 USER-CREATED OR ABSTRACT DATA TYPES ARE SUPPORTED WITH DOMAIN RULES:
  - 18.5 TIMESTAMPS ARE SUPPORTED:

OLCP OBJECTIVES: FLEXIBILITY IV

- 19.5 ERROR MESSAGES CAN BE LOGGED ALONG WITH USER-ID, APPLICATION-ID, STATEMENT, AND A TIMESTAMP:

- 19.6 A SHARABLE MEANS IS PROVIDED TO AUTOMATICALLY INCREMENT A COLUMN DURING A SET UPDATE OR INSERT (A SEQUENCER):

- 19.7 THE RESULTS OF A PARTICULAR QUERY MAY BE PROCESSED BY SUCCESSIVE QUERIES WITHOUT EXPLICITLY CREATING A TEMPORARY TABLE:

OLCP OBJECTIVES: FLEXIBILITY V

- *RULE 20. BATCH SUPPORT

  - 20.1 FAILED, IDLE, COMPLETED, OR ABORTED PROCESSES ARE AUTOMATICALLY DETECTED AND LOGGED OFF WITH APPROPRIATE RECOVERY:

    1. USER SETTABLE TIMEOUT:
    2. ERROR CONDITIONS:

  - 20.2 BUFFERPOOL SIZES MAY BE DYNAMICALLY CONFIGURED TO OPTIMIZE INPUT, INTERMEDIATE, OR OUTPUT PROCESSING:

  - 20.3 BATCH INSERT (PAGE APPEND) IS SUPPORTED:

ON-LINE COMPLEX PROCESSING

- THE POTENTIAL OF OLCP
  - CLOSING THE INFORMATION LOOP
    1. BETWEEN CUSTOMER AND FACTORY
    2. REDUCED TIME TO MARKET
    3. EFFICIENT CONTROL OF PRODUCTION VOLUMES
    4. FIELD MAINTENANCE AND REPAIR
    5. CUSTOMER SERVICE
    6. CUSTOMER SATISFACTION VIA DESIGN INFLUENCE
  - NO OFF-THE-SHELF EASY SOLUTION
  - REQUIREMENT WILL CONTINUE TO EXCEED CAPABILITIES
  - RELATIONAL DBMSs ARE A REQUIREMENT

BIOGRAPHY

- David McGoveran, Principal, Database Associates International

David McGoveran co-founded and is a principal of Database Associates, along with Colin White, Richard Finkelstein, and Paul Winsberg. He is President of Alternative Technologies, a Santa Cruz, California firm which has specialized in complex relational database applications for over ten years and which is an affiliate of Database Associates.

Mr. McGoveran designed and developed the first commercial CIM system using a relational database. He has authored numerous technical articles and has lectured around the world. Author of the forthcoming Guide to Sybase SQL Server with Chris Date, he serves as an associate editor of InfoDB magazine. Mr. McGoveran is listed in Who's Who in the World and Who's Who in the Computer Industry.