ALGEBRAIX DATA

PRODUCT VISION

INTRODUCTION

Imagine data-management products that could actually enable you to manage all enterprise data from a single, efficient platform or products that managed data equally well, regardless of its structure. Imagine products that could automatically organize your data for efficient processing and rapid response time or products built on a strong technological foundation that could support not just currently prevalent data types, but the data of tomorrow as well. Most database users have lived with the limitations of today's four-decades-old relational technology for so long that they cannot imagine such a scenario. At ALGEBRAIX™, we not only envision such a future, but we are building it today.

Algebraix Data Corporation is developing a new generation of revolutionary transrelational data-management software products that provide more-cost-effective, high-performance real-time access to all enterprise data -- no matter how or where it is stored. Patented ALGEBRAIX™ technology, based on a pioneering application of extended set mathematics, enables:

- Transparent, efficient and integrated access to all forms of enterprise information (i.e., structured and unstructured) through legacy applications
- Unlimited access to any size local or distributed data in any format
- Real-time complex-event reporting from perpetual queries processing both streaming input data and dynamically changing databases
- 10 times better cost-performance than the most commonly used products available today
- Rapid access to data through the elimination of time-consuming indexing and load operations associated with traditional databases
- Continuous performance improvement through automated self tuning
- Speed and ease of installation, integration, operation and use -- achieved via automated self-configuration and adaptive self-tuning
- Significant reductions in equipment, development, operating and maintenance costs

By modeling information with mathematical rigor, ALGEBRAIX technology brings the precision and power of mathematics to the problem of data management. The consequent boost in simplicity, performance and value represents a breakthrough in computer data management.

In this whitepaper, we will:

- Highlight the reasons why current technology can no longer meet the challenges of enterprise data management
- Explain ALGEBRAIX technology and how it enables a revolutionary new generation of enterprise data management
- Introduce our vision for a family of data-management products based on ALGEBRAIX technology

ENTERPRISE DATA-MANAGEMENT CHALLENGES

Enterprises face data-management challenges on two major fronts:

- Managing relational data
- Managing XML documents

Managing Relational Data

When the relational database model was introduced, it was a major step forward in data management. The relational data model enabled users and programmers to access a database without having to know how the data were physically stored on the computer. The system presented the data to the user in tabular form and allowed the data to be manipulated using relational operators. The relational database essentially shielded users and programs from the vagaries of how the data were actually stored in the computer.

Now, 40 years later, relational database management systems (RDBMSs) are by far the dominant data-management technology. But four decades have passed with virtually no fundamental technological advances; and, as valuable as the relational data model has proven to be, it is encumbered by a number of fundamental drawbacks that are at the source of the most daunting data-management issues facing enterprises today.

Two major problems with the relational data model are:

- The ambiguousness of the underlying mathematics
- The lack of a model for mapping data to the computer system

The relational data model was loosely based on classical set theory. However, the ambiguousness of results obtained from applying the classical relational algebra to relational tables has made modeling and optimizing RDBMSs largely ad hoc processes. The true power of mathematics can't be brought to bear, because the mathematical model is incomplete. Optimization decisions about how to structure the data, how to retrieve the data and how to partition the data cannot be readily performed by the system and instead fall largely upon the shoulders of programmers and system administrators.

While the relational data model provided a model for presenting data to the user, it did not provide a model for operating on data physically stored on disk. The model that has been adopted by all RDBMSs for storing data is a rigid row and column schema that mimics the relational user view. This organization works fairly well for updating transactional data, but is less well suited to performing analysis against it.

In a relational database, data are "trapped" in a rows and columns schema -- data loses all meaning when removed from the tables. Therefore, when one wants to manipulate data, the system must do so always in the context of tables. An enormous percentage of the resources expended by an RDBMS is spent managing these tables -- creating them, loading them, joining them, reading them

into memory, scanning them, sorting them, storing them, reorganizing them, etc. The inefficiency, overhead and excessive I/O associated with all of this table housekeeping is the fundamental reason that RDBMSs are failing to keep up with the enterprise data-management challenge.

Another important consequence of the limitations of the classical relational data model is that mappings of logical data models to physical data models, and of physical data models to machine operations and storage, must be predefined by applications programmers and database administrators (DBAs). In other words, someone must decide in advance how to arrange the data into tables, because how the data is laid out logically in tables and how those tables are then stored and manipulated by the computer system both have a massive impact on how the relational database will perform. Consequently, conventional relational databases are inherently application-specific as their design, implementation and use are dependent on the specialized schemas, table structures, table relationships and indexing schemes contrived to support the data relationships and operations required for particular applications. The lack of a mathematical model for mapping relational data and operations to physical storage render the system virtually unable to optimize itself.

Over the past three years, Fortune 1,000 companies have on average seen their data grow from 190TB to 1 petabyte (1 million gigabytes), and data at America's 9,000 midsize companies has grown from 2TB to 100TB during the same period.

"A zettabyte by 2010: Corporate data grows fifty-fold in three years," by Lucas Mearian, COMPUTERWORLD, 6 March 2007

For example, a database that is organized for efficient handling of transactions (where much of the processing occurs a row at a time) is probably not well organized for efficient data analysis (where one wants to inspect specific columns of data from numerous rows). A database that is optimized to look at one set of columns (e.g., those related to marketing) is probably not going to be efficient for doing analysis against a largely different set of columns (e.g., those related to finance or HR). Therefore, the DBA who designs and manages the system must ideally understand and anticipate how the data are going to be used before the table schema is created. The DBA must understand what kinds of transactions will be done and the kinds of analyses to which the data will be subjected. The DBA must also understand how the RDBMS will utilize system resources (CPU, memory, I/O channels and disk) to process the tables. In short, enormous amounts of human resources are devoted to designing, optimizing, tuning, indexing and reorganizing relational databases on a continuous basis. Moreover, many special-purpose schemas of the same enterprise data must typically be implemented, since a single database cannot be optimized to efficiently handle multiple divergent uses of the data. The problem becomes even more daunting when users want to do ad hoc analyses of the data, slicing and dicing the data in unpredictable ways. How does one create an environment that performs well when one does not know in advance how the data will be accessed? This same limitation impacts the ability of RDBMSs to scale with the computer systems on which they run. A related problem is that as computer systems are expanded to accommodate growth in transaction and data volumes, predetermined data schemas set up at the outset by DBAs for far smaller volumes frequently become less and less optimal.

Another area of concern is real-time data analysis. Increasingly, enterprises are looking for competitive advantage by analyzing and reacting to business data in real time. Financial institutions use real-time event processing to optimize trading and manage risk, as well as to detect fraud as it is being perpetrated. Telecommunications companies seek to optimize their rate plans in the light of up-to-the-minute analyses of the actual usage of their networks. Intelligence and law-enforcement agencies are charged with "finding the needle in the haystack" in massive real-time data flows that are gathered around the clock. As you can imagine, the relational data model presents a challenge when attempting to analyze and react to data flows in real time. Traditional RDBMSs require that data first be stored in tables before it can be queried. Obviously this is a hindrance to efficient, low-latency analyses of real-time data.

In summary, the inherent limitations of the mathematics underlying the relational data model result in databases that are:

- Unable to self-optimize the underlying data structures and, therefore, require extensive human resources devoted to schema design and data prestructuring and loading, as well as ongoing tuning and manual data reorganization;
- Highly likely to be application-specific;
- Inefficient for data analytical tasks;
- Unable to optimize utilization of computer system resources, and
- Do not scale with the computer systems on which they run.

The limitations of RDBMSs are well attested by the proliferation of special-purpose database hardware and software being introduced into the data-management marketplace to address the questions and issues raised previously:

- Object-oriented databases
- XML-"enabled" relational databases
- XML document servers
- Proprietary data appliances
- Column-oriented relational databases
- In-memory databases
- Massively parallel relational databases
- Relational database query appliances
- Special-purpose real-time databases
- Near-real-time relational databases
- Multidimensional relational databases for analytical processing
- In-house custom software

In every case, these special-purpose systems result in yet another "island of information" in the enterprise, optimized to some particular aspect of data

management but doing little or nothing to assist in integrating the information assets of the enterprise. Also, very little in the way of new technology is employed in these products -- they merely optimize certain aspects of the existing technology in order to overcome constraints encountered when making a particular use of the data.

Managing XML Documents

Vast amounts of enterprise data reside outside of relational databases. Increasingly, enterprises are looking to XML as the means of accessing the data stored in unstructured form within office documents and other formats. The rapid adoption of XML brings with it the realization that mushrooming islands of XML documents must be effectively integrated with the vast and growing islands of legacy data managed by RDBMSs. In fact, most e-commerce applications and the much-heralded and much-sought-after enterprise information portals demand integration of XML documents and relational data.

> A database that could truly marry the two worlds so that they could be managed and accessed uniformly is, to many, a database nirvana. ... John Parkinson, CTO of the Americas at Capgemini, agrees that there is potentially a big market: "There is increasing demand for such things because while the amount of SQL data goes up linearly, the amount of non-SQL data grows exponentially."
>
> As reported in "IBM Looks to 'Viper' Database to Combat Oracle, Microsoft," by Barbara Darrow, CRN, 20 May 2005

Given the enormous infrastructure and investment in enterprise RDBMSs, the natural first inclination of enterprise information executives is to extend the role of RDBMSs to that of managing all unstructured data in addition to structured data. Unfortunately, the limitations of the relational data model once again stand in the path of progress. If an RDBMS is to be a universal data-management system, then all enterprise information, including XML documents, must be mapped into relational data tables before it can be processed. However XML documents cannot efficiently be mapped into two-dimensional row-and-column structures. Such mappings require programmed conversions and redundant storage, which add substantial costs and overhead. Consequently, relational database technology is not cost-effective for managing XML documents.

Even after spending more than $300 billion for data-management software and perhaps a trillion more over a 10-year period for programmed format conversions and redundant data storage, fully 88% of unstructured enterprise information would not be centrally accessible or managed. In fact, given these assumptions, a 10-fold increase in data-management costs would be required to manage all enterprise information with relational databases, which is clearly not practical.

ALGEBRAIX TECHNOLOGY

Enterprises simply can't afford to manage all enterprise information without significant reductions in data-management costs from current levels. Clearly, a new and superior data- management technology is needed -- one that is capable of managing all forms of enterprise information while dramatically reducing data-management costs. ALGEBRAIX technology, based on a uniquely powerful application of extended-set processing (XSP) algebra, is just such a technology.

> "Mathematics is the tool specially suited for dealing with abstract concepts of any kind, and there is no limit to its potential. " -- Paul Adrien Maurice Dirac, English physicist

Algebraix Data Corporation is the only company that has taken a holistic approach to enterprise information management by redefining the techniques used to model, access and process enterprise information. The ALGEBRAIX approach is made possible by the ingenious use of extended set mathematics and XSP to formally model, process and transform all forms of enterprise information. Algebraix Data Corporation is unique in its use of this revolutionary technology and has filed numerous patent applications to date with many more to be filed. Algebraix Data products exploit XSP to address the important data-management challenges facing enterprises today.

Conventional relational data-management systems are only partially and loosely based on mathematics; only vaguely describing the relationships between rows, columns and tables. The primitive relational algebra upon which conventional relational data-management systems are based says nothing about set membership, ordering or the changes in information that occur with time.

By contrast, ALGEBRAIX technology has a rigorous mathematical foundation which enables accurate, complete and unambiguous representations of all types of computer data, as well as the software and hardware platforms in which they exist. Rigorous mathematics enable ALGEBRAIX databases to implement several key technologies, each of which delivers the specific advantages identified in the following table.

ALGEBRAIX Technology	Advantages
Data and System Models The systems accurately models both data and computer systems, enabling query responses that are optimized for both data-access patterns and supporting hardware platforms.	• Dynamically optimizes query performance on a continuous basis • Eliminates manual tuning • Adapts to changing data, workloads and system resources • Minimizes processing of unneeded data

ALGEBRAIX Technology	Advantages
Adaptive Data Restructuring The system continuously adapts the way data is stored, based upon how it is being used. The system utilizes spare resources for background processing to optimize performance for recognized query patterns.	• Adds a whole new dimension of query optimization. Traditional systems can only optimize how they store and retrieve data in rigid, fixed structures. ALGEBRAIX technology optimizes data structures as well. • Eliminates the need to manually predetermine data structures and indexes • Eliminates the need for ongoing tuning, manual data restructuring and reloading • Minimizes data transfers between memory and storage -- the critical obstacle to better database performance
Expression Manipulation ALGEBRAIX technology enables the recognition of mathematically equivalent but computationally more efficient solutions to queries. The system utilizes spare resources for background processing to optimize performance for recognized query patterns.	• Faster query response • More efficient use of CPU resources • Reduced I/O bandwidth requirement
Algebraic Substitution ALGEBRAIX technology enables the recognition of mathematically equivalent substitutes for some are all of a query, enabling simplified calculations and/or the reuse of prior query results.	• Enables prior query results saved in secondary storage to be efficiently utilized, directly or indirectly, to answer related current queries. The system is able to unambiguously determine validity of such reuse. • Improves query response time and system throughput for queries closely related to queries already performed • Enables the system to deliver performance gains to equivalent those of multidimensional databases without the need to prestructure data into “cubes”

ALGEBRAIX Technology	Advantages
Temporal Invariance ALGEBRAIX technology employs nondestructive recording of data updates to ensure that, once it is stored, data remains in the system unaltered unless it is intentionally removed.	• Enables the ability to query the data as it existed at some specific point in time in the past • Streamlines housekeeping associated with making changes to the database
Information Equivalence ALGEBRAIX technology recognizes when two data structures represent the same information and are mathematically equivalent.	• Enables highly-structured representations of both relational data and XML documents • Enables efficient real-time dynamic access to both relational data and XML documents • Simplifies process of providing dynamic bidirectional bridging of relational and XML data

The technological advantages outlined in the previous table enable Algebraix Data Corporation to deliver radically superior data-management products with clear business benefits, including the following.

- **More-flexible access to enterprise information**
  ALGEBRAIX technology is equally adept at managing relational data and XML documents and greatly simplifies the management of data across those two domains.
- **Rapid and simple Implementation**
  ALGEBRAIX technology eliminates the need for experts to prestructure and index data before users can begin to access it.
- **Lower total costs of ownership**
  ALGEBRAIX technology vastly reduces the human resources required to organize and manage data repositories, both up-front and throughout the system life. Also, the improved efficiency of ALGEBRAIX data management enables more data to be managed with fewer human and system resources than in existing technologies.

ALGEBRAIX DATA PRODUCT VISION

ALGEBRAIX technology provides the foundation for a new generation of superior data-management products, beginning with our first offering, the Advanced Analytic Database (A2DB). This discussion of product vision is organized in two sections, the first describing A2DB and the second providing a high-level overview of the multiple dimensions of product evolution that are planned for the future of ALGEBRAIX technology.

Advanced Analytic Database

The first product to leverage ALGEBRAIX technology for data management is the Algebraix Data Advanced Analytic Database (A2DB). A2DB is a relational database designed specifically for business intelligence (BI) and data-mining applications. It runs on standard 64-bit Windows and Linux servers and storage. A2DB provides rapid access to data via industry standard SQL and ODBC programming interfaces. Users interact with A2DB using their existing BI and SQL query tools.

Benefits of the Advanced Analytic Database

A2DB exploits the power of ALGEBRAIX technology to eliminate the critical limitations of traditional RDBMSs in data-mining environments:

- The time and DBA resources required to plan, design, implement or modify a data mart – i.e., a dedicated RDBMS for handling queries and analyses vs. transactions
- The time required to load data into an RDBMS prior to being able to query it
- The need for ongoing resources to manage the data mart and tune its performance
- Inefficient query processing
    - Poor price performance
    - Inability to produce results in a reasonable timeframe, particularly for databases of 1TB or more. This situation is often alleviated only through use of complex, proprietary hardware configurations.
- Unacceptable performance to users with dissimilar queries or to users doing unpredictable, ad hoc analyses

A2DB interacts with users and applications just like every traditional relational database on the market, but without these critical limitations. ALGEBRAIX technology eliminates the need for prestructuring, manual indexing and loading relational tables, making data immediately accessible for query. A2DB generates and accumulates mathematical metadata about the queries presented and the data being mined to optimize response to similar queries, with no need for manual restructuring or indexing. When new queries are encountered and used repeatedly, A2DB generates and accumulates additional metadata which are combined with all previously accumulated metadata to further optimize response.

As described in the previous technology discussion, this increased efficiency results from the fact that A2DB:

- Adapts data structures to queries, in addition to adapting queries to data structures
- Reuses any applicable parts of previous responses to similar queries
- Selects mathematically optimal processes for query processing
- Employs background processing to further optimize responses to anticipated queries
- Minimizes processing of unneeded data

Over time, this means that, much like today's "multidimensional" databases, A2DB can almost instantly answer commonly asked questions, such as, "What were the total sales of Product B in the Western Region during the first quarter?" However, in multidimensional databases, the commonly asked questions must be identified in advance and the structure of the multidimensional "cube" has to be predetermined to support them. If the commonly asked questions change, the multidimensional database will not have the answers at the ready, whereas A2DB will -- without DBA intervention.

The bottom line is that by massively reducing the need for configuration and tuning, and by optimizing the use of computer resources, A2DB delivers faster, more flexible access to enterprise analytical data at significantly lower total costs of ownership.

Potential Applications

The ability of A2DB to optimize the use of system resources through algebraic optimization and adaptive data restructuring make it an excellent platform for enterprises who are challenged by the size and complexity of their data-mining and business-intelligence queries. A few examples of the many business applications include:

- Telecommunications
  - Rate and routing optimization
  - Customer retention
  - Billing quality
- Government
  - Intelligence analysis
  - Fraud detection
- Marketing
  - Behavior-based targeting
  - Clickstream analysis
  - Customer loyalty

- Health sciences
  - Research and discovery
  - Brand management
  - Compliance and regulatory

b) Information Technology
  - Data warehouse prototyping
  - Rapid deployment of data marts, particularly in dynamic environments with many and frequently changing sources of data

The ability of ALGEBRAIX technology to eliminate the labor-intensive processes associated with prestructuring data and ongoing performance tuning make it a natural choice for enterprises that desire to broaden the implementation of BI into new business areas, but have not yet done so due to:

- Constraints in availability of database technical expertise
- High life-cycle costs of maintaining data warehouses and data marts

Future Products

Algebraix Data Corporation will introduce a series of progressively more powerful data-management products that exploit ALGEBRAIX technology to address specific and critical market needs. As shown in Figure 1, the inherent capabilities of ALGEBRAIX technology enable three primary dimensions of product enhancement and development:

- Additional capabilities and features for the enterprise relational data environment
- Capabilities and features which embrace unstructured (e.g., XML documents) as well as relational data
- Additional platform implementations for scalability

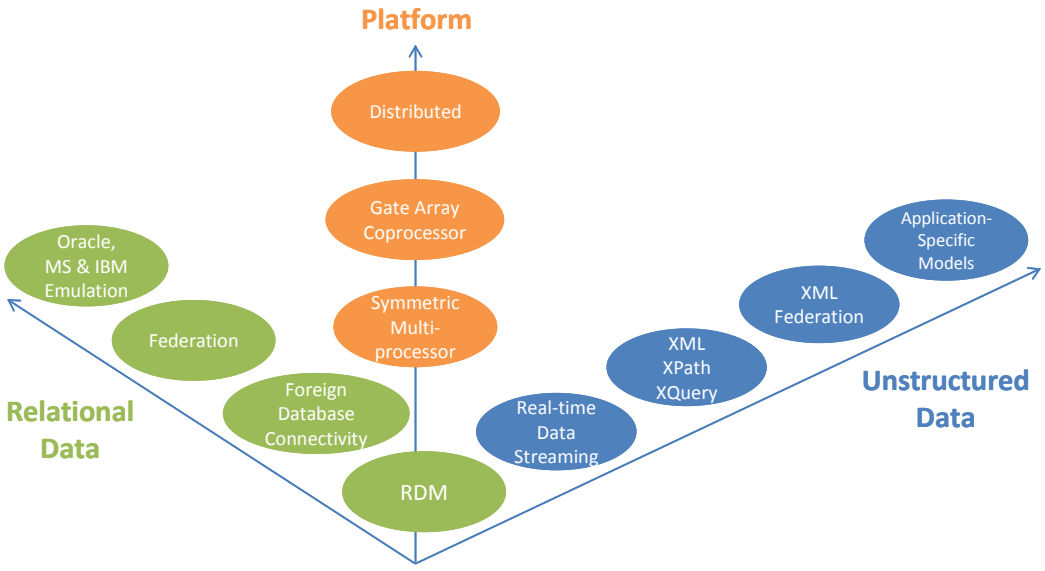

Figure 1. ALGEBRAIX Technology Directions

Future products, based upon market needs, will leverage improvements in multiple dimensions. Any given product may include enhancements in platform implementations, relational data handling and support for unstructured data. Particularly exciting are the unique advantages of ALGEBRAIX technology in enabling products that provide simultaneous real-time access to both dynamically changing relational data and XML documents.

Additional Capabilities and Features for the Enterprise Relational Data Environment

In the dimension of relational data, there are numerous enhancements that will be introduced. Three of the most-important will be:

- Direct database-database connections
- Oracle, Microsoft SQL Server and IBM DB2 emulation
- Distributed relational databases

Direct Database-Database Connections

The first release of A2DB will support what is essentially a file-oriented process for transferring data and data updates to A2DB from legacy databases. This highly-useful extension for integrating enterprise relational database will enable A2DB to issue SQL queries directly to other RDBMSs in the enterprise to download data as required. This feature will enable federation of enterprise databases at greatly reduced cost. Assuming that A2DB users have the needed authorizations to access data on other enterprise databases, A2DB will be able to retrieve the needed data from the other databases and optimally organize it in its repository based upon users' queries -- both initially and continuously.

Oracle, Microsoft SQL Server and IBM DB2 Emulation

The RDBMS market is dominated by Oracle, Microsoft and IBM. All except the most-disciplined of IT shops have made use of nonstandard features in the RDBMS products of these vendors. The ability of A2DB to emulate the popular proprietary extensions and features of the three leading RDBMS vendors will greatly simplify the process of integrating A2DB into the enterprise relational database environment and greatly enable a smooth transition from outdated relational databases to ALGEBRAIX technology. Oracle, SQL Server and DB2 emulation will enable an ALGEBRAIX database to interface directly with legacy applications and tools that employ nonstandard proprietary queries, calls and formats.

Distributed Relational Database

Patented ALGEBRAIX technology is highly amenable to physical distribution across physical nodes. The challenges of implementing a truly distributed database architecture are greatly diminished in comparison to traditional relational databases.

Capabilities and Features for Unstructured Data

In the dimension of unstructured data, there will be a series of enhancements for managing data streams and XML documents alongside relational data. Four key enhancements will be:

- Real-time data streams and complex-event processing
- XML, XPath and XQuery support
- XQuery federation
- Integration into the end-user platform

Real-Time Data Streams and Complex-Event Processing

The inherent technological advantages of ALGEBRAIX technology in handling real-time data will enable breakthrough performance, efficiency and flexibility in real-time event-processing applications. In fact, the primary effort associated with launching a real-time enabled ALGEBRAIX technology product to the market will be the development of an industry-standard real-time query language, such as StreamSQL (a query language that extends SQL with the ability to process real-time data streams). In addition to the benefits of any ALGEBRAIX data-management system (outlined above), a real-time enabled ALGEBRAIX product will provide the most cost-effective means possible for integrating real-time data feeds and large relational databases. Unlike other real-time systems, the Algebraix Data product will be capable of not only querying across both a data stream and a large relational database, but it could do so while continuously updating the database.

Real-time stream processing will enable Algebraix Data to address significant and growing applications in the following areas:

- Fraud detection

- Video surveillance
- Network security monitoring and protection
- Network health and performance management
- Intelligence analysis
- Securities trading
- Risk management

XML, XPath and XQuery Support

The time is fast approaching when the majority of enterprise information will reside not in relational tables but in XML documents. Enterprises need a way to dynamically access and update common data elements from either venue. Whether the underlying data are requested via an XML-enabled Web page or via an SQL query, the answer should be the same. Current approaches basically manage structured and unstructured data in separate domains (whether under the covers of a "single" data-management system or not). Users of such systems must go to great pain and expense to build bridges between their incompatible islands of relational data and XML documents. Adding support for data stored as XML documents will address the critical need for a unified data-management system for handling both structured and unstructured enterprise data.

ALGEBRAIX technology is inherently superior for handling unstructured data. As all data are represented mathematically, there is no inherent difference in manipulating relational data or XML documents, or any combination. Because ALGEBRAIX technology is based on a mathematically-rigorous data model, it can manipulate and interrogate any data structure.

ALGEBRAIX technology provides more-cost-effective, high-performance ***simultaneous*** real-time access to both dynamically-changing relational data and XML documents via industry-standard SQL and XQuery programming interfaces. By exploiting the unique capabilities of ALGEBRAIX technology, it is possible to store relational data and XML documents in a single integrated repository while providing automatic bidirectional run-time conversions between them — effectively, efficiently and without substantial programming effort. By eliminating most, if not all, of the programming now required, Algebraix Data products will be easier to use and dramatically more-cost-effective than the middleware and relational databases in use today.

A unitary system that is equally adept at handling relational data and XML documents, and facilitating the dynamic movement of information between these two structures, will be a natural alternative to current enterprise solutions that include:

- XML-only data repositories
- Enterprise information portals
- Home grown data transformation programs and scripts, which are expensive, time-consuming and difficult to maintain

XQuery Federation

Implementation of XQuery federation will enable users of the Algebraix Data XML-enabled data-management system to locate, retrieve and process XML documents stored on other XQuery-enabled data repositories (whether Algebraix Data or not) located anywhere in the enterprise.

The ability to interchange information with other XML document repositories will greatly speed implementation and enable a migration path to Algebraix Data as the authoritative repository for both XML documents and relational data in the enterprise.

Peer-to-Peer Implementations

The implementations of ALGEBRAIX technology discussed to this point have been server-based (whether centralized or distributed). Perhaps the most powerful breakthrough enabled by ALGEBRAIX technology will come when end users install ALGEBRAIX technology on their desktops to enable them to access any data, anytime, anywhere.

Peer-to-peer implementations are made possible by the inherent “distributability” of ALGEBRAIX technology and will enable all types of data on any Algebraix Data system to be accessible to authorized users and applications on any cooperating Algebraix Data system. No longer will vast amounts of enterprise information be “trapped” in end-user systems, inaccessible to all but the individuals who “own” those systems.

Scalability

The simplicity and mathematical optimizations that are inherent in ALGEBRAIX technology enable a high degree of parallelism. The elimination of the need to map and manage ad hoc data and index structures and the ability of ALGEBRAIX technology to accurately model system resources together enable Algebraix Data systems to make optimal use of all system resources. Because ALGEBRAIX technology adheres to a rigorous mathematical model, it is able to employ the distributive and commutative properties of its algebra to optimally distribute processing across system resources. Conceptually, multiple distributed Algebraix Data systems operate as a single system spanning multiple nodes of computing networks. This is in stark contrast to relational systems which, through lack of a mathematically-rigorous model, are relegated to ad hoc or procedural approaches to exploiting parallel system resources.

Consequently, ALGEBRAIX technology offers enormous scalability potential, exploiting multiple scaling approaches:

- **Clustering implementation**
  A cluster configuration is a simple means of scaling A2DB systems beyond the physical limitations of a single server (processor, memory and I/O processing).
- **Massively-parallel implementation**
  The efficiency of MPP integration allows for higher levels of processing efficiency than cluster configurations.
- **Exploitation of reconfigurable computing technology**
  In tandem with the approaches above, the exploitation of reconfigurable computing technology, such as the application of field-programmable gate arrays (FPGAs), promises orders of magnitude throughput gains as XSP operations are implemented by special-purpose coprocessors. Due to the underlying mathematics, the simplicity of the calculations lends itself readily to such implementations.

CONCLUSION

Algebraix Data Corporation is introducing a fundamentally new technology that has both a practical immediate benefit to enterprise RDBMSs users and a revolutionary long-term promise for streamlining enterprise data management. By delivering products that coexist in the legacy environment, Algebraix Data makes the adoption of its products virtually pain-free. Ultimately, Algebraix Data will deliver more-cost-effective, high-performance real-time access to all enterprise data -- no matter how or where it is stored. This Algebraix Data vision stands in stark contrast to the lack of innovative solutions being proposed by the incumbent giants of the enterprise data-management industry. The soundness, simplicity and superiority of ALGEBRAIX technology enable Algebraix Data to pursue this vision with confidence.◊

Algebraix Data Corporation's patented software is disrupting the entire BI complex. Its **A²DB** advanced analytic database enables real-time data access, eliminates performance tuning, and runs on affordable commodity hardware.

For more information about Algebraix Data Corporation, call 858.200.7215 or visit our Website at www.algebraixdata.com.

4350 Executive Drive, Suite 308
San Diego, CA 92121-2118
858.200.7215
www.algebraixdata.com
2.8.2010