# Flexible Integration of Manufacturing Databases Using a Model-Based Architecture Capable of Wide-Area Networking

Yicheng Tao[1]
Gilbert Babin[2]
Cheng Hsu[3]

July 1996

[1] Department of Decision Sciences and Engineering Systems, Rensselaer Polytechnic Institute, Troy, NY, 12180-3590.
[2] Assistant Professor, Départment of d'informatique, Université Laval, Québec City, Québec, Canada, G1K 7P4.
[3] Associate Professor, Department of Decision Sciences and Engineering Systems, Rensselaer Polytechnic Institute, Troy, NY, 12180-3590.

## ABSTRACT

Manufacturing enterprise systems entail both vertical and horizontal integration. Vertical integration involves different time frames of three basic cycles: the work order-oriented real time processes, the customer/demand-oriented aggregate systems, and the business strategy-oriented enterprise functions. Horizontal integration, on the other hand, is concerned with geographically distributed organizational systems, which typically require heterogeneous computing environments distributed not only over local-area networks but likely also over wide-area/global networks. Both integration require scalable building and adaptive configuration of systems, where interoperability must be accomplished without jeopardizing these systems' autonomy and concurrent processing. Furthermore, these systems may require active capabilities such as event-based information handling and fault-tolerant coordination among conventional database servers or data warehouses.

A model-based architecture providing flexible integration of manufacturing databases beyond local area networks is proposed in this paper. The concept of metadata independence for multiple systems is presented as the basis for the conceptual design of the architecture, while the ROPE (Rule-Oriented Programming Environment) method is developed to execute the design. Particular designs are developed to incorporate client-server technology (including CORBA) wherever appropriate.

## 1 AN ENTERPRISE INFORMATION MANAGER: THE OBJECTIVE

Multiple (heterogeneous, distributed, and autonomous) systems characterize distributed manufacturing enterprises. The enterprise information manager that integrates these systems (in part or in whole) needs to provide interoperability at three levels: file transfer (work flows), global query, and data and knowledge/processes management. At each level, the systems involved may require fault-tolerant, real-time, or even active capabilities in addition to the traditional database and file system operations. Furthermore, these multiple systems may have to operate over wide-area or even global networks where reliable and controllable enterprise-wide network transactions management capabilities are non-existent. To compound the complexity even further, distributed enterprises always involve dynamic changes that not only update the contents (data and operating rules) of these multiple systems, but also add or delete member systems and modify the configurations associating them. The enterprise information manager (EIM) technology must be assessed in this context. Integration which is done on a lesser scale is only accomplished by constructing special software programs which translate between individual proprietary systems. These translators are typically complex and expensive to modify in order to respond to changing business needs or improved information technologies.

The objective of an EIM is to enable manufacturing enterprises to have not only the disparate information integrated within each factory, but also to have access to relevant information across all factories, design centers, and other functional divisions. Further, the enterprises would be capable of making changes to data and have those changes automatically disseminate throughout the entire distributed organizations, which will result in having all related processes brought instantly into line. This objective can be demonstrated by using an Army logistics example which has the metadatabase installed into the configuration of Army depots and arsenals.

In the scenario, the EIM (which may be distributed) would provide a comprehensive linkage between each Army locale, thereby giving personnel the ability to search the global Army environment for particular pieces of data, and control of global events. The strength of global query capability would be matched by the ability to control events from a distance. Personnel may modify data on a global scale, such as changing the deadline in the research prototype and manufacture of a weapon system. Normally, any such scheduling change would begin a cascading effect; one Army arsenal would inform another by telephone until all were informed. Each in turn would have to determine which information systems required updating, locate the affected databases and implement the changes piecemeal, until al pertinent processes are realigned to ensure consistency. In contrast, the EIM would perform all of the above repetitious and error-prone tasks on-line. Ideally, installation of the EIM does not have to replace existing Army data structures but rather be expandable; its structure provides for computing growth within all Arsenals and Depots. This objectivity can only be achieved in steps, of course. An immediate step that this research attempts is to develop a flexible integration method applicable to EIM's under wide-area-networking (WAN) conditions. In particular, it proposes a model-based architecture with fault-tolerance to support global query, data management, and event-based processes.

The work is based on the authors' previous results [1,2,13,14,15,17]. New results which constitute the focus of this paper include a communication requirements model and a new design of the model-based architecture. These results respond to the specific needs of real-time process control [23] as well as the emerging client-server tecnologies, such as CORBA (Common Object Request Broker Architecture) and ANSI/ISO SQL3, so as to provide sufficient capabilities for active event-based processing. In addition, methods and algorithms are developed to execute the architecture. Fault tolerance is a built-in feature of the methods. Data replication is supported as a mechanism for both fault tolerance and real time process control. A prototype is implemented to verify the new execution methods. We now discuss the research problem in Section 2, and present the solution in Sections 3, 4 and 5. Conclusions are provided in Section 6.

## 2 AN ARCHITECTURE FOR FLEXIBLE INTEGRATION: THE PROBLEM

Industrial solutions tend to either be broad-based open systems, or be proprietary "one vendor" technology. The broad-based DAE (Distributed Application Environments)-class technology deals with API (Application Program Interface) and other inter-system interfacing tasks to provide file transfer-level of integration. A promising example is CORBA [21], which supports direct interoperability among various object-oriented systems and some gateway services for connecting relational or other traditional systems into an object environment. This class of enterprise information managers lacks a global model for the enterprise and typically does not include online knowledge bases (e.g., intelligent interfacing agents or rule-based shells) in its inter-system interface configuration design; both of which are strong suits of technologies that feature database-class capabilities at the expense of flexibility and openness. In between the DAE-class and the database-class technologies there exists data warehouse-class enterprise information managers. Prime examples include IBM's Information Warehouse and Hewlett-Packard's WorkManager.

Some global query functionalities (typically with build-in latency on the contents of information and limited flexibility on the types of queries supported) are provided in addition to file transfers.

The database-class technology - in particular, the various HDDBMS (Heterogeneous Distributed Database Management Systems) and multidatabases [3,6,8,9,11,22,24] - is somewhat antithetic to the other two classes of enterprise information managers in philosophy and practice. While being more analytically elegant and complete, available HDDBMS's tend to be limited to local-area networked environments and are not as robust as the other classes in dealing with tasks such as large-scale file transfers (e.g., engineering drawings) and work flow management. Moreover, due to the basic control methods employed, they tend to focus on the integration of information systems at a fixed point of time (the development time) and leave out such non-static aspects as (1) dynamically changing the (inter-system) processes and (intra-system) structures and (2) incrementally building or grouping systems (e.g., adding systems to an HDDBMS after the development time or connecting an HDDBMS to another). The two properties are referred to as **flexibility** in this work.

There are a few major concerns facing the database-class technology striving to extend flexibility for manufacturing over WAN's. We recognize the following problems from the perspective of this research:

1. Global serialization. Serialization across multiple databases has long been recognized as impractical if there is a higher probability of deadlock due to the network involved. The network problem of serialization in LAN environments might be solvable under certain conditions, but the same reliability issues become intractable in WAN and global network environments, where overall network management is totally lacking and instantaneous processing (for instantaneous correctness) is a physical impossibility due to delays.
2. Fault-tolerance. This is an acute concern for WAN-based interoperations. Since the global information processing among distributed systems might be interdependent, disruptions at a system would take place when any of those processings at other related systems fail or fail

to send results back and reach the system. Many systems could therefore be rendered inoperable frequently.

3. Real time process control. Manufacturing databases must support real time control processes to be implemented on shop floor. The key concern for this requirement is two-fold: performance (time frame) and storage. The architecture should be amenable to work with systems that do not possess mass storage capabilities. To satisfy stringent time constraints, local structure should always maximize parallel processing using self-reliant information.

A model-based architecture was provided by the metadatabase model to effect flexible integration of multiple systems. The previous results [13,14,15] do not include sufficient details to solve all of the problems discussed above. This research extends these results.

The metadatabase model proposes that (1) the enterprise metadata include contextual knowledge and be structured and managed as an independent full-fledged database - thus, the term metadatabase, (2) the inter-system operating knowledge be distributed into a concurrent shell system to effect interoperability at all three levels (file transfer, global query, and data and processes management), (3) a new criterion of correctness be employed without requiring a global synchronizer (which is optionally employed only for some mission-critical data that are designated dynamically and controlled by rules), (4) the metadatabase be commissioned as the enterprise information manager which directly manages the metadata instances at the enterprise level (global models and configuration logic), but only indirectly (through metadata) facilitates raw data processing at the local systems level, and (5) the metadatabase-defined data replication is used as a mechanism for both fault tolerant and real time process control over WAN/LAN environments. The metadatabase employs a distributed, rule-oriented shell system to interoperate with local systems, embed necessary knowledge of integration (at all three levels) into local environments, and let these local systems operate thereafter with this knowledge. It intervenes only when the enterprise metadata are changed, so as to update the localized knowledge. Based on these extended characteristics, the notion of **metadata independence** is defined to refer to the capability of an enterprise information manager being structurally neutral to multiple databases yet functionally supportive of their adaptive and scalable integration. This notion is directly comparable to the classical concept of data independence in single database with respect to multiple files.

The research extends the previous rule-oriented shell system [1,13] to support the overall architecture integrating WAN and LAN with reasonable performance by incorporating fault-tolerant capability. The major challenges and tasks include: (1) defining the rule-oriented shell system based on WAN and LAN technologies such as client-server architecture and CORBA and (2) developing execution methods with fault-tolerance and real time processing to support the inter-operation of shells across LANs and WANs.

# 3 THE MODEL-BASED ARCHITECTURE

## 3.1 The Concept

The structural model of the architecture is depicted in Figure 1 [13,14,16]. Note that the metadatabase can be distributed both horizontally and vertically [12]. Thus, the structural model shown in Figure 1 may be expanded recursively into a system of clusters of metadatabases and their local application systems. Each clusters might be based on an LAN, and the clustering of these clusters constitute the WAN environment. The major points from these citations can be summarized as follows. The metadatabase itself (a rigorously constructed collection of enterprise metadata) provides an integrated enterprise model for the multiple information systems, their databases, and the interactions among the different systems; i.e. the information contents and their contextual knowledge. The metadatabase approach through the architecture (1) uses the enterprise model to assist end-users performing global queries free of both technical details and a hierarchy of integrated schemata; (2) distributes the contextual knowledge to empower these local systems to update data and communicate with each other without central database control; and (3) incorporates legacy, new or changed local models into its generic structure of metadata to support evolution without system redesign or recompilation. The shells in the architecture, therefore, implements the distributed (localized) knowledge which, in turn, is managed by the metadatabase.

Figure 2 illustrates the laboratory environments at Rensselaer in which the model-based architecture prototype was developed.

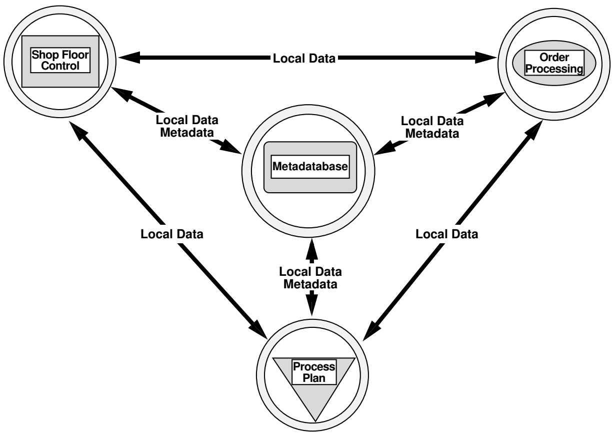


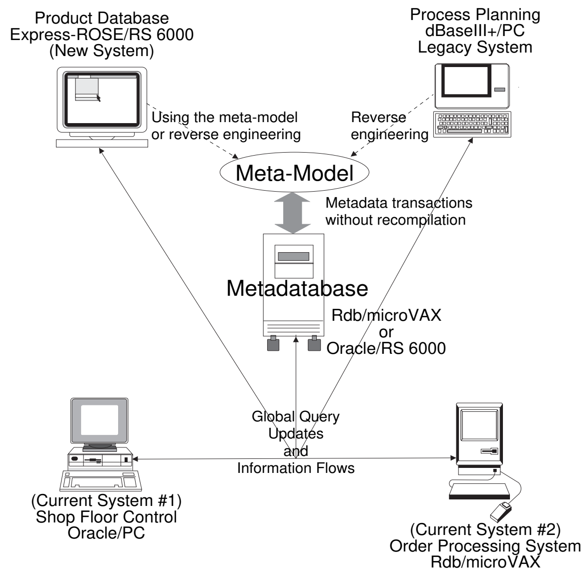


Figure 1 : The Model-Based Architecture Using a Metadatabase

Figure 2 : The Operation of the Model-Based Architecture

The metadatabase model employs a new criterion, event/usage correctness [13], for control and thereby removes the need for global serialization, which is infeasible in WAN environments, and fosters full concurrent processing among local systems. The event/usage correctness criterion is defined as follows: a data value only needs to be updated when it is needed by an application according to its contextual knowledge. Therefore, the consistency is established on a need-to-know basis, using both operating and integrity rules (also referred to as data management rules) defined in the enterprise model. This criterion eliminates the *requirement* of serialization across multiple systems but entails knowledge methods in its place. One important point to note is that the event/usage correctness criterion defines generally how to synchronize data across different platforms and databases; but does not prohibits serialization from being applied as the operating rules for certain (or all) data items in the enterprise. Concurrency control at the local level remains the responsibility of the local DBMS or the local database server at an LAN.

The new criterion reduces or even eliminates the need for a central controller, depending on the rules used. The rules can be tailored to differentially provide a variable level of consistency for different data items; while, in contrast, serialization presumes all data be equally critical and need complete synchronization at any point in time. Moreover, the rule-based schema allows for a traditional mechanism to be employed at a particular application site of the multiple systems. For data items that require traditional correctness, operating rules can be developed to route them to a serialization controller which can be included in the system on an optional basis. The global consistency as measured against the new criteria is therefore achieved in two stages: the correct modeling of the control rules to represent the logic correctness, and then the execution model to ensure the accurate implementation of these rules and hence the logical correctness.

The architecture facilitates flexibility. Instead of having the applications changed to fit a control method (e.g., serialization), we use the knowledge about the model (i.e., operating rules and integrity rules) in the metadatabase to build customized control shells around each application. These shells can then operate independently and concurrently with their own control knowledge according to the event/usage criteria; the control knowledge itself is coordinated at the metadatabase. The functionality of each shell depends on the specific knowledge of particular applications, therefore, each shell is in essence different. However, we can create these shells in a manner that allows their basic structure to be identical, with the differences only coming from the specific knowledge they process. Application Systems will be added, deleted or altered through the metadatabsae and the corresponding additions, deletions, or modifications of shells.

The shells must also be able to (1) communicate with each other in a concerted global behavior, (2) react to changes in the local applications that have global repercussions, and (3) process the knowledge they receive from the metadatabase. In addition, the shells should be efficient to implement, change and execute; otherwise, why should one have shells if it is easier to modify the application system? ROPE - Rule-Oriented Programming Environment - is the method developed to satisfy these requirements.

### 3.2 The Model and Functionality

The ROPE method develops the shells technology needed for the model-based architecture of multiple systems. It defines: (1) how the rules are stored in the local shells or database servers (representation and format used), (2) how the rules are processed, distributed, and managed, (3) how each shell interacts with its corresponding application system, (4) how the different shells interact with each other, and (5) how the shells are designed (i.e., their architecture). Furthermore, the ROPE approach prescribes three principles: (1) rules representing processing logic are separated from the program code by placing them in a distinct rulebase section for easy modifiability, (2) communications among shells are conducted through a message system (typically in an WAN environment) or a database server (typically in an LAN environment), and (3) the rule processing and the communicating methods are combined into the shells. As such, the local shells are invisible to the users, but control the "external" behavior of the local systems. The functionality of ROPE is described below, which serves as the design goals for the ROPE structure discussed in the next two sections.

**System Integration**

In ROPE, the information management capabilities are achieved in three ways. First, ROPE processes global queries at local nodes. Figure 3 shows how data and queries flow across the networked shells. The global query system of the MDBMS (metadatabase management system shell; circle in the center) sends requests for local queries to be executed by the different shells. The resulting data are assembled by the global query system. The different local shells can also initiate a request for a global query to be processed at the metadatabase management system shell. The general process is as follows: a user initiates a global query, either from MDBMS or from a global query interface located in the local application environment, using the Metadatabase Query Language (MQL) (see [7] for a full documentation of MQL). This query is sent to MDBMS, where it is processed; its syntax and semantics are validated, and local queries and result integration information are generated. The local queries are then sent to their respective shells, where they will be executed by the local DBMS. The results are then sent back to MDBMS for assembly and post-processing.