

# A Mediated Approach to Open, Large-Scale Information Management

Michael N. Huhns and Munindar P. Singh

Microelectronics and Computer Technology Corporation  
3500 West Balcones Center Drive, Austin, TX 78759-5398, U.S.A.  
{huhns,msingh}@mcc.com  
(512) 338-{3651,3431}

**ABSTRACT** The explosive expansion of the communications infrastructure, both within and among computing organizations, poses new challenges to the design of distributed information systems. Software systems are required that are open, that can be scaled to extremely large sizes, and that can be constructed and used effectively. The huge amount of information being made available through network technology is still incomprehensible and inconsistent. Previous attempts to address this problem of information management led to local solutions that were not maintainable or scalable, and were limited to retrieving, and not updating, data. We propose a mediator-based approach that supports interoperation of information resources, applications, and interfaces. Unlike previous approaches, mediator-based or otherwise, our solution embodies techniques that specially enable it to exploit modularity, abstraction, and adaptation. This solution enables *open, large-scale* information management, and emphasizes *updates* as well as retrievals.

## 1. Introduction

The extensive development of the communications infrastructure across the world poses special challenges to the design and implementation of software systems. Networking allows reliable access to widely distributed and disparate computing resources, both within large enterprises and globally. However, proper usage of this infrastructure presupposes that the different components that interact over it—user interfaces, applications, and information resources—are at least able to comprehend each other. Successful interoperation in large information spaces requires the participating components to be (a) locally autonomous, and (b) aware of remote components and the semantic constraints they embody and enforce. Each of these requirements is crucial in an open environment, to which service consumers and service providers are continually being added and removed. “Open” not only means that services can easily be added to a network, but also that they need not be captive to one enterprise. This underlines the need for local autonomy, which however is also important for technical reasons within an enterprise—if a system preserves local autonomy, then the various components need not be redesigned when they are integrated. Also, in an open system finding, accessing, and updating information entails integration into a semantic space that spans wide-area networks, so that the requirements imposed by restrictions such as integrity constraints and resource preferences can be met.

Our main contribution in this paper is an approach to make diverse components interoperate without disabling existing applications—a feature that ensures runtime scalability. A related problem is how the semantic links between different components can be “acquired.” To be scalable, an approach must not require

human effort to encode all possible relationships in advance. Our previous research has developed a solution to this problem [Huhns et al. 93] and we will not focus on it here.

Much research activity has recently been geared toward achieving interoperation within large information spaces. Our work is similar in spirit to some of that performed under the ARPA I<sup>3</sup> initiative, whose philosophy is described in [Wiederhold 92]. In a broad sense, our goals are also similar to those of the Open Distributed Processing (ODP) community [de Meer et al. 94]. However, our research focuses more on the semantics of information, than most ODP research. Further, we believe that our system is far more advanced than of most other related approaches. Whereas some of our software modules are prototypes others have been successfully deployed at customer sites. We address three major problems that extant approaches do not fully handle:

1. **Updates.** Updates are qualitatively more complex than retrieval. You can never violate a semantic constraint by just reading data (though an attempt to do so may violate some security or other policy constraint). Updates can, however, cause integrity constraints to be violated. Consequently, they call for a much more elaborate framework. Even in traditional homogeneous databases, updating views can be tricky; in heterogeneous systems, this problem is greatly exacerbated. However, most related research addresses information retrieval only (e.g., [Arens et al. 93]).
2. **Scalability.** How does one add a new component (e.g., a database, user interface, or application) to a large information space and have it interoperate effectively? This is the scalability problem, and there are multiple facets to it:
  - a. **Modularity.** We attain modularity via a novel combination of distribution and mediation.
  - b. **Abstraction.** We define a hierarchical structuring of mediators, where mediators at higher levels in the structure maintain abstract views of lower-level components of the information space.
  - c. **Incremental expansion.** We enable the coexistence of new applications with old ones, thereby facilitating a graceful introduction of technology.
3. **Heterogeneity of data and transaction models.** True interoperation requires resources designed with different data and transaction models to be coherently integrated. Instead of proposing yet another such model [Elmagarmid 92], we have built the functionality by which several of them can be uniformly integrated. Different transaction models are usually closed in that they cannot be extended with the features of other models. Our approach does not

commit to any single such model, but can express any of them succinctly and combine them in any way desired. These aspects of our approach are described in [Attie et al. 93; Huhns et al. 93], so we do not discuss them here.

To ground our discussions of scalability and updates, imagine the following scenario: a large brokerage house has 1000 brokers selling securities. The brokers each have a workstation or PC where they have created personal databases to record information about their clients. They store additional data in spreadsheets that they create locally. It is reasonable to assume that the computers are interconnected, but it is likely that the semantics and syntax of each database and each spreadsheet are different. The database systems might, for example, be Paradox, Oracle, and DB2. The tables might refer to “customers,” instead of “clients.” The data might be stored “last-name-first” or “last-name-last.” In such a diverse environment, how can

- a manager obtain information about the overall progress of the salesforce?
- the manager download new incentive tables for the brokers, when they each might represent the tables differently?
- a broker coordinate with another broker to avoid interference? It would be embarrassing if two brokers simultaneously gave conflicting recommendations.
- the manager remove from the databases all accounts that have been inactive for more than six months. This requires processing temporal interresource constraints.

We base our approach on mediators—for knowledge bases, applications, interfaces, and databases [Wiederhold 92]. We present a way of structuring mediators individually and collectively so that the three desiderata given above can be met. Of all the different functions that mediators can perform, we focus on workflow management as an interesting and important one well-suited to our approach. Section 2 discusses this problem and our overall approach. Section 3 presents our mediator architecture in some detail. Section 4 gives an example session of applying our approach to integrate resources within and among enterprises. Lastly, section 5 discusses applications in transaction processing in heterogeneous environments.

## 2. Technical Approach

### 2.1. The Problem: Workflow Management

The problem of information management can be cast in terms of the general problem of managing work—in government and corporate offices, design labs, factories, and so on. No one accesses information without a purpose and the purpose is best understood in the context of the overall enterprise activities. Therefore, instead of discussing information resource integration in the abstract, we approach it with a view to enabling business process integration. Business processes are the myriad activities that occur in typical enterprises. Frequently, these are informal and are not fully understood. Indeed, in certain cases their existence may not be known. This work focuses on the technical aspects of specifying and executing complex activities, and presupposes that the relevant business processes have been discovered, understood, documented, and formalized.

Workflows are the structured activities, in or abutting information systems, that enact business processes. Briefly, workflows consist of *tasks*, appropriately structured, where a task is some

useful computation. Thus, tasks can be information retrieval queries, database transactions, or other computations, e.g., those that generate visualizations. These activities frequently involve several database systems, user interfaces, and application programs. Traditional database systems do not support workflows to any reasonable extent. Usually, human beings must intervene to ensure the proper execution of workflows.

The activities that comprise a workflow of interest are typically already being carried out in the given organization. However, they are usually carried out by hand, with people intervening at crucial stages to ensure that the necessary tasks are done and that organization-wide consistency constraints are enforced. The semantics is supplied by the people or is implicitly encoded in different business procedures. For instance, when an order is received, it is first entered, and then numerous decisions are made about it. These decisions typically involve accessing several information resources within an enterprise and possibly some outside of it. For example, a request to buy stock against a money account requires that the authorization be verified, the account numbers be validated, and the source account be tested to have the required balance. External sources of information would be accessed for other requests, such as loan applications, where a credit bureau’s databases might be consulted to determine the credit worthiness of an applicant.

Integrating existing systems is much harder than designing them afresh. Many systems, especially those based on mainframe architectures, allow data to be accessed only through arcane interfaces. The systems and their interfaces cannot be easily modified, and our work assumes they cannot. This is for two main reasons: (1) the complexity of the programming effort that would be required to achieve any modifications, and (2) the constraint that older applications continue to run as before, since they typically have a wide user base. Thus, the integration must permit newly developed applications to coexist with previous applications. This proves to be a major requirement.

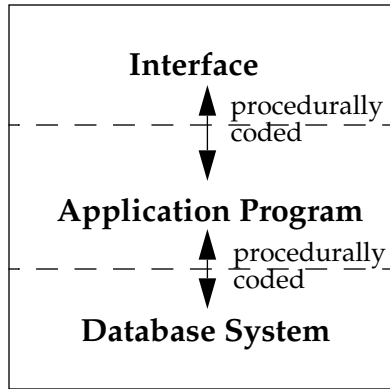
### 2.2. General Approach

Most importantly, our approach supports a new methodology for information systems design. Currently, information systems are constructed monolithically with a tight coupling among the interface, the application program, and the database (see Figure 1). In many cases, the three components are inseparable. This coupling makes the systems difficult to maintain, cumbersome to enhance, and expensive to reengineer as new technology becomes available. For example, it cannot benefit from advances in networking. Collectively, these problems contribute to the formation of what are pejoratively called “legacy systems.”

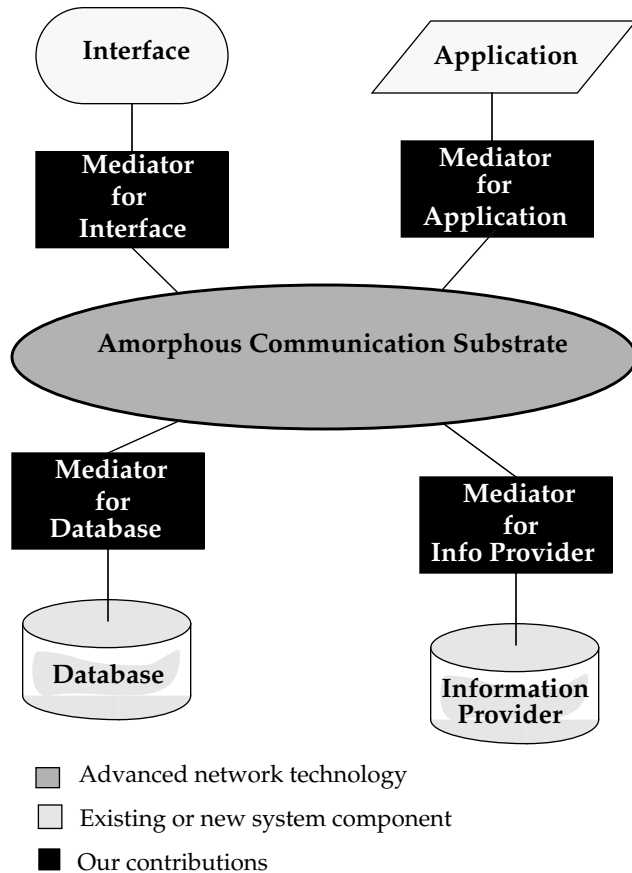
We loosen this coupling by introducing mediators between the three components (see Figure 2). Some of the mediators function as “wrappers,” i.e., protocol converters. But, in general, they can (1) update, create, and delete information, not just query it, and (2) learn models of their own component and of other mediators. This yields modular systems that can more readily be maintained, enhanced, or reengineered.

## 3. Mediator Architecture

In our approach, mediators translate semantics among different ontology domains, manage workflows, and enforce integrity and security constraints. Traditionally, mediators perform the func-



**Figure 1: Model (simplified for illustrative purposes) of the tight coupling among the three major components of most current information systems.**



**Figure 2: Information system components decoupled by the use of mediating agents.**

tions of semantic translation and security constraint enforcement. However, we see mediators as providing higher-level abstractions, and key among which are the semantic notions of workflows. Consequently, our mediators must also enforce integrity constraints, when the executed workflows may modify the data that might lead to an integrity violation.

As discussed above, the properties of modularity, abstraction, adaptation, and incremental expansion are highly desirable. Our

challenge, then, is to design a mediator architecture that

- is scalable,
- is extensible,
- is easy to implement and maintain,
- is reusable,
- does not violate the autonomy of integrated systems, and
- does not impose an unbearable performance penalty.

These might appear to be conflicting requirements. However, we motivate and present a design that meets the above criteria. We first discuss the structuring of each mediator and then the structuring of collections of them. The latter is particularly important in an open, extensible environment.

### 3.1. Layering of Functionality

Each mediator has a layered architecture, as shown in Figure 3.

|  |
|--|
| <b>Learning mechanism</b>                                    |
| <b>Cooperation and negotiation mechanism</b>                 |
| <b>Distributed truth maintenance</b>                         |
| <b>Security enforcement</b>                                  |
| <b>Workflow and interresource integrity constraints</b>      |
| <b>Semantic Translation: articulation axioms, value maps</b> |
| <b>Protocols: RDA SQL, KQML</b>                              |
| <b>Network Facilities: Actor treespace, Kerberos, etc.</b>   |

**Figure 3: The Layered Architecture of a Generic Mediator**

Not all layers may be present in each mediator. The lower layers handle the syntactic aspects of mediation and are closest to the communications infrastructure; the upper layers handle the semantic and adaptation-oriented aspects. Consequently, the lower layers are easily reused; the upper layers require human intervention, hopefully limited, in their construction. For reasons of reusability and testing, we implement the different layers of a mediator as different entities—actors in Carnot’s Extensible Services Switch. However, an entire set of such instantiated actors is one mediator because conceptually it performs as a logical unit—a mediator of a single resource.

The structuring aspects of our approach that we describe here apply to each layer of the mediators. Here we concentrate on the interresource aspects, which is the most interesting semantically.

### 3.2. Locality in Mediators

We hierarchically structure collections of mediators. The lowest level of mediators are associated with individual resources and applications. Semantics relating them to the rest of the information space are maintained by *local domain experts*. Higher level mediators manage interresource constraints among lower systems, and common ontologies to attach a semantics to those constraints. Local domain experts can construct these ontologies incrementally, using other parts of Carnot, thereby making those ontologies extensible. This approach is scalable, because it is based on the *local* construction and maintenance of mediators.

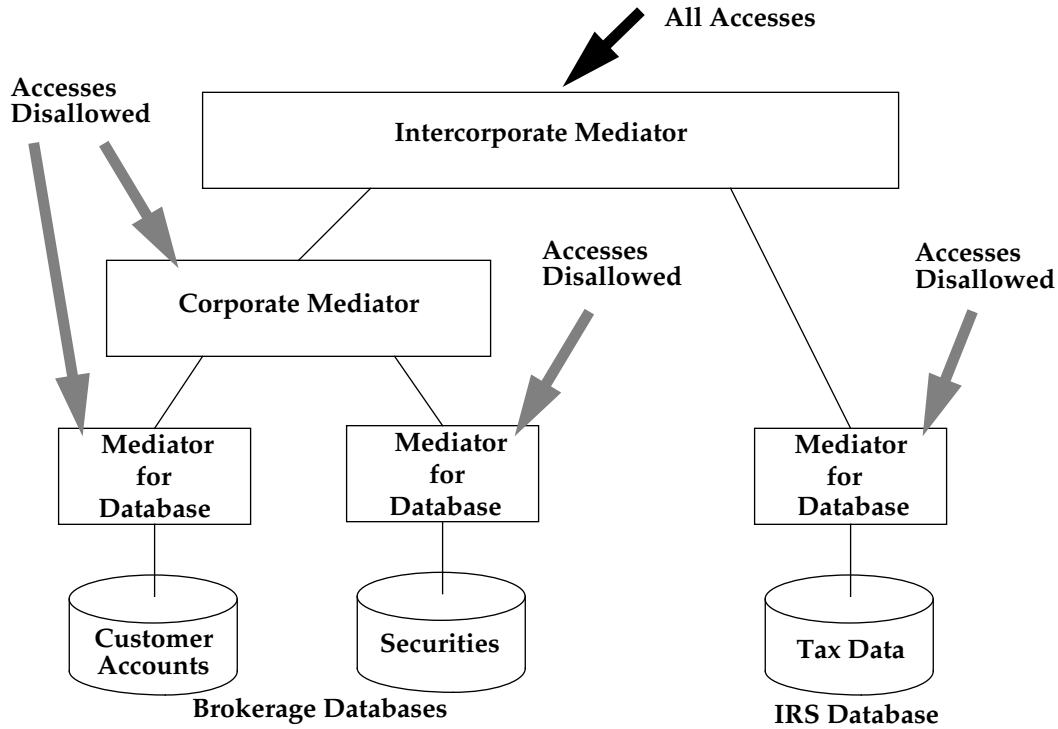


Figure 4: Control hierarchy

Our approach is bottom-up and resource-driven. It respects the organization of resources present in any enterprise: for example, individual databases, although usable by themselves, are often linked to other databases within a corporate division, to resources within an entire corporation, and increasingly nowadays to information resources in other corporations. The bottom-up approach assures not only that the system can be systematically constructed and tested, but also that constraints at a lower-level are never violated in favor of higher-level constraints. When higher-level constraints, such as those pertaining to intercorporate exchange of data, conflict with lower-level constraints, then the latter take precedence by default: however, such conflicts can and should be resolved through negotiation.

For each resource, we postulate a mediator that captures the constraints local to that resource. The schema of a database, its integrity constraints, and its transaction suite are in this category, as are security and access policies. In general, a mediator should be loyal to its own enterprise and not provide services to external entities that they are not authorized to access.

### 3.3. Scope Hierarchies

The requirements on mediators motivated above may appear to conflict. On the one hand, distribution is clearly crucial for scalability, extensibility, and to preserve local autonomy. On the other hand, the need to avoid significant performance or maintenance costs is suggestive of a hierarchical organization in which unnecessary interactions are avoided at runtime, and changes in the structure of a local system can be propagated as needed to other mediators. Control hierarchies are antithetical to distributed computing. They create bottlenecks and hot spots for performance, and make systems susceptible to the failure of the controlling node. These limitations subvert one of the key advan-

tages of distributed computing. Indeed, as we describe below, control hierarchies for mediators lead to additional complications that cannot be resolved without violating the autonomy of existing applications. We resolve this dilemma through a distributed architecture in which mediators can deal with other mediators on an even footing. Additionally, we also attach a hierarchical structure to these mediators, based not on control, but on *scope*.

The task of creating the initial local mediators can be carried out by different resource coordinators almost independently. Typically, however, there are additional resource constraints that pertain to multiple resources. Such constraints may be known from the beginning or may emerge later through experience. Rather than attempt to capture these constraints directly in one or some of the local mediators, we define another mediator whose scope includes the relevant resources. The scopes of the different mediators thus nest in a natural manner.

The new, high-level mediator can be seen as a descriptive construct, i.e., as a virtual mediator. Service requests from without can be directed to this virtual mediator as if it mediated for a monolithic resource with complex properties and constraints. This reflects the nature of business deals in general. Enterprises as a whole provide services and enter into contractual agreements. Their internal structure, though important, is irrelevant or hidden from without.

A control hierarchy creates higher-level mediators that control predefined lower-level ones. They determine the requests to be made to different lower-level mediators to maintain high-level constraints. Direct access to the predefined lower-level mediators must be disabled; otherwise, high-level constraints can be easily violated (see Figure 4). This debilitating limitation is not usually noted, since most approaches are restricted to information re-

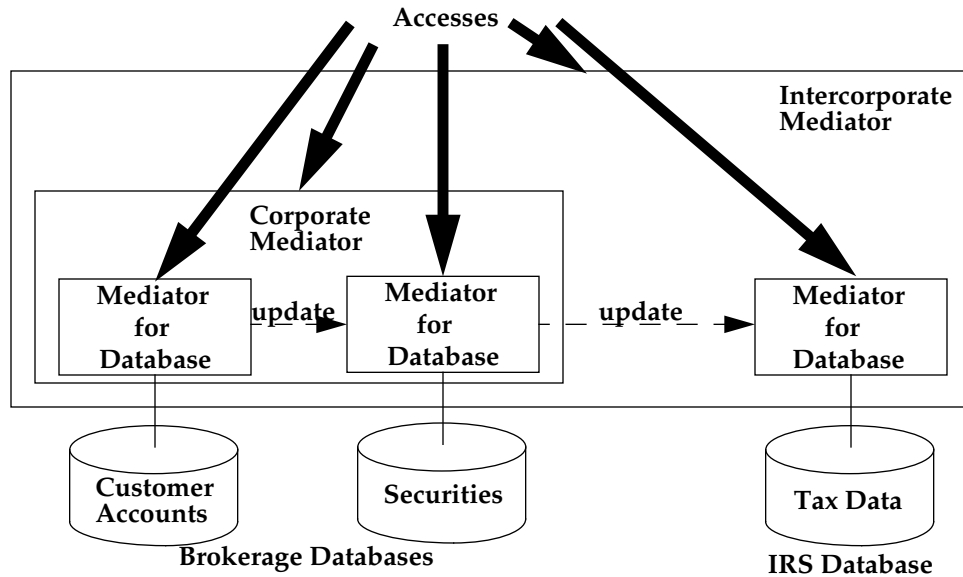


Figure 5: Scope hierarchy

trieval. As stated above, retrieval cannot cause semantic constraints to be violated. Hence, all read accesses can be safely permitted. In general, however, all requests must be funneled through the higher-level mediators.

By contrast, a scope hierarchy creates higher-level mediators that add constraints to the lower-level ones. All requests (physically) go to the lower-level mediators, which handle their old constraints as well as the new constraints. In this way the new higher-level mediator is distributed over the predefined lower-level mediators. Old applications can thus continue to issue requests directly to the lower-level mediators. Of course, if the requests issued by old applications violate new constraints, then such violations would be detected as easily as for new applications. Old applications would need to be disabled in such cases, as indeed they should be.

Hence, our architecture achieves the goals of scalability, extensibility, and local autonomy through distribution, in a manner in which no inessential constraints are imposed. Only the necessary interactions among mediators are required, and existing applications are affected just if they conflict with new constraints. Moreover, the mediators can be locally maintained and their different layers independently reengineered. Consequently, mediators can adapt through learning without wreaking havoc on the system as a whole.

#### 4. An Example Session

We now describe an example session of integration using our approach. Schematic and other semantic mismatches that may occur among the various databases must be resolved at the time of integration. Details of this are described elsewhere [Huhns et al. 93]. Here we describe the logical integration of databases at runtime, so that the proper integrity constraints are maintained, given that the mismatches have been taken care of already.

##### 4.1. Intracorporate Integration.

Assume that the customer account and the securities databases are already in use. The former gives the number of units of each


security owned by each account; the latter gives the total number of units of each security managed by the entire brokerage firm. The accounts database is accessed by different brokers; the securities database is accessed by a VP to determine how much of the firm's business is dependent on different securities and to adjust the risk for the mutual funds managed by the firm.

These databases are to be integrated to enforce the obvious semantic constraint between them. First, the administrators for the two databases must integrate their local schemas with a common ontology, which must be created if not already present. Then, the administrators are ready to capture the interresource integrity constraint. They begin by stating that constraint in terms of the common ontology and adding it to a new corporate mediator.

They then fold the constraint back into the local schemas of the two databases. The administrators also capture the temporal aspects of the constraint, namely, the frequency at which it must be validated. If it were required to be exactly maintained at all times, no feasible solution would exist in a heterogeneous environment. Fortunately, that is typically not the case. The administrators decide that whenever an update is performed on the accounts database, the mediator for the accounts database would send a corresponding transaction to the securities database. This update would eventually be executed. Thus, all existing applications continue to execute unchanged.

##### 4.2. Intercompany Integration.

Later, it is decided at the corporate level that certain information from the accounts and securities databases will be released to the Securities Exchange Commission. The SEC system administrators integrate their database with the common ontology, enhancing the common ontology as necessary. The additional constraint is added to a newly constructed intercompany mediator. The constraint is folded back into the schema of the corporate mediators and then on to the local schemas of the involved databases. The necessary workflow constraints are set up as before. Consequently, all applications continue to be usable. Periodically, updates are sent from the brokerage database to the SEC database to



**Figure 6: Task Dependency Example**

maintain consistency of the two databases. No new bottlenecks are introduced into the brokerage's information system.

In this example, after the interresource constraint was acquired, it was folded automatically into the applicable local schemas. However, the actual cross-system updates to be performed and when to perform them also need to be determined. These are a kind of flow constraint determined by reasoning about available actions and information flow in an integrated system. Such reasoning can be extremely complex. Our aim has been to design a tool that can assist humans in this activity, but not replace them. As different kinds of reasoning appear more and more useful, we augment the reasoning component of our tool so that it can be of increasing assistance.

## 5. Relaxed Transaction Processing

Transaction processing provides a related, but more complex, application for our approach. Classical transaction processing in databases deals with executing access and update tasks on a single database. Such tasks are traditionally assumed to have the so-called ACID properties:

- *atomicity*—each task happens either fully or not at all;
- *consistency*—each task takes the database from a consistent state to a consistent state;
- *isolation*—the intermediate results of a task are not visible to another task; and
- *durability*—the changes caused by a task are persistent.

These assumptions help simplify transaction management considerably [Gray & Reuter 93]. However, they are overly restrictive in loosely-coupled heterogeneous environments. For example, one of the ways in which ACID tasks may be coordinated is through mutual commit protocols, which ensure that either all of a given set of tasks commit or none do. Such protocols, e.g., the two-phase commit protocol, are notoriously inefficient when executed over networks. Further, to execute such a protocol, one requires access to the internal states of a transaction, such as the precommit states (when it is internally ready to commit, but is awaiting permission from the transaction manager to do so). Most commercial database systems do not provide access to such internal states, thereby making direct implementations of commit protocols difficult.

The ACID properties are naturally realized when the correctness of database transactions is characterized through some purely syntactic or structural criterion, such as serializability [Bernstein et al. 87]. However, serializability cannot be efficiently implemented in distributed systems whose component systems are autonomous. Following [Garcia-Molina & Salem 87], we characterize correctness criteria semantically, rather than syntactically. This simplifies coordination requirements, though at the cost of a deeper domain model. Thus, we can replace mutual commit by optimistic commit protocols. If need be, we undo the effects of incorrectly committed tasks by *compensating* them in a domain-dependent manner.

Figure 6 shows an example graph with tasks for updating brokerage accounts, updating the securities database, compensating for this update, and updating an alarm if some application-specific condition is met. Each task is modeled as a finite-state automaton whose nodes are events significant to the execution of the task. For example, the start of the accounts updates triggers the securities update. The conjunctive constraint (&) indicates starting the compensating transaction if the accounts update fails and the securities update succeeds.

Typically, to maintain the above constraints, a script would be required that starts the appropriate subtransactions to execute. This script might unnecessarily constrain executions. In any case, it would be difficult to modify. In our approach, the application, e.g., a user interface, that generates the accounts request does not need to be aware of the securities database or of the constraint. It generates the request as always. However, a higher-level mediator whose scope includes both databases introduces a constraint that at runtime triggers the securities subtransaction. Other constraints related to the commit or abort of the various subtransactions execute similarly. Thus the global semantics are enforced without affecting preexisting applications.

## 6. Conclusions

The concept of mediators is applied recursively within Carnot to provide applications with access to a variety of heterogeneous data and knowledge resources. Carnot has developed mediators for several kinds and varieties of database systems, including relational database systems, such as Oracle, Ingres, and Sybase, object-oriented database systems, such as Objectivity and Itasca,

and text retrieval systems, such as Topic. These mediators primarily handle protocol mediation, translate database schema semantics, and manage low-level concurrency [Woelk et al. 92]. Here we discussed how more sophisticated mediators that handle workflows are naturally incorporated in the same framework.

The ESS is constructed in Rosette, a high-performance implementation of an interpreter for the Actor model [Hewitt et al. 73] enhanced with object-oriented mechanisms. This interpreter has been developed over five years of research on parallel algorithms and control of distributed applications. Actors can be understood as concurrent and asynchronous units of computation [Agha 86]. Rosette is enhanced with remote evaluation, tree spaces, and virtual synchrony. We use cooperating groups of actors to implement application-driven mediated access to resources.

Scope hierarchies can be naturally implemented using ESS's treespaces [Tomlinson et al. 93a], which are a general-purpose mechanism for relaying messages anywhere on a network. Treespaces support symbolic addressing, and allow messages to be retrieved both by their content and by the recipient's name. The treespace facility pertains to the lower layers of the Carnot architecture. It is used for message addressing by the various database and expert system mediators we have already implemented. Thus, it provides just the right set of abstractions to implement scope hierarchies of mediators.

We presented a mediator architecture in which the individual mediators are layered by functionality and organized in a manner that yields the benefits of hierarchies without losing the benefits of distribution. Our approach yields local autonomy, easy maintainability, efficiency, and *scalability*. The architecture is embedded in an interpretive, dynamic development and execution environment, and provides capabilities ranging from protocol conversion to intelligent cooperation and learning. It enables the simple incorporation of new components into large information spaces. Our approach is also unique in that we develop mediation techniques for *transaction processing* on information resources, not just querying.

## References

- [Agha 86] Gul Agha, *Actors*, The MIT Press, Cambridge, MA, 1986.
- [Ahlsen & Johannesson 90] Matts Ahlsen and Paul Johannesson. "Contracts in Database Federations," in S. M. Deen, ed., *Cooperating Knowledge Based Systems 1990*, Springer-Verlag, London, 1991, pp. 293–310.
- [Ansari et al. 92] Mansoor Ansari, Marek Rusinkiewicz, Linda Ness, and Amit Sheth, "Executing Multidatabase Transactions," *Proceedings 25th Hawaii International Conference on Systems Sciences*, January 1992.
- [Arens et al. 93] Yigal Arens, Chin Y. Chee, Chun-Nan Hsu, and Craig A. Knoblock, "Retrieving and Integrating Data from Multiple Information Sources," *International Journal on Intelligent and Cooperative Information Systems*, Vol. 2, No. 2, 1993, pp. 127–158.
- [Attie et al. 93] Paul C. Attie, Munindar P. Singh, Amit P. Sheth, and Marek Rusinkiewicz, "Specifying and Enforcing Intertask Dependencies," *Proceedings of the 19th VLDB Conference*, 1993.
- [Bernstein et al. 87] Philip A. Bernstein, Vassos Hadzilacos, and Nathan Goodman, *Concurrency Control and Recovery in Database Systems*, Addison Wesley Publishing Co., 1987.
- [Bukhres et al. 93] Omran A. Bukhres, Jiansan Chen, Weimin Du, Ahmed K. Elmagarmid, and Robert Pezzoli, "InterBase: An Execution Environment for Heterogeneous Software Systems," *IEEE Computer*, Vol. 26, No. 8, Aug. 1993, pp. 57–69.
- [Buneman et al. 90] O.P. Buneman, S. B. Davidson, and A. Waters, "Querying Independent Databases," *Information Sciences*, Vol. 52, December 1990, pp. 1–34.
- [Cannata 91] Philip E. Cannata, "The Irresistible Move towards Interoperable Database Systems," *First International Workshop on Interoperability in Multidatabase Systems*, Kyoto, Japan, April 7–9, 1991.
- [de Meer et al. 94] J. de Meer, B. Mahr, and S. Storp, eds., *Open Distributed Processing II*, IFIP Transactions, C-20, North-Holland, 1994.
- [Elmagarmid 92] Ahmed K. Elmagarmid, ed., *Database Transaction Models for Advanced Applications*, Morgan Kaufman, San Mateo, CA, 1992.
- [Finin et al. 92] Tim Finin, Don McKay, and Rich Fritzson, "An Overview of KQML: A Knowledge Query and Manipulation Language," University of Maryland Computer Science Technical Report, March 1992.
- [Garcia-Molina & Salem 87] Hector Garcia-Molina and K. Salem, "Sagas," *Proceedings of ACM SIGMOD Conference on Management of Data*, 1987.
- [Gray & Reuter 93] Jim Gray and Andreas Reuter, *Transaction Processing: Concepts and Techniques*, Morgan Kaufman, San Mateo, CA, 1993.
- [Heimbigner & McLeod 85] Dennis Heimbigner and Dennis McLeod, "A Federated Architecture for Information Management," *ACM Transactions on Office Information Systems*, Vol. 3, No. 3, July 1985, pp. 253–278.
- [Huhns et al. 93] Michael N. Huhns, Nigel Jacobs, Tomasz Ksiezyk, Wei-Min Shen, Munindar Singh, and Philip Cannata, "Integrating Enterprise Information Models in Carnot," *International Conference on Intelligent and Cooperative Information Systems (ICICIS)*, Rotterdam, June 1993.
- [Klein 91] Johannes Klein, "Advanced Rule Driven Transaction Management," *Proceedings of the IEEE COMPCON*, 1991.
- [Singh & Huhns 94] Munindar P. Singh and Michael N. Huhns, "Automating Workflows for Service Order Processing: Integrating AI and Database Technologies," *IEEE Expert*, August 1994, in press. This is a special issue of selected papers from the *10th IEEE Conference on Artificial Intelligence Applications*, San Antonio, Texas, March 1994.
- [Tomlinson et al. 93b] Christine Tomlinson, Paul Attie, Philip Cannata, Greg Meredith, Amit Sheth, Munindar Singh, and Darrell Woelk, "Workflow Support in Carnot," *IEEE Data Engineering*, 1993.
- [Wiederhold 92] Gio Wiederhold, "Mediators in the Architecture of Future Information Systems," *IEEE Computer*, March 1992, pp. 38–49.
- [Woelk et al. 93] Darrell Woelk, Paul Attie, Philip Cannata, Greg Meredith, Munindar Singh, and Christine Tomlinson, "Task Scheduling Using Intertask Dependencies in Carnot," *ACM SIGMOD*, 1993.