

# Global Information Management via Local Autonomous Agents

Michael N. Huhns                      Munindar P. Singh  
Tomasz Ksiezuk  
Microelectronics and Computer Technology Corporation  
Research Division  
3500 West Balcones Center Drive  
Austin, TX 78759-5398  
{huhns,msingh,ksiezuk}@mcc.com

## Abstract

In this paper we describe how a set of autonomous computational agents can cooperate in providing coherent management of information in environments where there are many diverse information resources. The agents use models of themselves and of the resources that are local to them. Resource models may be the schemas of databases, frame systems of knowledge bases, or process models of business operations. Models enable the agents and resources to use the appropriate semantics when they interoperate. This is accomplished by specifying the semantics in terms of a common ontology. We discuss the contents of the models, where they come from, and how the agents acquire them. We then describe a set of agents for telecommunication service provisioning and show how the agents use such models to cooperate. Their interactions produce an implementation of relaxed transaction processing.

## 1 Introduction

Business operations, including sales, marketing, manufacturing, and design, can no longer be done in isolation, but must be done in a global context, i.e., as part of an enterprise. A characteristic of such enterprises is that their information systems are large and complex, and the information is in a variety of forms, locations, and computers. The topology of these systems is dynamic and their content is changing so rapidly that it is difficult for a user or an application program to obtain correct information, or for the enterprise to maintain consistent information.

Some of the techniques for dealing with the size and complexity of these enterprise information systems are modularity, distribution, abstraction, and intelligence, i.e., being smarter about how you seek and modify information. Combining these techniques implies the use of intelligent, distributed modules—a distributed artificial intelligence approach. In accord with this approach, we distribute and embed computational agents throughout an enterprise. The agents are knowledgeable about information resources that are local to them, and cooperate to provide global access to, and better management of, the information. For the practical reason that the systems are too large and dynamic (i.e., open) for global solutions to be formulated and implemented, the agents need to execute autonomously and be developed independently. To cooperate effectively, the agents must either *have models of each other and of the available information resources* or *provide models of themselves*. We focus on the latter in this paper.

For such an open information environment, the questions arise: what should be modeled, where do models come from, what are their constituents, and how should they be used? We discuss the types of models that might be available in an enterprise and how agents can acquire them. We use the ontology developed for the large knowledge-based system, Cyc, for semantic grounding of the models. This provides a common ontology. We

then describe a set of agents for telecommunication service provisioning—a scheduling agent, a schedule-repairing agent, a schedule-processing agent, and an interface agent—and describe their models and how they use them to cooperate. We also describe the use of actors [Agha 1986]—one per agent—who manage communications among the agents. Each actor independently maintains the relationship between its agent and the common ontology (in the form of articulation axioms), and updates that relationship as the ontology changes or the agent itself evolves.

## 2 Modeling

Enterprise information modeling is a corporate activity that produces the models needed for interoperability. The resultant models should describe all aspects of a business environment, including

- databases
- database applications
- software repositories
- part description repositories
- expert systems, knowledge bases, and computational agents
- business work flows, and the information they create, use, maintain, and own, and
- the business organization itself.

The models provide online documentation for the concepts they describe. They enable application code and data to be reused, data to be analyzed for consistency, databases to be constructed automatically, the impact of change on an enterprise to be assessed, and applications to be generated automatically.

An enterprise might have many models available, each describing a portion of the enterprise and each constructed independently. For example,

- the information present in a database is modeled by the schema for the database, which is produced through a process of logical data modeling
- the data values present in a database are modeled (weakly, in most cases) by data dictionary information, which is produced through data engineering
- the information present in an object-centered knowledge base is modeled by the ontology of the objects, which is produced through ontological engineering
- process models, possibly in the form of Petri nets or IDEFx descriptions, are produced through logical process modeling
- STEP (Standard for the Exchange of Product model data) schemas, written in Express, are produced from component and physical process modeling.

Although it might appear that interoperability would require all of these models to be merged into a single, homogeneous, global model, this is *not* the case in our approach. Instead, there are good reasons for retaining the many individual models: 1) they are easier to construct than a single large model; 2) enterprises may be formed dynamically through mergers, acquisitions, and strategic alliances, and the resultant enterprises might have inherited many existing models; 3) because enterprises are geographically dispersed, their resources are typically decentralized; and 4) as enterprises (and thus models) evolve, it is easier to maintain smaller models.

Unfortunately, the models are often mutually incompatible in syntax and semantics, not only due to the dif-

ferent things being modeled, but also due to mismatches in underlying hardware and operating systems, in data structures, and in corporate usage. In attempting to model some portion of the real world, information models necessarily introduce simplifications and inaccuracies that result in semantic incompatibilities. However, the individual models must be related to each other and their incompatibilities resolved [Sheth and Larson 1990], because

- A coherent picture of the enterprise is needed to enable decision makers to operate the business efficiently and designers to evaluate information flows to and from their particular application.
- Applications need to interoperate correctly across a global enterprise. This is especially important due to the increasing prevalence of strategic business applications that require *intercorporate linkage*, e.g., linking buyers with suppliers, or *intracorporate integration*, e.g., producing composite information from engineering and manufacturing views of a product.
- Developers require integrity validation of new and updated models, which must be done in a global context.
- Developers want to detect and remove inconsistencies, not only among models, but also among the underlying business operations that are modeled.

We utilize a mediating mechanism based on an existing common ontology to yield the appearance and effect of semantic homogeneity among existing models. The mechanism provides logical connectivity among information resources via a semantic service layer that automates the maintenance of data integrity and provides an enterprise-wide view of all the information resources, thus enabling them to be used coherently. This logical layer is implemented as a network of interacting agents. Significantly, the individual systems retain their autonomy. This is a fundamental tenet of the Carnot architecture [Woelk *et al.* 1992], which provides the tools and infrastructure for interoperability across global enterprises.

### 3 Semantic Integration via a Common Ontology

In order for agents to interact productively, they must have something in common, i.e., they must be either grounded in the same environment or able to relate their individual environments. We use an existing common context—the Cyc common-sense knowledge base [Lenat and Guha 1990]—to provide semantic grounding. The models of agents and resources are compared and mapped to Cyc but not to each other, making interoperation easier to attain. For  $n$  models, only  $n$  mappings are needed, instead of as many as  $n(n-1)$  mappings when the models are related pairwise. Currently, Cyc is the best choice for a common context, because of 1) its rich set of abstractions, which ease the process of representing predefined groupings of concepts, 2) its knowledge representation mechanisms, which are needed to construct, represent, and maintain a common context, and 3) its size: it covers a large portion of the real world and the subject matter of most information resources.

The large size and broad coverage of Cyc's knowledge enable it to serve as a fixed-point for representing not only the semantics of various information modeling formalisms, but also the semantics of the domains being modeled. Carnot can use models constructed using any of several popular formalisms, such as

- IRDS, IBM's AD/Cycle, or Bellcore's CLDM for entity-relationship models
- Ingres, Oracle, Sybase, Objectivity, or Itasca for database schemas, and
- MCC's RAD or NASA's CLIPS for agent models.

Cyc's knowledge about metamodels for these formalisms and the relationships among them enables transactions to interoperate semantically between, for example, relational and object-oriented databases.

The relationship between a domain concept from a local model and one or more concepts in the common con-

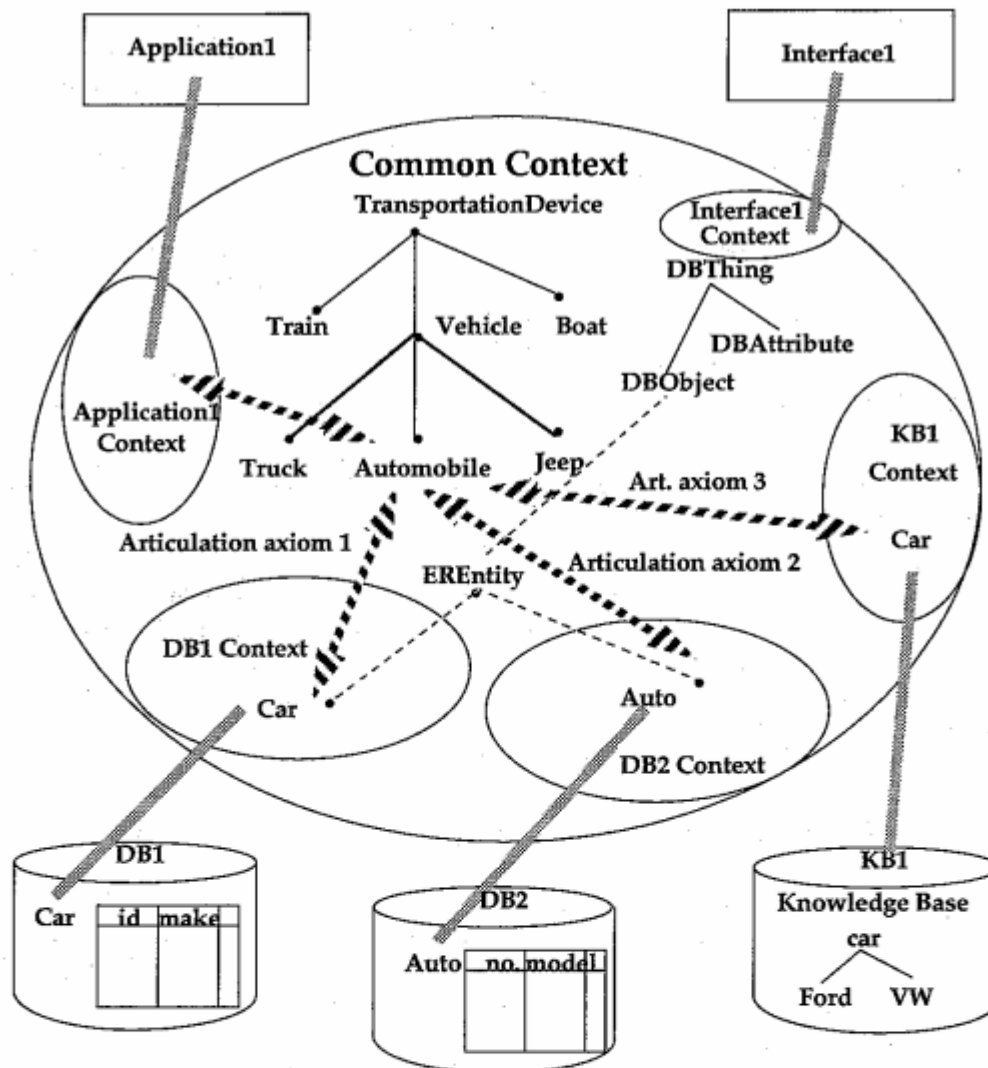


Figure 1: Concepts from different models are related via a common aggregate context by means of articulation axioms

text is expressed as an articulation axiom [Guha 1990]: a statement of equivalence between components of two theories. Each axiom has the form

$$ist(G, \varphi) \Leftrightarrow ist(C_i, \psi)$$

where  $\varphi$  and  $\psi$  are logical expressions and *ist* is a predicate that means "is true in the context." This axiom says that the meaning of  $\varphi$  in the common context  $G$  is the same as that of  $\psi$  in the local context  $C_i$ . Models are then related to each other—or translated between formalisms—via this common context by means of the articulation axioms, as illustrated in Figure 1. For example, an application's query about *Automobile* would result in sub-queries to DB1 about *Car*, to DB2 about *Auto*, and to KB1 about *car*. Note that each model can be added independently, and the articulation axioms that result do not have to change when additional models are added. Also note that applications and resources need not be modified in order to interoperate in the integrated environ-

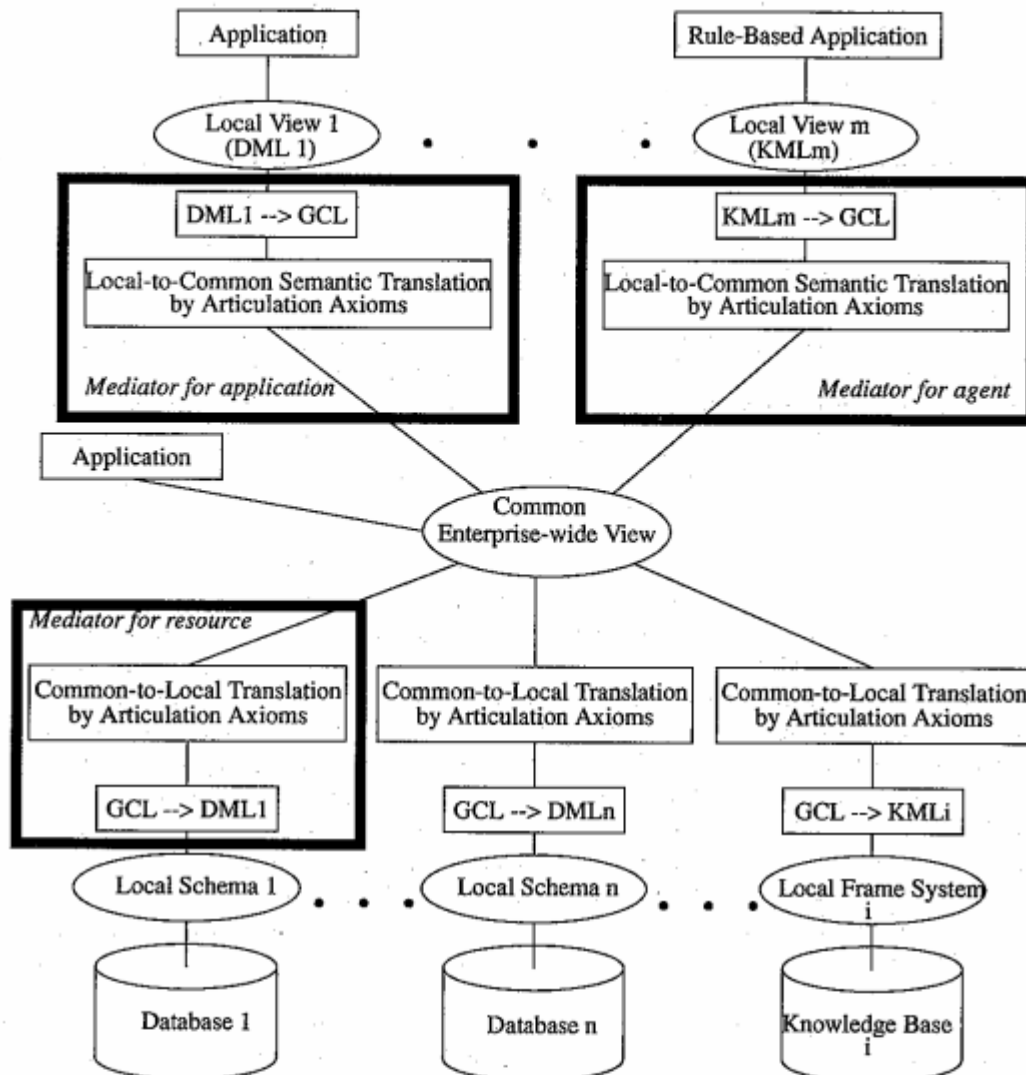


Figure 2: Logical view of the execution environment, showing how mediating agents apply articulation axioms to achieve semantic interoperability

ment. The Appendix contains a description of the graphical tool, MIST, that we have built to aid in the construction of articulation axioms.

Figure 2 shows a logical view of the execution environment. During interoperability, mediators [Wiederhold 1992], which are implemented by Rosette actors [Tomlinson *et al.* 1991], apply the articulation axioms that relate each agent or resource model to the common context. This performs a translation of message semantics. At most  $n$  sets of articulation axioms and  $n$  mediators are needed for interoperability among  $n$  resources and applications. The mediators also apply a syntax translation between a local data manipulation language,  $DML_i$ , and the global context language, GCL. GCL is based on extended first-order logic. A local data-manipulation language might be, for example, SQL for relational databases or OSQL for object-oriented databases. The number of language translators between  $DML_i$  and GCL is no greater than  $n$ , and may be a constant because there are only a small

number of data-manipulation languages that are in use today. Additional details describing how transactions are processed semantically through the global and local views of several databases can be found in [Woelk *et al.* 1992].

The mediators also function as communication aides, by managing communications among the various agents, databases, and application programs in the environment. They buffer messages, locate message recipients, and translate message semantics. To implement message transfer, they use a tree-space mechanism—a kind of distributed virtual blackboard—built on the OSI and TCP/IP protocols [Tomlinson *et al.* 1991].

## Application to Transaction Processing

We have applied our methodology to achieve relaxed transaction processing in the provisioning of telecommunication services, the task of providing communication facilities to customers. This task is executed in a heterogeneous multidatabase environment. It is an example of workflow control, in that it provides control and data flows among transactions executing on multiple autonomous systems [Jin *et al.* 1993; Tomlinson *et al.* 1993]. Service provisioning typically takes several weeks and requires coordination among many operation-support systems and network elements. Configuring the operation-support systems so that they can perform such a task often takes several months to complete.

We investigated ways to improve the provisioning of one type of communication facility—digital services (DS-1). Provisioning DS-1 takes more than two weeks and involves 48 separate operations—23 of which are manual—against 16 different database systems. Our goals were to reduce this time to less than two hours and to provide a way in which new services could be introduced more easily. Our strategy for accomplishing these goals was to 1) interconnect and interoperate among the previously independent systems, 2) replace serial operations by concurrent ones by making appropriate use of relaxed transaction processing [Attie *et al.* 1993; Bukhres *et al.* 1993; Elmagarmid 1992; Ansari *et al.* 1992], and 3) automate previously manual operations, thereby reducing the incidence of errors and delays. The transaction processing is relaxed in that some subsystems are allowed to be temporarily inconsistent, although eventual consistency is guaranteed. Relaxing the consistency requirements allows increased concurrency and, thus, improved throughput and response time.

The architecture of the agents used to implement relaxed transaction processing is shown in Figure 3. The agents operate as follows. The graphical-interaction agent helps a user fill in an order form correctly, and checks inventories to give the user an estimate of when the order will be completed. It also informs the user about the progress of the order.

The transaction-scheduling agent constructs the schedule of tasks needed to satisfy an order. The tasks are scheduled with the maximum concurrency possible, while still satisfying their precedence constraints. Some of the rules that implement the schedule are shown in Figure 4. These particular rules, when appropriately enabled, generate a subtransaction to update the database for customer billing. When executing such rules, the transaction-scheduling agent behaves as a finite-state automaton, as shown in Figure 5.

The schedule-processing agent maintains connections to the databases involved in telecommunication provisioning, and implements transactions on them. It knows how to construct the proper form for a transaction, based on the results of other transactions. The transactions are processed concurrently, where appropriate. If something goes wrong during the processing of a transaction that causes it to abort or fail to commit, the schedule-repairing agent provides advice on how to fix the problem and restore consistency. The advice can be information on how to restart a transaction, how to abort a transaction, how to compensate for a previously committed transaction, or how to clean-up a failed transaction. The integrity knowledge that is stored in the schedule repairing agent comes from a comparison of the models, as expressed in terms of the common ontology.

The agents, as described above, are simply expert systems whose expertise is in processing orders for telecommunication services. However, they have the additional abilities to interact and cooperate with each other via the mediators described above.

The agents cooperate, at the knowledge level [Newell 1982], via models of themselves. For example, a con-

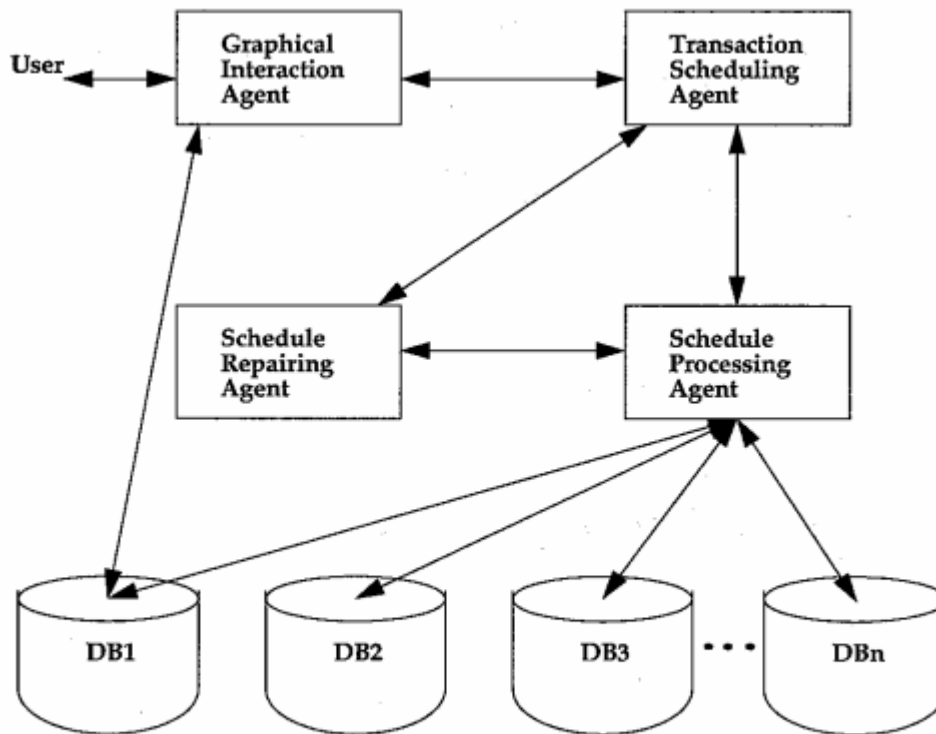


Figure 3: Agents used for relaxed transaction processing

ceptual domain model for the graphical-interaction agent is shown in Figure 6. An interface form that provides user access and modifications to the knowledge possessed by this agent is shown in Figure 7. Entries on the form, or the form's completion, cause queries and transactions to be sent to the other agents or databases in the environment. Note, however, that the model does not capture the procedural knowledge necessary to specify the queries and transactions; a technique for modeling processes is needed to capture such knowledge. In other words, the models represent the static knowledge of the agents, not (unfortunately) their dynamics. Nevertheless, they have proven useful in enabling the agents to interact coherently, as we describe next.

Conceptual models for two more of the agents are shown in Figures 8 and 9. Each model consists of organized concepts describing the context, domain, or viewpoint of the knowledge possessed by that agent, i.e., the knowledge base of each agent contains rules written in terms of these concepts.

The models in Figures 6, 8, and 9 are related to the common context, and thereby to each other, via articulation axioms. For example, the concept *Transaction* for the transaction-scheduling agent and the concept *DB-Transaction* for the schedule-repairing agent are each related to the common concept *DatabaseTransaction* via the axioms

$$ist(Cyc, DatabaseTransaction(T)) \Leftrightarrow ist(Scheduler, Transaction(T))$$

$$ist(Cyc, DatabaseTransaction(T)) \Leftrightarrow ist(Repairer, TDBTransaction(T))$$

The axioms are used to translate messages exchanged by the agents, so that the agents can understand each other.

```

;; This set of rules (1) executes an external program that
;; translates an Access Service Request into a command file
;; to update the database for customer billing, (2) executes
;; the command file, and (3) checks for completion. Note that
;; the scheduling agent, due to its truth-maintenance system,
;; stops processing this subtransaction whenever an abort of
;; the global transaction occurs.
;; ?gtid denotes the global transaction identifier.
Bill-Preparation:
  If (service-order(?gtid)
      new-tid(?subtid)
      unless(abort(?gtid)))
  then (do(,run-shell-program
          ("asr2bill"
           :input ("asr-?gtid.out")
           :output "bill-?gtid.sql"))
        bill(?gtid ?subtid)
        tell(GIAGENT "Task ?gtid BILLING ready"))

Bill-Execution:
  If (bill(?gtid ?subtid)
      logical-db(?db))
  then (tell(SchedProcAgent
            "task-execute ?subtid BILL ?db bill-?gtid.sql")
        tell(GIAGENT "Task ?gtid BILLING active"))

Bill-Completion:
  If (success(?subtid)
      bill(?gtid ?subtid))
  then (tell(GIAGENT "Task ?gtid BILLING done"))

Bill-Failure:
  If (failure(?subtid)
      excuse(bill(?gtid ?subtid)))
  then (abort(?gtid)
        tell(GIAGENT "Task ?gtid BILLING failed"))

```

Figure 4: Some of the rules used by the transaction-scheduling agent to generate a schedule for DS-1 workflow

In the above example, the two agents could use their axioms to converse about the status of database transactions, without having to change their internal terminology. Similar axioms describing the semantics of each of the databases involved enable the schedule-processing agent to issue transactions to the databases. The axioms also relate the semantics of the form shown in Figure 8 to the semantics of the other information resources in the environment. Such axioms are constructed with the aid of a graphical tool called MIST, for Model Integration Software Tool. The operation of MIST is described in the Appendix.

Operationally, the axioms are managed and applied by the mediators that assist each agent. They use the axioms to translate each outgoing message from their agent into the common context, and to translate each incoming message for their agent into its local semantics.



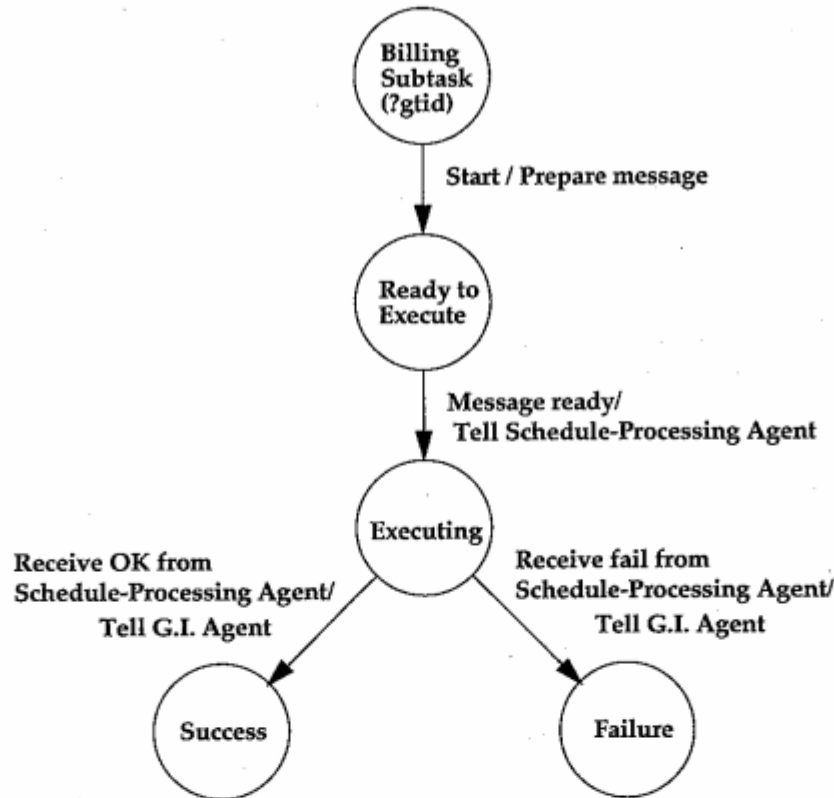


Figure 5: Representative finite-state automaton for a task to process a telecommunication service order, as implemented by the transaction-scheduling agent

## Background and Discussion

Integrating enterprise models is similar to integrating heterogeneous databases. Two approaches have been suggested previously for this [Buneman *et al.* 1990]. The *composite approach* produces a global schema by merging the schemas of the individual databases. Explicit resolutions are specified in advance for any semantic conflicts among the databases, so users and applications are presented with the illusion of a single, centralized database. However, the centralized view may differ from the previous local views and existing applications might not execute correctly any more. Further, a new global schema must be constructed every time a local schema changes or is added.

The *federated approach* [Heimbigner and McLead 1985, Litwin *et al.* 1990] presents a user with a collection of local schemas, along with tools for information sharing. The user resolves conflicts in an application-specific manner, and integrates only the required portions of the databases. This approach yields easier maintenance, increased security, and the ability to deal with inconsistencies. However, a user must understand the contents of each database to know what to include in a query: there is no global schema to provide advice about semantics. Also, each database must maintain knowledge about the other databases with which it shares information, e.g., in the form of models of the other databases or partial global schemas [Ahlsen and Johannesson 1990]. For  $n$  databases, as many as  $n(n-1)$  partial global schemas might be required, while  $n$  mappings would suffice to translate between the databases and a common schema.

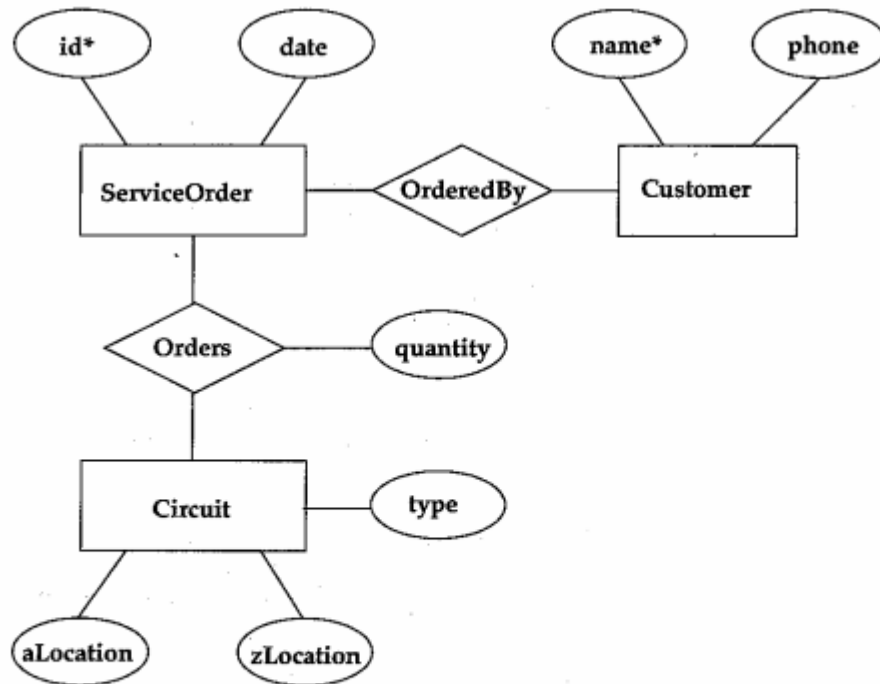


Figure 6: Semantic model (simplified) for the graphical-interaction agent

<b>DS-1 Access Service Request</b>			
Order ID	<input type="text"/>	Date	<input type="text"/>
Customer Name	<input type="text"/>	Phone	<input type="text"/>
Quantity	<input type="text"/>		
<b>Circuit Information</b>			
ALocation	<input type="text"/>	ZLocation	<input type="text"/>
		Type	<input type="text"/>

Figure 7: User interface form (simplified) corresponding to the declarative knowledge of the graphical-interaction agent

We base our methodology on the composite approach, but make three changes that enable us to combine the advantages of both approaches while avoiding some of their shortcomings. First, we use an *existing* common

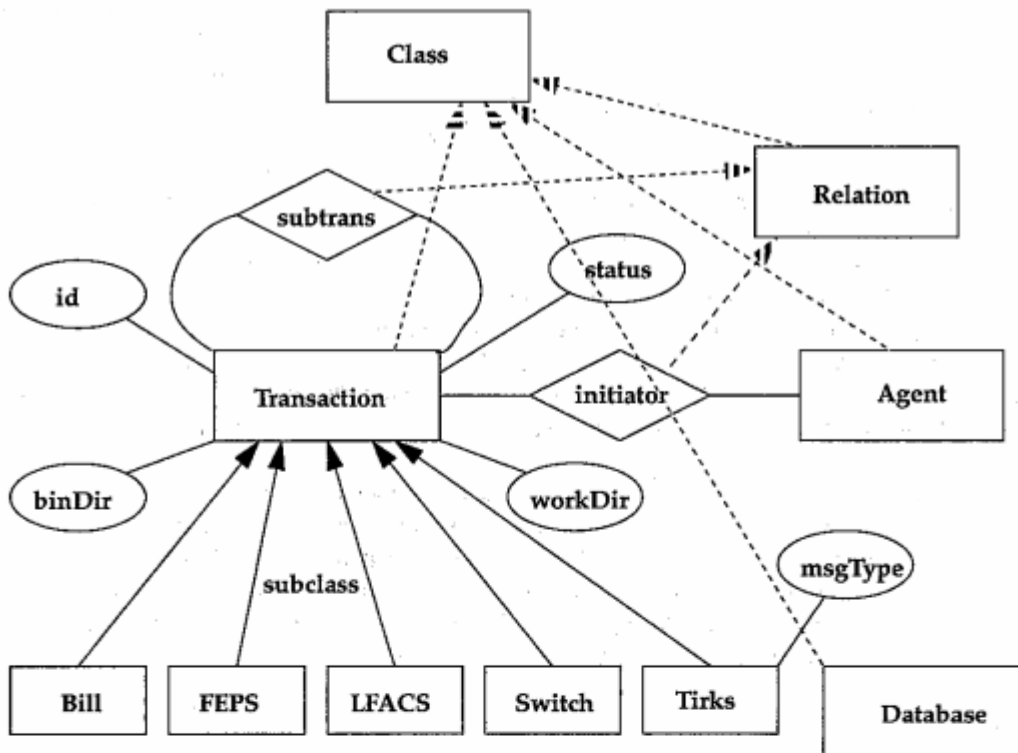


Figure 8: Semantic model for the transaction-scheduling agent (dashed arrows indicate instance relationships, and solid arrows indicate subclass relationships)

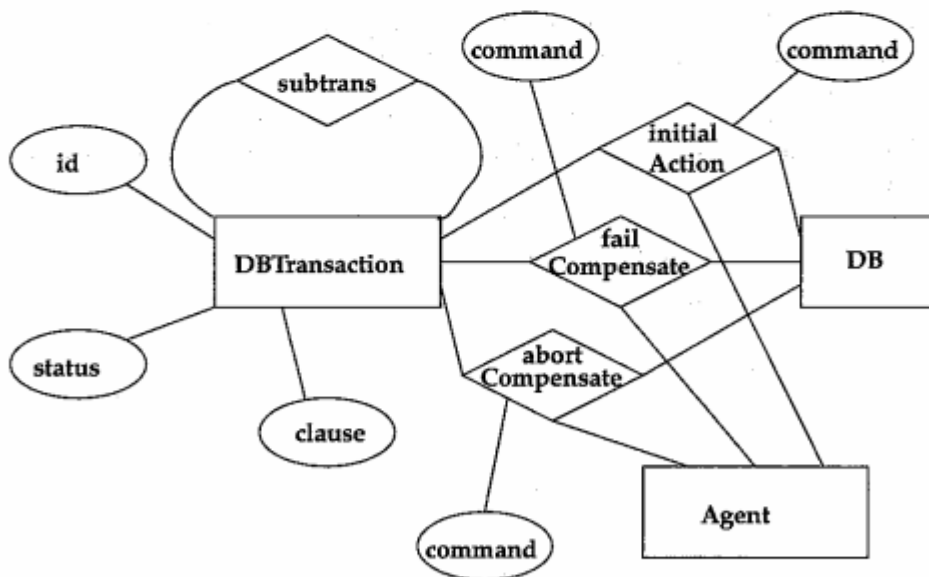


Figure 9: Semantic model for the schedule-repairing agent

schema or context. In a similar attempt, [Sull and Kashyap 1992] describes a method for integrating schemas by translating them into an object-oriented data model, but this method maintains only the structural semantics of the resources.

Second, we capture the mapping between each model and the common context in a set of articulation axioms. The axioms provide a means of translation that enables the maintenance of a global view of all information resources and, at the same time, a set of local views that correspond to each individual resource. An application can retain its current view, but use the information in other resources. Of course, any application can be modified to use the global view directly to access all available information.

Third, we consider knowledge-based systems (KBSs), interfaces, and applications, as well as databases.

Our use of agents for interoperating among applications and information resources is similar to the uses of mediators described in [Wiederhold 1992]. However, we also specify a means for semantic translation among the agents, as well as an implemented prototype. Other applications of similar agents, such as the Pilot's Associate developed by Lockheed *et al.* [Smith and Broadwell 1988], handcrafted their agents. This is not possible for large "open" applications: the agents must be such that they can be developed independently and execute autonomously.

Our architecture employs two kinds of computational agents: finer-grained, concurrent actors and coarser-grained, knowledge-based systems. The actors are used to control interactions among the components of the architecture. The knowledge-based agents are used where reasoning is needed, such as in deciding what tasks should be performed next or how to repair the environment when a task has failed. This seems to be a natural division of responsibilities for our example application. However, we took an engineering, rather than a scientific, approach, in that we did not investigate any alternative architectures.

## Conclusion

For years, information-system personnel managed corporate data that was centralized on mainframes. The data was kept consistent, but eventually the amount of data increased to the point that centralized storage was no longer viable. Also, users wanted a way to share data across applications and wanted more direct involvement in the management of the data. So, data then began proliferating onto workstations and personal computers, where users could manage it themselves. But this resulted in redundancy, inconsistency, and no coherent global view. Hence, there are now attempts to reintegrate data. Users still need to manage their own data, which remains distributed, but they and their applications need coherent global access and consistency must be restored.

This paper describes Carnot's approach to enabling interoperation among enterprise information objects, i.e., among suppliers and consumers of information. In this approach, an enterprise information object is integrated based on articulation axioms defined between two contexts: the context of a model of the object and a common context provided by the Cyc knowledge base. The methodology is based on the following principles:

- Existing information resources should not have to be modified and data should not have to migrate.
- Existing applications should not have to be modified.
- Users should not have to adopt a new language for communicating with the resultant integrated system, unless they are accessing new types of information.
- Resources and applications should be able to be integrated independently, and the mappings that result should not have to change when additional objects are integrated.

The above principles are incorporated in an integration tool, MIST, for assisting an administrator in generating articulation axioms for a model, and in a set of agents that utilize the resultant axioms to provide users and applications with access to the integrated resources. They can use a familiar local context, while still benefiting from newly added resources. These systems constitute part of the semantic services of Carnot [Cannata 1991],

under development at MCC. They help specify and maintain the semantics of an organization's integrated information resources.

Extensions of our work are focused on developing additional information-system applications for agents, including

- intelligent directory service agents
- negotiating electronic data interchange (EDI) agents
- database triggers—making passive databases active
- rule-based database applications
- database administration agents
- intelligent information retrieval agents.

Our most important future work is centered on ways in which agents can acquire and maintain models of each other in order to improve their interactions.

## References

- [Agha 1986] Gul Agha, *Actors: A Model of Concurrent Computation in Distributed Systems*, MIT Press, Cambridge, MA, 1986.
- [Ahlsen and Johannesson 1990] 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.* 1992] Mansoor Ansari, Marek Rusinkiewicz, Linda Ness, and Amit Sheth, "Executing Multidatabase Transactions," *Proceedings 25th Hawaii International Conference on Systems Sciences*, January 1992.
- [Attie *et al.* 1993] Paul C. Attie, Munindar P. Singh, Amit P. Sheth, and Marek Rusinkiewicz, "Specifying and Enforcing Intertask Dependencies," *Proceedings of the 19th VLDB Conference*, 1993.
- [Bukhres *et al.* 1993] Omran A. Bukhres, Jiansan Chen, Weimin Du, Ahmed K. Elmagarmid, and Robert Pezoli, "InterBase: An Execution Environment for Heterogeneous Software Systems," *IEEE Computer*, Vol. 26, No. 8, Aug. 1993, pp. 57–69.
- [Buneman *et al.* 1990] O. P. Buneman, S. B. Davidson, and A. Watters, "Querying Independent Databases," *Information Sciences*, Vol. 52, Dec. 1990, pp. 1–34.
- [Cannata 1991] Philip E. Cannata, "The Irresistible Move towards Interoperable Database Systems," *First International Workshop on Interoperability in Multidatabase Systems*, Kyoto, Japan, April 7–9, 1991.
- [Ceri and Widom 1992] Stefano Ceri and Jennifer Widom, "Production Rules in Parallel and Distributed Database Environments," *Proceedings of the 18th VLDB Conference*, Vancouver, British Columbia, Canada, 1992, pp. 339–351.
- [Collet *et al.* 1991] Christine Collet, Michael N. Huhns, and Wei-Min Shen, "Resource integration using a large knowledge base in Carnot," *IEEE Computer*, Vol. 24, No. 12, Dec. 1991, pp. 55–62.
- [Cutkosky *et al.* 1993] Mark R. Cutkosky, Robert S. Englemore, Richard E. Fikes, Michael R. Genesereth, Thomas R. Gruber, William S. Mark, Jay M. Tenenbaum, and Jay C. Weber, "PACT: An Experiment in Integrating Concurrent Engineering Systems," *IEEE Computer*, January 1993, pp. 28–38.

- [Elmagarmid 1992] Ahmed Elmagarmid, ed., *Database Transaction Models*, Morgan Kaufmann Publishers Inc., San Mateo, CA, 1992.
- [Guha 1990] R. V. Guha, "Micro-theories and Contexts in Cyc Part I: Basic Issues," MCC Technical Report Number ACT-CYC-129-90, Microelectronics and Computer Technology Corporation, Austin, TX, June 1990.
- [Heimbigner and McLeod 1985] 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.
- [Jin *et al.* 1993] W. Woody Jin, Linda Ness, Marek Rusinkiewicz, and Amit Sheth, "Executing Service Provisioning Applications as Multidatabase Flexible Transactions," Bellcore Technical Report (unpublished), 1993.
- [Lenat and Guha 1990] Doug Lenat and R. V. Guha, *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*, Addison-Wesley Publishing Company, Inc., Reading, MA, 1990.
- [Litwin *et al.* 1990] Witold Litwin, Leo Mark, and Nick Roussopoulos, "Interoperability of Multiple Autonomous Databases," *ACM Computing Surveys*, Vol. 22, No. 3, September 1990, pp. 267–296.
- [Newell 1982] Allen Newell, "The Knowledge Level," *Artificial Intelligence*, Vol. 18, No. 1, January 1982, pp. 87–127.
- [Sheth and Larson 1990] Amit P. Sheth and James A. Larson, "Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases," *ACM Computing Surveys*, Vol. 22, No. 3, Sept. 1990, pp. 183–236.
- [Smith and Broadwell 1988] David Smith and Martin Broadwell, "The Pilot's Associate—an overview," *Proceedings of the SAE Aerotech Conference*, Los Angeles, CA, May 1988.
- [Sull and Kashyap 1992] Wonhee Sull and Rangasami L. Kashyap, "A Self-Organizing Knowledge Representation Scheme for Extensible Heterogeneous Information Environment," *IEEE Transactions on Knowledge and Data Engineering*, Vol. 4, No. 2, April 1992, pp. 185–191.
- [Tomlinson *et al.* 1991] Chris Tomlinson, Mark Scheevel, and Vineet Singh, "Report on Rosette 1.1," MCC Technical Report Number ACT-OODS-275-91, Microelectronics and Computer Technology Corporation, Austin, TX, July 1991.
- [Tomlinson *et al.* 1993] Christine Tomlinson, Paul Attie, Philip Cannata, Greg Meredith, Amit Sheth, Munindar Singh, and Darrell Woelk, "Workflow Support in Carnot," *IEEE Data Engineering*, 1993.
- [Wiederhold 1992] Gio Wiederhold, "Mediators in the Architecture of Future Information Systems," *IEEE Computer*, Vol. 25, No. 3, March 1992, pp. 38–49.
- [Woelk *et al.* 1992] Darrell Woelk, Wei-Min Shen, Michael N. Huhns, and Philip E. Cannata, "Model-Driven Enterprise Information Management in Carnot," in Charles J. Petrie Jr., ed., *Enterprise Integration Modeling: Proceedings of the First International Conference*, MIT Press, Cambridge, MA, 1992.

## Appendix: The Development of Articulation Axioms

Carnot provides a graphical tool, the Model Integration Software Tool (MIST), that automates the routine aspects of model integration, while displaying the information needed for user interaction. The tool produces articulation

axioms in the following three phases:

1. MIST automatically represents an enterprise model in a local context as an instance of a given formalism. The representation is declarative, and uses an extensive set of semantic properties.
2. By constraint propagation and user interaction it matches concepts from the local context with concepts from the common context.
3. For each match, it automatically constructs an articulation axiom by instantiating axiom templates.

MIST displays enterprise models both before and after they are represented in a local context. MIST enables the Cyc knowledge base to be browsed graphically and textually to allow the correct concept matches to be located. With MIST, a user can create frames in the common context or augment the local context for a model with additional properties when needed to ensure a successful match. MIST also displays the articulation axioms that it constructs. The three phases of articulation axiom development are described next in more detail.

In the model representation phase, we represent the model as a set of frames and slots in a Cyc context created specially for it. These frames are instances of frames describing the metamodel of the schema, e.g., (for a relational schema) `Relation` and `DatabaseAttribute`.

In the matching phase, the problem is: given a (Cyc) representation for a concept in a local context, find its corresponding concept in the common context. The two factors that affect this phase are (1) there may be a mismatch between the local and common contexts in the depth of knowledge representing a concept, and (2) there may be mismatches between the structures used to encode the knowledge. For example, a concept in Cyc can be represented as either a collection or an attribute [Lenat and Guha 1990, pp. 339ff].

If the common context's knowledge is more than or equivalent to that of the local context's for some concept, then the interactive matching process described in this section will find the relevant portion of the common context's knowledge. If the common context has less knowledge than the local context, then knowledge will be added to the common context until its knowledge equals or exceeds that in the local context; otherwise, the common context would be unable to model the semantics of the resource. The added knowledge refines the common context. This does not affect previously integrated resources, but can be useful when further resources are integrated.

Finding correspondences between concepts in the local and common contexts is a subgraph-matching problem. We base subgraph matching on a simple string matching between the names or synonyms of frames representing the model and the names or synonyms of frames in the common context. Matching begins by finding associations between attribute/link definitions and existing slots in the common context. After a few matches have been identified, either by exact string matches or by a user indicating the correct match out of a set of candidate matches, possible matches for the remaining model concepts are greatly constrained. Conversely, after integrating an entity or object, possible matches for its attributes are constrained.

In the third phase, an articulation axiom is constructed for each match found. For example, the match between a relational attribute `phone` in model AAA and the Cyc slot `phoneNumber` yields the axiom

$$ist(Cyc, phoneNumber(L, N)) \Leftrightarrow ist(AAA, phone(L, N))$$

which means that the `phone` attribute definition determines the `phoneNumber` slot in the common schema, and vice versa. Articulation axioms are generated automatically by instantiating stored templates with the matches found.