



How Agents from Different E-Commerce Enterprises Cooperate

Qiming Chen, Meichun Hsu, Igor Kleyner
Software Technology Laboratory
HP Laboratories Palo Alto
HPL-2000-108
August 17th, 2000*

E-mail: qchen@hpl.hp.com

inter-enterprise
agent
cooperation,
cooperative
business process
management

Using agent technology to support ECommerce automation is a promising direction. However, the previous "proof-of-concept" efforts do not scale well in E-Commerce automation. An essential reason is that the conventional agent infrastructures are primarily designed for intra-enterprise, group-based agent cooperation, but most E-Commerce applications are based on *inter-enterprise* business partnership. Agents across enterprise boundaries are unlikely to be organized into the same "agent group" and under a centralized coordination.

We tackle the issue of scaling *inter-enterprise* agent cooperation from the following three angles. First, we have introduced the **Point of Presence (POP)** approach for integrating message-based agent communication with interface-based service invocation. This approach allows us to unify the messaging service interface for all the agents, and therefore greatly simplify both server-side and client-side interface implementation and maintenance. Next, we have developed the agent-embedded **cooperative process manager**, for elevating multi-agent cooperation from the conversation level to the business process level, and from centralized process management to peer-to-peer cooperative process management. Finally, we propose the conceptual **separation of agent cooperation messaging network from the bulk data network**. These emerging technologies are integrated with the ECarry agent infrastructure, an autonomous and decentralized system we developed at HP Labs. The feasibility of this approach has been demonstrated in a prototyping system.

* Internal Accession Date Only

Approved for External Publication

© Copyright Hewlett-Packard Company 2000

How Agents from Different E-Commerce Enterprises Cooperate

Qiming Chen, Meichun Hsu, Igor Kleyner

Hewlett Packard Laboratories
1501 Page Mill Road, MS 1U4
Palo Alto, California 94303, USA
+1-650-857-3060
qchen@hpl.hp.com

Abstract

Using agent technology to support E-Commerce automation is a promising direction. However, the previous “proof-of-concept” efforts do not scale well in E-Commerce automation. An essential reason is that the conventional agent infrastructures are primarily designed for intra-enterprise, group-based agent cooperation, but most E-Commerce applications are based on *inter-enterprise* business partnership. Agents across enterprise boundaries are unlikely to be organized into the same “agent group” and under a centralized coordination.

We tackle the issue of scaling *inter-enterprise* agent cooperation from the following three angles. First, we have introduced the **Point of Presence (POP)** approach for integrating message-based agent communication with interface-based service invocation. This approach allows us to unify the messaging service interface for all the agents, and therefore greatly simplify both server-side and client-side interface implementation and maintenance. Next, we have developed the agent-embedded **cooperative process manager**, for elevating multi-agent cooperation from the conversation level to the business process level, and from centralized process management to peer-to-peer cooperative process management. Finally, we propose the conceptual **separation of agent cooperation messaging network from the bulk data network**. These emerging technologies are integrated with the E-Carry agent infrastructure, an autonomous and decentralized system we developed at HP Labs. The feasibility of this approach have been demonstrated in a prototyping system.

Keywords: *Inter-Enterprise Agent Cooperation, Cooperative Business Process Management*

1. From Intra-enterprise Agent Cooperation to Inter-enterprise Agent Cooperation

E-Commerce applications operate in a distributed environment involving multiple parties with dynamic availability, and a large number of heterogeneous information sources with evolving contents. A business partnership is often created dynamically and maintained only for the required duration such as a single transaction. E-commerce activities typically rely on distributed and autonomous tasks for dealing with such operational dynamics. Today, they are initiated and executed primarily by humans. With the goal of automation, conducting them by software agents has been proposed [3,4,19,21]. However, the previous “proof-of-concept” efforts do not scale well in real E-Commerce applications.

Software agents have been studied for many years from all the aspects of software engineering: agent architecture [1,5,6,22,23,27], agent communication language (e.g. ACL, KQML [12,17,20]), agent coordination, conversation management (e.g. FIPA specification [13]), etc. However, previously agent technology primarily focused on autonomy, intelligence, ontology, language, etc, but rarely on scalability. To our knowledge, there is no sizable, inter-enterprise E-Commerce application using a commercial agent system ever deployed. It is time to examine the potential reasons for such a slow adoption.

1.1 Agent Cooperation across Enterprise Boundaries

Internet-based E-Commerce involves multiple enterprises separated by firewalls. One difficulty for the traditional agent technology to fit into this picture consists in the limitation of agent coordination model. Such a model assumes that agents are formed in groups, or **domains**, each group is provided with a *coordinator* for facilitating naming service, resource service, etc. Agents in a group rely on these services to communicate and cooperate. While it is possible for the agents belonging to the same enterprise, it is unlikely for the agents belonging to different enterprises, to form a single agent group, or domain. We see for example, a buyer agent for a retailer and a seller agent for a supplier might not be in the same agent group or under the same coordination (Figure 1). Organizing agent groups into a hierarchy may help, but does not eliminate the difficulty of coordination across enterprise boundaries.




Figure 1: *Group-based coordination does not support inter-enterprise agent cooperation*

One possible solution is to use a CORBA-like service bus [8] for agents to locate each other in peer-to-peer communications. Intuitively, any agent, A, can register a “send-message” service, making it possible for another agent in a foreign domain to send a message to A, using that service. However, if every agent has to register a messaging service in order to receive messages, and every agent has to maintain multiple client side messaging service implementations for all the agents it may need to have a contact with, it is not scalable. We call this “interface complexity” problem.

1.2 Agent Cooperation at Business Process Level

Next, agents cooperate through message exchanges, and this has led to a research topic: conversation management. However, many E-Commerce applications include complex business processes with a large number of concurrent, long-duration, long-waiting and nested tasks, and the flat conversation management lacks the scalability for handling and tracking such sizable applications. Instead, a more robust and scalable approach is to lift agent cooperation from the conversation level to the process level.

Conventionally, a workflow engine is used to provide a centralized scheduling, monitoring and

execution control of business processes, although the tasks that contribute to a process can be distributed (Figure 2) [7,9,15,18,26]. Such centralized process control is appropriate within a single enterprise but not across enterprise boundaries. *Intra-enterprise* process management differs from *intra-enterprise* process management significantly. Different enterprises are often separated by firewalls, have self-interests and individual data sharing scopes. When they are involved in a business process, they are unlikely to trust and rely on a centralized workflow server. Rather, they need support for peer-to-peer interactions [25]. This has become the major impendence for using the conventional centralized workflow systems for inter-enterprise E-Commerce automation. In fact, to our knowledge, there has been no such experience reported.




Figure 2: Centralized workflow control not suitable for cross-enterprise applications

1.3 Agent Cooperation with Bulk Data Exchange

As personalized, continuously running and semi-autonomous computational entities, software agents can be used to *mediate* users and servers to automate a number of time-consuming tasks in E-Commerce. However, agents communicate by exchanging messages, which may not be suitable for bulk data transfer. Routing and caching a large amount of data also impose a considerable burden for agents. For example, as shown in Figure 3, moving data between an operational database and a data warehouse via a software agent is unlikely.




Figure 3: Agents for mediating E-Commerce automation may not be suitable for bulk data transfer

A principle of business process management is the separation of tasks from the processes they contribute to, in the sense that a task does not necessarily know the entire process. Analogously, we can place bulk data transferring and agent cooperation at different conceptual layers, and we believe that is a necessary approach to scaling agent mediated E-Commerce automation.

1.4 Our Solutions

We envisage that the scalability issues introduced above are critical to agent-mediated E-Commerce automation, and these issues are inter-related. We tackle them by the following

approaches.

- We have developed the ***Point of Presence*** (POP) approach that supports inter-enterprise agent communication using a CORBA-like service bus, HP E-Speak [11], but provides a unified messaging interface to overcome the “interface complexity” problem. The use of E-Speak allows agents to communicate across enterprise boundaries, with fine-grained access control, firewall traversal and other infrastructure services. Under the proposed POP mechanism, each agent domain only registers the *messaging service* of the domain *coordinator* with E-Speak. This service then becomes the single gateway to the agent domain, and can be made standard for all the agent domains. Within a domain, the domain coordinator can forward messages to other agents through intra-domain agent communication. Thus on one hand, it is unnecessary for each agent to register an individual message service for receiving messages; on the other hand, an agent only needs a “standard” client-side interface for invoking the above messaging service to contact any agent in any foreign domain, using the domain name as a parameter.
- We have introduced the notion of ***cooperative business process*** and developed agent-embedded Cooperative Process Managers (CPMs), which elevates agent cooperation from the conversation-level to the process-level, and from centralized process management to cooperative process management with peer-to-peer interoperations. A cooperative business process is defined based on a commonly agreed operational protocol, such as the protocol for on-line purchase or auction. However, it is not executed by a centralized workflow engine, but by multiple agents with CPMs collaboratively. More specifically, each execution of a cooperative process, or a *logical process instance*, consists of a set of *peer process instances* run by the CPMs of participating agents. These peer instances share the same process definition, but may have private process data and sub-processes. The CPM of each agent is used to schedule, dispatch and control the tasks that agent is responsible for, and the CPMs interoperate through an inter-CPM messaging protocol to synchronize their progress in process execution. An XML-based Cooperative Process Definition Language, CPDL, extending the process definition language (PDL) [26], is developed for specifying cooperative business processes. This approach represents a technical emerging of, and significant extensions to both workflow technology and agent technology.
- We have developed the notion of ***two-tier agent cooperation infrastructure*** that is characterized by the conceptual *separation of cooperation message network and bulk data network*. This notion is analogous to the separation of signal network and voice trunks in modern telephone infrastructure, where the former is used for routing calls, setting-up and tearing-down connections, and the latter for transferring voice data. In our two-tier agent cooperation infrastructure, the message network is used for agents to exchange control information, meta-data and small size documents [10], and in addition, for setting-up and tearing-down connections between devices. The communication at this tier is message-based, asynchronous. Bulk data network is used for transferring large amount of data. The communication at this tier is generally synchronized.

The approaches proposed above are related. For example, peer-to-peer cooperative process

management requires inter-enterprise agent communication; the separation of cooperation messages and bulk data transfers ensures the required data throughput in agent-mediated E-Commerce applications, and allows the control information and data of business processes to be handled separately, for both efficiency and privacy.

Although the proposed mechanisms are independent of the underlying agent infrastructure, our experience has shown that implementing them using the *E-Carry agent infrastructure*, an autonomous and decentralized system we have developed at HP Labs, has many advantages due to its openness, scalability and flexibility. An E-Carry agent has the ability to load, maintain and start servers and applications dynamically. It also contains an embedded Web server with servlet functionality, enabling its state to be accessed or updated through a browser. Adding the proposed capabilities allows us to provide a migration from the traditional agent infrastructure to a dynamic and distributed middleware framework. The full details about E-Carry architecture will be reported separately.

The significance and feasibility of this work have been demonstrated in a prototype implemented at HP Labs.

Section 2 describes the POP approach for using E-Speak infrastructure to support inter-enterprise agent cooperation. Section 3 discusses peer-to-peer, agent cooperative process management. Section 4 illustrates the architecture for separating agent cooperation message network and bulk data network. Related work will be covered in Section 5 with conclusions.

2. Inter-Enterprise Agent Communication with Unified Messaging Interface

Agents in the same group, referred to as the agent **domain**, can communicate using the naming service provided by the *coordinator* of that domain. However, agents in different enterprises may not form a single agent domain. Instead, they need certain “service bus” to locate each other for peer-to-peer communication. Further, issues such as firewall, security, access control, and even billing, should be taken into account. We have adopted the HP E-Speak service bus, an interface based service provisioning and invocation framework with multiple interconnected *E-Speak Cores*. An E-Speak core provides a set of predefined and extensible *infrastructure services* including authentication, authorization, billing, etc. These infrastructure services represent the major difference between E-Speak and the traditional CORBA-like middleware. In this paper we do not intend to explain E-Speak in detail.

Conceptually, an agent, *A*, can register a “send message” service with E-Speak, making it possible for another agent in a foreign domain to send a message to *A* by invoking this service. In doing so, however, we want to avoid the following possible problems.

The first problem is the “interface diversity” problem we mentioned earlier, that is, if every agent has to register a messaging service in order to receive messages, and every agent has to implement and maintain multiple client side messaging service interfaces for all the agents it

may need to have a contact with, there would be too many interfaces for E-Speak to register, and for an agent to keep.

Furthermore, it is often the case that once an agent can reach a domain, it should be allowed to invoke certain services carried by the coordinator or other agents of that domain. If all those services must be registered, it is not scalable from service invocation point of view.

2.1 Inter-domain Agent Communication with Unified Messaging Interface

In order to unify the messaging interface for inter-enterprise agent communication, we integrate E-Carry and E-Speak in the following way.




Figure 4: Unify messaging service interface to simplify inter-enterprise agent communication

The *coordinator* of every agent domain carries a *messaging service*, and registers this service with E-Speak. This service then becomes the single entrance to the agent domain. We refer to it as the **Point of Presence (POP)**. Inside the domain, the coordinator can forward messages to other agents. Thus, it is unnecessary for each agent to register an individual message service, since the coordinator provides a gateway for any foreign agent to reach any agent in that agent-domain. Further, services provided either by the coordinator or by other agents in that domain, may be invoked through messages. This also eliminates the need of registering every individual service, maintaining therefore the single POP for service invocation as well (note that, however, there is no restriction on service registration, i.e. open more POPs, as needed).

Registering only the general messaging service also simplifies and unifies the client interface for sending messages. Every E-Carry agent only needs to be provided with a “standard” client-side stub for invoking the above messaging service, with the domain name encapsulated in the message envelop. By invoking this message service, the agent can contact any agent in any foreign domain, with messages routed by the coordinator of that domain.

Through the messages an agent sends to a foreign domain, services provided by the agents (including the coordinator) of that domain can be invoked, and such invocation is message-based without keeping continuous connection. Below we give some details about the messaging service used for inter-enterprise agent communication.

At “server-side”, the messaging service provided by the coordinator of an agent domain, D , is

registered with E-Speak. The interface of this service includes a single method

void sendMsg(String message)

The interface name, say *AgentMsgService*, plus a property “description” indicating the domain name, uniquely identify this service. In an intra-domain message, the destination is simply expressed by the receiver’s name. In an inter-domain message, the destination is expressed by

espeak:domain_name/agent_name

where *espeak* is the service bus, a concept at a higher-level than transport. For example, when our approach is extended to use http as the service bus, the logical address of an agent should be

http:domain_name/agent_name

In this case the Web server embedded in an E-Carry agent will be used. On the “client-side”, a standard implementation of the above interface is embedded to each E-Carry agent, as the “e-speak message dispatcher”.

As shown in Figure 4, for example, when agent A_1 in domain D_1 attempts to contact agent B_2 in a foreign domain D_2 , A_1 invokes function *sendMsg*, that is registered with E-Speak by the coordinator of domain D_2 as

Name='AgentMsgService' and Description=' D₂'

to send a message with destination *espeak: D₂/B₂*. The message is first received by the coordinator of D_2 , and then forwarded to B_2 by the coordinator. At the first step, the E-Speak infrastructure service is called; at the second step, the local naming service provided by the domain coordinator, is employed. If the sender intends to invoke a service provided by the coordinator or another agent of D_2 , the result of that service will be sent back to it via E-Speak as well.

2.2 Message Subscribing/Publishing

With E-Speak, agents from different domains can also communicate in the publish/subscribe mode. For example, when an agent intends to buy some electronic parts, instead of checking the vendor agents one by one, it can publish an availability-check message, and E-Speak can broadcast this message to all the vendor agents who subscribe this message.

The message publish server carried by an E-Carry agent and registered with E-Speak, implements the same interface as *AgentMsgService*, with a single method *sendMsg(String message)*. The agent represents a *virtual agent domain: MsgPublisher* (Figure 5). Therefore, when an E-Carry agent tries to publish a message, it sends the message to the *MsgPublisher* server, using *espeak:MsgPublisher* as the address, just like sending a message to an agent domain. To subscribe to message AvailabilityCheck, for instance, the subscribing agent should send the following message to *espeak:AgentMsgPublisher*.

```
<MESSAGE type="SUBSCRIBE" from="espeak:D2/A3" to=" espeak:MsgPublisher" interpreter="xml.default">
  <CONTENT> <MESSAGE_NAME> AvailabilityCheck </MESSAGE_NAME>... </CONTENT>
</MESSAGE>
```

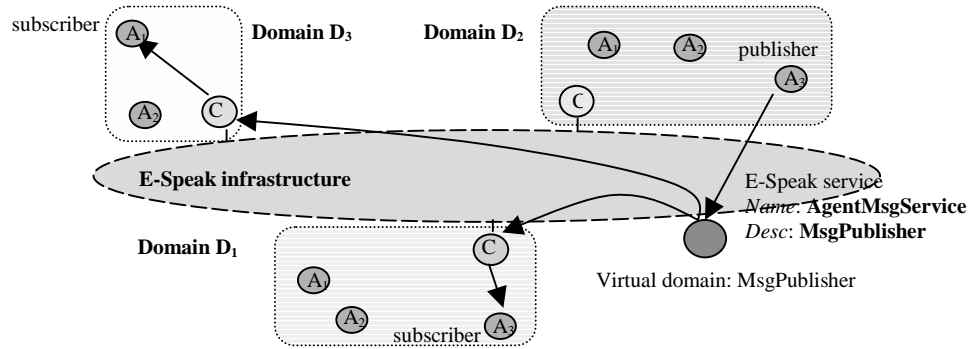



Figure 5: Publishing / Subscribing based communication

2.3 Message-based, Asynchronous Service Invocation

The POP approach has an additional advantage when used for agent service invocation, that is to keep the message-based, asynchronous client-server communication.

E-Commerce is a *plug and play environment*. Services need to be provided on demand. Business partnerships (e.g. between suppliers, resellers, brokers, and customers) need to be created dynamically and maintained only for the required duration such as a single transactional process. To support such dynamics, an E-Commerce infrastructure must support the cooperation of loosely coupled e-business systems.

Interface oriented, CORBA-like middleware is based on a technology for integrating tightly coupled local systems. It doesn't fit into the picture of e-business since it is too dependent on Remote Procedure Call (RPC), or Remote Method Invocation (RMI), a form of synchronous communications in which networked devices maintain continuous contact. Synchronous communications lack the flexibility for plug-and-play e-business, and don't cope well with firewalls. In the contrary, Web-based communication is asynchronous, where a message is sent when the line is available and contact is interruptible. This feature is also represented by IBM's MQ Series and Microsoft's MSMQ messaging middleware.

The E-Carry capability for carrying actions allows a domain coordinator to carry multiple services not registered with E-Speak, but invoked through messages. Services carried by other agents in the domain can also be invoked on the message basis. It is unnecessary (but optional) to register agent provided services with E-Speak. This well fits in agent cooperation where most of the services are kept message-enabled.

3. Inter-Enterprise, Peer-to-Peer Agent Cooperation at Business Process Level

Many E-Commerce applications include complex business processes, and involve multiple enterprises. To handle such applications through simple agent conversation is not scalable, and through a centralized server is unreasonable. This has motivated us to lift agent cooperation from the conversation level to the process level, and from centralized process management to peer-to-

peer cooperative process management.

3.1 From Centralized Process Management to Cooperative Process Management

A business process specifies the integration and synchronization of multiple steps, each step represents a logical piece of work that contributes to the accomplishment of the whole process. Although these tasks and the agents executing them can be distributed, they are scheduled and coordinated by a centralized workflow engine. Typically a business process includes a *data packet* containing the *process data* for control flow and data flow, and tasks can manipulate the process state by updating these data. However, consider a purchase process involving tasks belonging to different enterprises, e.g. the buyer and the seller. It is unrealistic to have the buyer and the seller coordinated by a single workflow engine, and it is unreasonable for them to put their private data (e.g. negotiation thresholds) into the common process data packet for flow control.

Our solution to the above problem is based on extending process management from the one-server model to the multi-server peer-to-peer model, a shift from *centralized process management* to *cooperative process management*.

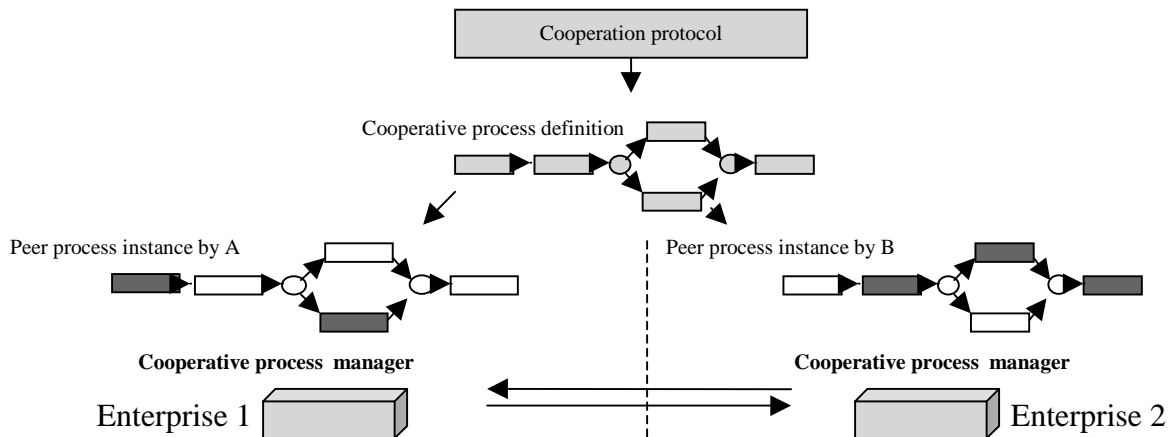


Figure 6: Peer-to-peer collaborative process management

We introduce the notion of *cooperative business process*. An inter-enterprise cooperative process involves multiple parties. It is defined based on the corresponding business protocols, and such a definition becomes the common template for all the participating parties to share. However, an execution of a cooperative process, viewed as a logical instance of the process, actually includes multiple peer instances that are not executed by a centralized workflow engine but by multiple CPMs and synchronized through peer-to-peer communication. These peer instances share the same process definition, but may have private process data and sub-processes. The CPM at each side recognizes its own share of the tasks (*shaded* in Figure 6) based on role-matching. For example, an on-line trading process, say *P*, is executed collaboratively by a seller and a buyer in such a way that each peer CPM runs an individual

process instance of P . For the CPM at buyer side, it is only responsible for (schedule and dispatch) the tasks to be executed by the buyer, such as preparing a purchase order and making a payment. Similarly the CPM at seller side is only responsible for the tasks belonging to the seller. The CPMs exchange task execution status messages for synchronization.

3.2 Agent Embedded Cooperative Process Manager

We have implemented CPM and integrated it into *E-Carry* agent platform. This novel integration achieves two purposes: on the one hand, it provides an implementation and execution platform for a CPM system; on the other hand, it elevates multi-agent cooperation in E-Carry from the conversation level to the process level for mediating E-Commerce applications.




Figure 7: *E-Carry agent with embedded CPM*

As shown in Figure 7, the functionality of CPM is *embedded* into the service tier of E-Carry,. The agent with CPM embedded can then be used as a CPM server. Since a CPM server can also be viewed as an agent, it is possible to consider the notion of *personalized CPM engine*. That is, each logical entity of an enterprise, say, an *electronic parts buying agent*, could have its (or his) own CPM engine to represent it (or him). When participating in inter-enterprise collaboration, it has its (or his) own CPM server executing peer process instances. Besides of acting as a CPM server, an E-Carry agent can also perform activities.

The CPM embedded in an E-Carry agent interacts with the hosting agent through a set of internal messages. The communication between agents is made through inter-agent message exchange. A set of agent messages specific to cooperative process management, are defined, and a corresponding message interpreter is provided for each agent. The E-Carry agent has the capability to load and switch interpreters based on message ontology types thus can easily handle applications in different contexts.

3.3 Cooperative Process Definition

To explain how the proposed cooperative process management approach extends the current workflow technology, we adopt the usual concepts of business process modeling in the following discussions. A *process* is modeled as a DAG with nodes representing the steps, or tasks, and arcs representing the links of those steps. A *work-node* represents a step (*task*) and associated with an activity, i.e. a piece of work that contributes to the accomplishment of the process, that may be executed either by a program (e.g. a software agent) or by a human worker. A process is associated with a packet of *process data*. When an activity is launched, a subset of the process

data, sub-packet, is passed to it; when it is completed, together with task status information, the sub-packet, possibly updated during the task execution, is sent back for reconciliation with the process data packet. A *route-node* specifies the rules and conditions for flow control, process data update, etc. Conventionally, a process execution creates a single process instance. However, for a cooperative process, the logical instance of each execution includes multiple peer process instances. Further, a cooperative process may have multiple concurrent executions. To support cooperative processes, the minimal extensions to process definition include the following.

A cooperative process has a list of **process-roles**, indicating the logical participants. For example, if a simple purchase process has two roles, “buyer” and “seller”, then there are two peer instances involved in its execution, one at the CPM for “buyer” and another at the CPM for “seller”. These two peer instances are assigned roles “buyer” and “seller” respectively.

A work-node has a **task-role**, and that must match one of the process-roles. In the above example, tasks can have roles “buyer” and “seller”. If the role of a task is “buyer”, it is only executed in the peer process instance with process-role “buyer”. Note that an activity also has a role called **activity-role**, such as “invoice-generator”, meaning that this task should be executed by (or dispatched to) an agent playing the “invoice-generator” role in this process. The notion of activity-role can be found in regular business process specifications.

In an inter-enterprise cooperative process execution, each party wants to keep some of the process data private. For example, the buyer in one enterprise and the seller in another enterprise do not want to expose their thresholds during price negotiation. In the process definition, **templates** for holding the definitions and initial values of process data objects can be specified. Furthermore, the **sharing scope** of the data objects is specified. A template may be *public*, i.e. sharable by all process-roles (and thus by all peer process instances) or *process-role specific*. A role-specific template is used by the peer process instances of the given roles (one or more) only, and such templates can be made different for different process-roles. Consider a cooperative process with roles “buyer”, “seller” and “bank”; some data are private to “buyer”; some are sharable by “buyer” and “seller”; some are public to all three roles. The initial data packet of a peer process instance consists of the appropriate templates, where the sharing scope of each data object is marked. This data packet can be updated or expanded at run time.

A task may represent a private sub-process with a private data packet. The sub-process binding is dynamic, that is, bound at run time. This allows a private sub-process to be designed separately from the host process. The process data of the internal sub-process is entirely private to the party executing the sub-process.

An XML[2]-based Cooperative Process Definition Language, CPDL, extending the process definition language (PDL) by capturing the above notions, is developed for specifying cooperative business processes.

3.4 Cooperative Process Execution

An execution of a cooperative process consists of a set of peer process instances run by the CPMs of the participating agents. These instances share the same process definition but they have additional properties and may have private process data and sub-processes.

Each peer process instance has a role that must match one of the process-roles. When a peer process instance is launched by a CPM at the seller side, for example, the process-instance-role is “seller”, and the CPM is only responsible for scheduling and dispatching the tasks with task-role “seller”. When executing a cooperative business process, the *player* of each peer process instance must be specified and bound to the corresponding process instance role. In addition, a *logical identifier* for this execution must be obtained. These pieces of information are captured as properties in every peer process instance. They are further described below.

The **players** of a cooperative process indicate the participating agents with embedded CPMs. A player is associated with four attributes.

- The role, e.g. “buyer” or “seller”, of the given process instance running at the CPM that represents this player. Note that without binding to a peer process-instance, a CPM does not have a fixed role.
- The domain name of the agent
- The local name of the agent within the domain to represent the player. For example, *corp.hp.com/buying_agnet* may be a player playing the *buyer* role in a *purchasing* business process, whose peer agent in this process might be *us.oracle.com/sales_agent*.
- The inter-domain messaging service infrastructure, such as HP E-Speak, that provides messaging services for inter-domain agent communication. The messaging service infrastructure is capable of delivering messages among multiple domains. As mentioned above, when peer agents participating a cooperative process execution rely on E-Speak to reach each other, the addressing structure is *espeak:domain_name/local_name*.

A **coop-key** is used to identify a logical instance of a cooperative process, that is, to correlate and synchronize the multiple peer instances of the execution of a single cooperative process. All the messages exchanged for that execution are marked by a unique coop-key. In our implementation, the CPM of each E-Carry agent can run multiple process instances concurrently, and each instance has a local ID. Each CPM engine maintains a mapping table between coop-keys and local process instance IDs. When a message relating to the execution of a cooperative process is received, the coop-key is used to identify the corresponding local process instance.

Let us use a simple example shown in Figure 8 for explanation purpose. The sample cooperative process for on-line purchase defined based on the OBI (Open Buying on Internet) protocol, *obi_process*, has process-roles “buyer” and “seller”. Each logical instance of *obi_process* has two peer-instances run at two peer CPMs, *A* and *B*, one at the buyer side and one at the seller side. It has several tasks (steps) including T_1 (make purchase order), T_2 (process purchase order), etc. T_1 is a step the buyer is responsible for, so its role is “buyer”, while the role of T_2 is “seller”. *A*, running the peer instance with role “buyer”, is responsible for executing T_1 , and *B*, running the

peer instance with the role “seller”, is responsible for executing T_2 . The initial data packet for process-role “buyer” and “seller” can be different. The execution of a cooperative process is carried out in the following way:

- CPM A, representing the process-instance-role of “buyer”, initiates a buyer-side process instance P_b and through messaging, tells CPM B to create a seller-side peer process instance P_s .
- A dispatches and executes T_1 , and upon receipt of the *task return message*, r_1 , forwards it to *all* other players of the process, in this case, simply B. Both A and B update their process state and schedule the possible next step of their own peer process instance based on that message.
- When A proceeds to activity T_2 , since the role represented by A does not match the role of T_2 , A simply waits for the peer server, that is B in this example, to handle it at the peer site.
- When B proceeds to activity T_2 , since the role represented by B matches that of T_2 , T_2 will be handled by B.
- The execution of peer process instances at both peer CPMs progress in this way, towards their ends.

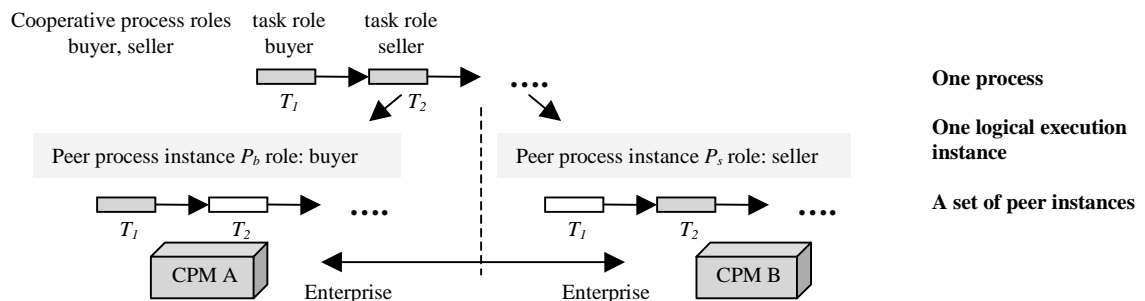


Figure 8: Peer-to-peer collaborative process management

An activity is dispatched to a software agent or a human user to perform, and upon its termination, a task return message is sent back and used to trigger the next step in process execution. Such a task return message contains the coop-key of the logical process instance, the handles (local Ids) of the process instance, task, and activity, the activity execution status, and the sub-packet, i.e. the subset of process data passed to the activity.

When a task return message comes back to the local CPM engine, it contains all the above information. Since the sub-packet of the process data passed to the activity may be updated during task execution, it must be reconciled with the process data packet after returning. However, before such a message is forwarded to a peer player, only the updated data elements that are shared with the player are retained. (Recall that the sharing scope of each data element is specified in the process definition.)

A key design issue is to maintain the right order of message processing. For various reasons the messages may not be delivered in the original order. *Queuing technique* and the knowledge drawn from process definitions are used to resolve the out-of-order message delivery problem. Each CPM has a **queuing server**, in addition to the regular message queue handler. This queuing server is *workflow specific* as it interfaces to the process definition handler and the process instance log handler, using process definitions and execution histories to make

operational decisions. When a message is received, the queue server checks if it is ready to be processed, if not, queue the message; and if it is, further checks if the change makes any queued message ready to be processed. It also responds to CPM internal events such as process instance status changes.

Turning agent cooperation from conversation-level to process-level is a natural and necessary move. In general, businesses collaborate following certain rules, such as “if you send me a price request then I will send you a quote”, and “if the quote I sent you is acceptable, then you will send me an order”. These rules include sequences of steps to form a business process. Such business collaboration usually involves multiple agents, each responsible for managing or performing certain tasks that contribute to the process. Adding inter-enterprise cooperative process management capability into agent-based systems is critical for these business collaborations.

4. Separating Bulk Data Transfer from Agent Cooperation Messaging Network

In previous agent cooperation frameworks, there is no explicit separation of control information and application data agents exchange. A business process instance is also combined with control flow and data flow. In some workflow systems such as Lotus-Notes, control flow and data flow are fully combined.

We have described early in this paper that for data privacy purpose, in cooperative process management we define the scopes of data visibility among peer CPMs and separate the non-sharable data from the sharable data. The messages for synchronizing cooperative process execution contain no sensitive private data, and data exchanges are handled at the task execution level, with explicit privacy control. Here we revisit this issue beyond process management.

In agent cooperative activities, agents communicate by exchanging messages. The content of agent conversation may be requests, inter-CPM messages, or moderate size business documents. However, transfer bulk data via software agents is neither efficient nor necessary. It is inefficient since the data transfer throughput of software agents is far less than large-scale data servers. It is unnecessary since the percentage of control data (to be handled by an agent) in a large data stream is fairly small. This has led to the idea of separating bulk data network from agent cooperation message network conceptually. We believe this is the key to the success of agent mediated E-Commerce automation, although it requires substantial work in standardization. In fact, such a system configuration is analogous to a modern telephone system.

In the telecommunication infrastructure, an important architecture feature is the separation of the voice network and the signal network. The voice trunks connect end-offices for voice data traffic. The end-offices are connected to and controlled by the SS7 signaling network (Figure 9) for switching, monitoring, traffic information gathering and so on. Such separation aims at enhancing the controlling and monitoring functionality over the whole infrastructure. With the separation of the signal network from the voice trucks, for example, the caller-id service can be

provided to allow a callee to know who the caller is, before picking up his phone.




Figure 9: Telephone network: separation of signal network from voice (data) trunks

Our proposed infrastructure is illustrated in Figure 10, where bulk data transfer is separated from the agent Cooperative Message Network (CMN). As mentioned above, the CMN still can be used for agents to exchange moderate size data as message contents. However, large data transfer such as database loading, must be separated from agent message delivery.

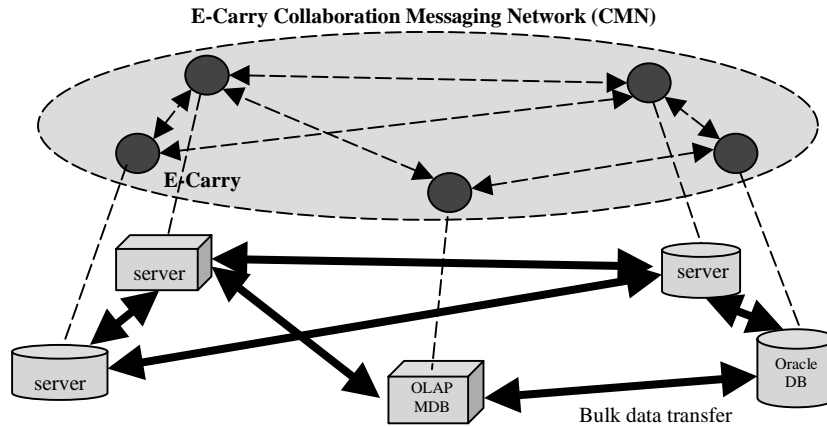


Figure 10: Separate collaboration message network from Data Network (Payload)

We can once again relate this concept to cooperative process management, where processes and tasks are already separated into two system layers. Bulk data transfer can be handled at task level without via process level messaging. Such separation can be introduced to the general agent mediated E-Commerce automation, that is, separate the tasks involving bulk data transfer, from the cooperation activities between agents. For instance, with the servlet capability, E-Carry agents can respond to certain messages by setting up corresponding services, including enacting the actions to set-up or tear-down connections between the systems they represent. For example, an E-Carry agent can execute a program that starts an Oracle Express OLAP script for loading data into its multi-dimensional database from a relational database.

5. Comparisons and Conclusions

We have developed several inter-enterprise agent cooperation approaches for supporting agent-mediated E-Commerce automation. These approaches represent an integrated solution based on emerging technologies and applications. The agent embedded CPMs enable peer-to-peer agent

cooperation at the business process level, and communicate using the unified messaging interface under the POP mechanism. This mechanism also supports message-based, asynchronous application invocation without requiring continuous connection. By separating bulk data transfer from cooperative message network, the data throughput limitation of agents can be overcome.

Compared with the group and group-hierarchy based multi-agent systems [12-14,16], we scale agent cooperation across enterprise boundaries by supporting peer-to-peer, non-coordinated inter-domain communication and collaboration.

Compared with RPC and RMI [8], the POP approach allows us to maintain, to the maximum, the benefits of asynchronous communication for Internet based applications. The regular use of RPC/RMI based middleware is to specify a service function in an interface class. Although multiple functions may be grouped into a single interface class, conceptually every service function must be specified and registered. In order to contact multiple interfaces, multiple client side implementations of those interfaces must be provided. Use this mechanism allows agents to reach out the local domains, but would require handling a lot of interfaces both for service provisioning and for service invocation. The proposed POP approach provides a unified interface for inter-domain agent communication, and therefore can reduce the above complexity. The use of E-Speak further offers the benefits of infrastructure services.

Compared with agent conversation management [13], the use of CPMs allows us to lift agent cooperation from the conversation level to the process level, and from centralized coordination to peer-to-peer collaboration.

Compared with existing workflow systems [26], the proposed cooperative process management can be used to enhance the collaboration of business partners and support inter-enterprise business process executions. This represents a novel extension to the workflow technology.

Compared with the conventional process federation and RosettaNet PIP approach [24], we conclude that our approach is capable of supporting PIPs. However, our approach goes beyond PIP in the following aspects. First, cooperative process management is based on *process-level business protocols* and PIP approach is based on *interface tasks*. PIPs expose individual “handshake” or conversation points of partner processes, but not a process level view to their cooperation. Second, we have a peer-to-peer execution model for cooperative processes but the PIP approach does not. In PIP approach, the execution of partner processes are not related and synchronized at process-level. Each party sees the trees, not the forest.

From the above comparison we can see the uniqueness of the proposed approaches in supporting inter-enterprise agent cooperation, the key to realize and scale agent-mediated E-Commerce automation. The significance and feasibility of this work has been demonstrated in a prototype implemented at HP Labs. We plan to further extend this system to a scalable, dynamic, inter-enterprise middleware framework.

REFERENCES

- [1] Aglets, "Programming Mobile Agents in Java", IBM, <http://www.trl.ibm.co.jp/aglets/>, 1997.
- [2] T. Bray, J. Paoli, C. M. Sperberg-McQueen, "Extensible Markup Language (XML) 1.0 Specification", February 1998, (<http://www.w3.org/TR/REC-xml>)
- [3] A. Chavez and P. Maes, Kasbah: An Agent Marketplace for Buying and Selling Goods, Proc. of PAAM96, 1996.
- [4] Q. Chen, Meichun Hsu, Umesh Dayal, Martin Griss, "Incorporating Multi-Agent Cooperation, Dynamic Workflow and XML for E-Commerce Automation", Proc. Fourth International Conference on Autonomous Agents, 2000, Spain.
- [5] Q. Chen, P. Chundi, Umesh Dayal, M. Hsu, "Dynamic-Agents", International Journal on Cooperative Information Systems, 1999, USA.
- [6] Q. Chen, P. Chundi, U. Dayal, M. Hsu, "Dynamic-Agents for Dynamic Service Provision", Proc. of 3rd Int. Conf. on Cooperative Information Systems (CoopIS'98), 1998, USA.
- [7] Q. Chen and Umesh Dayal, "A Transactional Nested Process Management System", Proc. of International Conference On Data Engineering, 1996.
- [8] CORBA, "CORBA Facilities Architecture Specification", OMG Doc 97-06-15, 1997.
- [9] U. Dayal and M. Hsu and R. Ladin, "A Transactional Model for long Running Activities", Proc. VLDB'91, 1991.
- [10] Document Object Model, <http://www.w3.org/DOM/>
- [11] E-Speak, <http://www.e-speak.net/>
- [12] T. Finin, R. Fritzson, D. McKay, R. McEntire, "KQML as an Agent-Communication Language", Proc. CIKM'94, 1994.
- [13] Foundation for Intelligent Physical Agents(FIPA)- FIPA97 Agent Specification, <http://www.fipa.org/>
- [14] R. J. Glushko, J. M. Tenenbaum and B. Meltzer, "An XML Framework for Agent-based E-Commerce", CACM 42(8), March, 1999.
- [15] N.R. Jennings, P. Faratin, MJ. Johnson, P O'Brien & ME Wiegand, "Using Intelligent Agents to Manage Business Processes". Proc. of PAAM96, U.K., 1996, pp. 245-360.
- [16] N. R. Jennings (1999) "Agent-based Computing: Promise and Perils" Proc. IJCAI-99, Sweden. 1429-1436.
- [17] T. John, Intelligent Agent Library/Factory, release 4 (<http://www.bitpix.com>)
- [18] N.Krishnakumar and A.Sheth"Specification of workflows with heterogeneous tasks in meteor", Proc. VLDB'94, 1994.
- [19] P. Maes, R. H. Guttman and A. G. Moukas, "Agents that Buy and Sell", CACM 42(8), March, 1999.
- [20] S.A. Moore, "KQML and FLBC: Contrasting Agent Communication Languages," proceedings. 32nd Hawaii international conference on system sciences, 1998,
- [21] A.G. Moukas, R. H. Guttman and P. Maes, "Agent Mediated Electronic Commerce: An MIT Media Laboratory Perspective", Proc. of International Conference on Electronic Commerce, 1998.
- [22] Odyssey, "Agent Technology: Odyssey", General Magic, <http://www.genmagic.com>, 1997.
- [23] B. Perry, M. Talor, A. Unruh, "Information Aggregation and Agent Interaction Patterns in InfoSleuth", Proc. of CoopIS'99, UK, 1999.
- [24] Rosetta-net, www.rosettaNet.org.
- [25] M. Rusinkiewicz, W. Klas, T. Tesch, J. Wasch, P. Muth, "Towards a Cooperative Transaction Model – The Cooperative Activity Model", VLDB'95, 1995.
- [26] Workflow Management Coalition, www.aiim.org/wfmc/mainframe.htm.
- [27] Voyager, "Voyager Core Package Technical Overview", Object Space, http://www.objectspace.com/voyager/technical_white_papers.html, 1997.