

## **Coordinating Distributed Objects with Declarative Interfaces**

Narinder Singh\*, Mark A. Gisi  
Software Technology Laboratory  
HPL-96-75  
June, 1996

object-oriented,  
distributed,  
enterprise-wide

This paper presents an architecture that supports coordination among loosely coupled distributed objects. The architecture has two components: objects that provide a declarative specification of their interface, and system programs that reason with these specifications to provide sophisticated interoperation services. Traditional object-oriented interoperation technologies rely on procedural interface specifications that do not address the semantics of the operations supported by the object. In addition, traditional approaches provide only limited support for automatic interoperation in a dynamic environment. For instance, a resource that is available at compile time may not be available at runtime, or a better resource may become available at runtime. Interoperation based on machine-processable specification of object interfaces reduces the coupling (interdependence) between a client and a server, and also shifts the burden of coordination from the programmer to the system.



## **Preface**

This paper describes one of several activities that was part of the Agent Based Software Interoperation (ABSI) investigation. The purpose of the investigation was to evaluate how simple inferencing technologies could be used to support application interoperation in an industrial setting. The activity discussed in this paper demonstrates how simple inferencing can support automatic coordination among interoperating distributed objects. Other ABSI activities explored additional work in coordination, application knowledge sharing, data translation and multicast communication. These activities are documented elsewhere.

## 1 Introduction

A number of object-oriented technologies have been developed to support inter-operation among applications distributed throughout an enterprise (e.g., CORBA [16, 17, 14, 15]). These technologies enable objects residing in one application to be accessed by objects residing in another. They also enable an entire application to be represented as a single object.

These technologies provide a number of system services that support a model of programming that is similar to programming in a single process (address) space. Although objects may be distributed over a network, conceptually they are viewed as belonging to a single application. Programs are constructed by one or a small group of programmers who have complete knowledge about the kinds of interactions that take place among the objects. The relationships between objects can be described as being tightly coupled and static.

However, the single-process space abstraction is not always appropriate for modeling interactions that occur among distributed objects supporting different functions of an enterprise. Distributed systems that support enterprises continually change and evolve over time, creating a dynamic environment. Furthermore, although objects that support different functions of an enterprise are required to interoperate from time to time, they often maintain high degrees of autonomy, that is, they do not relinquish total control to a central authority. Instead of viewing a collection of enterprise objects as belonging to a single application, it is more appropriate to view each object as maintaining a high degree of autonomy, interoperating with other enterprise objects to support different threads of activity. Examples of activities might include: to print a file, process a customer order, perform group scheduling, or obtain the latest temperature reading of a patient. The relationship between enterprise objects can be described as being loosely coupled and dynamic.

It is impossible for one or a small group of people to have complete knowledge about the different kinds of interactions that can take place during runtime. There is too much information and it changes too rapidly. A resource that is available at compile time may not be available during runtime, or a better resource may become available. Implementing objects that interact in dynamic runtime environments is a difficult task. We need to develop more sophisticated system services that assist object implementors in dealing with incomplete knowledge.

System services should be able to acquire, store, update, and reason about existing resources and object capabilities. In order for them to reason about such information, we need a precise and semantically rich way of describing an object's capabilities and properties. Unfortunately, conventional procedural interfaces provide little semantic information in a machine-processable form.

We address this limitation by providing objects with a second declarative interface so that system services can utilize simple inferencing techniques to reason about them. A declarative interface provides a semantic description of a server object's capabilities and properties. It also enables client objects to describe their requests declaratively. A client can state "what" it would like as opposed to specifying "how" to achieve it. That is, the client describes what

it would like; the system task determines how to satisfy the request. This reduces the complexity of a client's implementation in two ways. First, the steps required to satisfy a request can change as the environment changes; therefore this logic does not need to be embedded in the client's implementation. Second, the client doesn't have to know about the precise details of the interfaces of the servers that to satisfy the request, thereby reducing the coupling between client and server.

The ideas presented in this paper have been developed within a project focused on constructing distributed intelligent autonomous agents [8]. We believe a more restricted form of this work could be developed to extend existing distributed object models (e.g., CORBA and OLE). We describe the benefits of providing objects with a semantically rich declarative interface. We also describe services that utilize simple inferencing techniques to reason about existing resources that can assist object implementors in dealing with the complexities associated with large, evolving distributed environments.

## 2 Related Work

Conventional object models, e.g., CORBA [1, 16, 17] and the OLE/COM [14, 15], provide object communication infrastructures. Objects define their interfaces using a language referred to as IDL (Interface Definition Language). Figure 1 provides two simple interface definitions: print a PostScript file and translate a Microsoft WORD file to PostScript.

```
interface PrintServer
{
    void PrintPostscript(in file theFile)
}

interface WordServices
{
    void TranslateWordToPostscript (in file wordFile, out file postscriptFile)
}
```

Fig. 1. Interface Definition Language (IDL) Example.

IDL provides a syntax for object communication. In order to provide system services that support coordination, a mechanism is needed by which an object can provide semantic information that describes its behavior.

The ANSA project realized the limited expressiveness of IDL. They developed an architecture that supports interoperability among heterogeneous telecommunication services and distributed computer applications [10]. ANSA extended the expressiveness of IDL by enabling a server object to augment its specification

with a collection of properties (name/value pairs). Properties provide semantic information that describes different instances of a service. For example, different instances of the print server in Fig. 1 could have properties describing its location within an enterprise (e.g., building and floor), the different paper formats it can handle (e.g., 8x11, 11x17, A4) and whether it can print in color or not. Figure 2 illustrates this extension to IDL.

```

interface PrintServer
{
    void PrintPostscript(in file theFile)
    Properties:
    {
        Building : {A, ..., F}
        Floor : {1, ..., 5}
        size8by11 : {true, false}
        size11by17 : {true, false}
        color : {true, false}
    }
}

interface WordServices
{
    void TranslateWordToPostscript (in file wordFile, out file postscriptFile)
}

```

**Fig. 2.** IDL Extended to Include Attributes.

In addition to supporting more expressive interface specification, ANSA also provides a system service called a trader [3], that utilizes this information to facilitate interoperation. A server registers an interface specification (along with a collection of properties) with the trader, and clients make requests to the trader to help find a particular instance of a server. With each request, the client provides values for those properties it cares about. The trader uses this information to search its repository for an instance of a server that matches the client's request. The trader then passes a reference to the server back to the client. The advantage is that the reference is located at runtime, when the lookup is delegated to the trader. This reduces the complexity of a client's implementation and adds flexibility because binding takes place at runtime.

### **3 Declarative Interfaces and Intelligent System Services**

When an IDL specification is augmented with name/value pairs, an object's specification begins to resemble a declarative representation. Unfortunately, there are limits to how far IDL can be extended because it is inherently a syntactical

representation. In this section we describe declarative interfaces and present a sophisticated system service called the facilitator, which utilizes simple inferencing techniques to reason about an object's capabilities. The facilitator is a system service that bridges the gap between the objects that provide services and the objects that request those services.

Individual programmers can create objects without knowledge of the data structures and algorithms of other objects, and without knowledge of the hardware configuration in which the objects will run (this is similar to CORBA). However, in traditional systems, the object interface descriptions are in IDL, and the meaning of the IDL interface cannot be defined in IDL (e.g., the fact that a given interface returns the list price of a product). The meaning of the service provided by an object must be known by external means, and there is no support for standardizing or defining the semantics of the interfaces. This requires manual coordination between programmers in a standard CORBA-like environment to agree on the meaning of interfaces. In a large setting, no human can have such global knowledge, and the traditional coordination approaches break down. A declarative interface addresses these issues directly by providing a fixed semantics and by providing a mechanism for examining and defining the vocabulary used.

The system services provided by the facilitator are well suited to environments that are dynamic. In a dynamic environment, resources come and go, and it is impossible to make decisions at compile time. The services provided by the facilitator are based on explicit runtime dynamic data about available services, including other applications, resources, and so forth. Facilitation services use these dynamic specifications at runtime to determine the best way to service a request. The thesis of the paper therefore is: by providing an object with a second semantically rich machine-processable specification, new high-level coordination services, that utilize simple inferencing techniques, can be constructed.

In this section we describe (1) the language for declarative specifications and (2) a new system service, the facilitator, that utilizes simple inference to manage the coordination among client and server objects.

### 3.1 Agent Communication Language

The language we use for describing a declarative interface, developed in the artificial intelligence (AI) community to support interoperation among distributed autonomous agents [7, 8], is called Agent Communication Language (ACL). The use of ACL in the facilitator architecture is described in Sect. 3.2.

An object provides the facilitator with a declarative description of its capabilities and properties at runtime by sending it a collection of ACL messages. ACL has three components: (1) a vocabulary (a domain-specific semantic part), (2) a content language called KIF (Knowledge Interchange Format), and (3) a communication (wrapper) language called KQML (Knowledge Query and Manipulation). An ACL message is a KQML message which consists of a communication directive and a semantic content expressed in terms of the vocabulary in KIF. The communication directive instructs the facilitator on how to process

the content of the message. We describe the three components in the following subsections.

KIF and KQML were developed as part of the ARPA knowledge sharing effort, and both are being used by a number of different research groups. KIF has been standardized by ANSI and is currently under consideration for standardization by the ISO. KQML is being evaluated by the Object Management Group (OMG).

**Vocabulary.** The vocabulary of ACL is listed in a large and open-ended dictionary of words appropriate to common application areas (e.g., electronic commerce, medical industry, and so forth) [9]. Each word in the dictionary has an English description for use by humans in understanding the meaning of the word, and each word has formal annotations for use by programs. The dictionary is open ended to allow for the addition of both new words within existing areas and new application areas.

For example, consider an application that translates files from one format to another. We need to choose a vocabulary for file names, file types, file objects, and the translation of files. First, we need to select the names of individual files, e.g., we could choose "f1.tex" to be the name of a particular file. We need to choose symbols to refer to the types of files, e.g., `latex`, `ps`, `dvi`, `rtf`, `word`, etc. We also need to select a symbol for a function that maps the name of a file and its type to a file object, e.g., `file`. In addition we need to select a symbol for a relation that states that one file is the translation of another, e.g., `translation`.

Sharing information within a community demands agreement about the meaning of symbols, e.g., the meaning of the vocabulary for objects, functions, and relations in ACL. In a small setting (e.g., among a small group of programmers), it may be possible to mandate a single vocabulary for all objects to use in communication. However, in a large setting, this is impossible. More than likely, different communities in such an environment assign different meanings to the same symbols. Such an environment requires support for a collection of vocabularies. We have developed a framework for partitioning vocabularies (and defining the mappings between them) based on a Name-Space Context Graph [19].

**KIF.** In the collaboration architecture we use KIF [6] as the representation language to record facts and properties. This is similar to ANSA, with which it is possible to record the properties of objects. However, the KIF language is much more expressive. KIF is a prefix version of first-order predicate calculus with various extensions that enhance its expressiveness. The language has a well defined syntax and semantics.<sup>3</sup>

Sentences in KIF are composed of terms (similar to words in a natural language), which include words from a vocabulary, e.g., "f1.tex", `translation`,

---

<sup>3</sup> The semantics of KIF is based on the standard Tarski semantics for first-order logic. There are special semantics for the quote operator to prevent paradoxes [6].

latex, dvi, etc. Terms can be more complicated, as in the functional expression (file "f1.tex" latex). The function file maps a file name and a file type to a file object.

First and foremost, KIF provides for the expression of simple data. For example, the sentences shown below encode facts in a printer database. The first states that printer-1 is located in building 460. The second states that printer-1 is on the fourth floor. The last two state that printer-1 handles eight by eleven inch paper and eleven by seventeen inch paper.

```
(building printer-1 460)      (paper-size printer-1 8x11)
(floor printer-1 4)          (paper-size printer-1 11x17)
```

More complicated information can be expressed with the use of functional terms. For example, the first fact states that file "f1.tex" was modified on April 30, 1995. The second states that the size of "f1.dvi" is greater than one thousand bytes. The last fact states that the translation of a file "f2.tex" in tex format is the file "f2.dvi" in dvi format.

```
(modified (file "f1.tex" latex) 4/30/95)
(> (size (file "f1.dvi" dvi)) 1000)
(translation (file "f2.tex" tex) (file "f1.dvi" dvi))
```

KIF includes a variety of logical operators that assist in the encoding of logical information (such as negations, disjunctions, rules, quantified formulas, and so forth). The expression shown below is an example of a complex sentence in KIF. Note that all symbols beginning with ? are universally quantified variables that can be instantiated to match any expression (subject to the rules of unification). The following expression asserts that a printer is located in hplabs, if it is in building 460.

```
(<= (location ?x hplabs) (building ?x 460))
```

One of the distinctive features of KIF is its ability to encode knowledge about knowledge, using the ^ and , operators and related vocabulary. For example, the following sentence asserts that object scribe can handle printing any file that is in PostScript format. The use of commas signals that the variables should not be taken literally.

```
(handles scribe ^(print (file ,?x ps)))
```

KIF can also be used to describe procedures, i.e., to write programs or scripts for agents to follow. Given the prefix syntax of KIF, such programs resemble Lisp or Scheme. The following is an example of a three-step procedure written in KIF. The first step ensures that there is a fresh line on the standard output stream; the second step prints Hello! to the standard output stream; the final step adds a carriage return to get to a new line.

```
(sequence (fresh-line t) (print 'Hello!') (fresh-line t))
```

KIF defines a set of objects, functions, and relations whose meaning is fixed, e.g., numbers and arithmetic functions. However, it is important to note that KIF is open ended, i.e., users are free to define the meanings of new symbols that are not already predefined. The predefined set of symbols ensures that all objects can start communicating in a simple base language, which they are free to extend.

**KQML.** We use KIF to write sentences that define what is true in our application domain. In a community of objects this alone is not sufficient. When an object sends a KIF sentence to another object it also needs to indicate its attitude towards it, e.g., "I am telling you that  $x$  is true," "I am telling you that  $x$  is no longer true," "is  $x$  true?," "find one instance for which  $x$  is true," "find all instances for which  $x$  is true," "print the file  $f$ ," "perform action  $a$ ," etc.

The purpose of KQML is to provide this extra linguistic layer that describes the attitude of the sending object towards the embedded KIF expression. Intuitively, each message in KQML is one piece of a dialog between the sender and the receiver, and KQML provides support for a wide variety of such dialog types.

As used in ACL, each KQML message is a list of components enclosed in matching parentheses [4]. The first word in the list indicates the type of communication (`tell`, `ask-if`, `print`, etc). The subsequent entries are KIF expressions appropriate to that communication, in effect the "arguments."<sup>4</sup>

The expression shown below is the simplest possible KQML dialog. In this case, there is just one message— a simple notification. The sender is conveying the enclosed sentence to the receiver. In general, there is no expectation on the sender's part about what use the receiver will make of this information.

```
A to B: (tell (= (size (file "f2.tex" tex)) 12678))
```

The following dialog is a little more interesting. In this case, the first message is a request for the receiver to execute the operation of printing a string to its standard i/o stream. The second message tells the sender that the request has been satisfied.

```
A to B: (perform (print 'Hello!' t))
B to A: (reply done)
```

In the dialog shown below, the sender is asking the receiver a question in an `ask-if` message. The receiver then sends the answer to the original sender in a `reply` message.

```
A to B: (ask-if (> (size (file "f1.dvi" dvi)) (size (file "f1.ps" ps))))
B to A: (reply true)
```

In addition to the simple notifications, commands, and questions illustrated here, KQML also contains support for delayed and conditional operations, subscriptions, requests for bids, offers, promises, and so forth. Describing the complete list and its semantics is beyond the scope of this paper. Additional information can be found in [4].

KQML defines a set of performatives whose meaning is fixed, e.g., `ask-if`, `tell`, etc. However, it is important to note that, similar to KIF, KQML is open ended, i.e., users are free to define the meanings of any new performatives that are not already predefined. The predefined set of performatives ensures that all objects can start communicating in a simple base language, which they are free to extend.

<sup>4</sup> The current KQML manual [4] provides a procedural semantics for the various communication types. Developing a more formal semantics for KQML is an area of ongoing research [13].

### 3.2 Facilitators

The CORBA Object Request Broker (ORB) provides the communication infrastructure, but does not support coordination. The burden of coordination is placed entirely on the client programmer. Before a request can be satisfied, the client needs to identify the type of service, locate a reference to it, and explicitly invoke the appropriate method while passing it the required arguments. The ANSA trader reduces the client burden by providing a yellow page service. It assists a client in locating an "instance" of a server that is most appropriate. The client still needs to know a priori what type of service it will need, which method to invoke, and what the arguments for the method are. The facilitator reduces the burden of coordination further. It can determine the type of service needed, identify an instance of that service, and execute the appropriate methods. It can also create a new service by glueing together a number of existing services. The facilitator in effect reduces the interdependence of the client and server, since the client is not required to have explicit knowledge of the details of a server's interface.

A facilitator is a system service that can extend a distributed object model in a similar manner that a software trader can (see Fig. 3). Similar to the ANSA trader, a facilitator is a lightweight service that is visible to those objects that wish to utilize it, and transparent to those that do not. Objects are free to register services and/or make requests to a facilitator or access objects directly using conventional method invocations (e.g., CORBA ORB). In order to provide higher-level coordination services, objects will need semantically richer specifications than what ANSA's attribute extension of IDL can support. If an object chooses to utilize the facilitator they can communicate their ACL requests and/or specifications to the facilitator via the CORBA ORB.

In this section, we illustrate with examples services provided by a facilitator: the interface specifications of objects, and the processing of object requests by a facilitator.

Scalability is an important concern for the facilitator architecture. In a small setting, a single facilitator can service a collection of objects (on the same or different machines). However, this is impractical in a large setting, where the single facilitator will be a bottleneck. In a large setting, there can be a collection of facilitators, typically one facilitator per machine. Each facilitator is connected to one or more neighboring facilitators, thus forming a network of connected facilitators. Each facilitator also has a collection of objects directly connected to it. In order to support interoperation across the network, each facilitator informs each of its neighboring facilitators that it can directly handle all the capabilities both for its directly connected objects and for the other directly connected facilitators.

We next present a series of examples illustrating the coordination capabilities of the facilitator architecture (see [8] for details). In all the examples there is a single facilitator. Key elements of the architecture are the declarative specifications of object interfaces (a meta-description) and the declarative specifications of other background information. The power of the coordination services is made

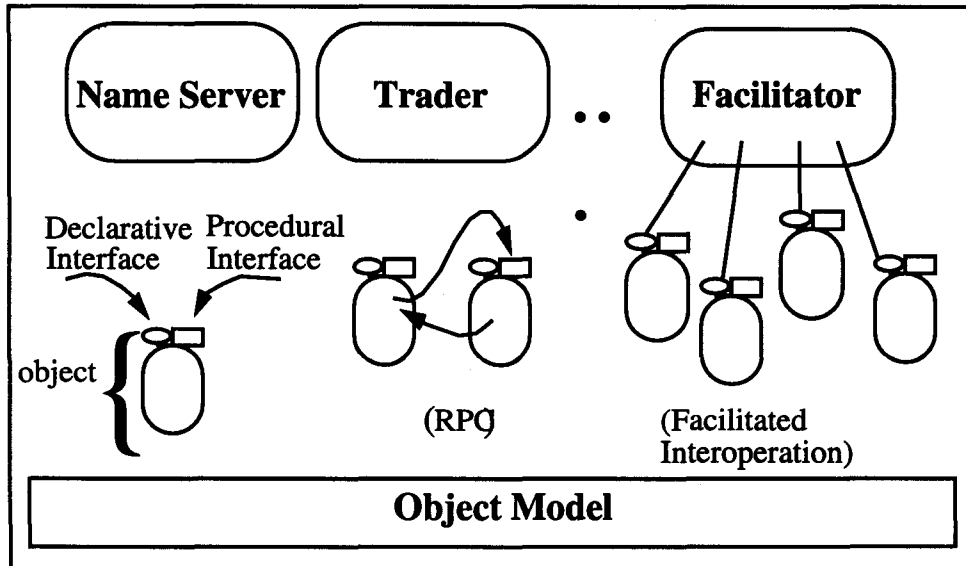


Fig. 3. Multiple system services extending an object model.

possible by the combination of these two types of information by the inference procedure of the facilitator.

It should be noted that simple computations are performed by the facilitator if simple interface specifications are given to it. The operations of the facilitator are very efficient in these situations. The clients and servers can make the trade-off on how much work they wish the facilitator to perform by controlling the facts communicated to the facilitator and by controlling the power of the inference method used by the facilitator, which is one of: simple pattern matching, backward chaining without recursion, backward chaining, or model elimination.

**Example: Manual Coordination.** This example illustrates a simple scenario where a client object manages all the coordination activities to satisfy its own request. In this example, there is a single facilitator running on a Unix workstation; there is an object named `textures` running on a Macintosh that can translate a file from Latex to DVI format; and there is an object,  $A_1$ , that accepts and processes requests on a user's behalf running on the same Unix workstation. Objects send their capabilities and requests to the facilitator. In this example, object  $A_1$  wants to translate a given file from Latex to DVI format.

Initially there are no services registered with the facilitator. The `textures` object registers with the facilitator, specifying its interface and any additional background facts to the facilitator. All communications between the objects and

the facilitator are restricted to be in ACL. The `textures` object must inform the facilitator that it can translate arbitrary Latex files into DVI format. It transmits the following message to the facilitator:

```
(tell (handles textures ~(ask-one ,?y (translation (file ,?x latex)
                                                    (file ,?y dvi))))))
```

This is the declarative meta-description of the capability of `textures`. The performative handled is the standard `ask-one` KQML performative, which takes two arguments: a pattern argument and a sentence. In order to handle this performative, `textures` must return the first argument instantiated by a binding for the variables in a proof of the second sentence argument. The pattern in this example is `?y`, and the sentence is `(translation (file ,?x latex) (file ,?y dvi))`. So the previous fact states that the `textures` object, when asked for one translation of a Latex file to a DVI file, will return the name of the DVI format file (`?y`).

The user interface object  $A_1$  next registers with the facilitator. It has no capabilities, so it does not transmit any interface specifications to the facilitator. However,  $A_1$  would like to find the translation of a file from Latex format to DVI format. In this example,  $A_1$  will do all the work to perform the translation. It does this by first asking the facilitator if there is any object that can perform the translation (using the yellow-pages services of the facilitator), and it then sends the request to an object that can handle it.

Object  $A_1$  first asks the facilitator if the following fact is true:

```
(handles ?a ~(ask-one ?result (translation (file 'f1.tex' latex)
                                           (file ?result dvi))))
```

The facilitator uses its repository of information, which includes object interface specifications, to find a proof of this fact using backward inference. The proof is direct, since this fact matches the previous specification for object `textures` (repeated below) with the variable `?a` bound to `textures`:

```
(handles textures ~(ask-one ,?y (translation (file ,?x latex)
                                                    (file ,?y dvi))))
```

The facilitator informs object  $A_1$  that `textures` can handle the request.  $A_1$  next sends the following request to `textures` (using the communications support provided by the architecture):

```
(ask-one ?result (translation (file 'f1.tex' latex) (file ?result dvi))))
```

`textures` translates the file `'f1.tex'` in Latex format to the file `'f1.dvi'` in DVI format, and returns this as the answer to  $A_1$ .

**Example: Simple Content-Based Routing.** This example is the same as the previous example, except that the object representing the user ( $A_1$ ) does not manage the coordination, but leaves this up to the facilitator. The example is the same as the previous example up to the point where  $A_1$  makes the request. In the previous example,  $A_1$  used the yellow-pages service of the facilitator to find

an object that could handle a request. In this example, the facilitator gives  $A_1$  the impression that it is directly achieving the request, and  $A_1$  is unaware that textures will actually be used in the translation.

$A_1$  makes the following request to the facilitator:

```
(ask-one ?result (translation (file 'f1.tex' latex) (file ?result dvi)))
```

The facilitator tries to find an object that can handle the quoted request:

```
(handles ?a ~(ask-one ?result (translation (file 'f1.tex' latex)
                                           (file ?result dvi))))
```

The facilitator uses backward inference to try to prove this fact. The proof is direct, since this fact matches the previous specification for textures (repeated below) with the variable  $?a$  bound to textures:

```
(handles textures ~(ask-one ,?y (translation (file ,?x latex)
                                             (file ,?y dvi))))
```

The facilitator passes the request from  $A_1$  to textures. textures translates the file 'f1.tex' in Latex format to the file 'f1.dvi' in DVI format. The textures object returns the answer 'f1.dvi', and this is the answer returned to  $A_1$ .

Note that in this example,  $A_1$  makes the following request to the facilitator:

```
(ask-one ?result (translation (file 'f1.tex' latex) (file ?result dvi)))
```

while in the previous example,  $A_1$  made the following request to the facilitator:

```
(handles ?a ~(ask-one ?result (translation (file 'f1.tex' latex)
                                           (file ?result dvi))))
```

and the second request to textures:

```
(ask-one ?result (translation (file 'f1.tex' latex) (file ?result dvi)))
```

This example illustrates a small improvement with content-based routing; however, in general an arbitrary amount of intermediate effort can be saved by content-based routing, as illustrated by the next example.

**Example: Simple Problem Decomposition.** This example illustrates a case requiring problem decomposition. The scenario is the same as the previous example, except that there is a third object `dvi2ps` running on a different Unix workstation that can translate DVI format files to PostScript files. In this example, the user object  $A_1$  requests the facilitator for the translation of a Latex file into a PostScript file. This requires first translating the file from Latex to DVI, and then translating the file from DVI to PostScript.  $A_1$  is unaware of the sequence of operations taking place. It appears to  $A_1$  that the facilitator is performing all the translations directly.

The example is the same as before up to the point when textures registers with the facilitator and sends its capabilities. Following this, `dvi2ps` registers with the facilitator and sends the following message:

```
(tell (handles dvi2ps ^ (ask-one ,?y (translation (file ,?x dvi)
                                                    (file ,?y ps))))))
```

This is similar to the message sent earlier by `textures` to the facilitator. For this example, assume that the facilitator also knows the following background fact (some other object has informed the facilitator of this fact):

```
(tell (<= (translation (file ?x latex) (file ?y ps))
        (translation (file ?x latex) (file ?z dvi))
        (translation (file ?z dvi)   (file ?y ps))))
```

The general form for a conditional rule is `(<= A B C)`, which means that `A` is true if `B` and `C` are true. The fact argument to the `tell` states that the translation of file `?x` in Latex format is file `?y` in PostScript format, if the translation from file `?x` in Latex format is file `?z` in DVI format, and the translation of file `?z` in DVI format is file `?y` in PostScript format. This rule in effect defines a simple transitivity relation.

The user interface object `A1` next registers with the facilitator. It has no capabilities, so it does not transmit any interface specifications to the facilitator. However, `A1` would like to find the translation of a file from Latex format to PostScript format. It sends the following message to the facilitator:

```
(ask-one ?result (translation (file 'f1.tex' latex) (file ?result ps)))
```

The facilitator tries to find an object that can handle the quoted request:

```
(handles ?a ^ (ask-one ?result (translation (file 'f1.tex' latex)
                                           (file ?result ps))))
```

The facilitator uses backward inference to try to prove this fact. Unfortunately, no object can handle this request directly. However, the facilitator can use the transitive rule to decompose this request into the conjunction:

```
(and (translation (file 'f1.tex' latex) (file ?z dvi))
      (translation (file ?z dvi)        (file ?result ps)))
```

The first conjunct is handled by the `textures` object, and it binds the variable `?z` to the file `'f1.dvi'`. This binding for `?z` is plugged into the second conjunct, which is handled by the `dvi2ps` object, and it returns the variable `?result` bound to `'f1.ps'`.

The problem decomposition rule for this example is applicable only in the special case of translating Latex files to PostScript by first translating them to DVI format. The following, more powerful rule allows an arbitrary intermediate format (not just DVI):

```
(<= (translation (file ?x latex) (file ?y ps))
    (translation (file ?x latex) (file ?z ?f))
    (translation (file ?z ?f)   (file ?y ps)))
```

Note that this is a recursive rule, which can be applied an arbitrary number of times. Content-based routing, in the presence of such rules, can tremendously reduce the complexity of problem solving compared with manual coordination.

**Example: Defining New Primitive Performatives.** In this example, we illustrate the definition of new performatives. The previous examples illustrated the use of the `tell` and `ask-one` primitive performatives of ACL. This example illustrates the definition of a conditional performative, though it possible to define simpler performatives without any conditions.

Performatives in ACL describe actions, as opposed to informational facts. It is important not to confuse informational facts with facts about actions. The purpose of informational facts is to define what is true, while the purpose of facts about actions is to define a procedure. For example, the fact `(and P P)` can be proved by proving `P` alone, while the action `(sequence P P)` cannot be satisfied by performing action `P` alone.

In this example, there is a single object called `Scribe` that can print files in PostScript format, as long as the size of the file is less than one megabyte. `Scribe` connects to the facilitator and sends it the following message:

```
(tell (<= (handles scribe ^(print (file ,?f ps)))
      (= (denotation ?f) ?fd) ; ?fd is the unquoted file name
      (< (file-size (file ?fd ps)) 1000000)))
```

The fact inside the `tell` in the above message states that `Scribe` can print any PostScript file as long as the size of the file is less than one million (bytes). The function `file-size` takes a file name as an argument and maps it to the size of the file in bytes. `Scribe` is defining a new performative `print`. It will handle all print requests of PostScript files (of a certain length), but how it does this is not specified or relevant.

The user object `A1` makes a request to the system to print a PostScript file by sending the facilitator the following message:

```
(print (file 'f1.ps' ps))
```

The facilitator first searches for an object that can handle the quoted form of the above expression. It uses backward inference to find a proof for the fact:

```
(handles ?a ^(print (file 'f1.ps' ps)))
```

The backward inference uses the interface description to conclude that the `Scribe` object can print the file `'f1.ps'` if the size of `'f1.ps'` is less than one megabyte. Assume that this is so, i.e., the following fact is true:<sup>5</sup>

```
(< (file-size "f1.ps" ps) 1000000)
```

The facilitator, therefore, concludes that `Scribe` can print the file `'f1.ps'`. Next the facilitator sends a message to `Scribe` to print the file. `Scribe` prints the file, and when it is done, the successful indication of this is reported to the user object `A1`.

---

<sup>5</sup> The size of a file is computed using procedural attachments, similar to the attachments for built arithmetic functions like addition and division.

**Example: Defining Procedures from Existing Performatives.** This example illustrates defining procedures in ACL, and it focuses on facts that specify a sequence of actions. The primitive actions are the KQML performatives, that include `ask`, `tell`, and other user-defined primitive performatives. From these primitives, it is possible to define more complex performatives as compositions of these (just as it is possible to define complex procedures from primitive statements in a programming language).

Every performative returns a value, e.g., the performative `(ask-one <exp> <fact>)` returns `<exp>` instantiated by a binding for the variables which make `<fact>` true. Performative definitions can be nested, where the evaluation is inside-out and a component performative is replaced by its value. This is similar to the inside-out evaluation of `(+ (* 2 3) 4)` in programming languages.

This example is identical to the previous, except that the user object  $A_1$  wishes to print a Latex file. In this case  $A_1$  sends the following message to the facilitator:

```
(print (file 'f1.tex' latex))
```

Unfortunately the facilitator is unable to find an object that can handle this request, since the file is in Latex format, and the printer Scribe can only print files in PostScript format. One way of overcoming this is to define a procedure for printing that defines a sequence of actions: to print an arbitrary file, first translate it to PostScript format, and then print it. This is specified by the following fact describing a procedure:

```
(define ~(print (file ,?n ,?f) )
      ~(print (file (ask-one ?n2 (translation (file ,?n ,?f)
                                             (file ?n2 ps))) ps)))
```

The general format for defining procedures is:

```
(define ~<performative-defined> ~<performative>)
```

A sequence of actions is defined by `(sequence <a1> ... <an>)`, and conditional actions are defined by:

```
(cond <condition-1> <action-1> ... <condition-n> <action-n>)
```

In this example the definition of the new performative is nested. In this case, the name of the translated file is given by the inner performative:

```
(ask-one ?n2 (translation (file ,?n ,?f) (file ?n2 ps)))
```

The request from  $A_1$  can now be handled. The facilitator cannot find an object that can handle the request directly; however, the facilitator uses the procedural definition above to find that the file can be printed if: 1) the file can be translated from Latex to PostScript (the inner performative in the definition), and 2) some object can print the resulting PostScript file. The process of handling these two steps was illustrated in previous examples. The transitive translation rule is used to translate the file from Latex to DVI, and then from DVI to PostScript. The PostScript file `'f1.ps'` can be printed directly by the printer Scribe since it is less than one megabyte long. The two steps are performed directly by the facilitator, and the indication of success is passed back to  $A_1$ .

**Example: Vocabulary Translation.** Hiding details about implementations reduces the coupling between a client and server. Unfortunately, with procedural interfaces, interaction between a client and server requires the client to have specific knowledge about the details of a server's interface. When a client makes a declarative request to the facilitator, the client is not required to know specific details about the interface of the server(s) that satisfies it. This further reduces the coupling between the client and server, which is particularly important when supporting interoperation among loosely coupled objects interacting in a dynamic environment.

This example illustrates the use of vocabulary translation to match a request with a server that can handle it.<sup>6</sup> A request is made with one vocabulary, while a service is provided with a different vocabulary. A translation rule is used automatically by the facilitator to map the request to the service. Such translation is not possible with a simple syntactic matching scheme.

Suppose that previously there was a different vocabulary for file translations, which is given below:

```
(file-translation <file1> <file2> <format1> <format2>)
```

For example, previously we would have used:

```
(file-translation "f1.tex" "f1.ps" latex ps)
```

while the current use is:

```
(translation (file "f1.tex" latex) (file "f1.ps" ps))
```

The new vocabulary makes explicit the concept of a file by using the function `file`. When the vocabulary was updated, the interface for `textures` was also updated (as illustrated in the previous examples). Unfortunately, there is an object `A0` (among others) that continues to use the old vocabulary. `textures` would like to continue supporting the old vocabulary. It does this by communicating the following message to the facilitator:

```
(tell (<= (file-translation ?file1 ?file2 ?format1 ?format2)
         (translation (file ?file1 ?format1) (file ?file2 ?format2))))
```

The fact embedded in the `tell` defines the mapping between the old and new vocabularies. Consequently, when `A0` makes the request to the facilitator:

```
(ask-one ?result (file-translation "f1.tex" ?result latex ps))
```

The facilitator uses the previous translation rule in the process of its backward reasoning to translate the request. The request is translated to:

```
(translation (file "f1.tex" latex) (file ?result ps))
```

---

<sup>6</sup> Vocabulary/interface translation should not be confused with file translation, which is what the vocabulary translation is being used to accomplish.

which is handled by textures. The processing from this point on is identical to that in the previous example.

This example illustrates that it is possible to modify the interface of an object, and still have the object support its old interface by using the translation capabilities of the facilitator.

These examples have illustrated the complex coordination behavior that is provided by the facilitator automatically in response to a request from a client. The client is not aware of the complicated sequence of events used to process the request. More generally, they also illustrate how semantically rich specifications significantly increase the ability of construct more sophisticated coordination services. That trying to add semantics to IDL by including name/value pairs is a good initial step but it is not sufficient.

## 4 Status

There is an API for the coordination architecture in C, C++, and Lisp. It hides all the details associated with network communication from the programmer. Each API uses TCP to communicate with the facilitator, but changing the communication protocol (e.g., to DCE) is a trivial task. The API has been used to support interoperation among applications running on UNIX, Macintosh System 7, and Windows NT. A prototype of the facilitator is implemented in Lisp and runs on Sun, Silicon Graphics, Hewlett-Packard Series 300/700 workstations and the Apple Macintosh.

The Logic Group in the Stanford University Computer Science Department, in conjunction with HP Labs in Palo Alto California, has been developing the facilitation architecture over the last five years. This system has been used in a collection of interoperation experiments, including an integrated design, manufacture, and diagnosis system for digital circuits[5], a multi-domain simulation of a robotic arm [2], and integrated CAD tools for civil engineering [12]. It is currently being used in the CommerceNet project [11], which provides smart search for product information using heterogeneous on-line catalogs, ordering, billing, etc.

## 5 Conclusion

In this paper we have presented an interoperation architecture based on declarative specification of object interfaces. Given machine-processable object specifications, it is possible to develop new system services that utilize simple interfering techniques to facilitate interoperation among loosely coupled distributed objects. These services provide additional capabilities (content-based routing, problem decomposition, translation, etc.), and the automation of these services enables using the best information available at runtime.

The reduction in the coupling between a client and server is a consequence of describing requests declaratively. That is, a client describes a request in terms of what it would like, and not how to achieve it. This has the following advantages:

- The client is not required to know about which servers are needed to satisfy the request. This logic has been factored out of the client's implementation and placed in the facilitator, thereby reducing the complexity of the client's implementation.
- The client is no longer responsible for obtaining references to particular instances of the different servers.
- The client is not required to know the exact details about a server's interface.

Facilitated interoperation provides considerable flexibility, but at the expense of runtime efficiency. In small distributed environments, this approach may not be acceptable or needed. As the distributed systems become larger, more dynamic, and more decentralized, this approach may become not only reasonable but necessary.

## 6 Acknowledgements

We would like to thank Mike Genesereth and Reed Letsinger for participating in discussions on some of the ideas presented in this paper. We would also like to thank Hewlett Packard Laboratories funding this work.

## References

1. Betz, M., "OMG's CORBA," *Dr. Dobb's Special Report*, Winter 1994/95.
2. Cutkosky, M. et al., "PACT: An Experiment in Integrating Engineering Systems," *Computer* 26, 1(1993), 28-37.
3. Deschrevel, J. P., "The ANSA Model of Trading and Federation," Architecture Projects Management, Cambridge, 1993.
4. Finin, T., and Wiederhold, G., et al., "Specification of the KQML Agent-Communication Language," available from the WWW with the URL <http://www.cs.umbc.edu/kqml/kqmlspec/spec.ps>, June, 1993.
5. Genesereth, M., "Designworld," in the *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*, Sacramento, California, April 1991, pp. 2785-2788.
6. Genesereth, M. R., Fikes, R., et al., "Knowledge Interchange Format Version 3 Reference Manual" Logic92-1, Stanford University Logic Group, 1992.
7. Genesereth, M. R., and Ketchpel, S. P., "Software Agents," *Communication of the ACM*, Vol. 37, No. 7 July 1994.
8. Genesereth, M., Singh, N., and Syed, M., "A Distributed Anonymous Knowledge Sharing Approach to Software Interoperation," in the *Proceedings of the International Symposium on Fifth Generation Computing Systems*, 1994, pp. 125-139.
9. Gruber, T., "Ontolingua: A Mechanism to Support Portable Ontologies," KSL-91-66, Stanford Knowledge Systems Laboratory, 1991.
10. Herbert, A., "An ANSA Overview," *IEEE Network*, January/February 1994, pp. 18-23.
11. Keller, A., "Smart Catalogs and Virtual Catalogs," in *USENIX Workshop on Electronic Commerce*, August 1995.

12. Khedro, T. and Genesereth, M., "The Federation Architecture for Interoperable Agent-Based Concurrent Engineering Systems," *International Journal on Concurrent Engineering, Research and Applications*, pp. 125-131, 1994.
13. Labrou, Y. and Finin, T. "A Semantic Approach for KQML – a general purpose communication language for software agents," in *Third International Conference on Information and Knowledge Management*, 1994.
14. Microsoft Corporation, *OLE 2 Programmer's Reference: Creating Programmable Applications with OLE Automation*, Volume 2, Microsoft Press, Redmond, Washington, 1994.
15. Microsoft Technical Backgrounder OLE 2.0, 1994.
16. The Object Management Group, "The Common Object Request Broker: Architecture and specification," Revision 1.1, TC Document Number 91.12.1, December 1991.
17. Mowbray, T. *Essential Corba: Systems Integration Using Distribute Objects*, Wiley, John & Sons, Inc., 1995.
18. Raj, R., Tempero, E., Levy, H., Black, A., Hutchinson, N., and Jul, E. "Emerald: A General Purpose Programming Language," *Software-Practice and Experience*, Vol 21, No. 1, January 1991.
19. Tawakol, O., and Singh, N. "A Name Space Context Graph for Multi-Context Systems," *Proceedings of the 1995 AAAI Fall Symposium Series*, Cambridge, Massachusetts, November, 1995.