

Derivation of Glue Code for Agent Interoperation*

Mark Burstein Drew McDermott
GTE/BBN Yale University
burstein@bbn.com drew.mcdermott@yale.edu

Douglas R. Smith
Stephen J. Westfold Kestrel Institute
{smith,westfold}@kestrel.edu

Abstract

Getting agents to communicate requires translating the data structures of the sender (the *source* representation) to the format required by the receiver (the *target* representation). Assuming that there is a formal theory of the semantics of the two formats, which explains both their meanings in terms of a neutral *topic* domain, we can cast the translation problem as solving higher-order functional equations. Some simple rules and strategies apparently suffice to solve these equations automatically. The strategies may be summarized as: decompose complex expressions, replacing topic-domain expressions with source-domain expressions when necessary. A crucial issue is getting the required formal theories of the source and target domains. We believe it is sufficient to find partial formalizations that grow as necessary.

1 Introduction

This paper is about getting agents to communicate with each other. By “agent” we mean programs that operate at a high enough semantic level that they can form new connections to other programs in order to get a job done. To make such a connection, an agent must find other agents that might carry out a task on its behalf, establish communication with them, send them tasks and receive results back. If this sort of connection is to happen without careful preprogramming by human coordinators, it must be

*This work was supported by the Air Force Research Lab (Rome, NY), and the DARPA CoABS program, under contract number F30602-98-C-0168.

the case that the agent has an explicit representation of its goals, and that the agents it talks to maintains an explicit representation of their abilities. Otherwise, it is difficult to see how the agents could conclude that agent *B* can help agent *A*.

There are many different ways of getting agents together dynamically. We spoke as though each agent solicited help from others on its own, but of course there could be other *middle agents* involved [13]. The information about the needs and abilities of the agents could be stored in any convenient place, not necessarily maintained by the agents themselves. The connection between agents *A* and *B* could be direct, or could be mediated in multiple ways by middle agents. It is also not important for our purposes who “takes the initiative,” *A*, *B*, or some third party, or whether there are other agents *C*, *D*, and so forth that also contribute.

Notational issues arise at multiple points in the process of hooking agents up. Assuming there are many agents, or an open-ended collection in an amorphous medium such as the Internet, it is impractical for an agent to find useful partners by enumerating all candidates and discarding those that don’t work. Instead, there has to be a *registration* service that allows agents with abilities to offer to make their presence known. When the agent signs in, it must specify what capabilities it offers. This specification is in an *advertising language* [3, 13]. The best-known example is the LARKS notation of [24] (although, as we shall see, LARKS does more than that). It is reasonable to suppose that an agent has to cope with just one advertising language. One plausible idea is that a registration service would be maintained by a consortium of companies in some industry. There would be little point in standardizing the registry without also standardizing the advertising language. For instance, suppose the registry serves the travel industry. There would be a vocabulary for describing airline flights, fare prices, number of passengers involved, dates involved, and so forth. An agent might advertise for someone who could provide flights for two people from London to Orlando sometime in May, and an airline that can provide flights for more than one person from London to anywhere in the U.S. any time might be worth talking to. See [24] or [15] for alternative models of how the linkage might work between advertisers and their customers.

Once agents are connected (whether or not there are helpers in between), they begin sending messages back and forth. One issue that arises here is what sorts of “conversation” they have. *Agent communication languages* (ACLs) [14] such as KQML [16, 6] and FIPA ACL [8, 7] offer the ability for agents to send messages such as (TELL *P*) or (REQUEST *A*). TELL and REQUEST are examples of *performatives*, or *communicative acts*, and tend

to be the focus of papers on ACLs. Unfortunately, it is hard to picture exactly what P 's and A 's two agents would send each other in the abstract, so descriptions of agent behavior at the communicative-act level are often unsatisfying. Let's look at a particular example.

Suppose agent R wants to know the address and phone number of each of a group of people. Suppose agent P knows phone numbers of individual people. It seems clear that P has something to offer R , that will allow it to come closer to its goal. But what it does it mean to say that agent P "knows" someone's phone number? How is a group of people represented? Abstractly we might think of a data structure as representing a person, but when two agents communicate they have to know what the pieces of the structure actually are, and how they relate to the person represented. It may turn out that, although they both reason about people's phone numbers, they represent people in different ways. For example, R may represent a person as a pair (*name, date-of-birth*), whereas P may represent the same person as *Social-Security-number*.

Note that this problem arises even if R and P share the same *ontology*, that is, vocabulary and logical framework. They may agree that there is a class called **Person** and that each person has a name, a Social Security number, and a date of birth. They can communicate, however, only if they both identify people the same way, or a translation mechanism can be found.

Another example appears in [24]. Here R wants to know the set of all air combat missions launched between two dates. P can, given two dates, produce the set of all AWAC missions launched between those two dates. Using the LARKS notation (and simplifying a bit), R 's request is

Types	Date = (mm: Int, dd: Int, yy: Int) DeployedMission = ListOf(mType: String, mID: String Int)
Input	sd: Date, ed: Date
Output	missions: Mission
InConstraints	sd ≤ ed
OutConstraints	deployed(mID), launchedAfter(mID, sd) launchedBefore(mID, ed)

and P 's advertisement is

Types	Date = (mm: Int, dd: Int, yy: Int) DeployedMission = ListOf(mType: String, mID: String Int, mStart: Date, mEnd: Date)
Input	start: Date, end: Date
Output	missions: DeployedMission
InConstraints	start \leq end
OutConstraints	deployed(mID), mt = AWAC, launchedAfter(mID, mStart) launchedBefore(mID, mEnd)

One can see how the sorts of filters employed by [13] could zero in on P as a good candidate for providing information to R . However, there are several issues that arise in completing the connection between the two. For instance: How do you know that the Ints mentioned in the Dates for R 's request line up properly with the Ints mentioned in P 's Dates? (the labels "mm," "dd," and "yy" are presumably just mnemonic, and have no content.) How do we transform the list produced by P into the list wanted by R ? Intuitively, if you have a list whose elements are tuples that are "too big" (have unwanted elements), you can just throw away the unwanted stuff. It appears from published descriptions [10, 25] that Sycara et al.'s RETSINA system has standard transformations for common cases such as this one.¹

We are interested in automating the generation of transformations such as that used by RETSINA to send dates from R to P and lists of missions back, given slightly more detailed descriptions of the data structures. In the rest of this paper, we will use the term *glue code* to refer to such a transformation. The problem is to generate glue code automatically.

The work most closely related to ours is the work by database and information-retrieval community on integrating disparate information sources. The term *schema* is in standard use to refer to description of the different record types in a database and the constraints among them. The schema

¹Our use of their example is not meant to suggest inadequacies in [13]'s treatment of agent communication. The published description makes it hard to be sure what's going on in their RETSINA system. It is difficult to tell what the scoping rules are, or how one could decipher the relationship between mStart and start (or mEnd and end) in P 's advertisement. It's possible that we are misconstruing the notation.

plays roughly the role of the InConstraints and OutConstraints in LARKS. Every database has its own schema, so putting two databases together requires *schema integration*, with subproblems of *query translation* and *value translation*. [20, 11, 9, 18].

Database researchers assume that the results of a query are tables in a standard format (or, equivalently, collections of answers to Datalog queries [26, 27]). When information sources are being integrated, the goal is to present the end user with an abstract database that behaves like the union of all the component databases. To make this happen, there must be axioms that allow the database *mediator* to map predicates in the source databases to predicates in the abstract database. If the component information sources don't look like databases, they can be "wrapped" with programs that make them look like databases [21, 28, 2, 3, 19, 1].

We don't think this model fits the case of agent communication particularly well. Couplings between agents may be less permanent than integration of information sources. As the examples in this paper will show, the I/O behavior of an agent may be most naturally described in terms of data structures, whose meaning is specified as a mathematical function of abstract domain entities.

A related issue has to do with the how information sources are compared. A date may be specified as a tuple $(year, monthday)$, but what makes it a date? In most database-integration papers, the answer ultimately is that someone told the mediator: "To infer `date_of(x, y, m, d)` in this database, solve the query `time_of(x, s, n, d, m, y)` in that one." If the same component is used elsewhere, another axiom will have to be found, and there is no way to get the human out of the loop. For agent communication, this is unsatisfactory.

The only solution to this problem is to describe each information source, or, more generally, the inputs and outputs of each *agent*, in terms of shared, global *ontologies*. This must be done just once for each agent (we hope), after which it can be coupled to other agents by figuring out the relationships between their outputs and its inputs. In this vision, what makes a data structure a *date* is that it is stipulated to be a tuple of the form

$$((time_instant_year\ t\ c), (time_instant_month\ t\ c), (time_instant_day\ t\ c))$$

for some time instant t and calendar c .

We don't want to overemphasize the difference between our approach and that of the database community. It is likely that we can make use of results from database researchers, and that the fields of agents and databases will eventually converge on answers to the questions we raise here.

2 The Setting

To explain our approach, we will start with an almost trivial problem. Suppose that we have two agents that use data structures that represent the same population of people, but that record different information about them. In particular, let S record the name, identification number, and age of each person, and let T record the name and the id number of the person if he/she is over age 30 and 0 otherwise. S is the output of agent 1, and T is the input to agent 2. If agent 1 is to feed information to agent 2, then we must derive a translator f such that $f(S) = T$. We call S the *source* domain, and T the *target* domain.

This problem is solvable only if S and T capture “the same” information, or, more precisely, if all the information to be included in T occurs in S . But by itself a database captures no information at all. It’s just a meaningless pile of symbols unless its users understand what the symbols are supposed to mean, and use the database in accordance with that understanding. Such an informal notion of meaning is sufficient if the users are all humans, but not in the case we’re concerned with, where the glue-code generator is supposed to be automatic.

Hence we must supply every agent with a formal theory glossing its input and output representations. This may seem like a big commitment, but we see no way to evade it. We’ll avoid being distracted by this issue here, but return to it in Section 7.

A key requirement of the formal theories about agents is that if two agents are to be connected, then their theories must *overlap*. If one of the agents contains the symbol **person**, and the other contains the symbol **employee**, then there must be a semantic domain whose topic is people, and the symbols must be defined in roughly the same terms from that domain. It must be possible to infer, for instance, that if a data structure in one domain represents information about **employees**, it can be translated into a data structure in the other that represents information about **persons**. This requirement raises many tough issues, which again we defer for later discussion. We will simply assume that there is an abstract common theory P of the application domain with the property that we can define S and T in terms of P . The common theory provides symbols for the concepts, operations, and relationships in the domain. Its axioms constrain the meaning of the vocabulary. In our example, we specify an abstract theory, called a *spec*, of persons P of sort² $set(person)$ where each *person* is an object that

²In this paper, we use the word “type” informally, reserving the word *sort* for technical

has a *name*, *id*, *age*, and other attributes, as follows:

```

spec Person-DB is
  sort Person
  op P : set(Person)
  op name : Person → string
  op id : Person → nat
  op age : Person → nat
  ...
end-spec

```

We will call the domain common to both S and T the *topic domain* and the specification above the *topic theory*. Sorts like *nat* and *string* correspond to manipulable data objects, which can be used as concrete representations of objects of sorts like *Person*. As we will see, agents cannot manipulate pure topic-domain entities directly, but can only perform computations over their representations.

We use a λ -calculus notation for functions. As usual in the functional-programming community, we will assume that a λ expression has just one argument, but this argument will often be a tuple. The i 'th element of a k -tuple $t = (x_1, \dots, x_k)$ will be written $\Pi_i^k(t) = x_i$. Application of function f to argument x will be written $f x$. So if $f : A_1 \times A_2 \times \dots \times A_k \rightarrow B$ (i.e., f 's sort is "function from k -tuples to B , where the i 'th element of its argument is from domain A_i "), then rather than write $f = (\lambda t g(\Pi_1^k(t), \dots, \Pi_k^k(t)))$, we will write $f = (\lambda(x_1, \dots, x_k) g(x_1, \dots, x_k))$. In some circumstances, when we write $(\lambda x \dots)$, we intend to cover the case where x is a tuple, but we don't currently care how long it is.

Now let's suppose that S and T both represent aspects of the set P . S is a set of triples $(name, id, age)$, one per person, so we can express it as

$$S = image((\lambda p (name p, id p, age p)), P)$$

where $image(f, X) = \{f(x) | x \in X\}$. We can express S more concisely using a couple of abbreviatory devices. *Function tupling* allows us to write

$$(name, id, age) \text{ for } (\lambda p (name p, id p, age p))$$

and we write $F \star X$ for $image(F, X)$. So we can write S as

$$S = (name, id, age) \star P$$

contexts.

Similarly,

$$T = (\textit{name}, (\lambda p \textit{ if } \textit{age } p > 30 \textit{ then } \textit{id } p \textit{ else } 0)) \star P$$

The problem is to find a function f such that $f S = T$. This type of equation will occur repeatedly in what follows, so we write it as

$$\begin{aligned} \textit{Find} : f(\textit{name}, \textit{id}, \textit{age}) \star P = \\ (\textit{name}, (\lambda p \textit{ if } \textit{age } p > 30 \textit{ then } \textit{id } p \textit{ else } 0)) \star P \end{aligned}$$

It seems clear that we can construct f by creating a function that does something to each element of the input dataset S . In other words, if we could solve:

$$\begin{aligned} \textit{Find} : i((\textit{name}, \textit{id}, \textit{age})(x)) = \\ (\textit{name}, (\lambda p \textit{ if } \textit{age } p > 30 \textit{ then } \textit{id } p \textit{ else } 0)) x \end{aligned}$$

where x is an arbitrary element of P , then the solution to the original equation would be

$$f = (\lambda X i \star X)$$

We express this as a subgoal-formation rule:

Image-Decomposition Rule

Given $g : A \rightarrow B$,

and $h : A \rightarrow C$,

To find: $f(g \star Y) = h \star Y$

Solution : $f = (\lambda X i \star X)$

$\Leftarrow \textit{Find} : i(g x) = h x$ for arbitrary $x \in X$.

We use the symbol “ \Leftarrow ” to indicate subgoaling.

Because we’re trying to find an i that works for all elements x of P , we can treat x as an arbitrary constant from now on. The subgoal may now be rewritten as

$$\begin{aligned} \Leftarrow \textit{Find} : i(\textit{name } x, \textit{id } x, \textit{age } x) = \\ (\textit{name } x, \textit{if } \textit{age } x > 30 \textit{ then } \textit{id } x \textit{ else } 0) \end{aligned}$$

To make further progress, we need the concepts of *imitation* and *projection* from higher-order matching theory [12], embodied in the following rules:

Imitation Rule

Given $g : B_1 \times \dots \times B_n \rightarrow C$

To *Find*: $f u = g(v_1, \dots, v_n)$

Solution: $f = (\lambda x g(h_1 x, h_2 x, \dots, h_n x))$

\Leftarrow *Find*: $h_i u = v_i$ for $1 \leq i \leq n$

Projection Rule

To *Find*: $f(a_1, \dots, a_k) = a_i$

Solution: $f = \Pi_i^k = (\lambda(x_1, \dots, x_k) x_i)$

That is, we let the structure of f mirror the desired structure on the right-hand (“target”) side of the equation.

In our example, imitation with $g = (\text{name}, (\lambda p \text{ if age } p > 30 \text{ then id } p \text{ else } 0))$ allows us to solve for i thus:

$$\begin{aligned} i &= (\lambda(n, d, a) (h_1(n, d, a), h_2(n, d, a))) \\ \Leftarrow & \quad \text{Find} : h_1(\text{name } p, \text{id } p, \text{age } p) = \text{name } p \\ & \quad \text{and Find} : h_2(\text{name } p, \text{id } p, \text{age } p) = \\ & \quad \quad \text{if age } p > 30 \text{ then id } p \text{ else } 0 \end{aligned}$$

The projection rule then gives us $h_1 = (\lambda(n, i, a) n)$. Imitation applied to h_2 gives

$$\begin{aligned} h_2 &= (\lambda(n, i, a) \text{ if } h_{21}(n, i, a) \\ & \quad \text{then } h_{22}(n, i, a) \\ & \quad \text{else } h_{23}(n, i, a)) \\ \Leftarrow & \quad \text{Find} : h_{21}(\text{name } p, \text{id } p, \text{age } p) = \\ & \quad (\text{name } p > 30) \\ & \quad \text{and Find} : h_{22}(\text{name } p, \text{id } p, \text{age } p) = \text{id } p \\ & \quad \text{and Find} : h_{23}(\text{name } p, \text{id } p, \text{age } p) = 0 \end{aligned}$$

A few more rounds of imitation and projection eventually yield:

$$f = (\lambda X (\lambda(n, i, a) (n, \text{if } a > 30 \text{ then } i \text{ else } 0)) \star X)$$

and $T = f(S)$. In English, to translate source dataset S into target dataset T , translate all the triples in S , leaving the name alone, and combining the age and id as shown.

3 More Complex Examples

We now introduce an example that we will refer to frequently in the rest of this paper. Suppose we are given a scheduling agent S that produces an airlift schedule. For each aircraft, the schedule gives the sequence of flights that it is scheduled to make. We also have an agent R that maintains what is known as a “commitment matrix,” a table that specifies, for each time slot, the number of each type of aircraft that are committed to scheduled flights (i.e. not free for allocation) in that time slot. The problem is to derive a translator f from schedules to commitment matrices, so that R is able to accept the information derived from the scheduler according to its declared input specification format.

This approach to the problem can be considered as a form of the *automatic programming* problem, but one that we believe is simpler than the general case. It “feels like” an exercise in moving data around, with a bit of condensing, summarization, and totaling thrown in. On the other hand, problems like this one are not trivial. The reader may wish to stop and try to produce the glue code for S and R by hand.

To proceed, we will need some more notational machinery. Let $x \mapsto y$ be the ordered pair (x, y) , considered as a component of a finite map; it records the association between element x and the element y it maps to. Let $P \triangleright U = \text{filter}(P, U) = \{u \in U | P(u)\}$, i.e., the subset of U satisfying predicate P . Let $(\text{size } U)$ be the cardinality of finite set U .

The topic-domain spec describes schedules:

```
spec Schedule-DB is
  import set, cargo, integer, port
  sort Reservation, time
  sort-axiom time=integer
  op R : set(Reservation)
  op aircraft : Reservation → string
  op flt-no : Reservation → string
  op depart : Reservation → time
  op arrive : Reservation → time
  op manifest : Reservation → set(cargo)
  op origin : Reservation → port
  op destination : Reservation → port
  op aircraft-class : Reservation → string
  ...
  axiom  $\forall(r1, r2 \in R)$ 
```

$$\begin{aligned}
& r1 \neq r2 \wedge (\text{aircraft } r1) = (\text{aircraft } r2) \\
& \supset ((\text{arrive } r1) \leq (\text{depart } r2) \\
& \quad \vee (\text{arrive } r2) \leq (\text{depart } r1)) \\
\text{end-spec}
\end{aligned}$$

where integers are used to represent time. The last axiom states that no two flights by the same aircraft overlap.

The source agent is a mission planner MP that produces a description of all reservations in the form:

$$S = (\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest}) \star R$$

The target agent requires a summarization of the schedule that indicates how many aircraft are flying at a given time over a given time interval. This summary, known as a *commitment matrix*³, is

$$\begin{aligned}
T(i, j) = & \\
& (\lambda t \\
& \quad t \mapsto \text{size}((\lambda r \ (\text{depart } r) \leq t < (\text{arrive } r)) \\
& \quad \triangleright R)) \\
& \star \{i..j\}.
\end{aligned}$$

(The notation $\{i..j\}$ denotes the set of all integers between i and j inclusive.)

The goal is to derive a translator from an arbitrary schedule $Sched$ to a commitment matrix⁴:

$$Find : f(S, i, j) = T(i, j).$$

A translator from S (domain DB) to T (domain MP) might be derived as follows⁵:

$$\begin{aligned}
Find: & f((\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest}) \star R, i, j) \\
& = (\lambda t
\end{aligned}$$

³A commitment matrix is supposed to give the number of each *type* of aircraft that is in use. For clarity, we have simplified this to a “commitment vector” that gives the total number of aircraft flying for each time period.

⁴Note that the parameters i and j control how much of S must be used to generate T . For technical purity, we could write this equation as $f(i, j)(S)$, but here and elsewhere we will move freely between expressing a multiargument function with currying and expressing it with tupling, generally favoring the latter.

⁵Technically, the first subgoal produced should be of the form $Find: g((\dots) \star R, i, j) = (\lambda t \ t \mapsto \dots)$, but we move the t into the arguments of g by the “decurrying” rule alluded to above.

$$t \mapsto \text{size}((\lambda r \text{ (depart } r \leq t < \text{(arrive } r))} \\ \triangleright R))$$

$\star\{i..j\}$

Solution: $f = (\lambda (X, i, j) (\lambda t (g(t, X, i, j))) \star h(X, i, j))$
by imitation.

\Leftarrow

$$\begin{aligned} \textit{Find: } & g(t, \\ & \text{(aircraft, depart, arrive, manifest)} \\ & \star R, \\ & i, j) \\ & = t \mapsto \text{size}((\lambda r \text{ (depart } r \leq t < \text{(arrive } r))} \\ & \triangleright R)) \end{aligned}$$

and *Find:* $h(t,$
 $\text{(aircraft, depart, arrive, manifest)}$
 $\star R,$
 $i, j)$
 $= \{i..j\}$

By straightforward imitation we derive

$$h = (\lambda (t, X, i, j) \{i..j\})$$

The other subgoal is attacked using imitation again:

Solution: $g = (\lambda (t, X, i, j) g_1(t, X, i, j) \mapsto g_2(t, X, i, j))$
 \Leftarrow

$$\begin{aligned} \textit{Find:} \\ & g_1(t, (\text{aircraft, depart, arrive, manifest}) \star R, i, j) \\ & = t \end{aligned}$$

and

$$\begin{aligned} [G^*] \textit{Find:} \\ & g_2(t, \\ & \text{(aircraft, depart, arrive, manifest)} \\ & \star R, \\ & i, j) \\ & = \text{size}((\lambda r \text{ (depart } r \leq t < \text{(arrive } r))} \\ & \triangleright R) \end{aligned}$$

The g_1 part succumbs to projection, with solution

$$g_1 = (\lambda (t, X, i, j) t)$$

But if we attempt to apply imitation to g_2 , we will get something useless:

$$\begin{aligned}
 \text{Solution attempt: } g_2 &= (\lambda(t, X, i, j) \text{ size}(h_{bug}(t, X, i, j))) \\
 &\Leftarrow \\
 \text{Find: } h_{bug}(t, (\text{aircraft, depart, arrive, manifest}) \star R, i, j) \\
 &= (\dots) \triangleright R
 \end{aligned}$$

The problem is that $(\dots) \triangleright R$ is of sort $set(Reservation)$, a topic-domain sort, whereas h_{bug} has access only to the source data, $(\dots) \star R$. There is no reason why h_{bug} could not exist as a mathematical abstraction, but it can't be the denotation of a computable process, because computations output data structures that *represent* reservations, not the reservations themselves. We must find a way to transform the target side of goal $[G^*]$ so that it is expressed in terms that can be found in the source representation.

4 Sourcing Rules

The key heuristic we employ is to replace occurrences of target-domain terms with their images under functions mentioned in the source sides of goals. This heuristic is licensed by the following rule:

★-Introduction Rule

For any $m : A \rightarrow B$ and $C \subseteq A$

$$C = m^{-1} \star m \star C$$

$$\Leftarrow m \text{ is injective}$$

If we have a problematic occurrence of a topic term C , and we know that $m \star C$ occurs in the source term S , then we try replacing C with $m^{-1} \star m \star C$. The hope is that the m^{-1} part can be simplified away by further transformations, leaving us with a term $(m \star C)$ that is closer to what is available on the source side. A rule of this kind that pushes a target-side term closer to being identical with a source-side term is called a *sourcing* rule.

All the previous rules matched an entire goal of the form “*Find: f ... = ...*,” producing a solution (and possibly some subgoals). We call them *solution rules*. The image-introduction rule matches a particular subterm of a goal, allowing us to produce a new goal by substituting another term for that subterm (again possibly accompanied by new subgoals). We call rules with this behavior *rewrite rules*. Unless stated otherwise, they are applied left-to-right. That is, $\alpha = \beta$ is used to rewrite terms matching α with terms

matching β . Rewrite rules that do not “sourcify” are labeled with the catch-all phrase *simplification rule*. Although in the end we will require dozens of simplification rules, for the purposes of our commitment-matrix example we need only two:

Image-Filter Rule

For any $p : A \rightarrow \text{boolean}$
 and any $f : A \rightarrow B$
 $p \triangleright (f \star X) = f \star (p' \triangleright X)$
 $\Leftarrow \text{Find: } p' x = p(f x) \text{ for all } x \in X$

“You can switch the order of a filter and image operation if you adjust the type of the filtering predicate.”

Size-Image Rule

For any $f : A \rightarrow B$
 $\text{size}(f \star X) = \text{size}(X)$
 $\Leftarrow f \text{ is injective}$

“If f is injective, it preserves the size of any set it operates on.”
 We also require one further solution rule:

Inverse-Composition Rule

For any $f : B \rightarrow C$
 and injective $g : B \rightarrow A$
 To *Find: $h x = f(g^{-1} x)$*
 $\Leftarrow \text{Find: } h(g y) = f y \text{ for arbitrary } y \in B$

In other words, to find an h that is the composition of f and g^{-1} , find an h such that $f = h \circ g$.⁶ To understand the rule, note that, since g is injective, for any $x \in A$, there is a unique $y \in B$ such that $x = g y$. Then $h x = h(g y) = f y = f(g^{-1} x)$.

Now we derive the commitment-matrix generator by proceeding as shown in figure 1 with our problematic subgoal.

Putting it all together, we have:

$$f = (\lambda(S, i, j) \quad ((\lambda t t \mapsto (\text{size}(\lambda(a, d, r, m) \ d \leq t < r) \triangleright S)) \star \{i..j\}))$$

⁶This looks a little different from the previous solution rules, because the h found in the subgoal happens to be identical to the h the rule finds. So rather than write *Solution: $h = h'$* , and $\Leftarrow \text{Find: } h'(g y) = \dots$, we just write the rule as shown.

$$\begin{aligned} \text{Find: } & g_2(t, (\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest}) \star R, i, j) \\ & = \text{size}((\lambda r (\text{depart } r) \leq t < (\text{arrive } r)) \triangleright R) \end{aligned}$$

Use \star -Introduction to get:

$$\begin{aligned} \text{Find: } & g_2(t, (\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest}) \star R, i, j) \\ & = \text{size}((\lambda r (\text{depart } r) \leq t < (\text{arrive } r)) \\ & \quad \triangleright (\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest})^{-1} \\ & \quad \star (\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest}) \\ & \quad \star R) \end{aligned}$$

\Leftarrow (by Image-Filter Rule and the definition of S)

$$\begin{aligned} & g_2(t, S, i, j) \\ & = \text{size}((\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest})^{-1} \\ & \quad \star p' \triangleright S) \\ & = (\text{by Size-Image Rule}) \\ & \quad \text{size}(p' \triangleright S) \\ & \Leftarrow \text{Find} : p'(a, d, r, m) \\ & \quad = (\lambda r (\text{depart } r) \leq t < (\text{arrive } r)) \\ & \quad \quad ((\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest})^{-1} \\ & \quad \quad (a, d, r, m)) \end{aligned}$$

Using the Inverse-Composition Rule:

$$\begin{aligned} \Leftarrow \text{Find} : & p'((\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest}) r) \\ & = ((\text{depart } r) \leq t < (\text{arrive } r)) \end{aligned}$$

which succumbs to imitation and projection, yielding:

$$p' = (\lambda (a, d, r, m) d \leq t < r)$$

So that the solution for g_2 can be found easily using imitation and projection:

$$g_2 = (\lambda (t, X, i, j) \text{size}((\lambda (a, d, r, m) d \leq t < r) \triangleright X))$$

Note: We have left out the required proof that $(\text{aircraft}, \text{depart}, \text{arrive}, \text{manifest})$ is injective, which follows from the “non-overlap” axiom in the Schedule-DB spec. However, it would be an unusual database that didn’t satisfy injectivity in this way.

Figure 1: Derivation of problematic subfunction

In this example, the Size-Image Rule served as “*deus ex machina*” to cause the function

$$(aircraft, depart, arrive, manifest)^{-1}$$

to vanish completely. In general, the process isn’t so abrupt or tidy. Typically several rounds of simplification are necessary to eliminate all such vestiges of the topic domain.

5 Search Management

Of course, there is a big gap between a bunch of higher-order equations and an algorithm for applying them. It is out of the question to fill this gap with a general-purpose higher-order theorem prover, because the search space for even small problems is enormous. However, for this application we believe that there is considerable guidance available. The strategy we followed in the examples above may be summarized thus:

To solve $f S = T$:

Use imitation and projection whenever the result would consist of operations on source- or target-domain entities.

When it wouldn’t, eliminate a topic-domain term t in the target T by replacing it with a term that may be schematized as $\sigma^{-1}(\sigma t)$. Here σt stands for a functional expression that appears in the source specification S , and σ^{-1} stands for whatever “upstream adjustment” is necessary to preserve the meaning of T .

Continue simplifying with imitation and projection. During this process, mark $\sigma(t)$ as *encapsulated*, so that it can be incorporated into other terms, but cannot be decomposed.

We use the term *Source-Substitution Principle* for this idea of replacing t by $\sigma^{-1}(\sigma t)$.⁷ The principle suggests arranging sourcifying rules by what sorts of σ they handle. Here are two more. The \triangleright -Introduction Rule is used when $\dots \triangleright t$ occurs in S and t occurs in T . The Representative-Element Rule is used when a set X occurs in S , and only an element $x \in X$ occurs in T .

⁷We have noticed the resemblance of this technique to the “rippling” technique of inductive theorem provers [5], in which the object is to transform a deductive goal into a form that matches an induction hypothesis. Our “encapsulations” correspond roughly to the “wave fronts” of the rippling mechanism. We don’t know how deep the resemblance is.

▷-Introduction Rule

For any $p : A \rightarrow \text{boolean}$ and $B : \text{set}(A)$
 $B = (p \triangleright B) \cup (\text{neg } p \triangleright B)$

where $\text{neg } p = (\lambda x \neg p x)$. This is a rewrite rule used to replace a topic-domain term B with the more complex term on the right. The subterm $p \triangleright B$ is encapsulated, and the other terms must be further transformed.

Representative-Element Rule

For any $f : A \rightarrow B$
and any $g : A \rightarrow C$
 $g(x) = h(\text{only_elt}((\lambda y (d y) = (d(f x))) \triangleright f \star X))$
 $\Leftarrow x \in X$
and $d \circ f$ is injective over X
and *Find*: $h(f y) = g y$ for all $y \in X$.

The function *only_elt*: $\text{set}(A) \rightarrow A$ denotes the only element of its argument if its argument has one element, and is undefined otherwise: $(\text{only_elt}\{x\}) = x$. In the context of this rule it will always be defined. What the rule says is that, to perform a computation on a topic-domain element x , you may transform it to the element of the source representation $f \star X$ that “represents” it.

This rule applies very generally, and hence is used only as a last resort. It applies to any formula at all with one or more occurrences of a term x such that $x \in X$, provided that $f \star X$ occurs in the source, for some function f , and a function d can be found that allows us to find x in $f \star X$. For example, if the source contains $(a, b) \star X$, and b is a unique identifier for every object, then the rule matches the target $\text{size}((\lambda u u > x) \triangleright Z)$ with $g = (\lambda v \text{size}((\lambda u u > v) \triangleright Z))$, and suggests rewriting the target as

$$h(\text{only_elt}((\lambda (y_a, y_b) y_b = b x) \triangleright ((a, b) \star X)))$$

after finding an h such that $h(a y, b y) = \text{size}((\lambda u u > y) \triangleright Z)$ for all $y \in X$. The idea is to use b to find the element $(a x, b x)$ of $((a, b) \star X)$ that represents x , and hope that there is enough information in that element so that $g x$ may be computed by computing $h(a x, b x)$.

We have worked several examples by hand using our strategy. For instance, consider the following source-data descriptor:

$S = (S_1, S_2)$
where $S_1 = (\text{id}, (\lambda x (\lambda (k, n) n > 0))$

$$\begin{aligned} & \triangleright (cid, (\lambda c \ qty(c, x))) \\ & \star C)) \\ & \star F, \\ \text{and } S_2 = & (cid, wcl) \star C \end{aligned}$$

S consists of two substructures, S_1 and S_2 that specify properties of a set of airplane flights, F , and a set of cargo-item classes C . S_1 is a table giving the number of items of each class aboard each flight as a set of tuples

$$(\text{flight-id } f, ((cid \ c_1, qty(c_1, f)), (cid \ c_2, qty(c_2, f)), \dots))$$

one per flight in the set F . We use the notation $\langle \dots \rangle$ to indicate a list of an indefinite number of items. The pair $(cid \ c_j, qty(c_j, f_i))$ specifies the number of items of cargo type c_j on board flight f_i . $(cid \ c)$ is the id of cargo type c . The entry is omitted if $qty(c_j, f_i) = 0$. The table S_2 gives the weight of each item class as a set of tuples $(cid \ c, \text{cargo-type-weight } c)$ for every cargo type c in the set C . In addition to cid , we use idf to refer to the id number of a flight; and $wcl \ c$ to refer to the weight of an item of class c .

The “Absent Items” problem is to derive a table showing, for each item class, whether it occurs on any flight. The table consists of a list of tuples $(\text{cargo-id}, \text{boolean})$:

$$T = (cid, (\lambda c \ \vee / ((\lambda x \ qty(c, x) > 0) \star F))) \star C$$

(\vee / X is **true** if and only if some element of X is **true**.)

This problem seems quite a bit harder than the examples we described above, but our method requires no search at all; every possible application of the Source-Substitution Principle is necessary.

We foresee two places where the method will break down:

1. If it is required to chain together an indefinite number of tuple elements, then some induction will be required. For example, given a relational database that shows the immediate successors of each airplane flight, transforming it to an explicit chain requires a recursive “collection” algorithm. Our technique works, but requires some insight, and so will be hard to mechanize.
2. When term rewriting becomes true deduction rather than mere simplification, the number of possible search paths could become large.

For instance, suppose you wanted to produce the first component S_1 of the source for the absent-items problems, which describes the *classes* of

items on board, given a data structure that lists the *individual items* on board. S_1 is the target in this problem, and the source is composed of two tables (R_1, R_2), where R_1 is a set of tuples (*flight-id*, \langle *item-id* \rangle), whose second component is a list of the ids of the items on the given flight; and R_2 is a set of tuples (*item-id*, *item-class-id*), which specifies for each item which class it belongs to. More formally:

$$\begin{aligned}
R &= (R_1, R_2) \\
\text{where } R_1 &= (id, (\lambda f iid \star \{x | aboard(x, f)\})) \\
&\quad \star F \\
\text{and } R_2 &= (iid, cid \circ icl) \star C
\end{aligned}$$

where $id(f)$ is the id of flight f , $iid(x)$ is the id of item x , $icl(x)$ is the class of item x , and $cid(c)$ is the id of class c . The problem is to find g such that $g R = S_1$; in other words:

$$\begin{aligned}
\text{Find: } g &((id, (\lambda f iid \star \{x | aboard(x, f)\})) \\
&\quad \star F, \\
&\quad (iid, cid \circ icl) \star C) \\
&= (id, (\lambda x (\lambda (k, n) n > 0) \\
&\quad \triangleright (cid, (\lambda c qty(c, x)))) \\
&\quad \star C)) \\
&\quad \star F
\end{aligned}$$

Our method can handle this, but it requires partitioning the set of all items in R_1 (i.e., $iid \star \{x | aboard(x, f)\}$) by item class and applying the following *deductive rule*:

Sum-partition Rule

For any $f : A \rightarrow Real$

and any $Y : set(A)$

$$+/\{ : f x | x \in Y \} = +/\{ : +/\{ : f x | x \in s \} | s \in P \}$$

$\Leftarrow P$ is a partition of Y

In this rule, $\{ : f(x) | \dots \}$ is the bag (multiset) of all $f(x)$ for x satisfying the condition "...". \oplus/b applies associative and commutative function \oplus to the elements of bag b . So $+/\{ : 3, 3, 4 \} = 10$.

The rule says that to add up $f(x)$ for every element of a set, find a partition and add up the subtotals for all the subsets the partition defines. Exactly when to apply deductive rules instead of imitation is the subject of our current research.

6 Glue Code Middle Agents

We have implemented a prototype glue code generation agent in a demonstration system simulating a military disaster relief scenario. In doing so, we have begun to discover and address issues of how such agents would be utilized and deployed.

In our demonstration, a glue code agent is deployed that receives requests for translation services and produces specific translation agents that the sender forwards messages through when the message is of a specific kind, and the ultimate recipient is of a specific agent type. In general, the indexing of these translators may become an issue, and various alternative schemes for, say, adding translation functions to a single middle agent responsible for interactions between two specific classes of agents might be preferable. Indeed, even for our simple case, there is a need to develop two different translation functions, one for the request and one for the response.

In one example from our scenario, an agent *S* uses a matchmaking service to find another agent *T* capable of delivering quantities of relief supplies to different locations by specified times. *S* standardly requests such services with a message like this:

```
(distribute :cargo-destination (:type <cargo-type>
                                :location <location-id>
                                :quantity <quantity-in-lbs>
                                :required-delivery-time <time-gmt>)*
           :cargo-source (:type <cargo-type>
                          :location <location-id>
                          :quantity <quantity-in-lbs>
                          :available-time <time-gmt>))
```

That is, a single source load of cargo, available at the stated time, is to be divided and distributed to a number of sites.

The agent that is found to do these deliveries uses a different request form, where the source and destination for a single delivery is specified together with the type and quantity, but where the quantity is in a different unit of measure.

```
(deliver :items (<cargo-type>
                <source-location-id>
                <destination-location-id>
                <quantity-in-kilograms>)
```

<available-time-local>)*)*

The glue code generator is established as a middle agent that accepts requests to “spawn” translator agents for particular messages between agents or classes of agents. In this case, agent *S* uses the description of the service found by the matchmaker to request a translator from the glue code agent. This request expresses the above two forms in terms of topic domain sets, variables and concepts:

Topic Variables: *REQ-GOALS* : *set(Move-Cargo-Goal)*
 CARGO-SOURCE : *Cargo-Availability*

SourceForm =

```
(distribute
  :cargo-destination
    (λ(g)
      (make cargo-goal-spec
        :type (goal-type g)
        :location (location-id (goal-destination g))
        :quantity (weight (goal-amt g) lbs)
        :required-delivery-time
          (date (goal-deadline g) zulu-time)))
    * REQ-GOALS
  :cargo-source
    (make cargo-source-spec
      :type (load-type CARGO-SOURCE)
      :location (load-loc CARGO-SOURCE)
      :quantity (weight (load-amt CARGO-SOURCE) lbs)
      :available-by (date (load-available-time
        CARGO-SOURCE) zulu-time)))
```

TargetForm =

```
(deliver
  :items (λ(g)
    (make delivery-spec
      (goal-type g)
      (weight (goal-amt g) kg)
      (load-loc CARGO-SOURCE)
      (goal-destination g)
      (date (load-available-time CARGO-SOURCE)
        (local-time (load-loc CARGO-SOURCE)))))
    *REQ-GOALS))))
```

This example raises certain mundane issues that are going to be quite important in realistic applications. A function like *date* or *weight* produces a computational value only when provided with a “unit system” such as *kg* or *zulu-time*.⁸ A unit system is not itself a computational object, and can’t be manipulated as such.

Handling these examples requires a simple—almost trivial—sort of source-coding rule. For instance:

Weight-Translation Rule

$$(weight\ x\ lb) = 2.2046225 \times (weight\ x\ kg)$$

If a term matching the left-hand side occurs in the target, and a term matching the right-hand side occurs in the source, then we rewrite and encapsulate just as for more profound rules.

The LISP code currently produced for this, after running through an optimizer, looks as follows, with variable names changed for readability:

```
(LAMBDA (MSG-FORM)
  (LET((CARGO-SRC (NTH 4 MSG-FORM)))
    (LIST 'DELIVER :ITEMS
          (MAPCAR #'(LAMBDA (REQ)
                    (LIST (NTH 1 REQ)
                          (* 2.2046225 (NTH 5 REQ))
                          (NTH 3 CARGO-SRC)
                          (NTH 3 REQ)
                          (ZULU-TO-LOCAL-TIME
                           (NTH 7 CARGO-SRC))))
                  (NTH 2 MSG-FORM))))))
```

A larger open issue for our approach to inter-agent translation arises from a difference between the potential differences between *requests* and *responses*. It is largely taken for granted in the agent literature that advertisements of agent services will need to include the specification of message formats for messages *to* the agent so advertising. However, there is little or no mention of the need for advertisements to include the specifications of allowable responses to requests generated by the agent. To translate responses, the glue code generation system will need to know the permissible

⁸A time system used by the U.S. military.

forms those responses can have, or have direct access to the full ontology of the requesting agent, so that appropriate terms and formulas can be chosen. This is a topic under consideration in our current work extending the approach outlined in this paper.

7 Summary and Prospectus

We have sketched an approach to automatic derivation of “glue code,” programs for translating the output of a *source* agent to the input representation of a *target* agent. We have argued that

1. Generating glue code requires a theory of a *topic domain*, the common subject matter of the source and target representations, and mappings from the source and target to the topic.
2. Given these mappings, the problem of glue-code generation can be solved by solving a higher-order equation for the unknown transformation function.

It might be objected that formalizing the mappings from source and target to the topic domain is just as difficult as writing glue code, and so we have traded a hard problem for another problem that might be even harder. We disagree, for a variety of reasons. For one thing, other researchers have proposed independent reasons for wanting a common “ontological” framework for agents. Advertising and matchmaking are two obvious applications of such shared notations. So this problem is already fairly urgent. (For a survey, see [17].)

Second, if it is impossible to supply formal theories of agent semantics, then we believe it is essentially impossible to glue agents together automatically. The only exceptions will be agents whose outputs are to be processed directly by humans. For instance, a search engine might look for agents that can give weather reports, and if it finds five such agents might send the output of all five to a screen for a human to look at. If one of the weather reports is a marine forecast, and the human is looking for an upper-atmosphere forecast, then the human will skim through it and reject it. However, we don’t believe this paradigm is going to apply to the majority of agent interactions. If a human has to be in the loop at “agent-composition time,” then agent composition will be unworkable.

Third, if formalization is done right, it has to be done only once per agent. Once the formalization is in place, it can support creation of glue code for multiple partner agents.

Fourth, we don't believe a commitment to formalizing domain theories is a commitment to finding a "formal theory of everything," which is fortunate, because there are good reasons to doubt that there is such a theory. We need to encourage agent producers to formalize the representations of their agents within some convenient framework. Such frameworks must be public and well advertised, but they don't have to be completely global. All that is necessary is that, when two agents are to be glued together, they are formalized within the same framework, or within two frameworks that are intertranslatable.

The third and fourth points in this list are in conflict. To avoid having to repeat work, we want the formal framework to be as general as possible, but to keep the formalization intuitive and clear, we want to neglect as much as possible. For instance, suppose one thing an agent represents is a list of available airplanes. On the "simple" end of the formalization spectrum we have theories in which `available` is a unary predicate. On the "complex" end we have a theory of what it means for something to be available to someone for some purpose over some time interval subject to a set of use constraints. It must be possible to pick a point on the spectrum that seems to make sense, and possibly generalize it later if that point fails to support liaison with some unanticipated class of agents.

We believe that these "meta" ideas can themselves be formalized to some degree. We need the notion of a *lattice of theories*, in which $T_1 \sqsubset T_2$ if T_2 can express everything T_1 can express but not vice versa. In other words, T_2 makes distinctions that T_1 doesn't make. We think this idea can be developed using the concept of *theory morphism* [23].

The other main unknown in this line of research is how tightly the search for a transformation can be controlled. The derivation strategy that has emerged from our examples is to perform second-order matching modulo the laws of the domain theory. Matching rules are applied unless doing so would expose the terms of the topic-domain theory. At that point we use *sourcifying* rules to wrap topic terms inside constructs from the source domain. The implementation mentioned in Section 6 used the Specware general-purpose theorem prover to find glue code [23]. It is worth seeing if a more specialized algorithm would be more efficient.

A key issue is acquisition of the common domain theory and the domain-specific rules. The rules of the domain theory are used to reformulate various terms to facilitate the matching process. This suggests that a means-end analysis might be useful in detecting the need for rules that translate between concepts on the left and right-hand sides of the translation equation. On a related note, one possibility is to analyze the rules of the domain theory

to determine how to orient them, and to construct a proof plan. There is an opportunity here to use our (meta)knowledge of the derivation strategy (or proof plan) to query the domain-expert/user for rules that link certain terms.

Another issue, to be dealt with in the future, is how to transform glue-code programs into a more efficient form once they have been created. This goal is addressed by our past work on program transformation [22].

8 Acknowledgements

Tao Gui and Duško Pavlović provided helpful comments on drafts of this paper. The referees made suggestions that helped us reorganize the discussion of agent communication in general.

References

- [1] J. L. Ambite, C. A. Knoblock, I. Muslea, and A. Philpot. Compiling source descriptions for efficient and flexible information integration. *J. Intelligent Information Systems*, 16(2):149–187, 2001.
- [2] N. Ashish and C. A. Knoblock. Semi-automatic wrapper generation for internet information sources. In *Proc. Second IFCIS Conference on Cooperative Information Systems (CoopIS)*, 1997. Charleston, South Carolina.
- [3] R. J. Bayardo, W. Bohrer, R. Brice, A. Cichocki, J. Fowler, A. Helal, V. Kashyap, T. Ksiezyk, G. Martin, M. Nodine, M. Rashid, M. Rusinkiewicz, R. Shea, C. Unnikrishnan, A. Unruh, and D. Woelk. Infosleuth: Agent-based semantic integration of information in open and dynamic environments. *ACM SIGMOD '97*, pages 195–206, 1997. ACM.
- [4] J. Bradshaw, editor. *Software Agents*. AAAI Press/MIT Press, 1997.
- [5] A. Bundy, A. Stevens, F. van Harmelen, A. Ireland, and A. Smaill. Rippling: a heuristic for guiding inductive proofs. *Artificial Intelligence*, 62(2):185–253, 1993.
- [6] T. Finin, Y. Labrou, and J. Mayfield. Kqml as an agent communication language. In Bradshaw [4]. .

- [7] FIPA. Fipa acl message structure specification, 2001. <http://www.fipa.org/specs/fipa00061/XC00061E.html>.
- [8] FIPA. Fipa communicative act library specification, 2001. <http://www.fipa.org/specs/fipa00037/XC00037H.html>.
- [9] D. Florescu, L. Raschid, and P. Valduriez. A methodology for query reformulation in Cis using semantic knowledge. *Int. J. of Cooperative Information Systems*, 1996.
- [10] J. A. Giampapa, M. Paolucci, and K. Sycara. Agent interoperation across multagent system boundaries. In *Proc. Fourth International Conference on Autonomous Agents (Agents 2000)*, 2000. New York: Association for Computing Machinery.
- [11] J. Hammer, H. Garcia-Molina, J. Cho, R. Aranha, and A. Crespo. Extracting semistructured information from the web. In *Proc. Workshop on Management of Semistructured Data*, 1997. Tucson, Arizona.
- [12] G. Huet and B. Lang. Proving and applying program transformations expressed with second-order patterns. *Acta informatica* , 11:31–55, 1978.
- [13] M. Klusch and K. Sycara. Brokering and matchmaking for coordination of agent societies: a survey. In A. Omicini, F. Zambonelli, M. Klusch, and R. Tolksdorf, editors, *Coordination of Internet Agents: Models, Technologies and Applications*. Springer Verlag, 2001. .
- [14] Y. Labrou, T. Finin, and Y. Peng. The current landscape of agent communication languages. *IEEE Intelligent systems* , 14(2):45–52, 1999.
- [15] D. L. Martin, A. J. Cheyer, and D. B. Moran. The open agent architecture: A framework for building distributed software systems. *Applied Artificial Intelligence* , 13:91–128, 1999.
- [16] J. Mayfield, Y. Labrou, and T. Finin. Evaluation of kqml as an agent communication language. In M. Wooldridge, J. P. Mueller, and M. Tambe, editors, *Intelligent Agents II: Agent Theories, Architectures, and Language*. Springer-Verlag Lecture Notes in AI, Volume 1037, 1996. .
- [17] T. Menzies. Cost benefits of ontologies. *Intelligence* , 10(3):27–32, 1999.

- [18] T. Milo and S. Zokar. Using schema matching to simplify heterogenous data translation. In *Proc. Conf. on Very Large Data Bases*, pages 122–133, 1998.
- [19] I. Muslea, S. Minton, and C. A. Knoblock. Wrapper induction for semistructured, web-base information sources. *Conf. on Automated Learning and Discovery*,, 1998.
- [20] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. Object exchange across heterogeneous information sources. In *Proc. Eleventh International Conference on Data Engineering (Taipei)*, pages 251–260, 1995.
- [21] Y. Papakonstantinou, A. Gupta, H. Garcia-Molina, and J. Ullman. A query translation scheme for rapid implementation of wrappers. In *Proc. DOOD'95*, 1995.
- [22] D. R. Smith. Kids: A semi-automated program development system. *IEEE Transactions on Software Engineering* , 16(9):1024–1043, 1990. Special Issue on Formal Methods, September.
- [23] Y. Srinivas and R. Jülig. Specware: Formal support for composing software. In *Proc. Conf. on Mathematics of Program Construction*, pages 399–422. Springer-Verlag, 1995. B. Moeller, Ed. LNCS 947.
- [24] K. Sycara, M. Klusch, S. Widoff, and J. Lu. Dynamic service match-making among agents in open information environments. In *J. ACM SIGMOD Record* (1), pages 47–53, 1999. (Special issue on Semantic Interoperability in Global Information Systems,).
- [25] K. Sycara, M. Paolucci, M. V. Velsen, and J. A. Giampapa. The Retsina Mas Infrastructure. Technical report, Carnegie Mellon University Robotics Institute, 2001. CMU-RI-TR-01-05.
- [26] J. D. Ullman. *Principles of Database and Knowledge-Base Systems*. Computer Science Press, 1988.
- [27] J. D. Ullman. *Principles of Database and Knowledge-Base Systems*. Computer Science Press, 1988.
- [28] J. D. Ullman. Information integration using logical views. In *Proc. 6th Intl. Conf. on Database Theory (ICDT-97), Lecture Notes in Computer Science*, pages 19–40, 1997. Springer-Verlag.