Matchmaking Multi-Party Interactions Using Historical Performance Data

David Lambert and David Robertson
School of Informatics
University of Edinburgh
Edinburgh, Scotland
d.j.lambert@sms.ed.ac.uk, dr@inf.ed.ac.uk

ABSTRACT

Matchmaking will be an important component of future agent and agent-like systems, such as the semantic web. Most research on matchmaking has been directed toward sophisticated matching of client requirements with provider capabilities based on capability descriptions. This is a vital mechanism for conducting matchmaking, but ignores the likelihood that in practice, and for various reasons, capability descriptions will not fully characterise the interaction behaviour of agents. This problem is further compounded in systems with many interacting agents, all of which have idiosyncrasies. As in everyday life, some groupings of agents will be more effective than others, regardless of their individual competencies or suitability to the task. The quality of the interaction between agents is a crucial factor. Using the incidence calculus and the lightweight coördination calculus, we show that we can easily implement matchmaking agents that will learn from experience how to select those groups known to inter-operate well for particular tasks.

General Terms

Algorithms, Reliability

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Correlation and regression analysis; I.2.11 [**Distributed Artificial Intelligence**]: Coherence and coordination; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

Keywords

incidence calculus, service composition

1. INTRODUCTION

The issue of matchmaking is key to the deployment of many types of multi-agent systems [1]. It first appeared as a research problem at the dawn of agent systems [2], and has resurfaced as a fundamental problem in newer areas like the semantic web and Grid computing.

Most of the previous research into agents has addressed interactions with only two-parties, service -requester and -provider. This bias continues on the semantic web. OWL-S [3], for instance, imagines that any interaction will be two-party. One can imagine interactions that are inherently multi-agent and which would thus require any matchmaker to find an appropriate set of agents. Indeed, recent developments in web services choreography [4] reveal the growing realisation that many real-world processes require multiple participants.

The technique presented here developed out of work on the lightweight coördination calculus (LCC) [5], a language for specifying and executing multi-party dialogues, where we found it necessary to provide a middle-agent to supply collaborators for the fulfilment of protocols.

We do not commit to any particular deployment environment. While we have implemented it for LCC, we believe the approach would have value any setting where service capabilities are specified in a flexible, semantic encoding. Consequently, we see applications not only in multi-agent systems, but on the Grid [6] and semantic web, and in work-flow and peer-to-peer systems.

The remainder of this paper is organised thus: We first examine the problem in greater detail in section 2, before introducing the two key ideas we use to tackle it—the lightweight coördination calculus and the incidence calculus—in section 3. Section 4 shows how we use the incidence calculus to select an appropriate team of agents for a task. Section 5 surveys related work, and section 6 concludes.

2. MOTIVATION

The motivation for matchmaking is clear: in any environment with large numbers of agents, the only feasible mechanism for connecting those requesting a service with those willing to provide it is via middle agents [1, 7]. However, most previous work has examined only the case of selecting a single agent (or in some cases, a range of agents, leaving the final selection of the agent to the client) for a single role. To our knowledge, [8] is the only exception.

While current interactions are primarily client-server, we can imagine a future where matchmade agent interactions are more distributed, involve many agents, and operate in a more peer-to-peer manner. It can be expected that these newer forms of dialogue will make even greater use of, and demands upon, matchmaking services than do current modes of employment.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

2.1 Capability description is not enough

In [8] the assumption is that agent implementations would vary in their *intrinsic ability* to complete a task, even if the task were completely specified and accurately advertised. We adopt this view, and extend it: we think that, as well as the agent's specific task-specific talent, there is also reason to believe that issues of ontological mismatch and other 'social' effects will reduce the efficacy of purely logical agent capability descriptions. Regardless of how carefully service providers specify the behaviour of their systems, there will always remain some level of 'semantic slack' such that matchmaking based on purely semantic service description can be improved by statistical or other learning techniques. In these circumstances, an approach like ours would be beneficial. Below, we examine several reasons for this encoding gap. Note that only the first is a technical problem: in the others, non-technical reasons prevent additional semantic information being embedded in profiles.

- **The capability description language lacks expressiveness.** This does not imply a criticism of the language: it is unreasonable to expect any general purpose capability description language to allow the communication of arbitrarily complex constraints in every imaginable domain. However, it would frequently be possible for matchmakers, especially domain-specific ones, to discover such constraints.
- **User ignorance of ability of the language to express a constraint, or of the effect of declaring the constraint.** As constraints become more complex, and services more common, it becomes increasingly likely that a user would be unaware of her ability to aid the matchmaker.
- **User expectation that the information will not be used by clients or matchmakers.** Service providers may refrain from supplying this kind of data until they observe a significant portion of the agent ecosystem using it.
- **Not wanting to express particular information.** In some instances, there is an incentive for service providers to keep the description of their services as general as possible, though not to the extent of attracting clients they have no possibility of pleasing. Alternatively, the provider may not wish to be terribly honest or open about her service's foibles.
- **The inter-relationship is not known to the provider.** Some of the dependencies may be subtle, even if they have a large effect.
- **Comprehensive constraints too expensive to generate or use.** Even if none of the above hold, it would often simply not be worthwhile for the service provider to analyse and encode the information. Further, in the case of web, semantic web and Grid services, it is reasonable to expect that users are discouraged by standards flux from investing much time in this endeavour.

But what underlying causes are there for these problems that are so difficult to encode? Below are some reasons why such issues might arise:

- **'Social' reasons.** For instance, different social communities, or communities of practice, may each cluster around particular service providers for no a priori reason, yet this would result in improved performance on some tasks if agents were selected from the same social pool.
- **Inter-business partnerships.** An airline may have a special deal with other airlines or car-hire companies.
- **Components designed by same group.** Organisations that seem to have nothing in common may well be using software created by a single group. Such software would be more likely to inter-operate well than software from others.
- **Different groups of engineers held differing views of a problem, even though the specification is the same.**Thus, the implementations are subtly incompatible, or at least do not function together seamlessly.
- **Particular resources or constraints shared between providers.** In a Grid environment, a computation server and a file store might share a very high bandwidth connection, leading to improved service. This particular case underlies our example scenario, detailed in figure 3.

3. TECHNICAL PRELIMINARIES

Our matchmaker relies on two techniques developed previously: the lightweight coördination calculus, and the incidence calculus. We briefly explain these here.

3.1 Lightweight coördination calculus

The lightweight coördination calculus (LCC) [5] is a method for specifying agent interaction protocols. A generalisation of the Electronic Institutions [9] model, it is based on process calculi, specifically CCS [10]. It provides a simple message passing framework (denoted $\Rightarrow$ for sending, and $\Leftarrow$ for receiving) with the operators *then* (sequence), *or* (choice), *par* (parallel execution), and $\leftarrow$ (if). The grammar is shown in figure 1. An LCC protocol framework is interpreted in a logic-programming style, using unification of variables which are gradually instantiated as the conversation progresses. Along with this dialogue framework, which specifies the various messages that can be transmitted and when, a protocol carries 'common knowledge' (data, specific to a conversation, that every participant in the dialogue can access and modify) and clauses which indicate to the agents where they have reached in the dialogue.

This style of protocol definition is flexible, and allows us to easily capture, and perform matchmaking for, multi-agent interactions, the primary contribution of this paper. Conveniently, LCC allows us to use the same language to express various degrees of decentralisation. For example, we can convert a protocol from one in which the middle-agent functions as a broker (routing all communication through itself, delivering only the final result to the client), to a matchmaker (once the middle-agent has identified agents for the required roles, it informs the client of the decisions and plays no further part in the protocol's execution).

The protocols can be created dynamically, but for this paper we choose to situate the matchmaker system in an environment where a library of standard protocols exists. Each protocol functions as a pre-defined plan. An agent, wishing to accomplish some task (such as creating an auction, or calling a meeting, finding participants, and arranging a mutually suitable time) will select a protocol from a library, and ask its matchmaker to suggest service-provider agents for each role. In some sense, each role is equivalent to an atomic service capability, although a role carries also the responsibility of participating in a specific pattern of on-going conversation.

Figure 1: Grammar for the LCC dialogue framework

$$
\begin{array}{rcl}
\mathit{Framework} & ::= & \mathit{Clause}, \mathit{Clause}+ \\
\mathit{Clause} & ::= & \mathit{Agent} :: \mathit{Def} \\
\mathit{Agent} & ::= & a(\mathit{Role}, \mathit{Id}) \\
\mathit{Def} & ::= & \mathit{Agent} \mid \mathit{Message} \mid \mathit{Def}\ \mathit{then}\ \mathit{Def} \mid \\
 & & \mathit{Def}\ \mathit{or}\ \mathit{Def} \mid \mathit{Def}\ \mathit{par}\ \mathit{Def} \\
\mathit{Message} & ::= & M \Rightarrow \mathit{Agent} \mid M \Rightarrow \mathit{Agent} \leftarrow C \mid \\
 & & M \Leftarrow \mathit{Agent} \mid C \leftarrow M \Leftarrow \mathit{Agent} \\
C & ::= & \mathit{Term} \mid C \wedge C \mid C \vee C \\
\mathit{Role} & ::= & \mathit{Term} \\
\mathit{Id} & ::= & \mathit{Term} \\
M & ::= & \mathit{Term}
\end{array}
$$

Figure 2: Incidence calculus rules

$$
\begin{array}{llll}
i(\top) & = \mathit{allworlds} & i(\bot) & = \{\} \\
i(\neg\alpha) & = i(\top) \backslash i(\alpha) & & \\
i(\alpha \wedge \beta) & = i(\alpha) \cap i(\beta) & i(\alpha \vee \beta) & = i(\alpha) \cup i(\beta) \\
i(\alpha \rightarrow \beta) & = i(\neg\alpha \vee \beta) = (i(\top) \backslash i(\alpha)) \cup i(\beta) & &
\end{array}
$$

Probabilities are derived from incidences thus:

$$p(\phi) = \frac{|i(\phi)|}{|i(\top)|} \qquad p(\phi|\psi) = \frac{|i(\phi \wedge \psi)|}{|i(\psi)|}$$

3.2 Incidence calculus

The incidence calculus [11] is a truth-functional probabilistic calculus in which the probabilities of composite formulae are computed from intersections and unions of the sets of worlds for which the atomic formulae hold true, rather than from the numerical values of the probabilities of their components. The probabilities are then derived from these incidences. The rules are given in figure 2. The key feature of the calculus is that, in general, $p(\phi \wedge \psi) \neq p(\phi) \cdot p(\psi)$. This raises the fidelity of derived probabilities in comparison to other probabilistic logics. The conditional probabilities of incidences drive our matchmaking, as we see in the next section.

As an illustration, consider the following set of incidences describing the weather in a given week:

$$
\begin{array}{ll}
i(\top) & = \{mon, tue, wed, thu, fri, sat, sun\} \\
i(rain) & = \{mon, wed, thu, fri, sat, sun\} \\
i(wind) & = \{mon, thu, fri, sun\} \\
i(sun) & = \{tue, wed, sat, sun\} \\
i(snow) & = \{sat\} \\
i(rain \wedge sun) & = i(rain) \cap i(sun) = \{wed, sat, sun\} \\
i(\neg rain \vee snow) & = i(\neg rain) \cup i(snow) = \{tue, sat\}
\end{array}
$$

To illustrate the computation of probabilities of compound formulae from the incidences, note that the probabilities of *wind* and *sun* are both $\frac{4}{7}$, but their conjunctions with *rain* (probability $\frac{6}{7}$) are different, at $\frac{4}{7}$ and $\frac{2}{7}$ respectively.

The incidence calculus is not frequently applied, since one requires exact incidence records to use it. For the application at hand, however, we have detailed information about each matchmaker invocation, and the calculus provides a simple, intuitive way of dealing with the problem. More sophisticated mechanisms exist in the calculus for dealing with situations where knowledge is incomplete, though we do not exploit them in this paper.

4. THE LCC MATCHMAKER

In using LCC for matchmaking, we must ask how we arrive at a protocol. A client agent has a task or goal it wishes to achieve. Using either a pre-agreed lookup mechanism, or by reasoning about the protocols available, the agent will select a protocol: more than one might be suitable. This done, it must recruit a matchmaker to propose other agents to fill the various roles in the protocol. These other agents we term 'collaborators' and denote $col(Role, AgentId)$.

The success of a protocol and the particular team of collaborators is decided by the client: on completion or failure of a protocol, the client informs the matchmaker whether the outcome was satisfactory to the client.

Each completed brokering session is recorded as an incident, represented as an integer. Our propositions are ground predicate calculus expressions, e.g. $col(astronomy\_database, greenwich)$, or $outcome(good)$. Each proposition has an associated list of worlds (incidents) for which it is true. Initially, the incident database is empty, and the broker selects agents at random. As more data is collected, a threshold is reached, at which point the matchmaker begins to use the probabilities. For our example scenario (figure 3), the database might look like this:

$i(protocol(\text{BLACK\_HOLE\_SEARCH}), [1, 2, \ldots, 25])$
$i(outcome(good), [1, 2, 3, 4, 6, 10, 11, 12, 16, 22, 23, 24])$
$i(col(astronomy\_database, greenwich), [18, 19, 20, 21, 22, 23, 24, 25])$
$i(col(astronomy\_database, herschel), [10, 11, 12, 13, 14, 15, 16, 17])$
$i(col(astronomy\_database, keck), [1, 2, 3, 4, 5, 6, 7, 8, 9])$
$i(col(black\_hole\_finder, barcelona\_sc), [8, 9, 16, 17, 24])$
$i(col(black\_hole\_finder, ucsd\_sdsc), [1, 2, 3, 4, 10, 11, 12, 13, 18, 19, 20])$
$i(col(black\_hole\_finder, uk\_hpcx), [5, 6, 7, 14, 15, 21, 22, 23])$
$i(col(visualizer, ncsa), [1, 2, \ldots, 25])$

For instance, the Barcelona supercomputer is rarely successful:

$$i(col(black\_hole\_finder, barcelona\_sc) \wedge outcome(good)) = \{16\}$$

not because it is a worse supercomputer than UCSD-SDSC or UK-HPCX, but because its network connections to the databases required for this task present a bottleneck, reducing client satisfaction.

4.1 Algorithms

We have developed three algorithms for choosing agents, though others are possible. The first, called MATCHMAKE-JOINT, fills all the vacancies in a protocol at the outset. It works by computing the joint distribution for all possible permutations of agents in their respective roles, selecting the grouping with the largest probability of a good outcome. More concretely, MATCHMAKE-JOINT operates thus: extract the roles required in the protocol $\mathcal{P}$; compute the joint distribution for all agent permutations for these roles; select the agent set with greatest likelihood of success.

For MATCHMAKE-JOINT, *extendcollaborators* for our Grid example looks like this:

$$extendcollaborators(\mathcal{P}, \mathcal{C}, R) = \left\{ \begin{array}{l} col(astronomy\_database, AD) \wedge \\ col(black\_hole\_finder, BHS) \wedge \\ col(visualizer, V) \} \end{array} \right\}$$

where *AD*, *BHS* and *V* are chosen to maximise:

$$P \left(\begin{array}{l} col(astronomy\_database, AD) \wedge \\ col(black\_hole\_finder, BHS) \wedge \\ col(visualizer, V) \wedge outcome(good) \\ \mid protocol(\text{BLACK\_HOLE\_SEARCH}) \end{array} \right)$$

The second approach, MATCHMAKE-INCREMENTAL, is to select only one agent at a time, as required by the executing protocol. This is done by the matchmaker on demand. The various agents already engaged in the protocol, on needing to send a message to an as

Figure 3: Astronomy workflow scenario with LCC dialogue framework

We take a hypothetical Grid workflow for our example scenario. We name this protocol BLACK_HOLE_SEARCH. Astrid, our *astronomer*, is attempting to find and visualise a suspected black hole in a region of space around Cygnus-X1. The voluminous data about this segment of space is kept in the very large file *cygnus_x1*, which is stored at numerous repositories, all of which can fill the role *astronomy_database*. She uses a computationally intensive service called *black_hole_finder* to actually determine if there is a black hole present. The *black_hole_finder*, if successful, will send the data (now refined and significantly smaller) to a visualisation service, which will pass the final image to Astrid. The variables *AD*, *BHF*, and *V* represent the systems providing the services (*astronomy_database*, *black_hole_finder*, and *visualizer* respectively). Each of these will be selected by the matchmaker when the protocol is executed.

The conceit on which this example hangs is that network bandwidth between various pairs of *black_hole_finder* and *astronomy_database* will be different, largely unknown to the persons providing the individual services, and hence *not declared to the matchmaker.* Since the file *cygnus_x1* is particularly large, this network bandwidth is likely to be a strong determiner of the satisfaction of Astrid.

$$
\begin{array}{ll}
a(astronomer(File), Astronomer) :: & search(File) \Rightarrow a(black\_hole\_finder, BHF)\ then \\
& \left(\begin{array}{l} success \Leftarrow a(black\_hole\_finder, BHF)\ then \\ receive\_visualisation(Thing, V) \leftarrow visualising(Thing) \Leftarrow a(visualizer, V) \end{array} \right) \\
& or \\
& failed \Leftarrow a(black\_hole\_finder, BHF) \\
\\
a(black\_hole\_finder, BHF) :: & search(File) \Leftarrow a(astronomer(File), Astronomer)\ then \\
& grid\_ftp\_get(File) \Rightarrow a(astronomy\_database, AD)\ then \\
& \left(\begin{array}{l} grid\_ftp\_sent(File) \Leftarrow a(astronomy\_database, AD)\ then \\ success \Rightarrow a(astronomer, Astronomer) \\ \qquad\qquad \leftarrow black\_hole\_present(File, Black\_hole)\ then \\ visualize(Black\_hole, Astronomer) \Rightarrow a(visualizer, V) \end{array} \right) \\
& or \\
& failed \Rightarrow a(astronomer(File), Astronomer) \\
\\
a(astronomy\_database, AD) :: & grid\_ftp\_get(File) \Leftarrow a(black\_hole\_finder, BHF) \\
& grid\_ftp\_sent(File) \Rightarrow a(black\_hole\_finder, BHF) \leftarrow grid\_ftp\_completed(File, AD) \\
\\
a(visualizer, V) :: & visualize(Thing, Client) \Leftarrow a(black\_hole\_finder, Requester)\ then \\
& visualising(Thing) \Rightarrow a(astronomer(Filename), Client) \leftarrow serve\_visualisation(Thing, Client)
\end{array}
$$

Note that LCC is being used only to coördinate the interaction: when domain-specific protocols, such as Grid FTP, are available and more appropriate, they are used to perform the heavy lifting.

Figure 4: Rewrite rules governing matchmaking for an LCC protocol

These rewrite rules constitute an extension to those described in [5]. A rewrite rule

$$\alpha \xrightarrow{M_i,M_o,\mathcal{P},O,\mathcal{C},\mathcal{C}'} \beta$$

holds if α can be rewritten to β where: M_i are the available messages before rewriting; M_o are the messages available after the rewrite; $\mathcal{P}$ is the protocol; O is the message produced by the rewrite (if any); $\mathcal{C}$ is set of collaborators before the rewrite; and $\mathcal{C}'$ (if present) is the—possibly extended—set of collaborators after the rewrite. $\mathcal{C}$ is a set of pairs of role and agent name, e.g. $\{col(astronomy\_database, greenwich), col(black\_hole\_finder, ucsd\_sdsc)\}$

$$\begin{array}{lll}
A :: B \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} A :: E & \textit{if} & B \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} E \\
A_1 \textit{ or } A_2 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} E & \textit{if} & \neg closed(A_2) \wedge A_1 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} E \\
A_1 \textit{ or } A_2 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} E & \textit{if} & \neg closed(A_1) \wedge A_2 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} E \\
A_1 \textit{ then } A_2 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} E \textit{ then } A_2 & \textit{if} & A_1 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} E \\
A_1 \textit{ then } A_2 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O} A_1 \textit{ then } E & \textit{if} & closed(A_1) \wedge collaborators(A_1) = \mathcal{C}' \wedge A_2 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C}',O} E \\
A_1 \textit{ par } A_2 \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},O_1 \cup O_2} E_1 \textit{ par } E_2 & \textit{if} & A_1 \xrightarrow{M_i,M_n,\mathcal{P},\mathcal{C},O_1} E_1 \wedge A_2 \xrightarrow{M_n,M_o,\mathcal{P},\mathcal{C},O_2} E_2 \\
C \leftarrow M \Leftarrow A \xrightarrow{M_i,M_i \setminus \{M \Leftarrow A\},\mathcal{P},\mathcal{C},\emptyset} c(M \Leftarrow A, \mathcal{C}) & \textit{if} & (M \Leftarrow A) \in M_i \wedge satisfied(C) \\
M \Rightarrow A \leftarrow C \xrightarrow{M_i,M_i,\mathcal{P},\mathcal{C},\mathcal{C}',\{M \Rightarrow A\}} c(M \Rightarrow A, \mathcal{C}') & \textit{if} & satisfied(C) \wedge \mathcal{C}' = extendcollaborators(\mathcal{P}, \mathcal{C}, role(A)) \\
null \leftarrow C \xrightarrow{M_i,M_i,\mathcal{P},\mathcal{C},\emptyset} c(null, \mathcal{C}) & \textit{if} & satisfied(C) \\
a(R,I) \leftarrow C \xrightarrow{M_i,M_o,\mathcal{P},\mathcal{C},\emptyset} a(R,I) :: B & \textit{if} & clause(\mathcal{P}, C, a(R,I) :: B) \wedge satisfied(C)
\end{array}$$

$$\begin{array}{lcl}
collaborators(c(Term, \mathcal{C})) & = & \mathcal{C} \\
collaborators(A_1 \textit{ then } A_2) & = & collaborators(A_1) \cup collaborators(A_2) \\
collaborators(A :: B) & = & collaborators(A) \cup collaborators(B)
\end{array}$$

We can capture our various algorithms for matchmaking using the same rewrite rules: the issue is *when* the set of collaborators is actually decided. *extendcollaborators* determines the selection of a new agent: it is the matchmaking function. It varies slightly, depending on the exact matchmaking algorithm in use. By using the same rewrite rules regardless of the matchmaking policy, we make it easier to re-use model-checking [12] and other tools on the protocols.

yet unidentified agent, will ask the broker to find an agent to fulfil the role at hand. The process of MATCHMAKE-INCREMENTAL is: compute probability of successful outcome for each agent available for role R given $\mathcal{C}$, the collaborators chosen so far; select the most successful agent.

Consider the operation of MATCHMAKE-INCREMENTAL in the context of our Grid workflow scenario. Initially, Astrid contacts the matchmaker to request an agent for the *black_hole_finder* role. The *BHF* agent's first action is to request the data file from an astronomy database. It therefore returns the protocol to the matchmaker, which selects the *astronomy_database* most likely to provide a successful outcome, given that the *black_hole_finder* is already instantiated to *BHF*. That is, we select *AD* to maximise

$$P\left(\begin{array}{c} col(astronomy\_database, AD) \wedge outcome(good) \\ \mid protocol(\text{BLACK\_HOLE\_SEARCH}) \wedge \\ col(black\_hole\_finder, BHF) \end{array}\right)$$

The final method, MATCHMAKE-TREE is a mix of the first two. Like MATCHMAKE-JOINT, it runs only once, before the protocol executes. Like MATCHMAKE-INCREMENTAL, it selects only one agent at a time (that is, when a message is sent). This seeming paradox is resolved by considering that MATCHMAKE-TREE walks through the protocol, exploring each possible branch, and selecting an agent when necessary in the same manner as MATCHMAKE-INCREMENTAL. This tree of possible choices can be stored with the protocol that is sent to the client, and consulted as required.

All three algorithms support the pre-selection of agents for particular roles. An example of this might be a client booking a holiday: if it were accumulating frequent flyer miles with a particular airline, it could specify that airline be used, and the matchmaker would work with this choice. This mechanism also allows us to direct the matchmaker's search: selecting a particular agent suggests that the client wants similar agents, from the same social pool, for the other roles. For instance, in a peer-to-peer search, by selecting an agent you suspect will be helpful in a particular enquiry, the broker can find further agents that are closely 'socially' related to that first one.

4.2 Discussion

Having described the three algorithms, we must decided which to use, and when. MATCHMAKE-JOINT is preferable when one wishes to avoid multiple calls to the matchmaker, either because of privacy concerns, or for reasons of communication efficiency. MATCHMAKE-INCREMENTAL and MATCHMAKE-TREE would be more suitable in protocols where many roles go unfilled: total work on the broker would be reduced, and the results would probably be at least as good as for brokering all agents. Such a protocol, in which many roles are never used, could be viewed as a generalisation of a class of more specific protocols: it is worthwhile asking if it is better to have many very specialised protocols (which might not be used often and hence leave the matchmaker short of data), or fewer, more general-purpose protocols that offer plentiful, but less specific, information on how to select agents.

One cannot determine in general which of the algorithms will provide the optimal selection of agents. MATCHMAKE-JOINT appears to provide the 'optimal' solution, but there are some issues with it. The most immediate is that, unlike MATCHMAKE-INCREMENTAL and MATCHMAKE-TREE, agents can be unfairly blackballed for apparently under-performing in unsuccessful protocols in which they never actively participated. Secondly, we hope to

Figure 5: Agent selection improves as matchmaking database grows

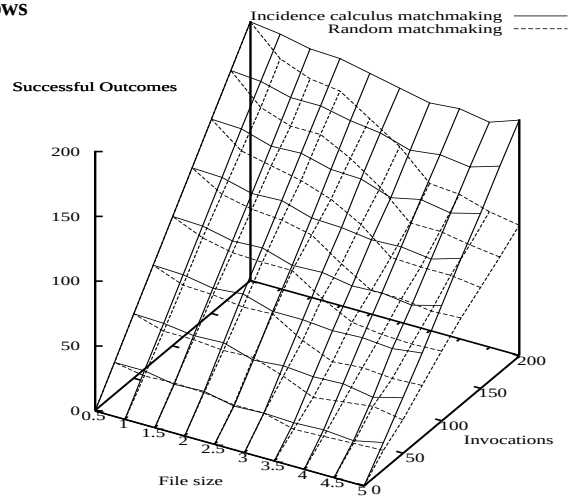

add backtracking to LCC, such that we might undo certain agent selections: this is more difficult if all roles are filled at the outset.

4.3 Results

The broker is currently implemented in Prolog. Performance, even in a naïve implementation, is adequate for tens of thousands of records, with response times well within the time frame that would be expected on the Internet.

We have so far executed simulations using agents constructed with plausible inter-agent relationships. A typical result (using MATCHMAKE-INCREMENTAL in the astronomy scenario) is shown in figure 5. We see that, as the file size increases, network bandwidth between nodes becomes more important in achieving a successful outcome. The incidence calculus matchmaker can detect and partially correct for this, while random matchmaking degrades much more quickly. This gives an impression of the benefit that can be expected, although actual performance will be sensitive to the scenario and the agents involved.

4.4 Inherent difficulties in the problem

We note here two significant problems that seem to be inescapable issues intrinsic to the problem: trusting clients to report honestly and in a socially 'normal' manner the outcome of protocol executions; and the problems of locating mutually co-operative agents in a large society.

Since individual client agents are responsible for the assigning of success metrics to matchmakings, there is scope for agents with unusual criteria, or downright malicious intentions, to contaminate the database.

Matchmaking is a social activity: clients wish to communicate with service providers that, by definition, they are unaware of. It is unclear how this aspect of agency will develop, and it will depend in many respects on companies' economic decisions, and the behaviour of individuals as to how many agents are deployed. There is a spectrum of possibilities, ranging from areas that are dominated by their 500lb gorillas (Google, Amazon, and E-Bay) through those that have dozens or hundreds of providers (insurers), to millions (personal calendar agents). We suppose that there will be a mix, and would presume that, for most purposes where one would use a matchmaker, we would be dealing with roles that supported numbers toward the lower end of the scale. Further, we must ask how many service types will be provided. Again, in each domain, we might have a simple, monolithic suck-it-and-see interface (Google again), or an interface with such fine granularity that few engineers ever fully understand or exploit it. Here, it is perhaps harder to predict the outcome.

While our technique handles large numbers of incidences, it does not scale for very large numbers of agents or roles. For any protocol with a set of roles R, and with each role having $|providers(r_i)|$ providers, the number of ways of choosing a team is

$$\prod_{r_i \in R} |providers(r_i)|$$

This is $O(m^n)$, making it an insuperable problem when considering protocols with large numbers of roles, or where each role has a large number of possible providers. This is not a specific criticism of our approach: no matchmaking system could possibly hope to examine all the various permutations of agents, although machine learning techniques might be helpful in finding non-obvious groupings of agents that simply could not be found by trial-and-error. How much of an issue this actually becomes in any particular domain will be heavily influenced by the outcomes to the social and economic issues discussed above.

5. RELATED WORK

The matchmaking/brokering problem arises in agent systems, semantic web, and grid environments. The matchmaking problem is discussed in [1, 7, 13]. We consciously ignored methods like those found in [14], though they would be crucial in any real-world deployment: we believe our technique would usefully augment such systems, and we intend to fuse the two approaches.

Our problem conception—matchmaking multiple roles for the same dialogue—is anticipated by the SELF-SERV system [15], although we believe our approach is novel in detecting emergent properties that are not known to the operator, and is more transparent, requiring less intervention (i.e. specification of service parameters) from the client. Our use of performance histories is predated by a similar approach found in [8], although that only examines the case of two-party interactions.

6. CONCLUSION AND FUTURE WORK

We have shown that, in plausible scenarios, the successful completion of a task may depend not only on the advertised abilities of agents but on their collective suitability and inter-operability. We presented a simple, but effective, technique for detecting successful groupings of agents. We highlighted the intractability of the problem in environments with large numbers of available provider agents and/or roles.

Despite talk of disintermediation and peer-to-peer systems, multi-agent scenarios that are amenable to matchmaking seem sparse. Typically, multi-party scenarios, where they do exist, have a client, several providers, and a central service which does the matchmaking itself—one common example is a travel agent service which arranges the various flight, car-hire, hotel bookings etc, without consulting a general purpose matchmaker. This paucity may be due to the emergence only recently of formalisms that can express such interactions easily, or it may reflect a deeper problem with conceiving problems as distributed sense as we adopted here. A standard library of such interactions would be helpful to research in the field.

Currently, only assignments of agents to roles are recorded. The utility of recording other events in the execution of the protocol will be investigated. For instance, 'partially satisfactory' protocol

executions could be interrogated to discover which agents are performing well, and which are proving to be a bottle-neck. Providing richer feedback from the client on its satisfaction with the outcome would be useful in itself. This could be extended to treating matchmaking as an interactive process: an agent requesting a service to satisfy a task might have various information that, unbeknownst to it, could aid the matchmaking in ascertaining the best protocol and collaborators to use. We also ignore issues of ontological compatibility and alignment. More sophisticated techniques for inferring compatibility will be examined, including those such as [14].

Previous work on LCC has examined the possibility of backtracking in protocols, in this case, allowing the broker to re-choose providers if an interaction failed. This would be easy for interactions, such as information gathering, where no commitments are made, but would require careful consideration of actions and state in, for example, a purchasing environment.

We think the intra-agent dependencies discussed here will be frequently encountered, but we have thus far only results for simulated systems. Current work is examining the applicability of the technique in the bioinformatics domain, using the same web services used by practising bioinformaticians.

The handling of older incidences is problematic: weighting more recent ones is not straight-forward in the incidence calculus. Issues of stability of solution, and settling on poor local maximums, remain unresolved. Although the incidence calculus provides a convenient framework for this model, it may not provide us with an optimal computational process, and the use of machine learning techniques may resolve some of these problems.

7. ACKNOWLEDGEMENTS

This research was sponsored by the Advanced Knowledge Technologies project, a six-year collaboration between groups at the universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University. AKT is funded by the United Kingdom's Engineering and Physical Sciences Research Council.

8. REFERENCES

[1] Decker, K., Sycara, K., Williamson, M.: Middle-Agents for the Internet. In: Proceedings of the 15th International Joint Conference on Artificial Intelligence, Nagoya, Japan (1997)

[2] R. G. Smith: The contract net protocol: high-level communication and control in a distributed problem solver. (1988) 357–366

[3] Martin, D., Burstein, M., Hobbs, J., Lassila, O., McDermott, D., McIlraith, S., Narayanan, S., Paolucci, M., Parsia, B., Payne, T., Sirin, E., Srinivasan, N., Sycara, K.: OWL-S 1.1 (2004)

[4] Kavantzas, N., Burdett, D., Ritzinger, G., Lafon, Y.: Web Services Choreography Description Language Version 1.0 (2004) W3C Working Draft 12 October 2004.

[5] Robertson, D.: A lightweight method for coordination of agent oriented web services. In: Proceedings of the 2004 AAAI Spring Symposium on Semantic Web Services, California, USA (2004)

[6] Blythe, J., Deelman, E., Gil, Y.: Planning for workflow construction and maintenance on the Grid. In: ICAPS03 workshop. (2003)

[7] Klusch, M., Sycara, K.: Brokering and matchmaking for coordination of agent societies: a survey. In: Coordination of Internet agents: models, technologies, and applications. Springer-Verlag (2001) 197–224

[8] Zhang, Z., Zhang, C.: An improvement to matchmaking algorithms for middle agents. In: Proceedings of the first international joint conference on Autonomous agents and multiagent systems, ACM Press (2002) 1340–1347

[9] Esteva, M., Rodriguez, J., Arcos, J., Sierra, C., Garcia, P.: Formalising Agent Mediated Electronic Institutions (2000)

[10] Milner, R.: Communication and Concurrency. Prentice Hall (1989)

[11] Bundy, A.: Incidence calculus: A mechanism for probabilistic reasoning. Journal of Automated Reasoning **1** (1985) 263–284

[12] Walton, C.: Model Checking Multi-Agent Web Services. In: Proceedings of the 2004 AAAI Spring Symposium on Semantic Web Services. (2004)

[13] Wong, H., Sycara, K.: A Taxonomy of Middle-agents for the Internet. (2000)

[14] Paulucci, M., Kawamura, T., Payne, T.R., Sycara, K.: Semantic Matching of Web Services Capabilities. In: The Semantic Web — ISWC 2002: Proceedings. (2002)

[15] Liangzhao Zeng and Boualem Benatallah and Marlon Dumas and Jayant Kalagnanam and Quan Z. Sheng: Quality driven web services composition. In: WWW '03: Proceedings of the twelfth international conference on World Wide Web, New York, NY, USA, ACM Press (2003) 411–421